

Output-Based Error Estimation and Mesh Adaptation for Steady and Unsteady Flow Problems

Krzysztof J. Fidkowski *
University of Michigan, USA

September 15, 2015

Contents

1	Introduction	5
1.1	Motivation	5
1.2	Finite Element Discretization	7
1.3	Background	8
2	Discretization	9
2.1	The Discontinuous Galerkin Method	9
2.1.1	Conservation Equations	9
2.1.2	Solution Approximation	9
2.1.3	Weak Form	10
2.1.4	Discrete System	11
2.1.5	Nonlinear Solver	13
2.2	Discrete Adjoint	13
2.2.1	Local Sensitivity Analysis	13
2.2.2	The Adjoint System	14
2.2.3	Adjoint Consistency	15
2.2.4	Adjoint Sensitivity Tests	18
3	Output Error Estimation	21
3.1	Two Discretization Levels	21
3.2	The Adjoint-Weighted Residual	21
3.3	Approximations	22
3.4	Error Effectivity	23
3.5	Examples	24
3.5.1	Drag Error for Euler Flow over a Bump	24
3.5.2	Drag Error for Viscous Flow over a NACA 0012 Airfoil	27

*Associate Professor, Aerospace Engineering Department, kfid@umich.edu

4	Mesh Adaptation	29
4.1	Error Localization	29
4.2	Adaptation Mechanics	29
4.2.1	Local Refinement	29
4.2.2	Global Re-Meshing	31
4.2.3	Targeting Strategies	32
4.2.4	Incorporating Anisotropy	33
4.2.5	Adapting in Order	34
4.3	Examples	34
4.3.1	Inviscid Flow over an Airfoil	35
4.3.2	Viscous Flow over an Airfoil	38
4.3.3	Which Output?	39
4.3.4	Transonic Turbulent Flow over an Airfoil	42
4.3.5	Transonic Turbulent Flow over a Wing	45
5	Mesh Optimization	49
5.1	Metric-Based Mesh Optimization Algorithm	49
5.1.1	Metric-Based Meshing	49
5.1.2	Error Convergence Model	50
5.1.3	Cost Model	52
5.1.4	Metric Optimization Algorithm	52
5.2	Element-Local Error Sampling	54
5.3	Examples	57
5.3.1	Inviscid Flow over a NACA 0012 Airfoil	57
5.3.2	High-Reynolds Number Flow over a Flat Plate	61
5.3.3	NACA 0012 Airfoil in Laminar Viscous Flow	62
5.3.4	RAE 2822 Airfoil in Turbulent Transonic Flow	69
6	Unsteady Systems	73
6.1	Primal and Adjoint Discretizations	73
6.1.1	Multi-Step Methods	73
6.1.2	Discontinuous Galerkin in Time	75
6.2	Deformable Domains	77
6.2.1	An Arbitrary Lagrangian Eulerian Treatment	77
6.2.2	Blended Analytical Mesh Motions	81
6.2.3	The Geometric Conservation Law	81
6.3	Error Estimation and Adaptation	83
6.3.1	The Adjoint-Weighted Residual and Error Localization	83
6.3.2	Incorporating Space-Time Anisotropy	84
6.3.3	Space-Time Mesh Adaptation	84
6.3.4	Implementation Notes	88
6.4	Examples	89
6.4.1	Static h -Refinement for an Impulsively-Started Airfoil	89
6.4.2	Dynamic-Order Adaptation for Pitching and Plunging Airfoils	93
6.4.3	Dynamic-Order Adaptation for a Three-Dimensional Wing	98

7	Conclusions	103
8	Acknowledgments	105
A	Compressible Navier-Stokes Equations	117
A.1	Euler Equations	117
A.2	Compressible Navier-Stokes	117
A.3	Reynolds-Averaged Compressible Navier-Stokes	119
B	Unsteady Adaptive Strategy	120
B.1	Greedy Fixed Growth Selection Algorithm	121
B.2	Temporal-Mesh Optimization	121

1 Introduction

1.1 Motivation

These notes present the theoretical background and practical implementation aspects of error estimation and mesh adaptation using output-based methods. The target application is the numerical simulation of aerodynamic flows, which are often convection-dominated and in which outputs of interest can be highly sensitive to small perturbations in the flow. The task of generating appropriate computational meshes for such simulations is therefore challenging, as inadequate resolution in a seemingly uninteresting area of the domain can have a large impact on engineering outputs of interest. It is in these situations that output-based adaptive methods, which specifically take into account such error propagation effects, can be advantageous, often by a margin that outweighs their (nontrivial) implementation cost.

When discussing adaptation, we are concerned with minimizing errors from the numerical discretization, which includes effects of mesh size, polynomial order, and time step. Such error is inherent to computational simulations because of the finite nature of discrete numerical calculations. As illustrated in Figure 1, discretization error is not the only source of error in computation, but its management tends to be the most elusive. In practice, with limited resources, numerical errors are managed by practitioners who are knowledgeable about the assumptions and limitations of the models, and/or by best-practice guidelines for mesh generation. However, even very experienced practitioners cannot reliably quantify the error in a discrete approximation of a complex flowfield – see for example the non-intuitive adaptive results in Figure 3. In addition, reliance on best-practice guidelines is an open-loop solution that leaves unchecked the possibility of large amounts of numerical error for computations on novel configurations.

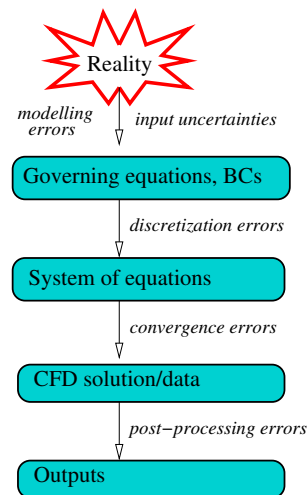


Figure 1: Various sources of error affecting computation.

Methods based on an *adjoint* analysis, where an adjoint associated with an output is the sensitivity of that output to residuals in the governing equations, help to close the loop in the control of numerical error via two ways. First, these methods provide output error estimates, which, although not bounds, can be used as confidence intervals to determine if an engineering output has been computed to sufficient numerical accuracy. Second, the error estimates can be localized, e.g. to elements of the computational mesh, in order to drive an adaptive method that reduces the error, as illustrated schematically in Figure 2. We view both of these uses of output-based analysis as improvements in the *robustness* of numerical simulations. Furthermore, note that the a posteriori nature of the error estimation and adaptation elicits a new paradigm for what constitutes a numerical solution – rather than treating properties of the discretization, such as the mesh, as givens, we propose to treat these properties as unknowns in the solution process.

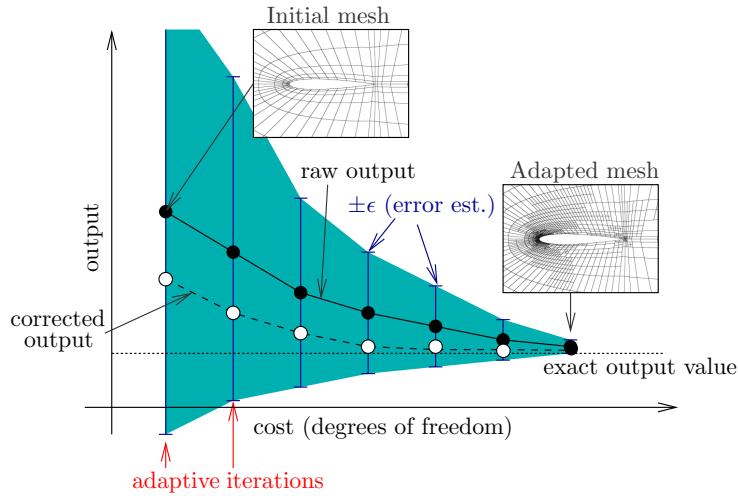


Figure 2: Typical adaptive convergence result obtained using output-based error estimates. The task of allocating degrees of freedom to regions of the domain is incorporated into the solver. Note, in many cases we can compute directional error estimates, i.e. corrections, and the resulting corrected outputs are illustrated by a dashed line. We can also often obtain conservative error estimates (though not strict bounds), and these are illustrated by the shaded regions at $\pm\epsilon$.

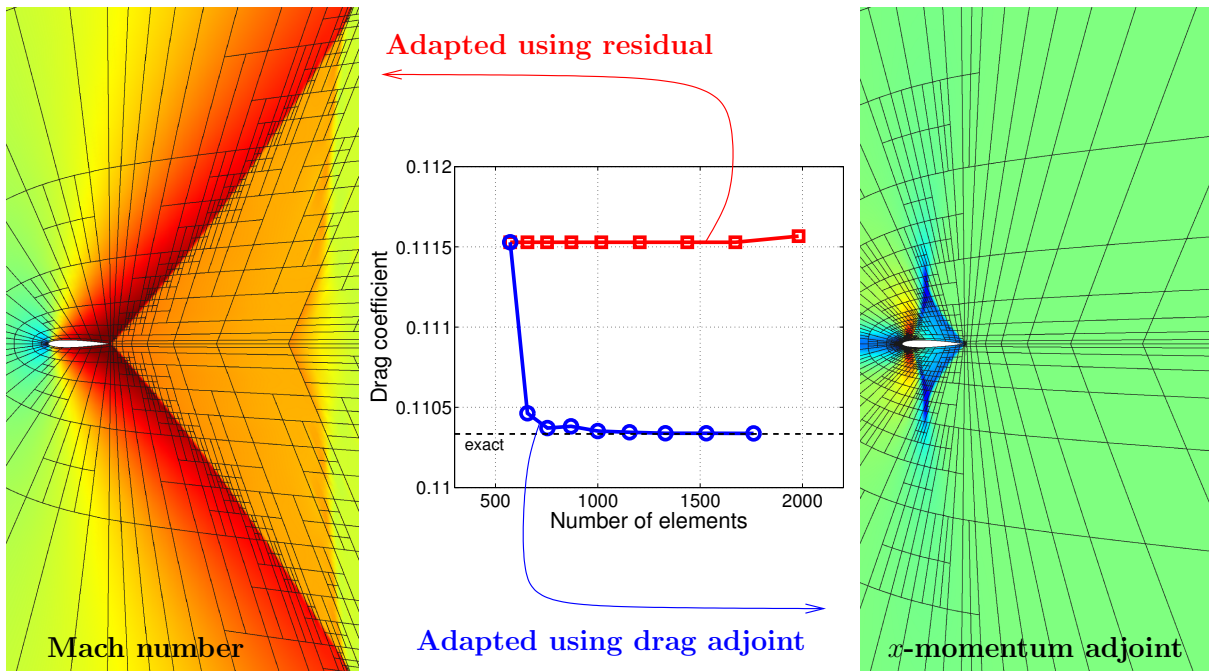


Figure 3: Two ways of adapting the mesh for a transonic, $M_\infty = 0.95$, Euler simulation around a NACA 0012 airfoil. The “fishtail” shock structure aft of the airfoil is a prominent feature, one targeted by a typical indicator such as an unweighted residual, which is actually not necessary for the prediction of drag. An indicator that weights the residual by the drag adjoint is not distracted by this feature and quickly hones in on the correct drag value.

1.2 Finite Element Discretization

In these notes we employ a finite element discretization of the governing equations of compressible fluid flow, i.e. the compressible Navier-Stokes equations. Many of the ideas presented extend to other discretizations, such as finite volume and finite difference methods. We choose finite elements for their variational formulation and flexibility in introducing high-order approximation. In particular, we employ the discontinuous Galerkin (DG) method, which has become a popular choice for convection-dominated conservation laws.

We employ a high-order discretization to allow for the possibility of order refinement during adaptation. The choice of low versus high-order methods is often a subject of debate, and a typical answer is that this choice depends on the problem: smooth flows are amenable to high order, whereas flows with singularities may be more efficiently resolved with low order. Most practical flows, however, exhibit both smooth and singular features, the locations of which are often not known a priori. In such circumstances, solution-based adaptation is critical for robustness and efficiency. High order applied selectively to parts of the domain becomes just another tool for efficient accuracy improvement. Other tools include mesh size and stretching optimization and mesh motion; one can imagine that more adaptive tools translate into a better use of degrees of freedom, and hence more efficient adaptation, as indicated in Figure 4. For our work, we focus on h (mesh), p

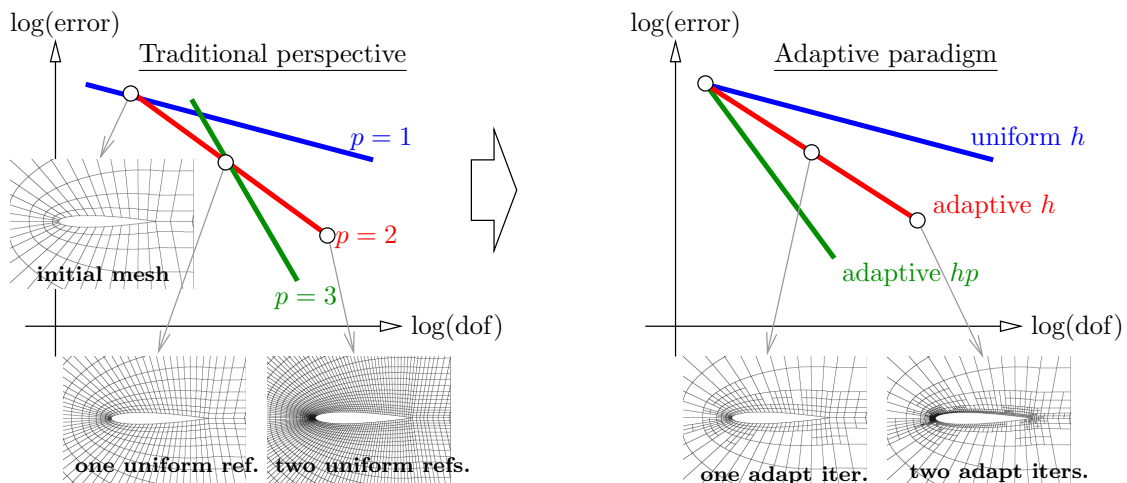


Figure 4: Schematic comparison of output convergence for uniform mesh refinement and for adaptive refinement. In aerodynamics, hp -refinement provides an important capability for addressing errors in flows with both smooth and discontinuous regions, specifically when targeting engineering outputs. Note that in the presence of singular features, the convergence rate of high-order approximations with uniform refinement (left) will generally be no better than that of low-order approximations, whereas this is not the case with hp -refinement.

(order), and hp (combined) adaptation. As aerodynamic flows generally exhibit both smooth and discontinuous features, hp -refinement is important in maximizing efficiency and robustness of CFD calculations, and the ability to locally enrich the order of the method is an key ingredient in making this possible.

1.3 Background

For flexibility in adaptive mechanics, we work with a discretization that admits high order, preferably in a locally-adaptive manner. Several such choices exist, and we turn to finite elements for their built-in variational framework, and to discontinuous approximation spaces for convective stability. Previous works have demonstrated the realizability of high-order accuracy [26, 38], error estimation [16, 60, 32], *hp*-adaptation [64], stable viscous discretization [11, 12, 4, 18, 92], and extension to variational space-time algorithms [6, 71, 117, 118, 68, 10, 79, 108, 37] using discontinuous Galerkin (DG) finite element methods.

Error estimation is a critical ingredient for improving the robustness of computational simulations. Efficient error control by systematic reduction requires the additional ability to identify regions in space and time contributing most to the error and the ability to change the discretization given such information. Table 1 lists several popular strategies for managing numerical error to improve simulation robustness. As indicated, only output-based methods provide both good error estimates and good adaptive indicators.

Table 1: Summary of existing methods for managing numerical error in engineering simulations.

Method and References	Description	Error Estimates	Adaptive Indicator
Resource exhaustion [45, 77]	Simulations run using the finest available numerical resolution.	None	None
Expert assessment [81, 76]	User employs previous experience, inspection of features, etc.	Poor	Poor
Convergence studies [100, 1]	Convergence rate estimated from a series of refined grids.	Good	Poor
Comparison to experiment [70, 84, 104]	Numerical error is low if experimental data is reasonably matched.	Medium	None
Feature-based adaptation [5, 8, 121]	Mesh is adapted in “interesting” areas, e.g. large gradients.	None	Medium
Residual methods [113, 103, 54, 102]	Nonzero residuals drive error propagation and mesh adaptation.	Medium	Medium
Output-based methods [14, 57, 119, 36]	Employ adjoint-weighted residual and localization for adaptation.	Good	Good

Output-based methods [96, 14, 57, 119, 82, 36] seek to quantify and minimize errors in scalar outputs of interest. Their robustness and efficiency arise from the fact that they specifically target for refinement those and only those areas of the computational domain that are important for predicting the output. They are well-suited for convection-dominated flows because through the use of adjoint solutions they properly account for error propagation effects, which can be troublesome for many heuristic adaptive indicators. Application to unsteady problems has been more limited [73, 74, 10, 79, 37, 34, 67] due largely to implementation challenges and computational expense associated with the solution of a fine-space adjoint equation, especially for nonlinear problems.

2 Discretization

2.1 The Discontinuous Galerkin Method

2.1.1 Conservation Equations

Consider a conservation law given by the partial differential equation (PDE)¹.

$$\partial_t \mathbf{u} + \partial_i \mathbf{H}_i(\mathbf{u}, \nabla \mathbf{u}) = \mathbf{0}, \quad (1)$$

where $\mathbf{u} \in \mathbb{R}^s$ is the state vector, $\mathbf{H}_i \in \mathbb{R}^s$ is the i^{th} component of the total flux, $1 \leq i \leq d$ indexes the spatial dimension d , and summation is implied on the repeated index i . We decompose the total flux into convective and diffusive parts,


$$\mathbf{H}_i = \mathbf{F}_i(\mathbf{u}) + \mathbf{G}_i(\mathbf{u}, \nabla \mathbf{u}), \quad (2)$$

$$\mathbf{G}_i(\mathbf{u}, \nabla \mathbf{u}) = -\mathbf{K}_{ij}(\mathbf{u}) \partial_j \mathbf{u}, \quad (3)$$

where $\mathbf{F}_i \in \mathbb{R}^s$ is the i^{th} component of the inviscid/convective flux, $\mathbf{G}_i \in \mathbb{R}^s$ is the i^{th} component of the viscous flux, and $\mathbf{K}_{ij} \in \mathbb{R}^{s \times s}$ is the (i, j) component of the viscous diffusivity tensor. We will initially focus on steady problems, so that $\partial_t \mathbf{u} = \mathbf{0}$.

2.1.2 Solution Approximation

DG is a finite element method in which the state \mathbf{u} is spatially approximated in functional form, using linear combinations of basis functions, usually polynomials, on each element. No continuity constraints are imposed on the approximations on adjacent elements. Denote by T_h the set of N_e elements in a non-overlapping tessellation of the domain Ω , as illustrated in Figure 5.



$$\mathbf{u}_h(\vec{x}) \approx \sum_{e=1}^{N_e} \sum_{n=1}^{N_{p_e}} \mathbf{U}_{en} \phi_{en}(\vec{x}), \quad (4)$$

Figure 5: Partition of a square domain into 14 triangular elements.

where

- N_{p_e} = number of basis functions needed for an order p_e approximation
- $\phi_{en}(\vec{x})$ = n^{th} order p_e basis function on element e (zero on all other elements)
- p_e = order of spatial basis functions on element e
- N_e = number of elements
- \mathbf{U}_{en} = vector of s coefficients on n^{th} basis function on element e

Formally, we can write that $\mathbf{u}_h \in \mathcal{V}_h = [\mathcal{V}_h]^s$, where

$$\mathcal{V}_h = \{u \in L_2(\Omega) : u|_{\Omega_e} \in \mathcal{P}^{p_e} \forall \Omega_e \in T_h\},$$

and \mathcal{P}^{p_e} denotes polynomials of order p_e on element e .² In many cases we use the same approximation order on all elements, so that $\forall e, p_e = p$. The lack of continuity in

¹See Appendix A for the PDEs used in this work

²A caveat here is that for elements that are curved, the polynomial approximation is usually performed on a master reference element, so that following the reference-to-global mapping, the state approximation on curved elements is not strictly of order p_e .

the solution approximation differentiates DG from continuous finite element methods, as illustrated in Figure 6. It adds computational expense, as we do not expect to approximate solutions that are discontinuous at every element interface, but it also provides convective stability and simplifies hanging-node mesh refinement and local order enrichment.

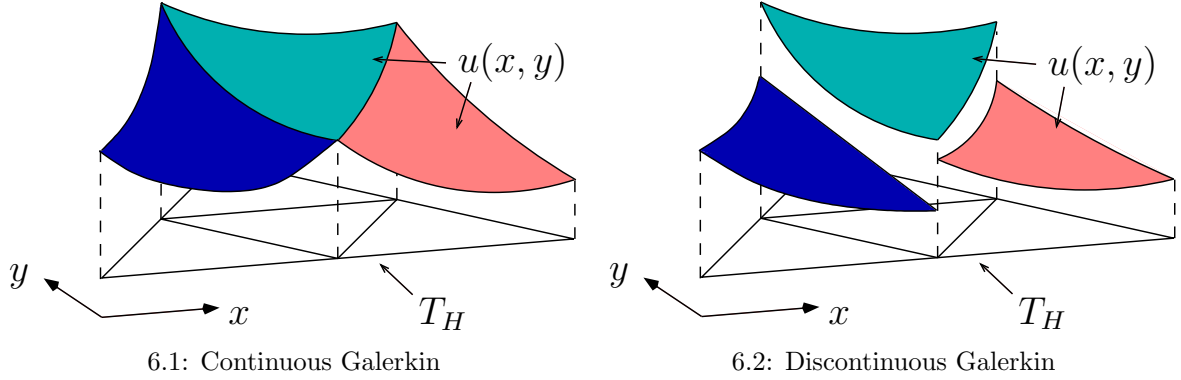


Figure 6: Solution approximation using continuous and discontinuous basis functions. Though the solution is discontinuous in DG methods, the inter-element flux is single valued, as in finite volume methods.

2.1.3 Weak Form

We obtain a weak form of (1) by multiplying the PDE by test functions $\mathbf{v}_h \in \mathbf{V}_h$ and integrating by parts to couple elements via fluxes. The convective fluxes on element faces are handled via a traditional finite-volume (approximate) Riemann solver, but the diffusive treatment is trickier and requires stabilization. Multiple diffusion formulations exist [4], and we employ the second form of Bassi and Rebay (BR2) [11].

We can write the final semilinear weak form as

$$\mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_h) = 0, \quad \forall \mathbf{v}_h \in \mathbf{V}_h, \quad (5)$$


which, by linearity of the second argument, we can decompose into contributions from each element,

$$\mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_h) = \sum_{e=1}^{N_e} \mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_h|_{\Omega_e}) = 0, \quad \forall \mathbf{v}_h \in \mathbf{V}_h. \quad (6)$$

Integrating by parts and applying the BR2 diffusion treatment, we find that the semilinear form associated with each element is

$$\begin{aligned} \mathcal{R}_h(\mathbf{u}_h, \mathbf{v}_h|_{\Omega_e}) &= \int_{\Omega_e} \mathbf{v}_h^T \partial_t \mathbf{u}_h \, d\Omega - \int_{\Omega_e} \partial_i \mathbf{v}_h^T \mathbf{H}_i \, d\Omega \\ &+ \int_{\partial\Omega_e \setminus \partial\Omega} \mathbf{v}_h^{+T} \left(\widehat{\mathbf{F}} + \widehat{\mathbf{G}} \right) \, ds + \int_{\partial\Omega_e \cup \partial\Omega} \mathbf{v}_h^{+T} \left(\widehat{\mathbf{F}}^b + \widehat{\mathbf{G}}^b \right) \, ds \\ &- \int_{\partial\Omega_e \setminus \partial\Omega} \partial_i \mathbf{v}_h^{+T} \mathbf{K}_{ij}^+ \left(\mathbf{u}_h^+ - \widehat{\mathbf{u}}_h \right) n_j \, ds \\ &- \int_{\partial\Omega_e \cup \partial\Omega} \partial_i \mathbf{v}_h^{+T} \mathbf{K}_{ij}^b \left(\mathbf{u}_h^+ - \mathbf{u}_h^b \right) n_j \, ds \end{aligned} \quad (7)$$

where $(\cdot)^T$ denotes transpose, and on the element boundary $\partial\Omega_e$ the notations $(\cdot)^+$, $(\cdot)^-$, $(\cdot)^b$ respectively denote quantities taken from the element interior, neighbor element, and boundary. The last two terms symmetrize the semilinear form for adjoint consistency. A schematic showing key quantities and the variables affecting the fluxes is given below. The unique state on an interior face is $\hat{\mathbf{u}}_h = (\mathbf{u}_h^+ + \mathbf{u}_h^-)/2$.



$$\begin{aligned}\hat{\mathbf{F}} &= \hat{\mathbf{F}}(\mathbf{u}_h^+, \mathbf{u}_h^-, \vec{n}) \\ \hat{\mathbf{G}} &= \hat{\mathbf{G}}(\mathbf{u}_h^+, \mathbf{u}_h^-, \nabla \mathbf{u}_h^+, \nabla \mathbf{u}_h^-, \vec{n}) \\ \hat{\mathbf{F}}^b &= \hat{\mathbf{F}}^b(\mathbf{u}_h^b, \text{BC}, \vec{n}) \\ \hat{\mathbf{G}}^b &= \hat{\mathbf{G}}^b(\mathbf{u}_h^b, \nabla \mathbf{u}_h^+, \text{BC}, \vec{n})\end{aligned}$$

In particular, on an interior face σ^f , the convective flux $\hat{\mathbf{F}}$ is computed using the Roe approximate Riemann solver [101], and the BR2-stabilized viscous flux is

$$\hat{\mathbf{G}} = \frac{1}{2} (\mathbf{G}_i^+ + \mathbf{G}_i^-) n_i^+ + \eta \frac{1}{2} (\delta_i^+ + \delta_i^-) n_i^+, \quad (8)$$

where the auxiliary variables $\delta_i^+, \delta_i^- \in [\mathbf{V}_h]^d$ have support on the two elements adjacent to the interior face σ^f and are obtained by solving $\forall \tau_{hi} \in \mathbf{V}_h$

$$\int_{\Omega_e^+} \tau_{hi}^T \delta_i^+ d\Omega + \int_{\Omega_e^-} \tau_{hi}^T \delta_i^- d\Omega = \int_{\sigma^f} \frac{1}{2} (\tau_{hi}^{+T} \mathbf{K}_{ij}^+ + \tau_{hi}^{-T} \mathbf{K}_{ij}^-) (\mathbf{u}_h^+ - \mathbf{u}_h^-) n_j^+ ds. \quad (9)$$

η is a stabilization factor that should not be less than the number of faces per element, and in our work is taken to be (at least) twice the maximum number of faces on adjacent elements.

On a boundary face σ^b , fluxes are typically computed directly from the boundary state, \mathbf{u}^b , which is a function (projection) of the interior state and the boundary-condition data, $\mathbf{u}_h^b = \mathbf{u}_h^b(\mathbf{u}_h^+, \text{BC})$. One exception is the convective flux for a boundary condition in which a complete exterior state is specified: in such a case, an approximate-Riemann solver is used to compute the boundary convective flux, $\hat{\mathbf{F}}^b$. The BR2-stabilized boundary viscous flux is

$$\hat{\mathbf{G}}^b = \Pi_G^{\text{BC}} [\mathbf{G}_i(\mathbf{u}_h^b, \nabla \mathbf{u}_h^+) n_i + \eta \delta_i n_i], \quad (10)$$

where the auxiliary variable $\delta \in [\mathbf{V}_h]^d$ has support on the element adjacent to the boundary face σ^b and is obtained by solving $\forall \tau_{hi} \in \mathbf{V}_h$

$$\int_{\Omega_e} \tau_{hi}^T \delta_i d\Omega = \int_{\sigma^b} \tau_{hi}^{+T} \mathbf{K}_{ij}^b (\mathbf{u}_h^+ - \mathbf{u}_h^b) n_j^+ ds. \quad (11)$$

The projection Π_G^{BC} in (10) incorporates boundary conditions on the viscous flux, such as a prescribed heat flux in the compressible Navier-Stokes equations.

2.1.4 Discrete System

For efficient implementation, the various contributions from all elements to the final semi-linear form in (5) are computed in three logical loops: one loop over all element interiors,

one loop over all interior faces, and one loop over all boundary faces. When computing these contributions, we choose as test functions the trial basis functions introduced in (4), ϕ_{en} , where

$$\text{span} \{ \phi_{en} \} = \mathcal{V}_h. \quad (12)$$

Recall that e is the element number and n indexes the polynomial basis functions inside element e . Each state equation is tested with the same set of functions, which means that we can define a size- s residual vector for the n^{th} test function in element e by

$$\mathbf{R}_{en} \equiv \{ \mathcal{R}_h(\mathbf{u}_h, \phi_{en} \mathbf{e}_r) \}_{r=1 \dots s} \in \mathbb{R}^s, \quad (13)$$

where $\mathbf{e}_r \in \mathbb{R}^s$, $r = 1 \dots s$, is a vector of all zeros except a 1 in position r . We now use the convention that dropping a subscript means considering the set of values over the entire range of that subscript. So \mathbf{R}_e is the set of residuals over all states and basis functions inside element e , and \mathbf{R} is the set of all residuals for all elements in the domain. Using the same convention for the state, \mathbf{U} , we can write the discrete system of equations (i.e. residuals) compactly as

$$\mathbf{R}(\mathbf{U}) = \mathbf{0}. \quad (14)$$

Note that both \mathbf{U} and \mathbf{R} lie in \mathbb{R}^N , where the total number of degrees of freedom including equation states is

$$N = \sum_{e=1}^{N_e} N_{pe} s. \quad (15)$$

When considering different discretization spaces, we will append a subscript h or H to the variables \mathbf{R} , \mathbf{U} , and N .

In practice, we store the state and residual vectors unrolled, as illustrated in Figure 7. Finally, the number of basis functions per element depends on the approximation space




Figure 7: Unrolled storage of the state vector $\mathbf{U} \in \mathbb{R}^N$.

and order p . We consider two approximation spaces that are given by the span of monomials in reference space coordinates ξ_i , $\prod_{i=1}^d \xi_i^{p_i}$: a full-order space in which $\sum_{i=1}^d p_i \leq p$, and a tensor-product space in which $\forall i, p_i \leq p$. The resulting dimensions of these spaces are

$$\text{full-order basis set: } N_p = \binom{p}{d} = \frac{p!}{d!(p-d)!} \quad \text{tensor-product basis set: } N_p = (p+1)^d.$$

2.1.5 Nonlinear Solver

To solve (14), we use a preconditioned Newton-Krylov method augmented with pseudo-transient continuation. Starting from an initial guess, usually the free-stream condition, and a conservative time step set using a Courant-Friedrichs-Lewy (CFL) number of approximately 1, the Newton solver iterates until steady state. The linear system at each Newton iteration is solved using an element-line preconditioned General Minimal Residual Krylov subspace method [106, 105]. As successful Newton updates are taken, the time step is increased according to an exponential growth formula for the CFL number. Details on the solver, including the incorporation of physical constraints on some of the variables, can be found in [24].

2.2 Discrete Adjoint

Suppose that we are interested in a scalar output computed from the solution to our PDE. Using the discrete state vector, we write

$$\text{output } J = J(\mathbf{U}). \quad (16)$$

The discrete adjoint, $\Psi \in \mathbb{R}^N$, is a vector of sensitivities of the output to the N residuals. That is, each entry of the adjoint tells us the effect that a perturbation in the same entry in the residual vector would have on the output J . One common source of residual perturbations is changes in input parameters for a problem, and so we ground the adjoint presentation in the context of a local sensitivity analysis.

2.2.1 Local Sensitivity Analysis

Consider a situation in which N_μ parameters, $\boldsymbol{\mu} \in \mathbb{R}^{N_\mu}$, affect the PDE in (14). For example, in aerodynamics, parameters could be used to set boundary conditions or the geometry, as illustrated schematically in Figure 8. We can then write the following chain of dependence,

$$\underbrace{\boldsymbol{\mu}}_{\text{inputs} \in \mathbb{R}^{N_\mu}} \rightarrow \underbrace{\mathbf{R}(\mathbf{U}, \boldsymbol{\mu}) = 0}_{N \text{ equations}} \rightarrow \underbrace{\mathbf{U}}_{\text{state} \in \mathbb{R}^N} \rightarrow \underbrace{J(\mathbf{U})}_{\text{output (scalar)}}. \quad (17)$$

We are interested in how J changes (locally for nonlinear problems) with $\boldsymbol{\mu}$,

$$\frac{dJ}{d\boldsymbol{\mu}} \in \mathbb{R}^{1 \times N_\mu} = N_\mu \text{ sensitivities}. \quad (18)$$

Note, if J depends directly on $\boldsymbol{\mu}$ we would add $\frac{\partial J}{\partial \boldsymbol{\mu}}$ to the above sensitivity, but for clarity of presentation we consider only the case when $J = J(\mathbf{U})$. Several options exist for computing these sensitivities. Two direct ones are finite differencing, in which the input parameters are perturbed one at a time, and forward linearization, in which the sequence of operations in (17) is linearized. Both of these become expensive when N_μ is moderate or large, because of the need to re-solve the system $\mathbf{R}(\mathbf{U}, \boldsymbol{\mu}) = 0$ for each parameter. A third choice is the adjoint approach, which requires an inexpensive residual perturbation

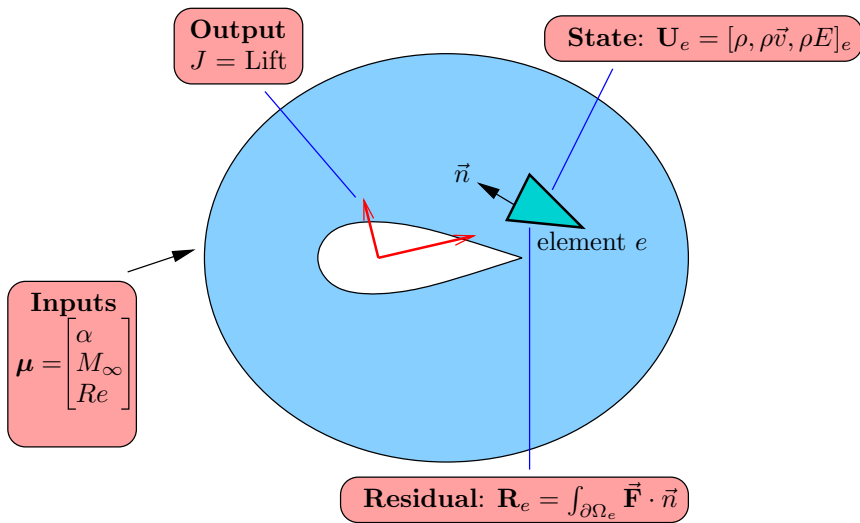


Figure 8: Schematic of typical inputs, state, residual, and output quantities in an aerodynamics simulation.

calculation followed by an adjoint weighting to compute the effect on the output. That is, we write

$$\frac{dJ}{d\boldsymbol{\mu}} = \boldsymbol{\Psi}^T \frac{\partial \mathbf{R}}{\partial \boldsymbol{\mu}}. \quad (19)$$

This approach is efficient for computing a large number of sensitivities for one output, as the cost is one residual perturbation calculation and one vector product per sensitivity.

The central idea in the adjoint approach is that we do not need to solve the forward problem each time we want a sensitivity. Suppose that for a given $\boldsymbol{\mu}$ we solve our discrete system and find \mathbf{U} such that $\mathbf{R}(\mathbf{U}, \boldsymbol{\mu}) = \mathbf{0}$. Now perturb $\boldsymbol{\mu} \rightarrow \boldsymbol{\mu} + \delta\boldsymbol{\mu}$... what is the effect on J ? We can re-solve for a perturbed state, but this is expensive. Instead, we can separate the effects of $\boldsymbol{\mu}$ on \mathbf{R} , and \mathbf{R} on J , as illustrated on the right. The adjoint method precomputes the effect of \mathbf{R} on J , which is the expensive step. The resulting N sensitivities are stored in the vector $\boldsymbol{\Psi}$.




Figure 9: Bypassing the forward solve via an adjoint approach to sensitivity calculation.

2.2.2 The Adjoint System

To derive an equation for the adjoint, we consider the chain of operations we would take in computing the sensitivities via a direct approach. In the following steps, we assume small perturbations.

1. Input $\boldsymbol{\mu} \rightarrow \boldsymbol{\mu} + \delta\boldsymbol{\mu}$
2. Residual $\mathbf{R}(\mathbf{U}, \boldsymbol{\mu} + \delta\boldsymbol{\mu}) = \delta\mathbf{R} \neq 0 \rightarrow \mathbf{R}(\mathbf{U}, \boldsymbol{\mu}) + \left. \frac{\partial \mathbf{R}}{\partial \boldsymbol{\mu}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \delta\boldsymbol{\mu} = \delta\mathbf{R}$
3. State $\mathbf{R}(\mathbf{U} + \delta\mathbf{U}, \boldsymbol{\mu} + \delta\boldsymbol{\mu}) = 0 \rightarrow \mathbf{R}(\mathbf{U}, \boldsymbol{\mu}) + \left. \frac{\partial \mathbf{R}}{\partial \boldsymbol{\mu}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \delta\boldsymbol{\mu} + \left. \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \delta\mathbf{U} = 0$
4. Output $J(\mathbf{U} + \delta\mathbf{U}) = J(\mathbf{U}) + \delta J \rightarrow \delta J = \left. \frac{\partial J}{\partial \mathbf{U}} \right|_{\mathbf{U}} \delta\mathbf{U}$

Subtracting step 2 from step 3, we obtain

$$\left. \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \delta\mathbf{U} = -\delta\mathbf{R} \Rightarrow \delta\mathbf{U} = - \left[\left. \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \right]^{-1} \delta\mathbf{R}. \quad (20)$$

Combining this result with the output linearization in step 4 gives the output perturbation in terms of the residual perturbation,

$$\delta J = \left. \frac{\partial J}{\partial \mathbf{U}} \right|_{\mathbf{U}} \delta\mathbf{U} = \underbrace{- \left. \frac{\partial J}{\partial \mathbf{U}} \right|_{\mathbf{U}} \left[\left. \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \right]^{-1}}_{\boldsymbol{\Psi}^T \in \mathbb{R}^N} \delta\mathbf{R}. \quad (21)$$

Taking the transpose of the equation $\boldsymbol{\Psi}^T = - \left. \frac{\partial J}{\partial \mathbf{U}} \right|_{\mathbf{U}} \left[\left. \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \right]^{-1}$ and moving everything to the left-hand side gives the *adjoint equation*,

$$\left(\left. \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \right)^T \boldsymbol{\Psi} + \left(\left. \frac{\partial J}{\partial \mathbf{U}} \right|_{\mathbf{U}} \right)^T = 0. \quad (22)$$

The n^{th} component of $\boldsymbol{\Psi}$ is the sensitivity of J to changes in the n^{th} residual.

Since $\mathbf{R}(\mathbf{U}, \boldsymbol{\mu}) = 0$, from step 2 above we have $\delta\mathbf{R} = \left. \frac{\partial \mathbf{R}}{\partial \boldsymbol{\mu}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \delta\boldsymbol{\mu}$ and (21) becomes

$$\delta J = \boldsymbol{\Psi}^T \left. \frac{\partial \mathbf{R}}{\partial \boldsymbol{\mu}} \right|_{\mathbf{U}, \boldsymbol{\mu}} \delta\boldsymbol{\mu} \Rightarrow \frac{dJ}{d\boldsymbol{\mu}} = \boldsymbol{\Psi}^T \left. \frac{\partial \mathbf{R}}{\partial \boldsymbol{\mu}} \right|_{\mathbf{U}, \boldsymbol{\mu}}. \quad (23)$$

Therefore, once we have $\boldsymbol{\Psi}$, no more solves are required for new sensitivities for the same output. Note that the calculation of $\left. \frac{\partial \mathbf{R}}{\partial \boldsymbol{\mu}} \right|_{\mathbf{U}, \boldsymbol{\mu}}$ is typically very cheap compared to a forward solve.

Although we have presented the adjoint in the context of a parameter sensitivity analysis, we will see in the next section that residual perturbations also arise when discrete solutions are viewed from an enriched space. This will be the motivation for using adjoint solutions in output error estimation.

2.2.3 Adjoint Consistency

The solution to (22) is a *discrete* adjoint, $\boldsymbol{\Psi}$, which at the simplest level we can think of as a vector of N numbers. However, the adjoint also has a continuous counterpart, call it $\boldsymbol{\psi}(\vec{x})$, and we can think of the N numbers in $\boldsymbol{\Psi}$ as expansion coefficients in an approximation of $\boldsymbol{\psi}$ using the same basis functions used for the primal problem. The accuracy of this approximation is of interest for various reasons, including error estimation.

Suppose that the exact primal solution, $\mathbf{u} \in \mathcal{V}$, satisfies

$$\mathcal{R}(\mathbf{u}, \mathbf{v}) = 0, \quad \forall \mathbf{v} \in \mathcal{V}, \quad (24)$$

for an appropriately defined space \mathcal{V} . The exact adjoint $\boldsymbol{\psi} \in \mathcal{V}$ then satisfies

$$\mathcal{R}'[\mathbf{u}](\mathbf{v}, \boldsymbol{\psi}) + \mathcal{J}'[\mathbf{u}](\mathbf{v}) = 0, \quad \forall \mathbf{v} \in \mathcal{V}, \quad (25)$$

where the primes denote Fréchet linearization about the arguments in square brackets, and \mathcal{R} and \mathcal{J} are the continuous versions of the semilinear form and output functional, respectively. While we have assumed that both \mathbf{u} and $\boldsymbol{\psi}$ lie in \mathcal{V} , this may not always be the case [72].

The exact adjoint can be regarded as a Green's function that relates source perturbations in the original partial differential equation to perturbations in the output[46, 48], as illustrated schematically in Figure 10. A sample adjoint solution is shown in Figure 11




Figure 10: Interpretation of the adjoint as the sensitivity of a scalar output, such as lift, to residual source perturbations.

for Reynolds-averaged compressible flow over an airfoil. While the adjoint solution often shares qualitative characteristics similar to the primal, such as the presence of a boundary layer in a high-Reynolds number flow, it also shows marked differences, such as the “wake reversal” seen in the far-field view in Figure 11. In this case, the output is drag, and upstream of the airfoil, residual perturbations on/near the stagnation streamline (the flow that is going to hit or come closest to the airfoil) will have a larger magnitude impact on the drag than residual perturbations elsewhere; hence we see an adjoint “reversed” wake telling us that there are large sensitivities to perturbations in front of the airfoil. The figure shown tells us that this is the case for residual perturbations in the conservation of x -momentum equations, but plots of the other adjoint components show similar behavior.

Figure 12 illustrates an adjoint field for another case: supersonic flow over a diamond airfoil. The output in this case is a pressure line integral, and nonzero values of the adjoint (y -momentum component shown) indicate where perturbations to the (y -momentum) residuals will affect the output.

The adjoint field depicted in Figure 11 is the discrete adjoint solution on a fine mesh. It can only be regarded as a faithful representation of the exact adjoint if the discretization





11.1: x -momentum state (near view)11.2: x -momentum adjoint (near view)11.3: x -momentum state (far view)11.4: x -momentum adjoint (far view)

Figure 11: Comparison of the primal solution (x -momentum component) and the adjoint solution (conservation of x -momentum equation component) for a drag output in Reynolds-averaged turbulent flow over an RAE 2822 airfoil. The color scales are clipped to show the interesting features of each quantity – in the adjoint plots, yellow is near zero.




Figure 12: Sample adjoint solution for a pressure integral output in supersonic flow over a diamond airfoil. Large regions of zero adjoint indicate areas where residuals have no affect on the output – these could also be identified via a characteristic analysis.

is in some manner consistent with the exact adjoint problem. *Primal consistency* in the variational problem requires that the exact solution \mathbf{u} satisfy the discrete variational statement,

$$\mathcal{R}_h(\mathbf{u}, \mathbf{v}) = 0, \quad \forall \mathbf{v} \in \mathcal{W}_h, \quad (26)$$

where $\mathcal{W}_h = \mathcal{V}_h + \mathcal{V} = \{\mathbf{h} = \mathbf{f} + \mathbf{g} : \mathbf{f} \in \mathcal{V}_h, \mathbf{g} \in \mathcal{V}\}$. Similarly, the combination of the discrete semi-linear form \mathcal{R}_h and the functional \mathcal{J}_h is said to be *adjoint consistent* if [72, 56, 86]

$$\mathcal{R}'_h[\mathbf{u}](\mathbf{v}, \psi) + \mathcal{J}'_h[\mathbf{u}](\mathbf{v}) = 0, \quad \forall \mathbf{v} \in \mathcal{W}_h. \quad (27)$$

Discretizations that are not adjoint consistent may still be *asymptotically adjoint consistent* if Eq. 27 holds in the limit $h \rightarrow 0$, by which we mean the limit of uniformly increasing resolution, over suitably normalized $\mathbf{v} \in \mathcal{W}_h$. For non-variational discretizations, the definition of consistency must involve an approximation operator to map exact solutions into discrete spaces [28].

Adjoint consistency has an impact on the convergence of not only the adjoint approximation but also the primal approximation [4, 53, 47, 52, 72, 56, 86]. In error estimation, an adjoint-inconsistent discretization can lead to irregular or oscillatory adjoint solutions that pollute the error estimate with noise and lead to adaptation in incorrect areas [72]. Enforcing adjoint consistency imposes restrictions on the output definition and on the interior and boundary discretizations that enter into the semi-linear form. These restrictions have been studied by several authors in the context of the discontinuous Galerkin method [4, 72, 56]. In general, discretizations that are found to be adjoint inconsistent can often be made adjoint consistent by adding terms to either the semi-linear form or the output functional.

2.2.4 Adjoint Sensitivity Tests


One way to test a discrete adjoint is to compare output sensitivities computed using the adjoint to those computed using finite differences. In this example we show this comparison for a viscous flow over a NACA 0012 airfoil at $M = 0.5$, $Re = 5000$, $\alpha = 2^\circ$. We are interested in the sensitivity of the lift coefficient to angle of attack. Since the lift is defined as the force perpendicular to the free-stream direction, the output depends directly on the input parameters, so that we need to augment (19) with the partial derivative of J with respect to the input angle of attack, α ,

$$\frac{dJ}{d\alpha} = \Psi^T \frac{\partial \mathbf{R}}{\partial \alpha} + \frac{\partial J}{\partial \alpha}. \quad (28)$$

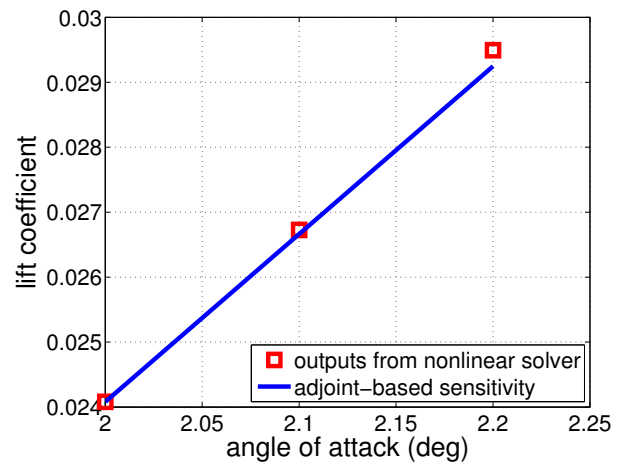
For the present test we compute $\frac{\partial \mathbf{R}}{\partial \alpha}$ using a simple forward difference,

$$\frac{\partial \mathbf{R}}{\partial \alpha} \approx \frac{\mathbf{R}(\mathbf{U}, \alpha + \Delta\alpha) - \mathbf{R}(\mathbf{U}, \alpha)}{\Delta\alpha}, \quad (29)$$

where $\Delta\alpha$ is a small value, e.g. 0.01rad. In Figure 13, we show a comparison of the sensitivity computed using (28), i.e. a local linearized sensitivity, to outputs obtained by actually perturbing the input parameter. The qualitative agreement is excellent, and a quantitative look shows convergence at the expected rate as the angle of attack perturbation goes to zero.



13.1: Mach number contours



13.2: lift coefficient sensitivity

Figure 13: Verification of the discrete adjoint solver using a parameter sensitivity test. The angle of attack is varied in a viscous flow over a NACA 0012 airfoil, and the resulting data points are overlaid on a line through the baseline, $\alpha = 2^\circ$, case, with a slope computed from the discrete lift coefficient adjoint. The agreement is excellent for small α perturbations; deviations at larger α are due to the nonlinear nature of the compressible Navier-Stokes equations.

3 Output Error Estimation

Numerical errors due to insufficient mesh resolution can affect outputs, often in seemingly subtle but significant ways. The latter point is especially true for the convection-dominated flows common to aerodynamics. Our goals are to quantify the effect of these numerical errors and to reduce them through mesh adaptation. In this section, we consider the first of these: estimation of output error.

3.1 Two Discretization Levels

Without access to infinite resolution, estimating the true numerical error in an output is practically out of reach for general nonlinear problems. We thus resign ourselves to estimating the output error between two finite-dimensional spaces: a coarse approximation space (\mathbf{V}_H) on which we calculate the state and output, and a fine space (\mathbf{V}_h) relative to which we estimate the error. The equations and output representations on these spaces are

$$\begin{aligned} \text{coarse space: } & \rightarrow \underbrace{\mathbf{R}_H(\mathbf{U}_H) = 0}_{N_H \text{ equations}} \rightarrow \underbrace{\mathbf{U}_H}_{\text{state} \in \mathbb{R}^{N_H}} \rightarrow \underbrace{J_H(\mathbf{U}_H)}_{\text{output (scalar)}} \\ \text{fine space: } & \rightarrow \underbrace{\mathbf{R}_h(\mathbf{U}_h) = 0}_{N_h \text{ equations}} \rightarrow \underbrace{\mathbf{U}_h}_{\text{state} \in \mathbb{R}^{N_h}} \rightarrow \underbrace{J_h(\mathbf{U}_h)}_{\text{output (scalar)}} \end{aligned}$$

We would like to measure the output error in the coarse solution relative to the fine space,

$$\text{output error: } \delta J \equiv J_H(\mathbf{U}_H) - J_h(\mathbf{U}_h). \quad (30)$$

The fine space is typically constructed by uniformly refining each element in the coarse space, or by increasing each element's approximation order. Figure 14 illustrates a fine space obtained by uniform refinement. We assume that the fine approximation space contains the coarse approximation space, so that the following lossless state injection, \mathbf{U}_h^H , is possible:

$$\mathbf{U}_h^H \equiv \mathbf{I}_h^H \mathbf{U}_H, \quad (31)$$

where \mathbf{I}_h^H is the coarse-to-fine state injection (prolongation) operator.

3.2 The Adjoint-Weighted Residual

On the fine space, the exact solution $\mathbf{U}_h \in \mathbb{R}^{N_h}$ would give us zero fine-space residuals,

$$\mathbf{R}_h(\mathbf{U}_h) = \mathbf{0}. \quad (32)$$

However, the state injected from the coarse space will generally not be a fine space solution and hence will not give us zero fine-space residuals,

$$\mathbf{R}_h(\mathbf{U}_h^H) \neq \mathbf{0}. \quad (33)$$




Figure 14: Sample fine approximation space obtained by uniform refinement of each element in the mesh of the coarse approximation space. An alternative fine space is obtained by incrementing the approximation order of each element.

Instead, the injected coarse state solves a *perturbed* fine-space problem,

$$\text{find } \mathbf{U}'_h \text{ such that: } \mathbf{R}_h(\mathbf{U}'_h) - \underbrace{\mathbf{R}_h(\mathbf{U}_h^H)}_{\delta \mathbf{R}_h} = 0 \quad \Rightarrow \quad \text{answer is: } \mathbf{U}'_h = \mathbf{U}_h^H. \quad (34)$$

As this is just the fine-space problem with a residual perturbation, the fine-space adjoint, Ψ_h , tells us to expect an output perturbation given by the inner product between the adjoint and the residual perturbation,

$$\underbrace{J_h(\mathbf{U}_h^H) - J_h(\mathbf{U}_h)}_{\approx \delta J} = \Psi_h^T \delta \mathbf{R}_h = -\Psi_h^T \mathbf{R}_h(\mathbf{U}_h^H). \quad (35)$$

This derivation assumes small perturbations in \mathbf{U} and \mathbf{R} when the output or equations are nonlinear. Calling the left-hand side δJ assumes $J_H(\mathbf{U}_H) = J_h(\mathbf{U}_h^H)$, which is true if the output definition (e.g. geometry) does not change between the coarse and fine spaces.

In summary, we have

$$\boxed{\delta J \approx -\Psi_h^T \mathbf{R}_h(\mathbf{U}_h^H)} \quad (36)$$

Note that this error estimate does not require the fine-space primal solution, \mathbf{U}_h .

3.3 Approximations

The error estimate in (36) uses the adjoint on the fine space, Ψ_h . Obtaining Ψ_h requires solving a large, possibly expensive, linear system. Suppose we have the *coarse*-space adjoint, Ψ_H , which injected into the fine space is $\Psi_h^H \equiv \mathbf{I}_h^H \Psi_H$. Define the adjoint perturbation as

$$\delta \Psi_h \equiv \Psi_h^H - \Psi_h.$$

We then re-write (36) as

$$\delta J \approx \underbrace{-\left(\Psi_h^H\right)^T \mathbf{R}_h\left(\mathbf{U}_h^H\right)}_{\text{computable correction}} + \underbrace{\left(\delta \Psi_h\right)^T \mathbf{R}_h\left(\mathbf{U}_h^H\right)}_{\text{remaining error}}. \quad (37)$$

The computable correction is tempting to use as the sole error estimate. It does provide important information on δJ for many discretizations, including reconstructed finite volume. However, it performs poorly as an adaptive indicator because it does not incorporate any new information from the fine space. Moreover, it is zero for finite element discretizations with Galerkin orthogonality. Therefore we need some estimate of Ψ_h . Several methods are used in practice,

1. Reconstruct Ψ_h from Ψ_H using information from neighboring elements.
2. Solve for Ψ_h approximately using a cheap iterative smoother on the fine space.
3. Solve for Ψ_h exactly on the fine space if the $N_h \times N_h$ linear system is tractable.

In this work we employ the third option when possible, as it gives the most accurate error estimates (albeit at the highest cost). When the fine-space solve becomes too expensive, we turn to one of the first two approximations; although we sacrifice some accuracy in the error estimate, the adaptive indicator obtained from the error estimate remains similar to that obtained from solving the fine-space adjoint exactly.

For nonlinear problems the leading term not present in (37) is quadratic in the state and adjoint errors. This “error in the error estimate” can be reduced to third-order in the state and adjoint errors by using [98]

$$\delta J \approx -\left(\Psi_h^H\right)^T \mathbf{R}_h\left(\mathbf{U}_h^H\right) + \frac{1}{2}\left(\delta \Psi_h\right)^T \mathbf{R}_h\left(\mathbf{U}_h^H\right) + \frac{1}{2}\left(\delta \mathbf{U}_h\right)^T \mathbf{R}_h^\psi\left(\Psi_h^H\right), \quad (38)$$

where $\mathbf{R}_h^\psi\left(\Psi_h^H\right)$ is the residual vector of the fine-space adjoint problem, (22), and $\delta \mathbf{U}_h = \mathbf{U}_h^H - \mathbf{U}_h$ is the state perturbation. In practice, both the state and adjoint perturbations could be approximated using one of the three approaches outlined above.

3.4 Error Effectivity

The error estimate calculated above is not a bound. We measure its accuracy by defining an *effectivity*,

$$\eta_H = \frac{J_H\left(\mathbf{U}_H\right) - J_h\left(\mathbf{U}_h\right)}{J_H\left(\mathbf{U}_H\right) - J}, \quad (39)$$

where J is the exact output, i.e. calculated from the exact solution. An effectivity close to 1 is desirable. In practice this value will depend on the choice of fine space (order enrichment versus element subdivision), and on approximations made in the fine-space adjoint error estimation. However, we can make some rough a priori estimates. If the output converges as $J_H\left(\mathbf{U}_H\right) - J = CH^k$, uniform element subdivision for the fine space yields an effectivity of $\eta_H = 1 - (1/2)^k$ – this does not approach unity as $H \rightarrow 0$. On the other hand, if order enrichment is used for the fine space, then the effectivity converges as $\eta_H = 1 - C_1 H^{\delta k}$, where δk is the increase in convergence rate of the fine-space output relative to the coarse-space output. In this case, the effectivity does approach unity with mesh refinement [36].

3.5 Examples

We now present some examples that test the error estimates described in this section. We use uniform mesh refinement studies to determine the rate, k , at which certain errors, notably the output error, asymptotically converge,

$$\text{error} \propto h^k, \quad \text{valid as } h \rightarrow 0, \quad (40)$$

where h is a measure of the size of the mesh elements. Even though we have many elements, not all of the same size, we use a single h value that we can think of as indicating a refinement level. When we do a uniform refinement study, we are not concerned with the particular value of h ; we just want to know how “uniform” changes in h for every element translate into changes in the error.

In the following studies we will uniformly subdivide mesh elements, by bisecting each edge³, and measure the resulting effect on the error. In this case, a reasonable definition of h is $\sqrt{1/N_e}$, where N_e is the number of elements. So a uniform refinement of a two-dimensional mesh would increase N_e by a factor of 4, and decrease h by a factor of 2. Taking the logarithm of (40), we obtain

$$\log(\text{error}) = C + k \log\left(\sqrt{\frac{1}{N_e}}\right). \quad (41)$$

3.5.1 Drag Error for Euler Flow over a Bump

In this example, we consider inviscid subsonic flow inside a channel that has a smooth Gaussian perturbation on the bottom wall. The compressible Euler equations govern the flow, and the output J is the drag (horizontal) force on the bottom wall. The solution on a fine mesh is shown in Figure 15.1.


Figure 15.2 shows the initial coarse mesh for the uniform refinement study. On this mesh, we perform the following steps:

1. Solve for the flow using $p = 1$. This gives us \mathbf{U}_H , which is the state in the coarse approximation space. Also compute $J_H = J_H(\mathbf{U}_H)$.
2. Solve the adjoint exactly using $p = 2$. This gives us Ψ_h , which is in the fine approximation space.
3. Compute the error estimate using (36) and find the associated corrected output,


$$\text{corrected output} = J_H - \delta J. \quad (42)$$

We then repeat these steps on three successive refinements of the coarse mesh. We also obtain the “exact” output, J , by solving using $p = 3$ approximation on a mesh that is uniformly refined once more compared to the finest mesh in the study.


³When using curved elements it is important to bisect curved edges as close to along the arc length as possible, and the same applies for element interiors; for example, bisecting a reference element when using a reference-to-global mapping with a lot of nonlinear stretching would not necessarily reduce the size of each element by the same amount.




15.1: Mach number contours; range is 0.27 to 0.43



15.2: Coarsest mesh for uniform refinement study: quartic curved triangles



15.3: Convergence of baseline and corrected output




15.4: Convergence of error effectivity

Figure 15: Error estimates and effectivity for inviscid flow in a channel with a Gaussian bump, at $M_\infty = 0.3$, using $p = 1$ approximation. The output of interest is the drag force coefficient on the bottom wall. The adjoint is solved exactly on the fine space.


Figure 15.3 shows the convergence of the errors in the coarse-space output, $J_H - J$, and in the corrected output, $J_H - \delta J - J$. Based on the formula in (41), we expect (asymptotically) straight lines on a log-log plot, and this is what we see. The slopes tell us the convergence rates. First, we observe a rate of $k = 3$ for the output error $J_H - J$, and for our $p = 1$ approximation, this is a super-convergent result ... approximation theory would lead us to expect a rate of $p + 1$ but we are actually seeing a rate of $2p + 1$ (we can verify this with data from higher p). Second, we observe a rate of 4, $2p + 2$, for the corrected output. So we see that the output correction buys us an extra order of convergence for the output, which is a reasonable result as we are using order enrichment for the fine space.

Figure 15.4 plots the error effectivity, (39), for the solutions on the 4 mesh refinements. We actually plot two effectivities: one as defined in (39) (this is relative to the exact output J), and a similar one but relative to the fine-space output J_h . The latter choice gives us an idea of the errors we make in estimating $J_H - J_h$, while the former tells us about $J_H - J$. We see that as expected (by design), we estimate the error better relative to the fine space than relative to the true output, but that both effectivities approach 1 as $H \rightarrow 0$.

We mentioned that in practice we sometimes solve the fine-space adjoint problem approximately, to avoid the cost of a full solve on the fine space. It turns out that this approximation degrades the accuracy of the error estimates, resulting in worse convergence rates for the corrected output compared to when using an exact fine space adjoint. Figure 16 shows this degradation when using ν iterations of an element-block Jacobi solver to smooth the fine-space adjoint after injection from the coarse space. We see that when



16.1: Convergence of baseline and corrected outputs



16.2: Convergence of error effectivities

Figure 16: Effect of inexact fine-space adjoint solve, using ν element-block smoothing iterations, on the error estimates for inviscid flow in a channel with a Gaussian bump, at $M_\infty = 0.3$, using $p = 1$ approximation. The output of interest is the drag force coefficient on the bottom wall.

ν is small (less work on the fine space), the error estimates are not great, although the corrected output is still better than the original coarse one. As ν increases, the error estimates approach the exact adjoint solve result ($\nu \rightarrow \infty$), and the effectivities close in on 1. The “adequate” number of smoothing iterations will depend on the case and on

the smoother, in addition to the desired accuracy levels, so these results should not be interpreted as a recipe for how to choose ν . Instead, one take-away message is that to get the best error estimates possible, we need the best possible fine-space adjoint; however, another one is that even approximate fine-space adjoints can lead to improved outputs. Furthermore, in an adaptive setting, the indicators obtained from these error estimates will turn out to be very good at driving mesh refinement.


3.5.2 Drag Error for Viscous Flow over a NACA 0012 Airfoil

In this example, we consider viscous subsonic ($M_\infty = 0.5$) flow over a NACA 0012 airfoil (closed-trailing-edge [32]) at zero angle of attack. The compressible Navier-Stokes equations, at $Re = 5000$, $Pr = 0.71$, and constant viscosity, govern the flow, and the output J is the drag (horizontal) force on the airfoil, at which an adiabatic no-slip boundary condition is imposed. The solution on a fine mesh is shown in Figure 17.1.


We follow the same steps as in the previous example for measuring the convergence rate of the output and the corrected output (i.e. the error estimate)⁴. Figure 17.3 shows the convergence of the errors in the coarse-space output, $J_H - J$, and in the corrected output, $J_H - \delta J - J$. The lines are not straight initially, indicating that the solution is not yet in the asymptotic regime. However, by the finer meshes we observe relatively straight lines on the log-log plot. We observe a rate of $k \approx 2.7$ for the output error $J_H - J$, and for our $p = 1$ approximation, this is still super-convergent relative to the expected rate of $p + 1$. For the corrected output, we observe a higher rate of 3.8, and so we see that the output correction buys us an extra order of convergence, which is again reasonable as we are using order enrichment for the fine space.

The effectivity story in Figure 15.4 is similar to that of the previous example. The error estimate does a better job at predicting the error between the coarse space and the fine space than at predicting the error relative to the exact output, but both of the effectivities converge to 1 as $H \rightarrow 0$.


⁴We also follow the remedy for alleviating the effects of p -dependence of the BR2 residual on the error estimates, by evaluating the fine-space residual with order p integration rules and stabilization approximation. [126]




17.1: Mach number contours; range is 0.0 to 0.6



17.2: Coarsest mesh for uniform refinement study: quartic curved quadrilaterals



17.3: Convergence of outputs



17.4: Convergence of error effectivity

Figure 17: Error estimates and effectivity for viscous $Re = 5000$, $M_\infty = 0.5$ flow over a NACA 0012 airfoil at $\alpha = 0$, using $p = 1$ approximation. The output of interest is the drag force coefficient on the airfoil, and the adjoint is solved exactly on the fine space.

4 Mesh Adaptation

The output error estimate derived in the previous section provides error bars on quantities of interest from numerical simulations. However, the benefit of these estimates extends beyond the error bars. Because the output error estimate takes the form of a weighted residual, and because local mesh refinement decreases residuals, the error estimate provides a means of targeting for refinement areas of the computational domain that give rise to the output error. This is the key idea of output-based mesh adaptation.

4.1 Error Localization

The adjoint-weighted residual error estimate in (36) can be localized to the elements by keeping track of the contributions from each fine-space element, indexed by e below,

$$\begin{aligned} J_H(\mathbf{U}_H) - J_h(\mathbf{U}_h) &\approx -\Psi_h^T \mathbf{R}_h(\mathbf{U}_h^H) = -\sum_e \Psi_{he}^T \mathbf{R}_{he}(\mathbf{U}_h^H) \\ \Rightarrow \epsilon_e &\equiv \left| \Psi_{he}^T \mathbf{R}_{he}(\mathbf{U}_h^H) \right|, \end{aligned}$$

where the subscript e indicates restriction to element e , and the adaptive indicator ϵ_e is obtained by taking the absolute value of the elemental contributions. When order enrichment is used for the fine space in error estimation, this ϵ_e is the adaptive indicator for each element in the current mesh and can be used directly to drive a mesh adaptation strategy. When element refinement is used for the fine space, then the indicators for the fine sub-elements of a coarse element need to be summed to obtain the coarse-element adaptive indicator. The steps involved in obtaining the adaptive indicator, ϵ_e are summarized graphically in Figure 18.

4.2 Adaptation Mechanics

Numerous strategies exist for translating the error indicator into a modified computational mesh. In CFD for aerospace engineering, the most popular adaptation strategy is h -adaptation, in which only the triangulation forming the mesh is modified. This modification usually consists of targeted refinement and coarsening, although pure node repositioning, sometimes called r -refinement, has also been investigated [21, 78]. For high-order methods, additional strategies include p -adaptation, in which the approximation order is changed on a fixed triangulation [116, 72], and hp -adaptation in which both the order and the triangulation are varied [50, 16, 97, 2, 62, 27, 59, 80, 61, 65]. For CFD applications, in which solutions often possess localized, singular features, h -adaptation is key to an efficient adaptation strategy. With the growing availability of high-order methods, however, hp -adaptation can now be used to further improve efficiency.

4.2.1 Local Refinement

Many approaches to adapting a mesh rely upon the application of local operators through which the mesh is modified incrementally. A simple example of a local operator is element sub-division in a setting that supports non-conforming, or hanging, nodes [15, 66, 98, 27, 114]. For triangular and tetrahedral meshes, local mesh modification operators consist of







18.1: $p = 1$ state, \mathbf{U}_H 18.2: $p = 1$ residual, $\mathbf{R}_H(\mathbf{U}_H)$ (zero as expected)18.3: $p = 1$ state injected to $p = 2$, \mathbf{U}_h^H 18.4: $p = 2$ residual, $\mathbf{R}_h(\mathbf{U}_h^H)$ 18.5: $p = 2$ lift adjoint, Ψ_h 18.6: Error indicator, $\epsilon_e = |\Psi_{he}^T \mathbf{R}_{he}(\mathbf{U}_h^H)|$

Figure 18: Quantities involved in the calculation of the error estimate and adaptive indicator for $Re = 5000$, $\alpha = 2^\circ$, $M = 0.1$ flow over a NACA 0012 airfoil.

node insertion, face/edge swapping, edge collapsing, and node movement, as illustrated in Figure 19.1. These operators have been studied extensively by various authors [43, 20, 22,

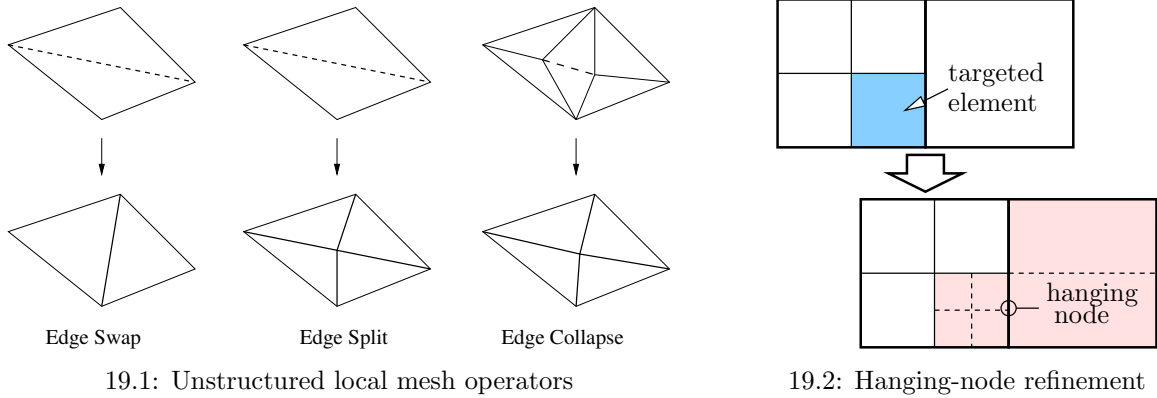


Figure 19: Local mesh modification adaptation operators in two dimensions.

122, 51, 123, 5, 88, 89] in different contexts. The primary advantage of local operators is their robustness: the entire mesh is not regenerated all at once, but rather each operator affects only a prescribed number of nodes, edges, or elements.

Another local operation, especially relevant for discontinuous Galerkin discretizations, is hanging-node mesh refinement, illustrated for quadrilateral elements in Figure 19.2. The non-conforming nature of a hanging node mesh does not significantly affect the DG discretization, which does not enforce solution continuity between elements. Hanging node refinement could be isotropic or directional [23]. It is particularly useful for initially-structured, stretched meshes for viscous flows, in which only a few refinements of key regions may dramatically improve the accuracy of an output.

4.2.2 Global Re-Meshing

Another approach to adapting a mesh is global re-meshing, in which a new mesh is generated for the entire computational domain. The original, or background, mesh is used to store desired mesh characteristics during regeneration. For applications to adaptation, the desired mesh characteristics are often described using a Riemannian metric, the idea being that in an optimal mesh, all edge lengths will have unit measure under the metric [22, 51]. In a Cartesian coordinate system of dimension d , an infinitesimal segment $\delta\mathbf{x}$ has length $\delta\Gamma$ under a Riemannian metric \mathbf{M} ,

$$\delta\Gamma^2 = \delta\mathbf{x}^T \mathbf{M} \delta\mathbf{x} = \delta x_i M_{ij} \delta x_j, \quad (43)$$

where δx_i are the components of $\delta\mathbf{x} \in \mathbb{R}^d$, M_{ij} are the components of the symmetric, positive definite metric, $\mathbf{M} \in \mathbb{R}^{d \times d}$, and summation is implied on the repeated indices $i, j \in [1, \dots, d]$.


The metric \mathbf{M} contains information on the desired mesh edge lengths in physical space. As \mathbf{M} is symmetric and positive definite, the unit measure requirement,

$$\mathbf{x}^T \mathbf{M} \mathbf{x} = 1,$$


describes an ellipsoid in physical space that maps to a sphere under the action of the metric. The eigenvectors of \mathbf{M} form the orthogonal axes of the ellipsoid – i.e. the principal directions. The corresponding eigenvalues, λ_i , are related to the lengths of the axes, h_i , via

$$\lambda_i = \frac{1}{h_i^2} \quad \Rightarrow \quad \frac{h_i}{h_j} = \left(\frac{\lambda_j}{\lambda_i} \right)^{1/2}$$

Physically, the h_i are the principal stretching magnitudes. A diagram of a possible ellipse resulting from the unit-measure requirement in two dimensions is given in Figure 20.1. Thus, the ratio of eigenvalues of \mathbf{M} can be used to define a desired level of anisotropy.



20.1: Riemannian metric ellipse



20.2: Sample mesh obtained by global re-meshing

Figure 20: Depiction of adaptation using global metric-based re-meshing. On the left, an ellipse representing requested mesh sizes implied by equal measure under a Riemannian metric \mathbf{M} , together with the principal directions, \mathbf{e}_i , and the associated principal stretching magnitudes, h_i . On the right, sample mesh for viscous flow generated using global re-meshing.

A successful approach for generating simplex meshes based on a Riemannian metric is mapped Delaunay triangulation, in which a Delaunay mesh generation algorithm [110] is applied in the mapped space, allowing for the creation of stretched and variable-size triangles or tetrahedra [75]. This method is implemented in the Bi-dimensional Anisotropic Mesh Generator (BAMG) [17, 58], which has been used in various finite volume [107, 119] and discontinuous Galerkin [35, 7, 85] applications requiring anisotropic meshes. An example of an output-adapted mesh obtained using BAMG is shown in Figure 20.2.

4.2.3 Targeting Strategies

In global re-meshing, all elements can be refined or coarsened based on their error indicators. However, when using local operators, such as hanging-node or order refinement, one must decide which elements to target. The determination of which elements to refine or coarsen has important implications, as too little refinement at each adaptive iteration may result in an unnecessary number of iterations, while too much refinement may result in an expensive solve on an overly-refined mesh. A useful tool for analyzing adaptation targeting strategies is an error distribution histogram [1], in which elements are binned

according to the error indicator. A typical assumption is that in an ideal mesh, the error is equidistributed among the elements [5], and this situation yields a “delta” histogram, in which all elements lie in the same bin. In contrast, the initial coarse mesh will generally have some distribution of error indicators, and the goal of an adaptation targeting strategy is to drive the histogram towards the ideal delta distribution. This characterization of adaptation targeting strategies also holds for runs in which a maximum element count is specified instead of an error tolerance. The ideal mesh in this case is one for which the error is equidistributed among a number of elements within the element budget [125, 126].

Most adaptation targeting strategies are based on a decreasing refinement threshold [83], in which elements with the highest error are targeted for refinement first so that the mesh size grows gradually. For example, a fixed-fraction approach prescribes a fraction of elements with the highest error indicator to be refined at each adaptation iteration, such that the decreasing threshold is a function of the shape of the error histogram. Then, the elements targeted for adaptation are typically refined in a locally uniform manner, e.g. by splitting all edges in half. This simple approach has been applied to output-based adaptation in several studies [13, 57, 9, 112, 65, 82, 25]. The fixed-fraction parameter is often chosen heuristically in a trade-off between an excessive number of iterations and a risk of over-refinement. Nevertheless, the method works quite well for practical problems.

4.2.4 Incorporating Anisotropy

An important ingredient in h -adaptation for aerodynamics is the ability to generate stretched elements in areas such as boundary layers, wakes, and shocks, where the solution exhibits anisotropy, which refers to variations of disparate magnitudes in different directions. Anisotropy can be created to a limited, albeit often sufficient, extent with hanging-node refinement [111, 115]; yet global re-meshing with unstructured triangular and tetrahedral grids offers the most flexibility in tailoring the required stretching.

For spatially second-order methods, the dominant method for detecting anisotropy involves estimating the Hessian matrix of second derivatives of a scalar (e.g. Mach number) computed from the solution [93, 75, 22, 51]. A mesh metric is obtained from the Hessian by requiring that the approximation error estimate of the scalar quantity u be the same in any chosen spatial direction. The Hessian matrix stores precisely this information, so that this requirement leads to a metric that is directly proportional to the Hessian. The proportionality constant can be tied to the error indicator to incorporate the output error estimate [119].

The definition of a metric tensor becomes difficult for high-order methods because the standard Hessian matrix approach assumes linear approximation of the scalar quantity. One possible extension is based on constructing a metric around the direction of maximum $p + 1$ st derivative [32, 31, 87]. Another extension involves the calculation of the order $p + 1$ derivative tensor that is analogous to the Hessian for $p = 1$, or the use of surrogate heuristics based on inter-element jumps for quadrilateral or hexahedral meshes [69].

The metric tensor may also be used to guide an adaptation procedure based on local operators, for example in an effort to make all edges approximately the same length when measured using the metric tensor [22, 29, 123, 89].

The methods mentioned so far rely on a priori analysis and heuristics to predict the desired mesh anisotropy. For example, the approximation error assumptions are made

without regard to the output of interest by using a single scalar, such as the Mach number, to control the anisotropy for a system of equations. However, the assumption that the directional approximation error must be equidistributed for one or more scalar variables at each point in the domain may not be valid, especially when only a scalar output is desired. This observation has motivated research into adaptation algorithms that more directly target the error indicator.

Hessian-based approximation error estimates have been combined with output-based a posteriori error analysis to arrive at an output-based error indicator that explicitly includes the anisotropy of each element [42, 41, 40]. Other works have considered directional output error estimates for quadrilateral and hexahedral meshes [99], and direct node position optimization using an auxiliary adjoint problem [109].

For general unstructured meshes, output error estimation has been incorporated into local mesh operators of element swapping, node movement, element collapse, and element splitting [90]. A direct optimization approach for output error reduction has also been applied to quadrilateral and hexahedral elements, using anisotropic discrete refinement options [63, 55, 23]. Finally, a recent work extends the ideas of direct mesh optimization to general unstructured meshes with global re-meshing, through optimization of a parametrized metric field via local refinement sampling [126].

4.2.5 Adapting in Order

In order/ p -adaptation [116], the approximation space is refined or coarsened by changing the order of approximation. With the discontinuous Galerkin method, changing the order is simple and can be done locally on each element [72, 34]. Advantages of p -adaptation are that the computational mesh remains fixed and that an exponential error convergence with respect to degrees of freedom (DOF) is possible for sufficiently smooth solutions. Disadvantages include difficulty in handling singularities and areas of anisotropy and the need for a reasonable starting mesh.

hp -adaptation strives to combine the best of both strategies, employing p -refinement in areas where the solution is smooth and h -refinement near singularities or areas of anisotropy. The motivation for this strategy is that, in smooth regions, p -refinement is more effective at reducing the error per unit cost, compared to h -refinement [120, 65]. Implemented properly, hp -adaptation can isolate singularities and yield exponential error convergence with respect to DOF. In practice, however, the difficulty of hp -adaptation lies in making the decision between h - and p -refinement, which typically requires either a solution regularity estimate or a heuristic algorithm [65, 19]. An approach that employs output error information to make the hp decision and to choose the appropriate element anisotropy for quadrilateral and hexahedral elements has also been presented [25].

4.3 Examples

The following sub-sections demonstrate applications of output-based refinement to various aerodynamic flows. We consider inviscid, viscous, and turbulent flows in two and three dimensions. Our figure of merit will generally be degrees of freedom, although we also present some results for computational time.

4.3.1 Inviscid Flow over an Airfoil

In this example we adapt a mesh to predict drag in inviscid flow over a NACA 0012 airfoil at $M_\infty = 0.5$ and $\alpha = 2^\circ$. Figure 21 shows the Mach number contours and the initial mesh. We compare uniform refinement to output-driven, isotropic, hanging-node, fixed-fraction refinement with $f^{\text{adapt}} = 0.05$.




Figure 21: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$: Contours of the Mach number and the initial mesh.

Figure 22 shows the results of one adaptation run, using $p = 2$. The drag coefficient error decreases as the number of elements increases, at a steady rate. At each iteration, both the estimated and actual errors are shown, and these agree very well, indicating that the error estimate is effective at predicting the true error. The figure also shows uniform refinement results, in which the error does not drop as rapidly.

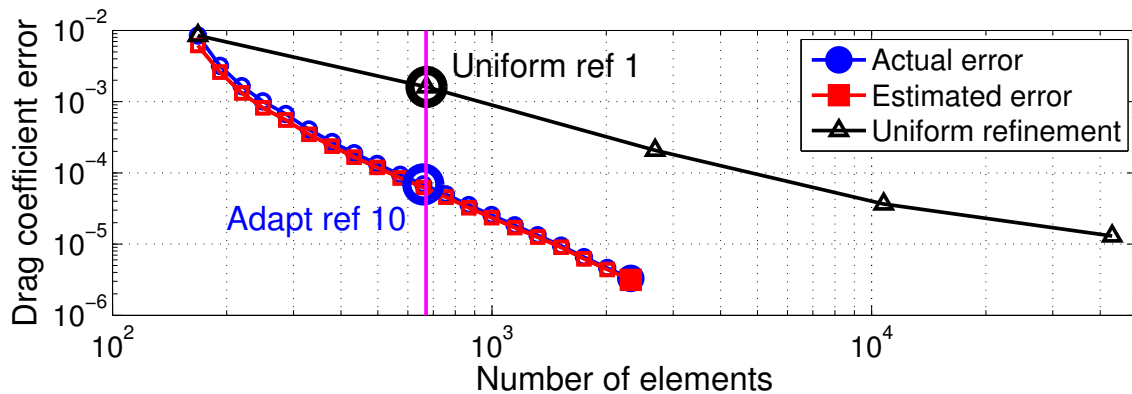


Figure 22: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$: actual and estimated drag coefficient errors obtained from $p = 2$ output-based adaptation, compared to uniform refinement.

We can gain insight into where adaptive refinement allocates degrees of freedom by comparing meshes from uniform refinement and adaptive refinement at a similar number of elements. Figure 22 identifies two such meshes, which are shown in Figure 23. We can see that for accurate drag prediction, the leading and trailing edges of the airfoil must be resolved with small elements.

Figure 24 shows the convergence of the drag coefficient error with a measure of the mesh size, $\text{dof}^{-1/2}$, for $p = 1$ and $p = 2$ approximation. We see that uniform refinement

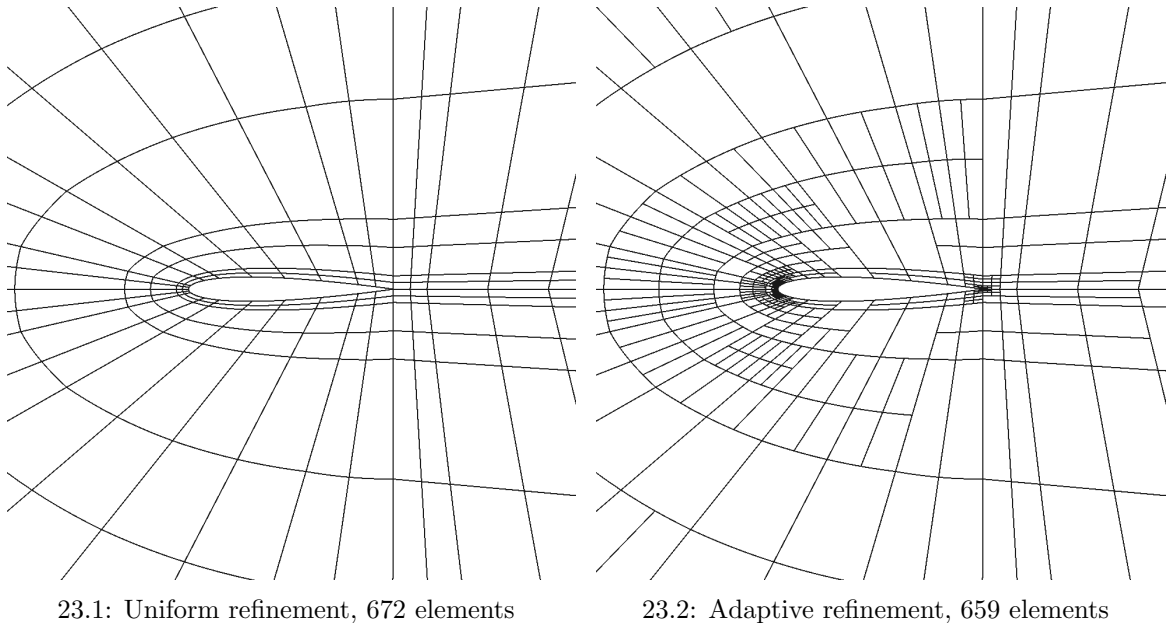


Figure 23: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$: comparison of meshes obtained from uniform refinement and adaptive refinement, at a similar total number of elements.

converges at a rate of 2.5 for $p = 1$ and $p = 2$. This rate is limited due to a singularity of the flow at the trailing edge, which then dominates the error convergence. On the other hand, we see that adaptive refinement produces meshes that make better use of degrees of freedom: the equivalent “convergence rate” for $p = 2$ adaptation is approximately 5, i.e. the expected super-convergent rate of $2p + 1$ for an inviscid simulation. We note that although we usually speak of convergence rates only for uniform refinement studies, this notion generalizes to adaptive mesh refinement [125]. For example, for a case in which the entire domain is important (e.g. a very smooth solution), adaptive refinement would (eventually, in turn) target the entire domain too so that its error versus degrees of freedom would converge at the same rate as for uniform refinement. More interestingly, when there are (isolated) singularities present, a situation in which uniform refinement would see its rate limited, adaptive refinement could “quarantine” these singularities with nominal resources and allocate the rest to resolving the remaining smooth regions at the optimal rate. This is indeed what we observe in Figure 24, where, while uniform refinement sees its rate limited by non-smooth solution features for high order, adaptive refinement attains a higher, super-convergent rate. We also note that correcting the output with the error estimate further increases the order of accuracy by 1 to $2p + 2$. In this example, the fine-space problem was solved exactly, but solving it approximately (e.g. with a few block-Jacobi smoothing iterations) does not appreciably change the results. For more complex problems, not solving the fine-space adjoint exactly can yield poor error estimates.

Adapting using an output adjoint consumes more resources, i.e. CPU time, than just solving the primal problem because of the need for a fine-space adjoint. To determine whether this additional cost pays off, Figure 25 plots the convergence of the drag output with CPU time, for both an exact and an approximate fine-space adjoint solution. We see that the cost of an exact adjoint solve pushes the output-based results close to the



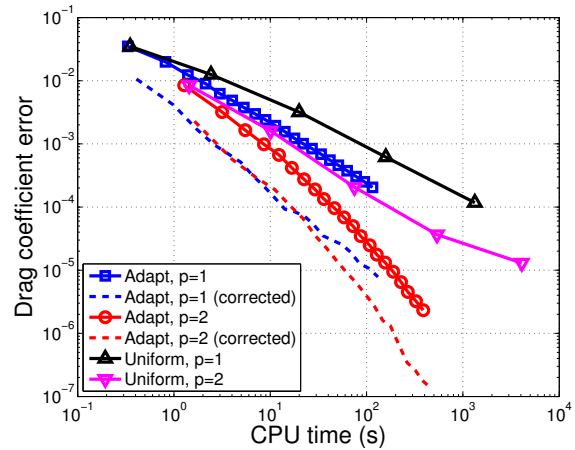


Figure 24: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$: drag coefficient convergence with mesh size, $\text{dof}^{-1/2}$, using uniform refinement and output-based adaptation. Orders of approximation $p = 1$ and $p = 2$ are shown. The convergence rate is identified for each run: we observe a rate of $2p + 1$ for the output-adapted results, and $2p + 2$ for the adapted results corrected with the error estimate. Uniform refinement yields suboptimal rates due to a singularity at the airfoil trailing edge.

uniform refinement results – the corrected results still remain more efficient, however. If we instead solve the fine-space adjoint problem approximately, we can reduce the CPU time of the output-based results by a factor of 2 or more.




25.1: Exact fine-space adjoint




25.2: Approximate fine-space adjoint

Figure 25: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$: output convergence with CPU time for different adaptation strategies. Using an approximate fine-space adjoint solve of five block-Jacobi smoothing iterations yields noticeable efficiency improvements over an exact (machine precision) adjoint solve. Correcting the output with the error estimate further improves efficiency.

Finally, Figure 26 shows the breakdown of the CPU time at each adaptive iteration when using the exact and approximate adjoint solves. The exact adjoint solve consumes more than half of the CPU time, whereas the approximate adjoint solve is much cheaper: less than 10% of the total CPU cost. In both cases, the cost of adapting the mesh is negligible compared to the primal and adjoint solutions.



26.1: Exact fine-space adjoint




26.2: Approximate fine-space adjoint

Figure 26: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$: timing breakdown, summed over all adaptive iterations for $p = 2$, when using exact and approximate adjoint solves. output convergence with CPU time for different adaptation strategies.

4.3.2 Viscous Flow over an Airfoil

We consider again the case of a NACA 0012 airfoil in viscous flow, introduced in Section 3.5.2. We use $p = 2$ solution approximation and fixed-fraction, $f^{\text{adapt}} = 0.05$, and

isotropic hanging-node adaptation. The output of interest is the drag coefficient, and the initial mesh is the same as in the previous example. Figure 27 shows contours of the primal solution (Mach number) and the drag adjoint solution (x -momentum component).



27.1: Mach number




27.2: x -Momentum drag adjoint

Figure 27: NACA 0012, $M_\infty = 0.5$, $\alpha = 0^\circ$, $Re = 5000$: contours of the Mach number and the x -momentum component of the drag adjoint. Note the wake reversal in the adjoint solution, which indicates high sensitivity of the output to residual perturbations made upstream of the airfoil.

Figure 28 shows the adaptive results when using a drag-adjoint weighted residual indicator, compared to uniform refinement. As in the previous example, we see that adaptive refinement produces meshes that make better use of degrees of freedom: the equivalent convergence rate for $p = 2$ is approximately 4.5, slightly above the expected rate of $2p$ for a viscous simulation. Uniform mesh refinement at $p = 2$ only attains a rate of 2.6. Thus, as in the previous example, adaptive refinement can “uncover” the optimal convergence rates for high-order discretizations, resulting in a more efficient use of degrees of freedom.

4.3.3 Which Output?

In this example, we consider a NACA 0012 airfoil with a closed trailing edge and a far field approximately 40 chord-lengths away. The initial mesh of cubic quadrilateral elements is illustrated in Figure 31. While the initial mesh appears structured, this structure disappears with the first adaptation iteration and the mesh storage is always fully unstructured. In the following results, quadratic solution approximation, $p = 2$, was used in the discretization, and isotropic h -adaptation was driven by a fixed-fraction


28.1: Final drag-adapted mesh for $p = 2$ 

28.2: Output convergence

Figure 28: Adaptation results for viscous $Re = 5000$, $M_\infty = 0.5$ flow over a NACA 0012 airfoil at $\alpha = 0$. The output of interest is the drag force coefficient on the airfoil, and, at each adaptive iteration, the adjoint is solved approximately with $\nu = 10$ element-block Jacobi smoothing iterations on the fine space. Using an exact adjoint solve does not perceptibly change the adaptive result.

strategy with $f^{\text{adapt}} = 0.1$, meaning that at each step of the adaptation, those elements lying in the top 10% of the error criterion were chosen for refinement.

Mach number contours for the airfoil in inviscid flow at $M_\infty = 0.4$, $\alpha = 5^\circ$ are shown in Figure 29. Three different engineering outputs are considered: drag coefficient, lift co-



29.1: Mach number contours



29.2: x -momentum moment adjoint

Figure 29: NACA 0012, $M_\infty = 0.4$, $\alpha = 5^\circ$: Contours of the Mach number and the x -momentum component of the moment adjoint.


efficient, and leading-edge moment coefficient. All of these outputs were computed using integrals of the inviscid momentum flux, that is, the pressure, on the airfoil surface. Adjoint solutions associated with these outputs were used to drive three different adaptation runs. One adaptation run was also performed using an “entropy-adjoint” indicator [39], in which the entropy variables are interpreted as a “free adjoint” for an output that expresses an entropy balance statement. For comparison, an unweighted residual indicator,

equivalent to summing the absolute values of the discrete fine-space residuals, was also tested.


Figure 30 shows the results of adaptation runs driven by the different indicators. Uniform mesh refinement results are given for comparison. The plots show the error in the engineering outputs versus degrees of freedom. Each “truth” output was calculated from a $p = 3$ solution on a mesh obtained by uniformly refining the finest output-adapted mesh. For all three outputs of interest, the adjoint-based adaptive strategies, including the en-



30.1: Drag output



30.2: Lift output



30.3: Moment output

Figure 30: NACA 0012, $M_\infty = 0.4$, $\alpha = 5^\circ$: Comparison of output convergence histories for various adaptation strategies.

trophy adjoint, perform similarly and are orders of magnitude better than uniform mesh refinement. The unweighted residual indicator performs well for the drag output, and similarly to the output adjoints for the lift and moment outputs. Interestingly, the refinement based on the entropy adjoint actually gives better predictions for lift and moment than the refinements that specifically target those outputs. These results are certainly surprising, but not actually paradoxical, because the procedure does have empirical elements. For example, the lift and moment adaptive indicators target the stagnation streamline in

front of the airfoil, perhaps excessively so. As shown in Figure 29 for the moment output, the adjoint varies rapidly across the stagnation streamline. This behavior was suggested in the analysis of Giles and Pierce who found that a square root singularity with respect to distance from the stagnation streamline exists for sources that perturb the stagnation pressure [46]. Intuitively, a force output on the airfoil should respond differently to perturbations that affect the flow over the upper surface of the airfoil versus to those that affect the flow over the lower surface of the airfoil. The singularity is strongest for the lift and moment outputs, and for these cases the performance of the output adjoint adaptation deteriorates the most. The noise created by polynomial approximation of the adjoint on discrete finite elements in this area may be responsible for the excessive refinement.


The meshes after eight adaptation iterations of each strategy are shown in Figure 31. The leading edge, trailing edge, and upper surface of the airfoil are consistently targeted for refinement by the adjoint indicators. The unweighted residual adaptation targets the vicinity of the leading edge and the trailing edge, but not the upper surface of the airfoil, leading to errors in the lift and moment outputs. Refinement of the stagnation streamline is evident in the adjoint-based runs, especially for the lift and moment adaptations. Finally, we remark that although residual-based refinement performs well in this example, this is not always the case for more complicated flows (e.g. see Section 6). In addition, residual-based adaptation does not provide an error estimate that could be used as a stopping criterion for the adaptation. The same may be argued for adaptation with the entropy adjoint, as that output is typically not of direct engineering interest. However, it turns out that the entropy adjoint output can be related to drag error [30] and hence could be used as an adaptive stopping criterion when targeting drag.

4.3.4 Transonic Turbulent Flow over an Airfoil


To demonstrate the importance of anisotropy, we present an example of output-based adaptation for turbulent, transonic flow over an NACA 0012 airfoil. The free-stream conditions are $M = 0.8$, $\alpha = 1.25^\circ$, $Re = 100,000$. In this case, the flow experiences a relatively strong normal shock on the upper surface and a weak shock on the lower surface. Element-wise constant artificial viscosity is used to capture the shocks [95]. $p = 2$ is used for solution approximation on all elements, and the initial mesh consists of 1740 cubic quadrilateral elements, as shown in Figure 32.

Convergence of the drag and lift coefficients is shown in Figure 33. The plots are similar in that the lift output converges as rapidly as the drag output for all of the adaptation schemes. Also, the drag adaptation performs well for the lift output and vice versa. The anisotropic adaptations, performed using a discrete-choice output-based approach [23], converge much more rapidly compared to the isotropic adaptations: the outputs do not change much after 40,000 degrees of freedom with the anisotropic adaptation, while changes are still observed after nearly 100,000 degrees of freedom when using isotropic adaptation.


Two of the drag-adapted meshes are shown in Figure 32: one from isotropic adaptation after six iterations and one from anisotropic adaptation after ten iterations. Differences are evident in the boundary layer and wake, where anisotropic adaptation is more efficient. The shock also appears to be more tightly resolved in the anisotropically-adapted mesh.




31.1: Initial mesh




31.2: Drag-adapted




31.3: Lift-adapted



31.4: Moment-adapted



31.5: Entropy-adjoint-adapted



31.6: Residual-adapted

Figure 31: NACA 0012, $M_\infty = 0.4$, $\alpha = 5^\circ$: Meshes after eight adaptation iterations for the tested adaptation strategies.

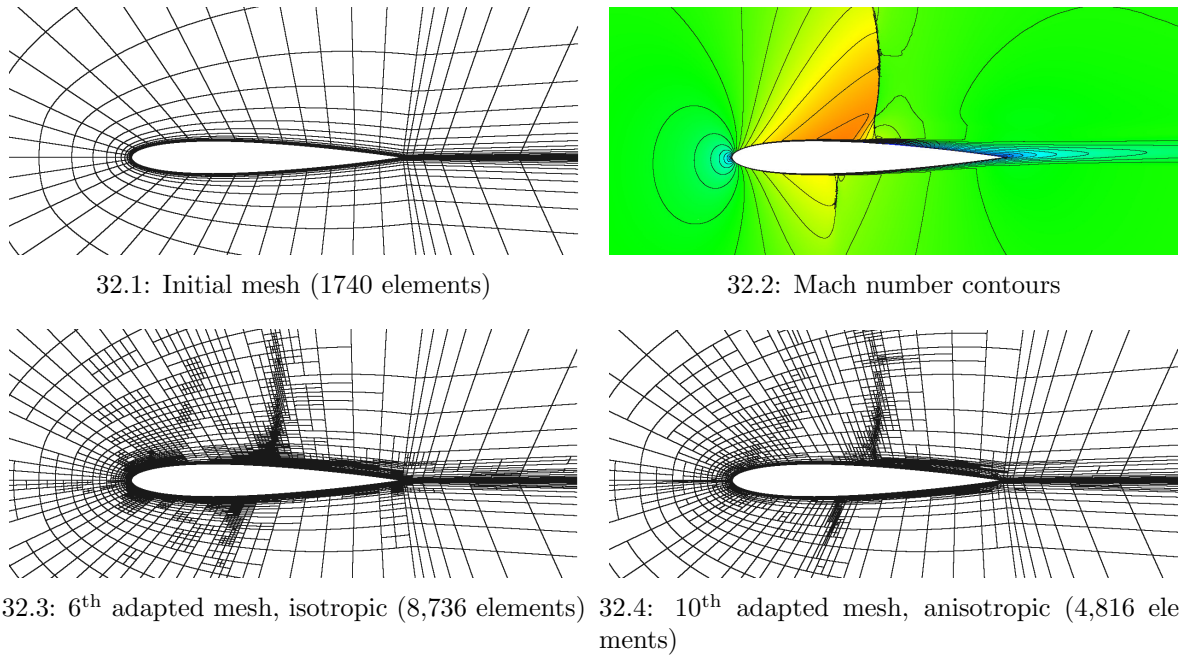


Figure 32: Turbulent NACA 0012, $M = 0.8$, $\alpha = 1.25^\circ$, $Re = 100,000$: Initial mesh, solution contours, and adapted meshes. Note, under isotropic adaptation, some elements are refined anisotropically, and these arise from additional refinements required to keep a maximum refinement ratio between adjacent elements below 2. of 2-to-1 ratio between is due to

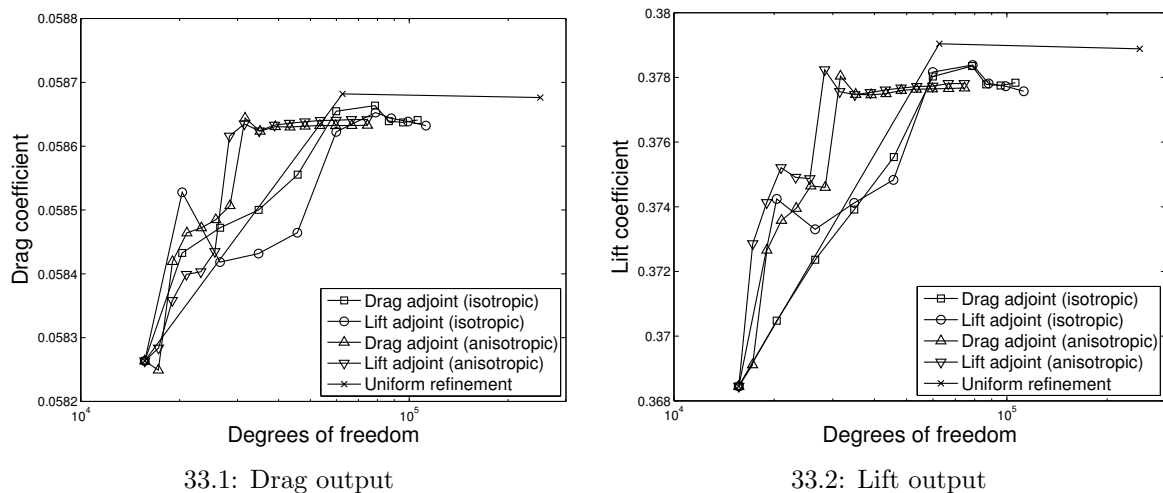



Figure 33: Turbulent NACA 0012, $M = 0.8$, $\alpha = 1.25^\circ$, $Re = 100,000$: Drag and lift convergence with degrees of freedom using isotropic and anisotropic refinement. Note, outputs from uniform refinement overshoot the exact values.

4.3.5 Transonic Turbulent Flow over a Wing

In this example, we demonstrate drag-based adaptation for a steady, three-dimensional, turbulent, transonic flow. We consider the baseline wing geometry (DPW-W1) from the third AIAA Drag Prediction Workshop [45]. The initial curved mesh, shown in Figure 34.1, was obtained through agglomeration of cells from a finer structured linear C-grid generated specifically for this purpose. In the agglomeration, each curved hexahedral element was obtained by merging twenty seven linear elements using a distance-based Lagrange interpolation of the nodal coordinates, resulting in cubic ($q = 3$) geometry interpolation. Also, the spacing of the linear mesh is such that the agglomerated mesh presents $y^+ \approx 1$ for the first element off the wall as recommended in the workshop guidelines [44] and the outer boundary is located at 100 mean-aerodynamic-chord-lengths away from the wing.

We use the Spalart-Allmaras turbulence model without trip terms, and element-wise constant artificial viscosity for shock capturing [95]. The baseline flow solution is obtained with linear ($p = 1$) approximation order, and we study anisotropic, output-based hp -adaptation using discrete choices that include order enrichment [25] – in this strategy, when deciding which discrete refinement option to choose, a figure of merit is used that takes into account the benefit (error addressed) and the cost (the increased Jacobian matrix size). The fine-space adjoint used for error estimation is obtained approximately by using $\nu^{\text{smooth}} = 5$ iterations of an element-block Jacobi smoother. All of the adaptive schemes start from the same initial solution. For the adjoint-based adaptation methods, the CPU time taken for the initial adjoint solve is also included in the initial starting time.




34.1: Initial pressure contours (29310 cubic elements, $p = 1$).


34.2: Pressure contours on the 7th drag-adapted mesh (85377 cubic elements).

Figure 34: DPW Wing 1, $M_\infty = 0.76$, $\alpha = 0.5^\circ$, $Re = 5 \times 10^6$: Initial and drag-adapted meshes with pressure contours.

We compare two mesh improvement strategies starting from the initial $p = 1$ solution shown in Figure 34.1. One of the strategies is uniform h -refinement, in which all hexahedra are divided into 8 elements. The other is drag-based hp -adaptation in which $f^{\text{adapt}} = 10\%$ of the elements is selected for refinement at each adaptation step. Additionally, we fix the overall budget of CPU wall-time for each of the three runs and the last converged solutions obtained within that budget are shown in Figure 34.



35.1: Convergence with degrees of freedom



35.2: Convergence with CPU time

Figure 35: DPW Wing 1, $M_\infty = 0.76$, $\alpha = 0.5^\circ$, $Re = 5 \times 10^6$: drag coefficient convergence for output-based adaptation compared to uniform refinement, using both degrees of freedom and CPU time.

Figure 35 shows the drag coefficient convergence for the mesh refinement strategies. Note that the dashed lines indicate the output corrected with the error estimate. The difference between these corrected values for the last two adaptation steps of the output-based strategy is within 0.15 counts of drag. Note that the performance in terms of degrees of freedom and CPU time of the output-based strategy is better compared to uniform refinement, especially when using the corrected output results.

Figures 36 and 37 show two cuts at representative span-wise positions for the output-based strategy. Note the presence of anisotropic cells along the shock and on the boundary layer, and the minimal use of p -refinement in this case – this is due in part to the under-resolved nature of the mesh for this flow, and to the relatively large cost ascribed to order enrichment in the definition of the merit function.




36.1: 7th adapted mesh with Mach contours.36.2: 7th adapted mesh p -order distribution; the range is $p = 1 \rightarrow 5$.


Figure 36: DPW Wing 1, $M_\infty = 0.76$, $\alpha = 0.5^\circ$, $Re = 5 \times 10^6$: cut at $y = 220\text{mm}$ of the drag-adapted mesh. Note, the reference span is 1524mm .



37.1: 7th adapted mesh with Mach contours. 37.2: 7th adapted mesh p -order distribution; the range is $p = 1 \rightarrow 5$.

Figure 37: DPW Wing 1, $M_\infty = 0.76$, $\alpha = 0.5^\circ$, $Re = 5 \times 10^6$: cut at $y = 620\text{mm}$ of the drag-adapted meshes.

An optimization-based mesh adaptation algorithm may offer insight on “best-practice” gridding guidelines. We notice that several regions of the flow are frequently targeted for refinement. One of these regions is near the leading edge where the flow accelerates through the sonic condition. This change in character of the flow causes strong variations in the adjoint solution, which are responsible for large error indicators. Another region is the edge of the boundary layer, where the turbulent working variable, $\tilde{\nu}$, transitions to zero rapidly. The other two regions are the location of shock-boundary-layer interaction and the trailing edge. These regions exhibit strong gradients in $\tilde{\nu}$ that contribute to the drag output. Figure 38 shows the interaction between the shock and the boundary layer. Note the concentration of cells in the boundary layer and the sharp variation of $\tilde{\nu}$. Further downstream, in the trailing edge region (Figure 39), the beginning of the turbulent wake is also adapted.



38.1: Mach contours for initial mesh.


38.2: Mach contours for the 7th adapted mesh.38.3: \tilde{v} contours for initial mesh.38.4: \tilde{v} contours for the 7th adapted mesh.

Figure 38: DPW Wing 1, $M_\infty = 0.76$, $\alpha = 0.5^\circ$, $Re = 5 \times 10^6$: interaction between shock and boundary-layer at $y = 620\text{mm}$.


39.1: \tilde{v} contours for initial mesh.39.2: \tilde{v} contours for the 7th adapted mesh.

Figure 39: DPW Wing 1, $M_\infty = 0.76$, $\alpha = 0.5^\circ$, $Re = 5 \times 10^6$: trailing edge at $y = 620\text{mm}$.

5 Mesh Optimization

Unstructured meshes offer flexibility in mesh generation and in adaptation, since resolution can be placed only where necessary. *Resolution* refers to the size and shape of an element, as both of these affect the approximation power of the mesh. This information can be encoded in a metric field [17, 91] over the computational domain. The goal in mesh optimization is to relate this metric field to a practical objective, such as minimizing error or computational cost.

Previous works have derived metric fields from consideration of such objectives through heuristic [22, 5, 20, 51], semi-heuristic [119, 89, 31, 125], and more recently, rigorous [126] ways. In this section we present a rigorous output-based approach [126] for determining the metric that gives the best possible mesh. Some modifications are made from the original presentation to simplify the implementation and reduce the computational burden of this method.

5.1 Metric-Based Mesh Optimization Algorithm

Our goal is to optimize the computational mesh, T_h , in order to minimize the output error at a prescribed computational cost. We follow the approach introduced by Yano [126], which iteratively determines the optimal change in the mesh metric field given a prescribed metric-cost relationship and a sampling-inferred metric-error relationship. In this section we briefly review the key elements of this method, and in Section 5.2 we present a modification to the error sampling procedure.

5.1.1 Metric-Based Meshing

A Riemannian metric field, $\mathcal{M}(\vec{x})$, is a field of symmetric positive definite (SPD) tensors that can be used to encode information about the desired size and stretching of a computational mesh. At each point in physical space, \vec{x} , the metric tensor $\mathcal{M}(\vec{x})$ provides a “yardstick” for measuring the distance from \vec{x} to another point infinitesimally far away, $\vec{x} + \delta\vec{x}$. This distance is

$$\delta\ell = \sqrt{\delta\vec{x}^T \mathcal{M} \delta\vec{x}}. \quad (44)$$

After choosing a Cartesian coordinate system and basis for physical space, \mathcal{M} can be represented as a $d \times d$ SPD matrix. The set of points at unit metric distance from \vec{x} is an ellipse, as illustrated in Figure 40: eigenvectors of \mathcal{M} give directions along the principal axes, while the length of each axis (stretching) is the inverse square root of the corresponding eigenvalue.

A mesh that *conforms* to a metric field is one in which each edge has the same length, to some tolerance, when measured with the metric according to (44), which can be integrated to obtain the distance between points that are not infinitesimally close together. An example of a two-dimensional metric-conforming mesher is the Bi-dimensional Anisotropic Mesh Generator (BAMG) [17], and this is used to obtain the present results.

BAMG generates a mesh given a metric field, which is specified at nodes of a background mesh – the current mesh in an adaptive setting. The optimization will determine *changes* to the current, mesh-implied, metric, $\mathcal{M}_0(\vec{x})$, which is obtained on each simplex





Figure 40: A metric, $\mathcal{M} \in \mathbb{R}^{d \times d}$ is a symmetric positive definite tensor field that provides a “yardstick” for measuring distances in different directions. In two dimensions, the set of points equidistant to a point O under the metric measure defines an ellipse.

element of the background mesh by solving a linear system for the $d(d+1)/2$ independent entries of $\mathcal{M}_0(\vec{x})$; the equations in this system enforce that each of the $d(d+1)/2$ edges has unit metric measure. The element-based mesh-implied metrics are then averaged to the nodes using an affine-invariant⁵ algorithm [91]. Given the desired metric at each node of the background mesh, BAMG generates a new, metric-conforming mesh, as illustrated in Figure 41.



41.1: Input: background mesh and desired metric

41.2: Output: Metric-conforming mesh

Figure 41: Input and output to/from the Bi-dimensional Anisotropic Mesh Generator [17]. The desired metric on the background mesh is visualized via ellipses at each node.

Affine-invariant changes to the metric field are made via a symmetric *step matrix*, $\mathcal{S} \in \mathbb{R}^{d \times d}$, according to

$$\mathcal{M} = \mathcal{M}_0^{\frac{1}{2}} \exp(\mathcal{S}) \mathcal{M}_0^{\frac{1}{2}}. \quad (45)$$

Note that $\mathcal{S} = 0$ leaves the metric unchanged, while diagonal values in \mathcal{S} of $\pm 2 \log 2$ halve/double the metric stretching sizes.

5.1.2 Error Convergence Model

The mesh optimization algorithm requires a model for how the error changes as the metric changes. We consider one element, Ω_e , with a current error \mathcal{E}_{e0} , the absolute value of the element’s contribution to (36), and a proposed metric step matrix of \mathcal{S}_e . What will the new error, \mathcal{E}_e , be over Ω_e following refinement with this step matrix? A typical a priori

⁵For example, just averaging matrix entries would be coordinate-system dependent and hence not affine invariant.

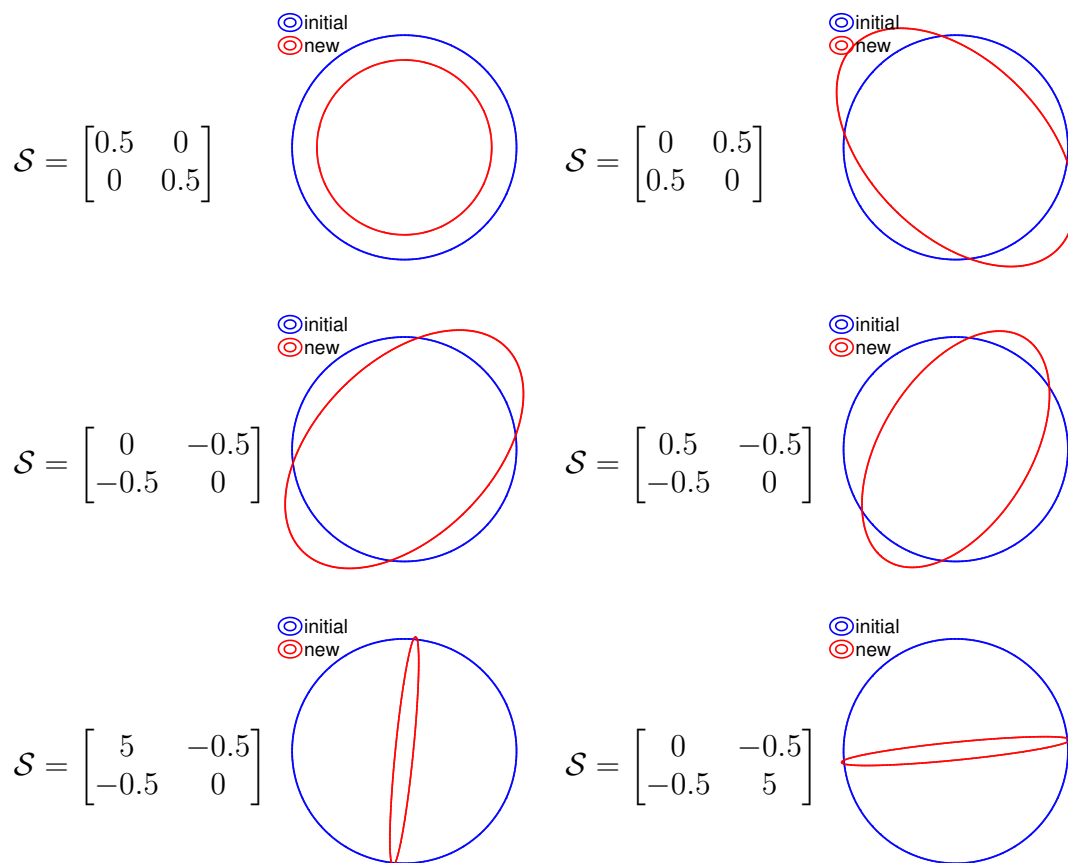


Figure 42: Examples of the effect of metric modification via (45) using $\mathbf{M}_0 = \mathcal{I}$ and various step matrices, \mathcal{S} . The ellipses show the loci of points of unit metric measure away from the center using $\mathbf{M}_0 = \mathcal{I}$ (blue) and $\mathcal{M} = \mathcal{M}_0^{\frac{1}{2}} \exp(\mathcal{S}) \mathcal{M}_0^{\frac{1}{2}}$ (red).

model may predict

$$\mathcal{E}_e = \mathcal{E}_{e0} \left(\frac{h}{h_0} \right)^r = \mathcal{E}_{e0} \exp [r \log(h/h_0)], \quad (46)$$

where h/h_0 are some consistent measures of the new/original element size, and r is the error convergence rate. A generalization to anisotropic metric tensors, for which the error may converge differently depending on the direction of change, is that S_e plays the role of $\log(h/h_0)$ and a symmetric *rate tensor* \mathcal{R}_e replaces the scalar r :

$$\mathcal{E}_e = \mathcal{E}_{e0} \exp [\text{tr}(\mathcal{R}_e \mathcal{S}_e)] \quad \Rightarrow \quad \frac{\partial \mathcal{E}_e}{\partial \mathcal{S}_e} = \mathcal{E}_e \mathcal{R}_e. \quad (47)$$

Note, we include a linearization with respect to \mathcal{S}_e for use in Section 5.1.4. The total error over the mesh is the sum of the elemental errors,

$$\mathcal{E} = \sum_{e=1}^{N_e} \mathcal{E}_e \quad (48)$$

During optimization we will want to keep \mathcal{E} small, and we will be able to change the step matrices at the mesh vertices, \mathcal{S}_v . The rate tensor, \mathcal{R}_e , will be determined separately for each element through the sampling procedure described in Section 5.2.

5.1.3 Cost Model

To measure the cost of refinement, we use degrees of freedom, \mathbf{dof} , which on each element just depends on the approximation order p , assumed constant over the elements. By (45) and properties of the metric tensor, when the step matrix S_e is applied to the metric of element e , the area of the element decreases by $\exp [\frac{1}{2} \text{tr}(S_e)]$. Equivalently, the number of new elements, and hence degrees of freedom, occupying the original area Ω_e *increases* by this factor. So the elemental cost model is

$$C_e = C_{e0} \underbrace{\exp \left[\frac{1}{2} \text{tr}(\mathcal{S}_e) \right]}_{\text{Area}_0/\text{Area}} \quad (49)$$

where $C_{e0} = \mathbf{dof}_{e0}$ is the current number of degrees of freedom on element e , and \mathcal{I} is the identity tensor. Again we include a linearization with respect to \mathcal{S}_e for use in Section 5.1.4. The total cost over the mesh is the sum of the elemental costs,

$$\mathcal{C} = \sum_{e=1}^{N_e} C_e. \quad (50)$$

5.1.4 Metric Optimization Algorithm

Given a current mesh with its mesh-implied metric ($\mathcal{M}_0(\vec{x})$), elemental error indicators \mathcal{E}_{e0} , and elemental rate tensor estimates, \mathcal{R}_e , the goal of the metric optimization algorithm is to determine the step matrix field, $\mathcal{S}(\vec{x})$, that minimizes the error at a fixed cost.

The step matrix field is approximated by values at the mesh vertices, \mathcal{S}_v , which are arithmetically-averaged to adjacent elements⁶:

$$\mathcal{S}_e = \frac{1}{|V_e|} \sum_{v \in V_e} \mathcal{S}_v, \quad (51)$$

where V_e is the set of vertices ($|V_e|$ is the number of them) adjacent to element e (see Figure 43). The optimization problem is to determine \mathcal{S}_v such that the total error \mathcal{E} is minimized at a prescribed total cost \mathcal{C} .




Figure 43: Definition of vertices surrounding an element and vice versa, used for expanding linearizations of the error and cost functions with respect to the step matrix at the nodes.

First-order optimality conditions require derivatives of the error and cost with respect to \mathcal{S}_v . From the equations in Sections 5.1.2 and 5.1.3, and using (51), these are

$$\frac{\partial \mathcal{E}}{\partial \mathcal{S}_v} = \sum_{e \in E_v} \underbrace{\frac{\partial \mathcal{E}_e}{\partial \mathcal{S}_e}}_{(47)} \underbrace{\frac{\partial \mathcal{S}_e}{\partial \mathcal{S}_v}}_{1/|V_e|}, \quad \frac{\partial \mathcal{C}}{\partial \mathcal{S}_v} = \sum_{e \in E_v} \underbrace{\frac{\partial \mathcal{C}_e}{\partial \mathcal{S}_e}}_{(49)} \underbrace{\frac{\partial \mathcal{S}_e}{\partial \mathcal{S}_v}}_{1/|V_e|}, \quad (52)$$

where E_v is the set of elements adjacent to vertex v (see Figure 43). We note that the cost only depends on the *trace* of the step matrix; i.e. the trace-free part of \mathcal{S}_e stretches an element but does not alter its area. We therefore separate the vertex step matrices into trace ($s_v \mathcal{I}$) and trace-free ($\tilde{\mathcal{S}}_v$) parts,

$$\mathcal{S}_v = s_v \mathcal{I} + \tilde{\mathcal{S}}_v. \quad (53)$$

Derivatives of the error with respect to s_v and $\tilde{\mathcal{S}}_v$ are

$$\frac{\partial \mathcal{E}}{\partial s_v} = \text{tr} \left(\frac{\partial \mathcal{E}}{\partial \mathcal{S}_v} \right), \quad \frac{\partial \mathcal{E}}{\partial \tilde{\mathcal{S}}_v} = \frac{\partial \mathcal{E}}{\partial \mathcal{S}_v} - \frac{\partial \mathcal{E}}{\partial s_v} \mathcal{I} \quad (54)$$

The optimization algorithm is then the same as presented by Yano [126]:

1. Given a mesh, solution, and adjoint, calculate $\mathcal{E}_e, \mathcal{C}_e, \mathcal{R}_e$ for each element e .
2. Set $\delta s = \delta s_{\max} / n_{\text{step}}, \mathcal{S}_v = 0$.
3. Begin loop: $i = 1 \dots n_{\text{step}}$

⁶There is no need for an affine-invariant average because entries of \mathcal{S} are coordinate system independent.

- (a) Calculate \mathcal{S}_e from (51), $\frac{\partial \mathcal{E}_e}{\partial \mathcal{S}_e}$ from (47), and $\frac{\partial \mathcal{C}_e}{\partial \mathcal{S}_e}$ from (49).
- (b) Calculate derivatives of \mathcal{E} and \mathcal{C} with respect to s_v and $\tilde{\mathcal{S}}_v$ using (54).
- (c) At each vertex form the ratio $\lambda_v = \frac{\partial \mathcal{E}/\partial s_v}{\partial \mathcal{C}/\partial s_v}$ and
 - Refine the metric for 30% of the vertices with the largest $|\lambda_v|$: $\mathcal{S}_v = \mathcal{S}_v + \delta s \mathcal{I}$
 - Coarsen the metric for 30% of the vertices with the smallest $|\lambda_v|$: $\mathcal{S}_v = \mathcal{S}_v - \delta s \mathcal{I}$
- (d) Update the trace-free part of S_v to enforce stationarity with respect to shape changes at fixed area: $S_v = S_v + \delta s (\partial \mathcal{E} / \partial \tilde{\mathcal{S}}_v) / (\partial \mathcal{E} / \partial s_v)$.
- (e) Rescale $S_v \rightarrow S_v + \beta \mathcal{I}$, where β is a global constant calculated from (49) to constrain the total cost to the desired dof value: $\beta = \frac{2}{d} \log \frac{\mathcal{C}_{\text{target}}}{\mathcal{C}}$, where $\mathcal{C}_{\text{target}}$ is the target cost.

Note, λ_v is a Lagrange multiplier in the optimization. It is the ratio of the marginal error to marginal cost of a step matrix trace increase (i.e. mesh refinement). The above algorithm iteratively equidistributes λ_v globally so that, at optimum, all elements have the same marginal error to cost ratio. Constant values that work generally well in the above algorithm are $n_{\text{step}} = 20$ and $\delta s_{\text{max}} = 2 \log 2$.

In practice, the mesh optimization and flow/adjoint solution are performed several times at a given target cost, $\mathcal{C}_{\text{target}}$, until the error stops changing. Then the target cost is increased to reduce the error further if desired. Figure 44 illustrates this process. For reporting the final error and cost at each adaptive iteration, the values are averaged over the last few solution iterations at each target cost.




Figure 44: Illustration of the typical error (left axis) and cost/dof (right axis) history during mesh optimization. At each solution iteration, the primal and adjoint problems are solved and the mesh optimization algorithm is applied to determine the metric for the mesh for the next solution iteration. The target cost is held fixed for multiple (e.g. 10) solution iterations, during which the error eventually stabilizes as each mesh “hovers” about the optimum. Following an allowable cost increase, the error drops over a few solution/optimization iterations and settles to a lower value. Note that the primal and adjoint solutions can be transferred from one mesh to the next, making the solves at each iteration much cheaper compared to solving from scratch (e.g. using the freestream state).

5.2 Element-Local Error Sampling

The mesh optimization algorithm requires that the error convergence rate tensor, \mathcal{R}_e be known for each element e . We estimate this rate tensor a posteriori by *sampling* a small

number of element refinements (cuts) for each element and performing a regression. This is also the approach taken by Yano [126], and our work differs in that we do not solve any primal or adjoint problems on the refined elements. Indeed, we never actually modify the mesh when considering the refinement samples, and this simplifies data structure handling in the implementation, especially in parallel.

For a triangular element, we consider four refinement options, as shown in Figure 45. We would like to know how much the error would decrease under each refinement option.




Figure 45: Four refinement options for a triangle. Each one is considered implicitly during error sampling, though the elements are never actually refined.

One expensive option is to refine the element with the proposed cut, re-solve the primal and fine-space adjoint problems globally, and re-compute the error estimate. Though accurate, this would be impractically expensive. Another option is to only solve the primal/adjoint problems on a subset of the original mesh: the current element and its neighbors. This approach, taken by Yano, is less accurate but still performs very well as globally-exact primal/adjoint states are not necessary to estimate the error rate tensor. In this work we further simplify the estimation by not solving additional problems, even on a local patch of elements. Instead, our proposed algorithm uses element-local *projections* of the fine-space adjoint to semi-refined spaces associated with each refinement option.

Consider one element e . The fine space adjoint $\boldsymbol{\psi}_h^{p+1}|_{\Omega_e}$ gives, via (36), an estimate of the output error in the current order p solution, as measured relative to a “truth” order $p + 1$ solution: this is \mathcal{E}_{e0} . If the fine-space adjoint were of order $p + 2$, we would have an estimate of the error relative to an even finer space. Now, suppose that we are looking at refinement option i in Figure 45. Refining an element with such an option creates a solution space that is finer than the original, though (we assume) not as fine as increasing the order to $p + 1$ – we still use $p + 1$ as the “truth” space. Suppose we have an order p adjoint on this refined space; call it $\boldsymbol{\psi}_{hi}^p|_{\Omega_e}$, where the i indicates that we are considering refinement option i . With this adjoint we can compute an error indicator $\Delta\mathcal{E}_{ei}$: this estimates the error between the coarse solution and that on refinement option i . The remaining error associated with refinement option i is then given by the difference,

$$\mathcal{E}_{ei} \equiv \mathcal{E}_{e0} - \Delta\mathcal{E}_{ei}. \quad (55)$$

Calculating $\Delta\mathcal{E}_{ei}$ requires an adjoint-weighted residual evaluation on the element refined under option i . We bypass this complication with an additional simplification: we project $\boldsymbol{\psi}_{hi}^p|_{\Omega_e}$ into the $p + 1$ space on the original element and evaluate the adjoint weighted residual there. That is, we perform

$$\Delta\mathcal{E}_{ei} \equiv |\mathcal{R}_h^{p+1}(\mathbf{u}_h^p, \tilde{\boldsymbol{\psi}}_{hi}^p|_{\Omega_e})|, \quad (56)$$

where $\tilde{\psi}_{hi}^p$ is ψ_{hi}^p projected from order p on refinement option i into order $p + 1$ on the original element. The final simplification is to not actually calculate $\psi_{hi}^p|_{\Omega_e}$ for refinement option i . Instead, we simply project the known fine-space order $p + 1$ adjoint onto order p in refinement option i .

In summary, the error uncovered by refinement option i , $\Delta\mathcal{E}_{ei}$, is estimated by the adjoint-weighted residual in (56), with all calculations occurring at order $p + 1$ on the original element. Figure 46 breaks down the key computation of $\tilde{\psi}_{hi}^p$ into individual projections. Using least-squares projections in reference space, the combination of the




Figure 46: Projection of the fine-space adjoint, ψ_h^{p+1} , down to the space created by refinement option $i = 1$, and then back up to order $p+1$ on the original element. This sequence of projections is encapsulated in a single transfer matrix in (57).

steps depicted in Figure 46 can be encapsulated into one transfer matrix that converts ψ_h^{p+1} into $\tilde{\psi}_{hi}^p$, both represented in the order $p + 1$ space on the original element:

$$\tilde{\psi}_{hi}^p = T_i \psi_h^{p+1}, \quad (57)$$

$$T_i = [M_0(\phi_0^{p+1}, \phi_0^{p+1})]^{-1} \sum_{k=1}^{n_i} T_{ik}, \quad (58)$$

$$T_{ik} = M_k(\phi_0^{p+1}, \phi_k^p) [M_k(\phi_k^p, \phi_k^p)]^{-1} M_k(\phi_k^p, \phi_k^{p+1}) [M_k(\phi_k^{p+1}, \phi_k^{p+1})]^{-1} M_k(\phi_k^{p+1}, \phi_0^{p+1}) \quad (59)$$

In these equations, n_i is the number of sub-elements in refinement option i , k is an index over these sub-elements, ϕ_k^p, ϕ_k^{p+1} are order p and $p + 1$ basis functions on sub-element k , ϕ_0^p, ϕ_0^{p+1} are order p and $p + 1$ basis functions on the original element, and components of the mass-like matrices are defined as

$$M_k(\phi_l, \phi_m) = \int_{\Omega_k} \phi_l \phi_m d\Omega, \quad M_0(\phi_l, \phi_m) = \int_{\Omega_0} \phi_l \phi_m d\Omega, \quad (60)$$

where Ω_k is sub-element k and Ω_0 is the original element. Note that the transfer matrix T_i can be calculated for each refinement option i once in reference space and then used for all elements, so that the calculation of $\Delta\mathcal{E}_{ei}$ consumes minimal additional cost – and most importantly, no solves or residual evaluations are needed on the refined element, as these generally require cumbersome data management and transfer.

Finally, after calculating \mathcal{E}_{ei} , the errors remaining after each refinement option i according to (55), we use least-squares regression to estimate the rate tensor \mathcal{R}_e . Note that for triangles, we have 4 refinement options and 3 independent entries in the symmetric \mathcal{R}_e tensor. Using (47), we formulate the regression to minimize the following error, summed

over refinement options,

$$\sum_i \left[\log \frac{\mathcal{E}_{ei}}{\mathcal{E}_{e0}} - \text{tr}(\mathcal{R}_e \mathcal{S}_{ei}) \right]^2. \quad (61)$$

In this equation, \mathcal{S}_{ei} is the step matrix associated with refinement option i , given by (from (45)),


$$\mathcal{S}_{ei} = \log \left(\mathcal{M}_0^{-\frac{1}{2}} \mathcal{M}_i \mathcal{M}_0^{-\frac{1}{2}} \right), \quad (62)$$

where \mathcal{M}_i is the affine-invariant metric average of the mesh-implied metrics of all sub-elements in refinement option i . Differentiating (61) with respect to the independent components of \mathcal{R}_e yields a linear system for these components.


5.3 Examples

5.3.1 Inviscid Flow over a NACA 0012 Airfoil

We first consider inviscid flow over a NACA 0012 airfoil at $M_\infty = 0.5$ and $\alpha = 2^\circ$. The output of interest is the drag coefficient. Figure 47 shows the Mach number contours and the initial coarse mesh for this case. Curved elements of geometry order $q = 4$ are used adjacent to the airfoil to represent the geometry, and the farfield boundary is over 2000 chord-lengths away.



47.1: Mach number contours



47.2: Initial mesh (356 triangles)

Figure 47: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, inviscid: Mach number contours and the initial coarse mesh for metric-based mesh optimization.

Starting from the coarse mesh, adaptive runs were performed using approximation orders $p = 1, 2, 3$, and at four `dof` targets: 2000, 4000, 8000, and 16000. Note that for a given `dof` target, we expect coarser meshes (fewer elements) at higher p , when each element consumes more `dof`. At each `dof` target, ten solution iterations were performed before moving to the next `dof` target. Figure 48 shows a sample run using $p = 2$. Note how quickly the error drops in the first ~ 3 iterations after increasing the `dof` target and after starting from the coarse mesh. Since the 356-element initial mesh has 2136 `dof` at $p = 2$, the `dof` plot shows little change in the first ten iterations, when the target is 2000




Figure 48: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, inviscid: output error and degrees of freedom histories for mesh optimization with $p = 2$ approximation and 15 solution iterations at each target dof value.

dof, even though the error drops as the ~ 350 elements are resized and repositioned to an optimal configuration for drag prediction.

Figure 49 compares the convergence of the error in the drag coefficient with a measure of the mesh size, $\text{dof}^{-1/2}$, when using the presented mesh optimization algorithm and when using uniform refinement of the initial mesh. The uniform refinement results exhibit stagnating convergence, even at high order, due to the presence of singular solution features (e.g. at the trailing edge) which limit the attainable rate. Indeed there is little benefit in using $p = 3$ compared to $p = 2$ on the uniform refinement sequence. On the other hand, the adaptive results with the presented optimization algorithm show greatly-improved convergence. Orders of magnitude error reductions are achieved compared to uniform refinement, and there is a clear benefit to higher-order approximation. Note that Figure 49 presents data points for every solution iteration in the mesh optimization, with larger symbols denoting the output and dof averaged over the five iterations at each dof target. As shown, the 2000 initial dof target required refinement of the initial mesh for $p = 1$ and coarsening of the initial mesh at $p = 3$; at $p = 2$ the total number of elements remained virtually unchanged at this target.

Figures 505152 show the adapted meshes for $p = 1, 2, 3$ approximation, at the four target dof values. The differences in the number of elements among the various meshes are evident: the $p = 1$ adapted meshes are the finest, whereas the $p = 3$ ones are the coarsest. The adaptation adds mesh resolution to the most interesting areas in the flowfield: the leading and trailing edges of the airfoil. Especially at the trailing edge, small elements persist even at $p = 3$, in order to resolve the strong singularity.







Figure 49: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, inviscid: comparison of output error convergence with dof for uniform refinement and mesh optimization at orders $p = 1, 2, 3$. In the optimization results, the smaller symbols denote the output at each solution iteration, whereas the larger symbols show the average over the last five iterations at each target dof.




50.1: dof = 2000



50.2: dof = 4000



50.3: dof = 8000



50.4: dof = 16000

Figure 50: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, inviscid: optimized meshes (last in sequence for each target dof) for $p = 1$.

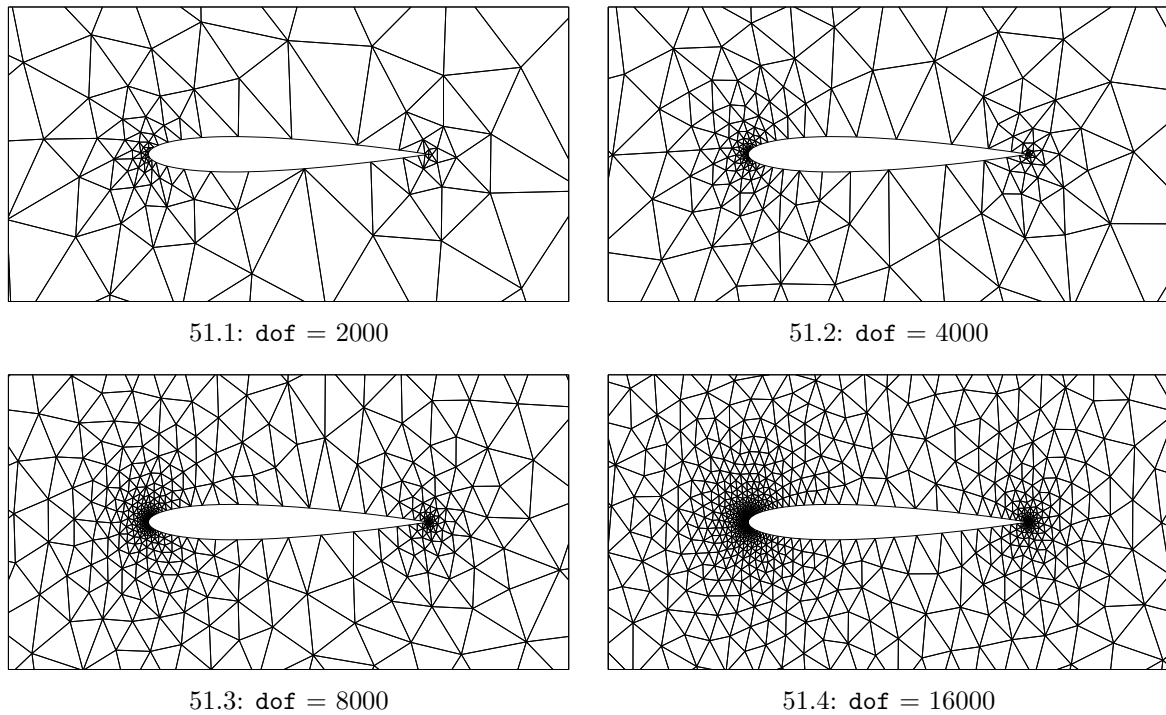


Figure 51: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, inviscid: optimized meshes (last in sequence for each target dof) for $p = 2$.

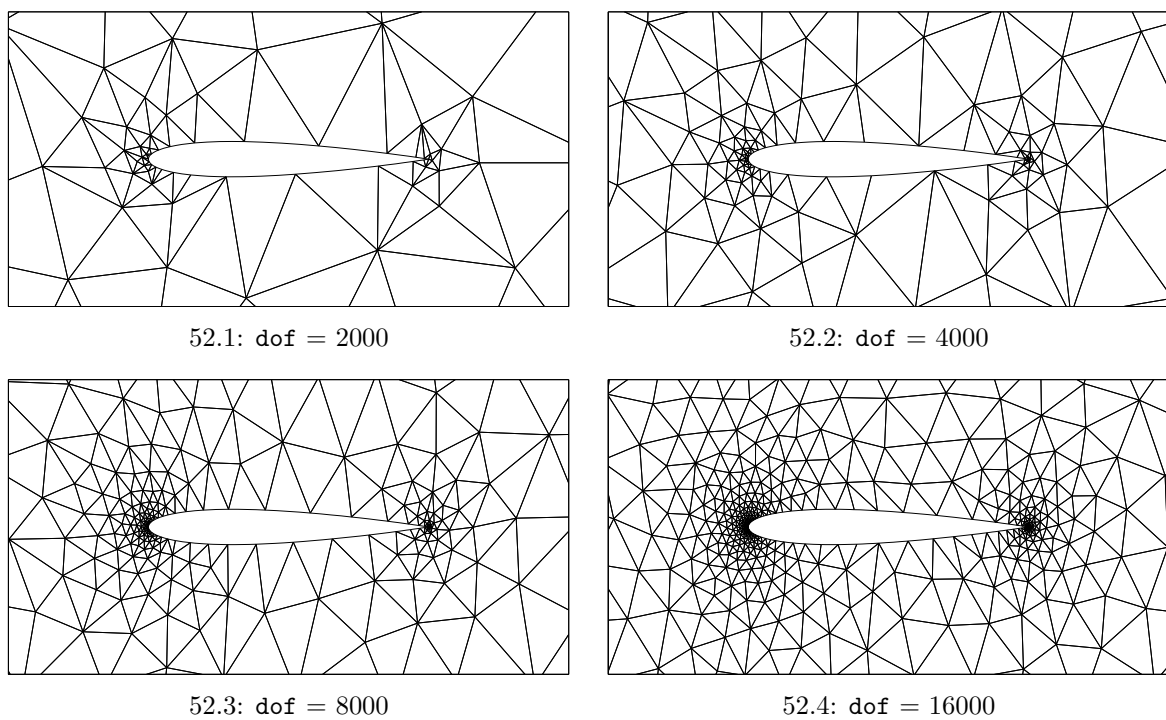




Figure 52: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, inviscid: optimized meshes (last in sequence for each target dof) for $p = 3$.

5.3.2 High-Reynolds Number Flow over a Flat Plate

This case is laminar flow over a flat plate at $M_\infty = 0.5$, $\alpha = 0$, $Re_L = 10^6$. The compressible Navier-Stokes equations are solved over the domain shown in Figure 53, where $L_H = L_V = L = 1$. The Prandtl number is $Pr = 0.72$ and the ratio of specific heats is $\gamma = 1.4$. The output of interest is the drag coefficient on the flat plate.



53.1: Domain and boundary conditions



53.2: Initial mesh (336 triangles)

Figure 53: Flat plate, $M_\infty = 0.5$, $Re_L = 10^6$: problem setup and the initial coarse mesh for metric-based mesh optimization.

Figure 53 also shows the initial mesh used for uniform refinement and mesh optimization runs. This mesh was obtained from subdivision of elements in an initially structured quadrilateral mesh that was generated with extra refinement in the boundary layer and leading-edge of the flat plate. In the uniform refinement runs, each triangular element was subdivided isotropically into four new elements to obtain the next mesh. In addition, “quasi-uniform” refinement was considered in which a hand-generated fine mesh, also with stretched spacing, was coarsened by removal of every other grid line. The resulting mesh sequence was no longer nested but did a better job at resolving the boundary layer and leading-edge regions.




Figure 54: Flat plate, $M_\infty = 0.5$, $Re_L = 10^6$: comparison of output error convergence with dof for uniform refinement, quasi-uniform refinement (coarsening of a graded fine mesh), and mesh optimization at orders $p = 1, 2, 3$. In the optimization results, the smaller symbols denote the output at each solution iteration, whereas the larger symbols show the average over the last five iterations at each target dof.

Figure 54 shows the results of the uniform refinement and adaptive runs for approximation orders $p = 1, 2, 3$. Note that among the uniform refinement runs, increasing the order does decrease the error, but that no higher-order convergence rate is observed due to insufficient resolution of the singular region of the flow (i.e. the leading edge of the flat plate). The quasi-uniform meshes show better results because they are not nested and the finer meshes intentionally have high resolution at the leading edge and in the boundary layer. The meshes optimized using the proposed method show, for a given order, the best results. Note that in this case four cost (**dof**) values were considered: 2000, 4000, 8000, 16000. For each cost value, fifteen iterations of mesh optimization were performed, and errors from the last five were averaged to produce the large symbols shown in Figure 54. The stall in convergence around error values of 10^{-10} is likely due to the residual tolerance used (eight orders of magnitude drop).

Figures 55 and 56 show the optimized meshes for $p = 2$ and $p = 3$, respectively. The meshes are shown over the entire computational domain and zoomed-in with a stretched vertical axis to see refinement in the boundary layer. Comparing $p = 2$ and $p = 3$, we see coarser meshes for $p = 3$ since the number of degrees of freedom is the same for both orders, and $p = 3$ consumes, with full-order approximation, 10 degrees of freedom compared to 6 for $p = 2$.

From the un-zoomed figures, we see that the optimized meshes show no sign of the initial-mesh refinement pattern: not much resolution is required away from the boundary layer. In the zoomed-in figures we see very high resolution in the leading-edge region, and in an arc-shaped region extending from the leading edge to the trailing edge. The boundary layer of course does not disappear to zero height at the trailing edge, but errors made away from the wall this far down the plate have little impact on the drag. The adaptive algorithm picks up on this. In addition, we see reasonably-high anisotropy of the elements in the boundary layer region – note that the ratio of horizontal to vertical axis scaling is 100-to-1 in the zoomed-in figures. This anisotropy is detected automatically by the error sampling procedure that yields the rate tensor for each element, as described in Section 5.2.

5.3.3 NACA 0012 Airfoil in Laminar Viscous Flow

In this case we consider viscous flow over a NACA 0012 airfoil at $M_\infty = 0.5$, $\alpha = 2^\circ$, and $Re = 5000$. The output of interest is the drag coefficient. Figure 57 shows the Mach number contours and the initial coarse mesh for this case. Curved elements of geometry order $q = 4$ are used adjacent to the airfoil to represent the geometry, and the farfield boundary is over 2000 chord-lengths away.

Starting from the coarse mesh, adaptive runs were performed using approximation orders $p = 1, 2, 3$ and four **dof** targets: 2000, 4000, 8000, and 16000. Figure 58 shows drag error and **dof** histories for a sample run using $p = 2$. Compared to the inviscid case the error does not drop as quickly after each **dof** target increase, but it does plateau after 10-15 solution iterations. The output and **dof** values reported are again averages over the last 5 solution iterations at each target **dof**.

Figure 59 compares convergence of the error in the drag coefficient with a measure of the mesh size, $\text{dof}^{-1/2}$, when using the presented mesh optimization algorithm, when using uniform refinement of the initial mesh, and when using an alternative adaptation

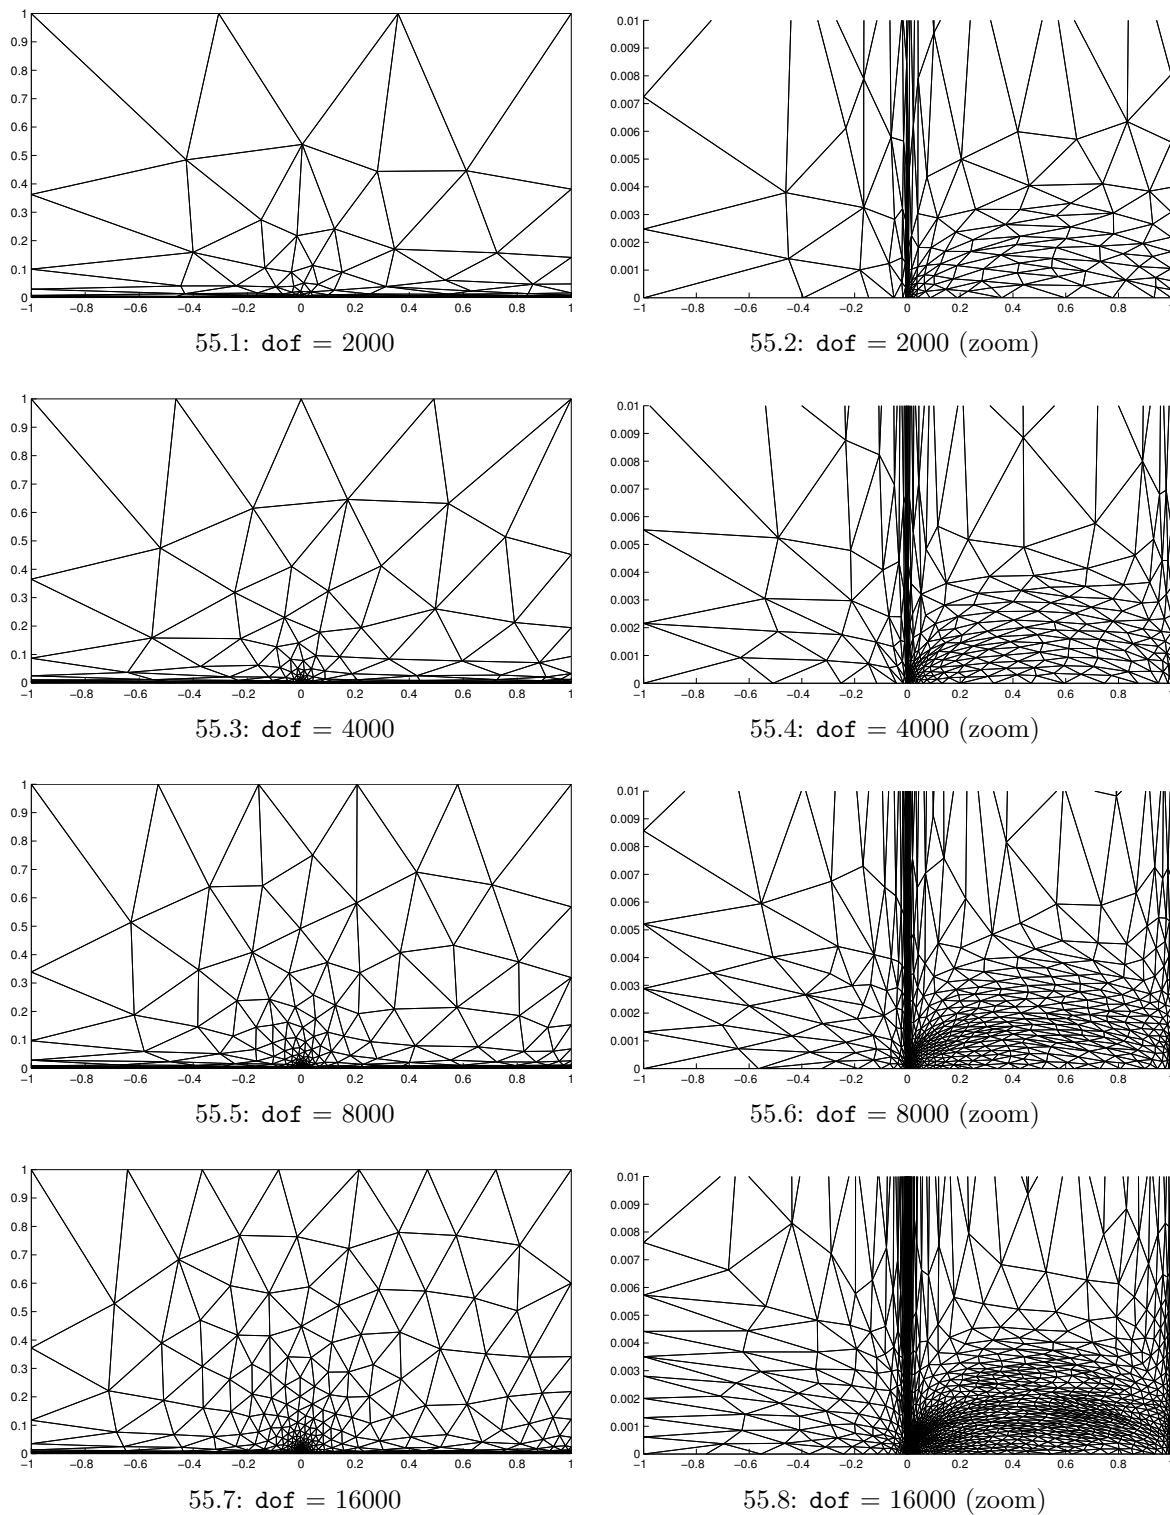


Figure 55: Flat plate, $M_\infty = 0.5$, $Re_L = 10^6$: optimized meshes (last in sequence for each target dof) for $p = 2$. Note the vertical axis stretching in the zoomed plots.

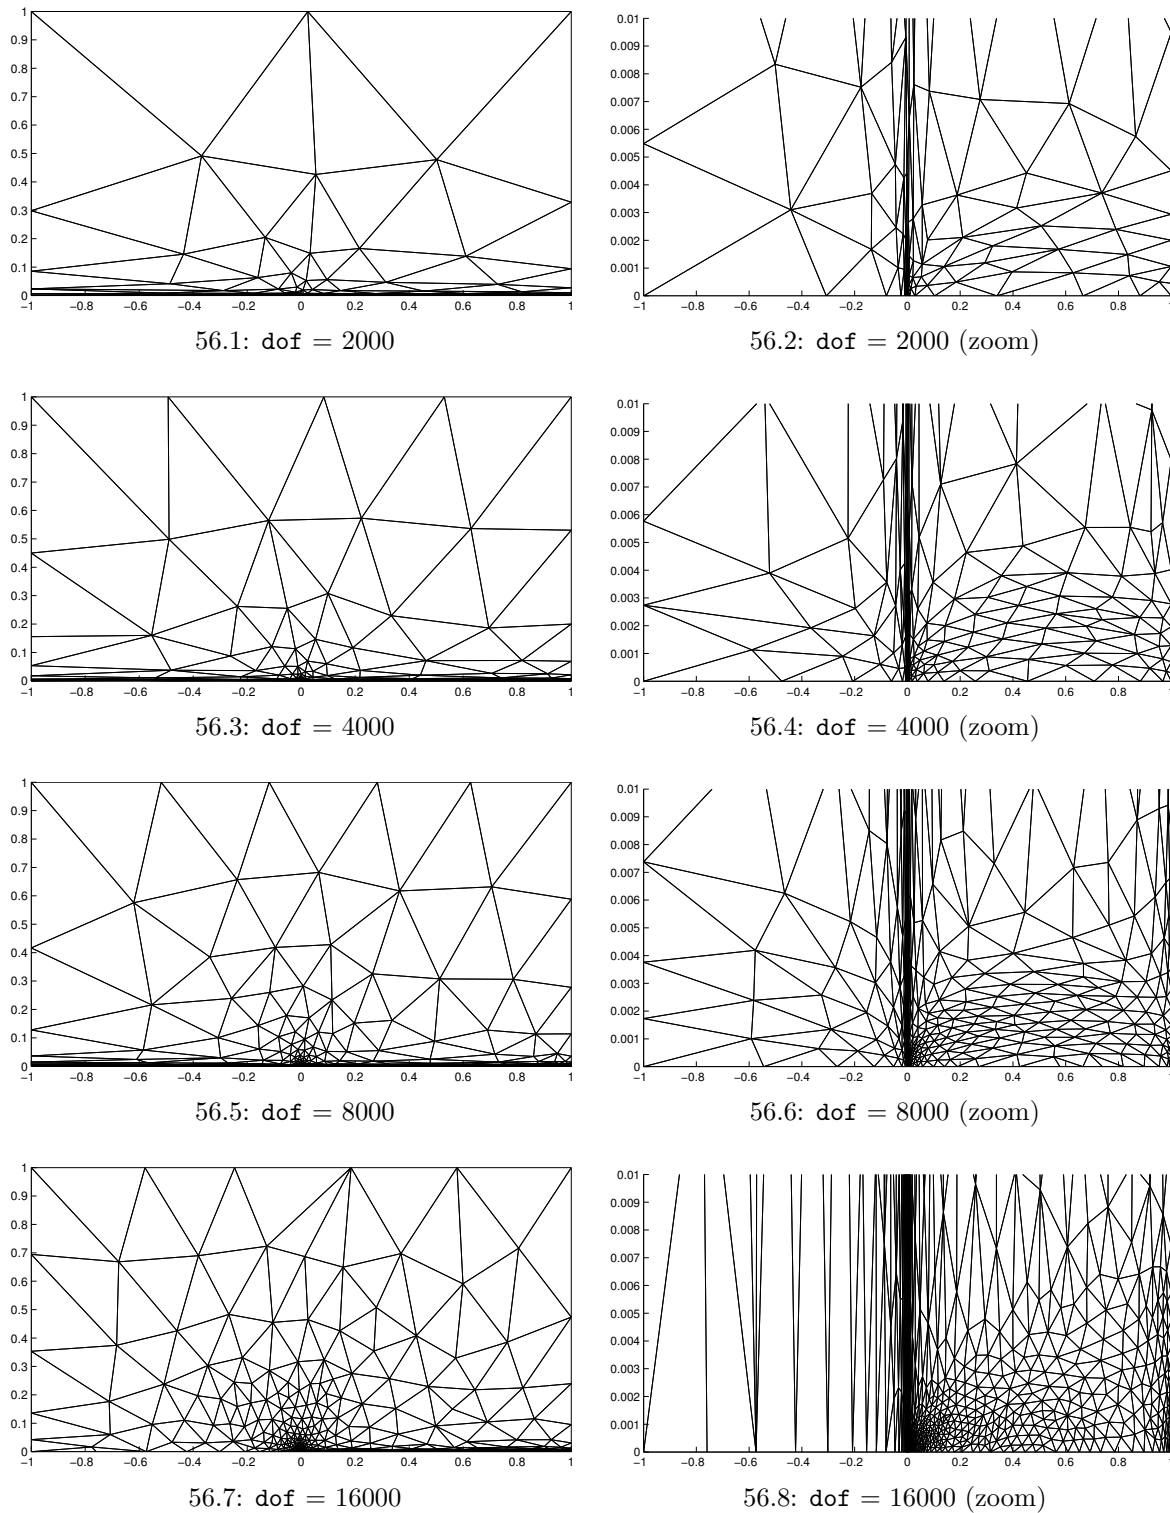


Figure 56: Flat plate, $M_\infty = 0.5$, $Re_L = 10^6$: optimized meshes (last in sequence for each target dof) for $p = 3$. Note the vertical axis stretching in the zoomed plots.




Figure 57: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, $Re = 5000$: Mach number contours and the initial coarse mesh for metric-based mesh optimization.




Figure 58: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, $Re = 5000$: output error and degrees of freedom histories for mesh optimization with $p = 2$ approximation and 15 solution iterations at each target dof value.

algorithm: one in which the output error indicator is used to set the mesh size while the anisotropy comes from the Hessian matrix associated with the Mach number [119, 31]. As in the inviscid case, the uniform refinement results exhibit stagnating convergence, even at high order, due to the presence of singular solution features – there is little benefit to high order in this scenario. On the other hand, the adaptive results show greatly-improved convergence, especially with the presented optimization algorithm. Using the Mach number Hessian in lieu of the metric optimization yields reasonable results for $p = 1$ and $p = 2$, but stalls out for $p = 3$. This should not come as a surprise, since: (1) the Hessian matrix approach was derived for linear ($p = 1$) solution approximation; and (2) the use of the Mach number Hessian is a heuristic that is not guaranteed to provide optimal meshes for the prediction of an output.






Figure 59: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, $Re = 5000$: comparison of output error convergence with dof for uniform refinement, adaptation using the Mach number Hessian to determine anisotropy, and adaptation using the mesh optimization algorithm described in this section. The approximation orders are $p = 1, 2, 3$. In the adaptive results, the smaller symbols denote the output at each solution iteration, whereas the larger symbols show the average over the last five iterations at each target dof .


Figures 60 and 61 show the adapted meshes for $p = 2$ and $p = 3$ approximation, respectively, at the four target dof values. The adaptation algorithm adds resolution in the boundary layer, wake, and in the vicinity of the leading-edge stagnation streamline. The two sets of meshes, one optimized using the error sampling procedure and the other generated using the Mach number Hessian, show resolution of similar regions but with different anisotropy. In the boundary layer and in the edges of the wake, using the Mach number Hessian leads to elements with somewhat larger anisotropy. However, on the leading-edge stagnation streamline, the sampling-based meshes exhibit flow-aligned anisotropy, a result of features in the adjoint solution that the Mach Hessian cannot identify. In general, the Mach number Hessian is a reasonable heuristic for identifying anisotropy, but it is not perfect, as the output convergence results show that the sampling-based optimized meshes are more efficient.




60.1: dof = 2000, optimized




60.2: dof = 2000, Mach Hessian




60.3: dof = 4000, optimized




60.4: dof = 4000, Mach Hessian




60.5: dof = 8000, optimized



60.6: dof = 8000, Mach Hessian




60.7: dof = 16000, optimized




60.8: dof = 16000, Mach Hessian


Figure 60: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, $Re = 5000$: adapted meshes (last in sequence for each target dof) for $p = 2$ using mesh optimization and, for comparison, the Mach number Hessian to set anisotropy.




61.1: dof = 2000, optimized




61.2: dof = 2000, Mach Hessian




61.3: dof = 4000, optimized




61.4: dof = 4000, Mach Hessian




61.5: dof = 8000, optimized



61.6: dof = 8000, Mach Hessian



61.7: dof = 16000, optimized




61.8: dof = 16000, Mach Hessian


Figure 61: NACA 0012, $M_\infty = 0.5$, $\alpha = 2^\circ$, $Re = 5000$: adapted meshes (last in sequence for each target dof) for $p = 3$ using mesh optimization and, for comparison, the Mach number Hessian to set anisotropy.

5.3.4 RAE 2822 Airfoil in Turbulent Transonic Flow

The final case is an RAE 2822 airfoil in transonic turbulent flow: $M_\infty = 0.734$, $\alpha = 2.79^\circ$, $Re = 6.5 \times 10^6$. The Spalart-Allmaras [3] closure is used to model the turbulent flow, and element-constant artificial viscosity [95] is used to capture shocks. The output of interest is the drag coefficient. Figure 62 shows the Mach number contours and the initial coarse mesh for this case. Note that the initial mesh is much coarser in the boundary layer than required for accuracy. Curved elements of geometry order $q = 4$ are used adjacent to the airfoil to represent the geometry, and the farfield boundary is over 500 chord-lengths away. A linear elasticity analogy is used to deform anisotropic linear meshes to curve them to the $q = 4$ boundary geometry approximation.



62.1: Mach number contours



62.2: Initial mesh (758 triangles)

Figure 62: RAE 2822, $M_\infty = 0.734$, $\alpha = 2.79^\circ$, $Re = 6.5 \times 10^6$: Mach number contours and the initial coarse mesh for metric-based mesh optimization.

Starting from the coarse mesh, adaptive runs were performed using approximation orders $p = 1, 2, 3$, and at four dof targets: 5000, 10000, 20000, and 40000. At each dof target, fifteen solution iterations were performed before moving to the next dof target. Figure 63 shows a sample run using $p = 2$. The error drops more slowly in this case compared to the inviscid and laminar cases, and the behavior of the error from mesh to mesh is more oscillatory. This is likely due to the highly nonlinear nature of the RANS equations and the sensitivity of the output to anisotropy in the boundary layer.

Figure 64 shows the convergence of the drag coefficient error with respect to the mesh size, $\text{dof}^{-1/2}$, when using the presented mesh optimization algorithm, and when using uniform refinement. Here we see a dramatic benefit from the optimized meshes. Uniform refinement of the coarse mesh does not appear to consistently reduce the error at any of the orders – the refined meshes are still not fine enough to resolve the boundary layer. On the other hand, the adapted meshes show excellent error reduction: almost four orders of magnitude lower error on the finest meshes compared to the error on the initial mesh. The differences in the adapted results among the three orders are minor, with slight efficiency gains for $p = 2$ and $p = 3$ at the lower error levels.

Finally, Figures 65,66,67 show the adapted meshes for the different orders and dof targets. These meshes are much different from the initial one: we observe high levels of anisotropy in the thin boundary layer adjacent to the airfoil, and in the wake and leading-edge stagnation streamline resolution. We also see some refinement in the foot of




Figure 63: RAE 2822, $M_\infty = 0.734$, $\alpha = 2.79^\circ$, $Re = 6.5 \times 10^6$: output error and degrees of freedom histories for mesh optimization with $p = 2$ approximation and 15 solution iterations at each target dof value.




Figure 64: RAE 2822, $M_\infty = 0.734$, $\alpha = 2.79^\circ$, $Re = 6.5 \times 10^6$: comparison of output error convergence with dof for uniform refinement and adaptation using the mesh optimization algorithm described in this section. The approximation orders are $p = 1, 2, 3$. In the adaptive results, the smaller symbols denote the output at each solution iteration, whereas the larger symbols show the average over the last five iterations at each target dof.

the shock on the upper surface, and resolution of the “ λ ” feature of the adjoint solution above the airfoil. Note that resolution of the entire shock is not needed for accurate drag calculation: the important regions are the foot of the shock and the characteristics that connect to the shock-boundary layer junction.







Figure 65: RAE 2822, $M_\infty = 0.734$, $\alpha = 2.79^\circ$, $Re = 6.5 \times 10^6$: optimized meshes (last in sequence for each target dof) for $p = 1$.




66.1: dof = 5000



66.2: dof = 10000




66.3: dof = 20000




66.4: dof = 40000


Figure 66: RAE 2822, $M_\infty = 0.734$, $\alpha = 2.79^\circ$, $Re = 6.5 \times 10^6$: optimized meshes (last in sequence for each target dof) for $p = 2$.




67.1: dof = 5000



67.2: dof = 10000



67.3: dof = 20000



67.4: dof = 40000

Figure 67: RAE 2822, $M_\infty = 0.734$, $\alpha = 2.79^\circ$, $Re = 6.5 \times 10^6$: optimized meshes (last in sequence for each target dof) for $p = 3$.

6 Unsteady Systems

Results in the previous section showed that adapting a mesh using an indicator obtained from an adjoint-weighted residual can improve robustness (through error estimates) and efficiency (through targeted refinement) of a steady-state CFD simulation. The motivation for such adaptation becomes even stronger in unsteady simulations, where we have control of both spatial and temporal resolution. Identifying the regions in both space and time that require adaptation is more challenging than identifying spatial regions alone. In addition, with an extra dimension, the percentage of the total space-time domain that is important for predicting an output is likely to be reduced compared to a steady problem, which improves the prospects of efficiency gains.

The idea of output-error estimation through an adjoint-weighted fine-space residual extends to unsteady problems without significant theoretical changes. However, unsteady problems do pose challenges for solver efficiency and adaptation mechanics. In this section we address some of these challenges and present adaptive results for reasonably-complex unsteady simulations.

6.1 Primal and Adjoint Discretizations

We discretize time using a semi-discrete approach, i.e. independently of the spatial discretization. Although our adaptive results employ a finite-element temporal discretization, we also present multi-step methods to simplify the presentation of the adjoint extension.

6.1.1 Multi-Step Methods

Primal Form: The discrete residual equation $\mathbf{R}(\mathbf{U}) = \mathbf{0}$, see (14), represents a steady system. A simple method of incorporating unsteady variations is to use a semi-discrete formulation, in which a PDE that is already discretized in space becomes

$$\mathbf{M} \frac{d\mathbf{U}}{dt} + \mathbf{R}(\mathbf{U}) = \mathbf{0}, \quad (63)$$

where $\mathbf{M} \in \mathbb{R}^{N \times N}$ is the mass matrix,

$$\mathbf{M}_{ij} = \mathbf{I}_s \int_{\Omega} \phi_i \phi_j d\Omega. \quad (64)$$

In the above expression, $\mathbf{I}_s \in \mathbb{R}^{s \times s}$ is the state identity matrix, and $1 \leq i, j \leq N$ index the global degrees of freedom. Note that in DG, ϕ_i will have support over only one element, which means that the mass matrix is element-wise block diagonal.

In a multi-step discretization, the time derivative in (63) is discretized using finite-differences. Denote by a superscript n or m the time nodes: e.g. in a uniform temporal subdivision, the time at each time node would be given by $t_m = m\Delta t$, where $\Delta t = T/N_t$ is the time step size for integration with N_t time intervals over a time $0 \leq t \leq T$. For example, a backward Euler method uses the approximation

$$\left. \frac{d\mathbf{U}}{dt} \right|_{t_m} \approx \frac{\mathbf{U}^m - \mathbf{U}^{m-1}}{\Delta t}, \quad (65)$$

6.1.2 Discontinuous Galerkin in Time

Primal Form: Instead of finite differences we could also consider finite elements for the temporal discretization. This is in fact a reasonable approach for output error estimation, which is most rigorous in a variational formulation. In particular, for our adaptive work, we use a discontinuous Galerkin method in time. This consists of time slabs on which the temporal solution variation is approximated with polynomials of order r . The term “time slab” is used only to emphasize that no local-in-space time stepping is performed, and that, at a given time, all elements advance at the same time step, Δt . This Δt is just the width of the current time slab, which can vary once adaptation is performed.

Figure 68 illustrates the concept of time slabs for a two-dimensional simulation. We




Figure 68: Illustration of time slabs and dynamic-order (depicted by color) refinement for a DG-in-time discretization.

enumerate elements in the spatial mesh by $1 \leq e \leq N_e$, and these are assumed fixed (not moving or refined) over the course of the simulation. We enumerate the time slabs by $1 \leq k \leq N_k$, where N_k is the total number of time slabs. Each space-time element is then identified by two indices, (e, k) . p_e^k denotes the order of approximation in element e , time slab k , and this can vary across elements and time slabs.

On each space-time element (e, k) , we approximate the state spatially using a standard DG-in-space approximation, as introduced in Section 2. The result of the spatial approximation is that on each element we have $N_{p_e^k}$ spatial unknowns, \mathbf{U}_{he}^k . In DG-in-time, we approximate the temporal variation of these unknowns using polynomials in time,

$$\mathbf{U}_{he}^k(t) = \sum_{n=1}^{r+1} \mathbf{U}_{he}^{kn} \varphi_h^n(t), \quad (70)$$

where $\varphi_h^n(t)$ are order r temporal basis functions. The subscript h on the above terms indicates the space-time mesh, i.e. the approximation space. When discussing error estimation, we will be dealing with coarse and fine spaces, which we will denote by H and h subscripts, respectively.

For compactness of notation, we lump all spatial degrees of freedom associated with time node n on slab k into one vector,

$$\mathbf{U}_h^{kn} = \{\mathbf{U}_{he}^{kn}\}_{\forall e} \in \mathbb{R}^{N_{hk}}. \quad (71)$$

As a shorthand, we will denote by \mathbf{U}_h^k the sets of unknowns over a whole time slab, k . Finally, the set of states over the entire space-time domain will be referred to simply as \mathbf{U}_h .

A nonlinear system of equations on each time slab is obtained by substituting the approximation from (70) into (63), multiplying by test functions in the same space as the approximation functions, and integrating by parts to incorporate discontinuities at time slab and spatial element interfaces. We group the resulting set of discrete equations by time nodes on each time slab into $(r + 1)$ unsteady residual vectors, enumerated by m ,

$$\mathbf{R}_h^{km} \equiv a^{mn} \mathbf{M}_h^{k,k} \mathbf{U}_h^{kn} - \varphi_h^m(t_{k-1}) \mathbf{M}_h^{k,k-1} \mathbf{U}_h^{k-1,r+1} + \int_{t_{k-1}}^{t_k} \varphi_h^m(t) \mathbf{R}_h(\mathbf{U}_h^k(t)) dt = \mathbf{0}, \quad (72)$$

where $\mathbf{M}_h^{k,l} \in \mathbb{R}^{N_{hk} \times N_{hl}}$ is the spatial mass matrix formed by spatial basis functions on neighboring time slabs k and l , which could be different in cases of dynamic spatial order refinement. $\mathbf{U}_h^{k-1,r+1}$ is the end-of-slab state from the previous time slab, as shown in Figure 69. For the first time slab, $k = 1$, this is $\mathbf{U}_h^{0,r+1} \equiv \mathbf{U}_h^0$, which is the initial condition.




Figure 69: DG in time for $r = 2$: depiction of temporal basis functions on one time slab and definitions of key quantities in (72). The vertical axis represents a “state” – from the spatial discretization, we actually have N_h states, and we can imagine having N_h such plots.

The values a^{mn} in (72) constitute a temporal stiffness matrix and are given by

$$a^{mn} = - \int_{t_{k-1}}^{t_k} \varphi_h^n \frac{d\varphi_h^m}{dt} dt + \varphi_h^n(t_k) \varphi_h^m(t_k). \quad (73)$$

Using a Lagrange temporal basis on equally-spaced nodes, we have,

$$\text{for } r = 1, a^{mn} = \begin{bmatrix} 1/2 & 1/2 \\ -1/2 & 1/2 \end{bmatrix}; \quad \text{for } r = 2, a^{mn} = \begin{bmatrix} 1/2 & 2/3 & -1/6 \\ -2/3 & 0 & 2/3 \\ 1/6 & -2/3 & 1/2 \end{bmatrix}. \quad (74)$$

The time integral in (72) is evaluated with $(r + 1)$ points of Gauss quadrature for order $2r + 1$ accuracy.

Adjoint Form: The discrete adjoint equation, (68), remains valid for a DG-in-time discretization. The number of time nodes is now $N_t = N_k(r + 1)$, where N_k is the number of time slabs, since each time slab has $r + 1$ temporal unknowns. Denoting the adjoint at time slab k , time node m , by Ψ_H^{km} , a more descriptive version of (68) is

$$\underbrace{\left(\frac{\partial \mathbf{R}_h^{km}}{\partial \mathbf{U}_h^{ln}}\right)^T \Psi_h^{km} + \left(\frac{\partial J_h}{\partial \mathbf{U}_h^{ln}}\right)^T}_{\mathbf{R}_{\psi h}^{ln}(\Psi_h^{km})} = 0, \quad (75)$$

where k, l index time slabs and n, m index time nodes. We have defined the adjoint residual, $\mathbf{R}_{\psi h}^{ln}(\Psi_h^{km})$, as the entire left-hand side of the adjoint equation. Linearizing the residual expressions in (72), the $r + 1$ adjoint residual vectors on time slab k are

$$\mathbf{R}_{\psi h}^{ln} = a^{mn} \mathbf{M}_h^{l,l} \Psi_h^{lm} - \varphi_h^n(t_l) \mathbf{M}_h^{l,l+1} \Psi_h^{l+1,1} + \int_{t_{l-1}}^{t_l} \varphi_h^n \frac{\partial \mathbf{R}}{\partial \mathbf{U}} \Big|_{\mathbf{U}_h(t)}^T \Psi_h^l(t) dt + \left(\frac{\partial J_h}{\partial \mathbf{U}_h^{ln}}\right)^T \quad (76)$$

In this equation, $\Psi_h^l(t) = \sum_m \Psi_h^{l,m} \varphi_h^m(t)$, and $\Psi_h^{l+1,1}$ is the adjoint vector associated with the start of the next time slab. When calculating on the last time slab, $\Psi_h^{l+1,1} = 0$.

Both the primal and adjoint equations are solved using a Newton iteration based on an approximate factorization. This solver requires solutions of systems that are the same size as a steady-state solution. Details are given in [37].

6.2 Deformable Domains

6.2.1 An Arbitrary Lagrangian Eulerian Treatment

In an Arbitrary Lagrangian Eulerian (ALE) method, the mesh can move at a velocity different from that of the flow. This is useful for analyzing problems in which the computational domain undergoes deformation (e.g. flutter or flapping flight in aerodynamics), as in these situations the mesh generally needs to “move with” the geometry so that much of the computational domain may be affected.

The key idea of an ALE formulation is to map the original PDE on the deforming physical domain to a modified PDE on a static reference domain, as solving on a static domain is something we already know how to do. This transformation is illustrated graphically in Figure 70, and Table 2 defines key quantities.

The expressions for the transformations of the normals are obtained using $dv = g dV$ for infinitesimal volumes and $d\vec{l} = \mathcal{G} d\vec{L}$ for infinitesimal vectors [94]. The system of conservation laws on the physical domain is, repeating (1),

$$\partial_t \mathbf{u} + \partial_i \mathbf{H}_i(\mathbf{u}, \nabla \mathbf{u}) = \mathbf{0}, \quad \mathbf{H}_i = \mathbf{F}_i(\mathbf{u}) - \mathbf{G}_i(\mathbf{u}, \nabla \mathbf{u}), \quad (77)$$

where both inviscid and viscous fluxes are included. Integrating over a time-varying volume $v(t)$ yields,

$$\int_{v(t)} \partial_t \mathbf{u} dv + \int_{\partial v(t)} \vec{\mathbf{H}} \cdot \vec{n} da = 0, \quad \vec{n} \text{ is outward-pointing on } \partial v(t), \quad (78)$$

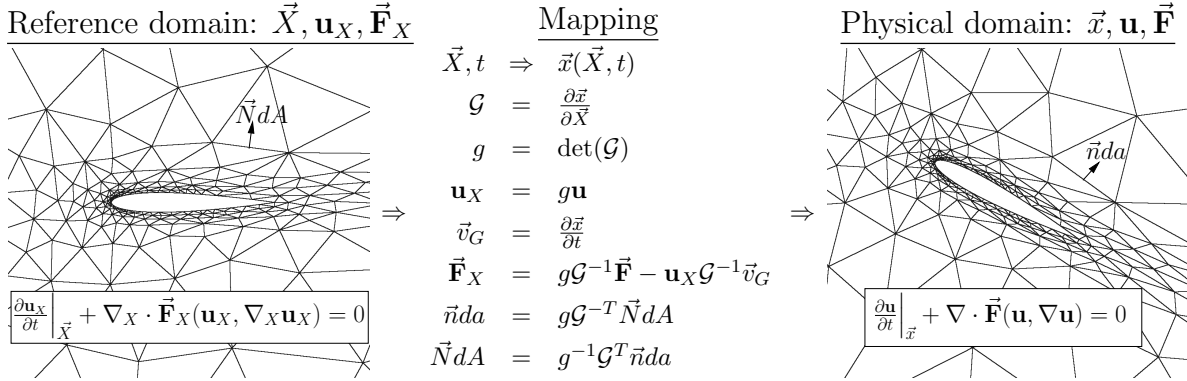


Figure 70: Summary of the mapping between reference and physical domains. The equations are solved on the reference domain, which remains fixed for all time. When denoting reference-domain quantities, we use a subscript X .

Table 2: Definitions of variables used in the ALE mapping. Bold indicates a state vector and an arrow indicates a spatial vector.

\vec{X}	=	reference-domain coordinates	\vec{x}	=	physical-domain coordinates
\mathbf{u}_X	=	state on reference domain	\mathbf{u}	=	physical state
$\vec{\mathbf{F}}_X$	=	flux vector on reference domain	$\vec{\mathbf{F}}$	=	flux vector on physical domain
dA	=	differential area on reference domain	da	=	differential area on physical domain
\vec{N}	=	normal vector on reference domain	\vec{n}	=	normal vector on physical domain
V	=	reference domain (static)	$v(t)$	=	physical domain (dynamic)
\mathcal{G}	=	mapping Jacobian matrix	\vec{v}_G	=	grid velocity, $\partial\vec{x}/\partial t$
g	=	determinant of Jacobian matrix			

where $\vec{\mathbf{H}}$ is a spatial vector with components \mathbf{H}_i . We now transform to the reference domain, V . The boundary integral of the flux is

$$\int_{\partial v(t)} \vec{\mathbf{H}} \cdot \vec{n} da = \int_{\partial V} \vec{\mathbf{H}} \cdot (g\mathcal{G}^{-T}\vec{N}) dA = \int_{\partial V} (g\mathcal{G}^{-1}\vec{\mathbf{H}}) \cdot \vec{N} dA. \quad (79)$$

The first integral in (77) transforms using Leibniz's rule,

$$\begin{aligned} \int_{v(t)} \frac{\partial \mathbf{u}}{\partial t} &= \frac{d}{dt} \int_{v(t)} \mathbf{u} dv - \int_{\partial v(t)} (\mathbf{u}\vec{v}_G) \cdot \vec{n} da \\ &= \frac{d}{dt} \int_V \mathbf{u}g dV - \int_{\partial V} (\mathbf{u}\vec{v}_G) \cdot (g\mathcal{G}^{-T}\vec{N}) dA \\ &= \int_V \frac{\partial(g\mathbf{u})}{\partial t} dV - \int_{\partial V} (g\mathbf{u}\mathcal{G}^{-1}\vec{v}_G) \cdot \vec{N} dA. \end{aligned} \quad (80)$$

Substituting (79) and (80) into (78) and applying the divergence theorem gives the PDE on the reference domain,

$$\left. \frac{\partial \mathbf{u}_X}{\partial t} \right|_{\vec{X}} + \nabla_X \cdot \vec{\mathbf{H}}_X(\mathbf{u}_X, \nabla_X \mathbf{u}_X) = 0, \quad (81)$$

where $\mathbf{u}_X = g\mathbf{u},$
 $\vec{\mathbf{H}}_X = g\mathcal{G}^{-1}\vec{\mathbf{H}} - \mathbf{u}_X\mathcal{G}^{-1}\vec{v}_G.$

∇_X denotes the gradient with respect to the reference coordinates. We break up the transformed flux, $\vec{\mathbf{H}}_X$, into inviscid and viscous fluxes by lumping the grid-velocity term into the inviscid flux,

$$\vec{\mathbf{H}}_X = \vec{\mathbf{F}}_X - \vec{\mathbf{G}}_X, \quad \vec{\mathbf{F}}_X = g\mathcal{G}^{-1}\vec{\mathbf{F}} - \mathbf{u}_X\mathcal{G}^{-1}\vec{v}_G, \quad \vec{\mathbf{G}}_X = g\mathcal{G}^{-1}\vec{\mathbf{G}}. \quad (82)$$

The gradient of the state transforms via the chain and product rules. Using implied summation,

$$\begin{aligned} \frac{\partial \mathbf{u}}{\partial x_j} &= \frac{\partial(g^{-1}\mathbf{u}_X)}{\partial X_d} \frac{\partial X_d}{\partial x_j} = \left(g^{-1} \frac{\partial \mathbf{u}_X}{\partial X_d} - g^{-2} \frac{\partial g}{\partial X_d} \mathbf{u}_X \right) \mathcal{G}_{dj}^{-1} \\ &= g^{-1} \left(\frac{\partial \mathbf{u}_X}{\partial X_d} - g^{-1} \frac{\partial g}{\partial X_d} \mathbf{u}_X \right) \mathcal{G}_{dj}^{-1}, \end{aligned} \quad (83)$$

where d and j index the reference and physical coordinates, respectively. We also have,

$$\mathcal{G} = \mathcal{G}_{jd} = \frac{\partial x_j}{\partial X_d}, \quad \delta_{ji} = \frac{\partial x_j}{\partial x_i} = \frac{\partial x_j}{\partial X_d} \frac{\partial X_d}{\partial x_i} = \mathcal{G}_{jd} \mathcal{G}_{di}^{-1} \Rightarrow \mathcal{G}^{-1} = \mathcal{G}_{di}^{-1} = \frac{\partial X_d}{\partial x_i}.$$

In a DG setting, discretization of the new reference-domain equation requires modifications to the numerical flux function on inter-element faces, to the boundary conditions, to the face normal vectors, and to the quadrature integration weights. These modifications are based on the reference-to-global mapping and its derivatives.

The weighted residual statement on the reference domain is obtained from the PDE by multiplying by test functions (defined in the reference domain) and integrating over

reference-domain elements. The discretization would be straightforward were it not for the fact that fluxes and boundary conditions are specified on the physical domain. A natural approach that minimizes intrusion into the code is to express the reference-space fluxes and boundary conditions in terms of the physical fluxes and boundary conditions.

For example, the inviscid flux on the reference domain includes the standard Galilean transformation expected from changing reference frames and also a multiplication by $g\mathcal{G}^{-1}$, which is done by post-processing the equation-set specific flux,

$$\vec{\mathbf{F}}_X = g\mathcal{G}^{-1}\vec{\mathbf{F}} - \mathbf{u}_X\mathcal{G}^{-1}\vec{v}_G = g\mathcal{G}^{-1}\left(\vec{\mathbf{F}} - \mathbf{u}\vec{v}_G\right). \quad (84)$$

To account for the Galilean transformation on element interfaces, the Riemann solver needs to operate on $\vec{\mathbf{F}} - \mathbf{u}\vec{v}_G$ instead of just $\vec{\mathbf{F}}$.

The reference-domain viscous flux is related to the physical viscous flux through

$$\vec{\mathbf{G}}_X = g\mathcal{G}^{-1}\vec{\mathbf{G}}. \quad (85)$$

Since the physical viscous flux is calculated using a diffusion matrix and the physical state gradient, $\mathbf{G}_i = \mathbf{K}_{ij}\partial_j\mathbf{u}$, then, using (83) for the physical gradient, the reference-domain viscous flux is

$$\begin{aligned} \mathbf{G}_{X,d} &= g\mathcal{G}_{di}^{-1}\mathbf{K}_{ij}\partial_{x_j}\mathbf{u} \\ &= g\mathcal{G}_{di}^{-1}\mathbf{K}_{ij}g^{-1}\left(\partial_{X_c}\mathbf{u}_X - \mathbf{u}_Xg^{-1}\partial_{X_c}g\right)\mathcal{G}_{cj}^{-1} \\ &= \underbrace{\mathcal{G}_{di}^{-1}\mathbf{K}_{ij}\mathcal{G}_{cj}^{-1}}_{\mathbf{K}_{X,dc}}\left(\partial_{X_c}\mathbf{u}_X - \mathbf{u}_Xg^{-1}\partial_{X_c}g\right), \end{aligned} \quad (86)$$

where c, d index the reference domain coordinates. $\mathbf{K}_{X,dc}$ represents the diffusion matrix on the reference domain. It can be re-written in a more symmetrical form as

$$\mathbf{K}_{X,dc} = \mathcal{G}_{di}^{-1}\mathbf{K}_{ij}\mathcal{G}_{cj}^{-1} = \mathcal{G}_{di}^{-1}\mathbf{K}_{ij}\mathcal{G}_{jc}^{-T}. \quad (87)$$

Boundary conditions also require modifications when simulating problems on deformable domains. In particular, the physical boundary flux must be aware of motion on the boundary, \vec{v}_G . For example, on a moving wall, the flow tangency boundary condition states that the normal component of the fluid velocity is equal to the normal component of the boundary motion velocity (which would be zero without mesh motion). This physical consideration is separate from the subtraction of $\mathbf{u}^b\vec{v}_G$ from the flux – both must be included.

Calculation of the viscous contribution on a boundary requires not only the boundary state, \mathbf{u}^b , but also the boundary flux. For pure Dirichlet boundary conditions, the state gradient information is taken from the interior. In other cases, the physical viscous flux is prescribed on the boundary (e.g. zero heat flux for an adiabatic wall), and in these cases, the viscous flux contribution is added directly to the residual. Note that no transformation needs to be applied to the viscous flux dotted with the normal, since

$$\vec{\mathbf{G}} \cdot \vec{n}da = \left(g^{-1}\mathcal{G}_{id}\vec{\mathbf{G}}_X\right) \cdot \left(g\mathcal{G}^{-1}\vec{N}\right) dA = \vec{\mathbf{G}}_X \cdot \vec{N}dA. \quad (88)$$

6.2.2 Blended Analytical Mesh Motions

The ALE method described above requires an analytically defined mapping between reference and physical domains. Therefore, the user must prescribe a motion (e.g. sinusoidal pitch/plunge) in a certain region of the domain. If only a portion of the domain needs to move, the mapping can be smoothly blended into the static domain outside the moving region. As Persson *et al* present in [94], a polynomial blending function is a simple way




Figure 71: Airfoil and wing undergoing analytical motions. The blue regions are those in which the prescribed inner motion is blended into the static outer mesh. The boundaries of these blending regions are circular in 2D and spherical in 3D.

to transition between deforming and static regions. A typical scenario is to have an inner disk in 2D (or sphere in 3D) undergo a prescribed rigid-body motion, and to then blend this motion into the static mesh via the polynomial blending function. Figure 71 shows an example of this blending for an airfoil and a wing.

6.2.3 The Geometric Conservation Law

The ability to preserve a free stream is a desirable property of numerical schemes. However, for schemes employing a finite-dimensional basis (say a set of polynomials in space-time), a constant state $\bar{\mathbf{u}}$ in the physical domain will generally not be a solution to the discrete form of (81) in the reference domain. This means that for an arbitrary motion of the mesh, an initially free-stream state will not be preserved.

This lack of free-stream preservation can be explained by noting that, for general mappings, the Jacobian g will be non-polynomial in both space and time. Hence, the reference state $\mathbf{u}_X = g\bar{\mathbf{u}}$ will likewise be non-polynomial, which means it cannot be represented exactly with the reference-domain bases. This inexact representation will introduce both spatial and temporal errors into an initially free-stream state, which manifest as conservation errors that accumulate in time.

To eliminate these conservation errors, a separate Geometric Conservation Law (GCL) is enforced alongside the governing equations. The idea of the GCL is twofold: (i) to address the representation issues mentioned above, replace the analytical g with a new variable, \bar{g} , which is a polynomial approximation to g in space-time; and (ii) to ensure that this \bar{g} actually allows for free-stream preservation, compute it from the following equation [94]:

$$\frac{\partial \bar{g}}{\partial t} - \nabla_X \cdot (g \mathcal{G}^{-1} \vec{v}_G) = 0. \quad (89)$$

This equation ensures that the change in element area (i.e. the change in \bar{g}) is directly linked to what the grid velocities on element boundaries claim it should be. Hence, there is no disagreement between grid velocities and Jacobians on what the geometry is, and in that sense we have “geometric conservation.”

The strategy then is to use this \bar{g} to define a new reference-domain state $\mathbf{u}_{\bar{X}} = \bar{g} g^{-1} \mathbf{u}_X = \bar{g} \mathbf{u}$, which is used instead of the original state $\mathbf{u}_X = g \mathbf{u}$. If \bar{g} is discretized using the same spatial basis as the state and is marched in time using the same unsteady solver, a free-stream state $\bar{\mathbf{u}}$ will be preserved. In the end, what we have done is replaced the original analytical g with a “best fit” space-time polynomial \bar{g} , which then makes the free-stream state $\mathbf{u}_{\bar{X}} = \bar{g} \bar{\mathbf{u}}$ exactly representable in the discrete space.

Once \bar{g} is obtained on each element, it is used instead of g to convert the stored reference state to the physical state. The final form of the reference-domain equation is then

$$\left. \frac{\partial \mathbf{u}_{\bar{X}}}{\partial t} \right|_X + \nabla_X \cdot \vec{\mathbf{H}}_{\bar{X}}(\mathbf{u}_{\bar{X}}, \nabla_X \mathbf{u}_{\bar{X}}, \bar{g}) = 0, \quad (90)$$

where $\vec{\mathbf{H}}_{\bar{X}}$ is just $\vec{\mathbf{H}}_X$ but with $\mathbf{u}_{\bar{X}}/\bar{g}$ replacing \mathbf{u}_X/g in the calculation of the physical state.

Figure 72 shows the effect of the GCL on a free-stream preservation test. In this case, we solve the Navier-Stokes equations on a rectangular domain with an analytical sinusoidal deformation mapping defined in the domain interior. The temporal and spatial discretization are both discontinuous Galerkin, with order $r = 1$ and p , respectively. Without the GCL, the free-stream solution is not preserved, though the L_2 error converges with both h - and p -refinement of the spatial mesh. With the GCL, the free stream is maintained to residual tolerance, which was approximately ten orders of magnitude for these runs. To achieve this level of accuracy, high order quadrature rules are required, with rules of order $6p$ used in this case to demonstrate the GCL’s full effect. In practical cases we use more modest rules, namely $2p + 5$ for the Navier-Stokes equations, since discretization errors tend to dominate and make high quadrature rules unnecessary.

Finally, note that (89) has introduced an additional numerical quantity, \bar{g} , which is subject to discretization errors just like the state quantities, and this is relevant for error estimation.




Figure 72: Free-stream errors with and without the GCL. A large number of time steps is used so that the spatial error dominates the temporal error in each case.

6.3 Error Estimation and Adaptation

6.3.1 The Adjoint-Weighted Residual and Error Localization

The adjoint-weighted residual output error estimate in (36) extends directly to unsteady problems,

$$\delta J \approx -\Psi_h^T \mathbf{R}_h(\mathbf{U}_h^H) = -\sum_{k=1}^{N_k} \sum_{m=1}^{r+1} (\Psi_h^{km})^T \mathbf{R}_h^{km}(\mathbf{U}_h^H). \quad (91)$$

The effect of the GCL can be incorporated into the above formula by extending the state, residual, and adjoint vectors to include the GCL variable and equation. The contribution of the GCL to the error estimate is important because, when the GCL is used, the GCL residuals tell us where the mesh or time step should be refined to more accurately represent the true motion [67].

To localize the error contributions to individual space-time elements in the mesh, we note that the output error estimate in (91) can be written as a sum over all space-time elements,

$$\delta J \approx \sum_{k=1}^{N_k} \sum_{e=1}^{N_e} \varepsilon_e^k, \quad (92)$$

where the error contribution of a given space-time element (e, k) is, assuming a temporal order of r ,

$$\varepsilon_e^k = \sum_{m=1}^{r+1} (-\Psi_{he}^{km})^T \mathbf{R}_{he}^{km}(\mathbf{U}_h^H), \quad (93)$$

This is just the adjoint-residual product restricted to the element (e, k) , with a sum taken over the intra-slab temporal degrees of freedom m . The error indicator is then taken as the absolute value of this elemental contribution to the output error,

$$\text{error indicator for element } e \text{ of time slab } k = \epsilon_e^k = |\varepsilon_e^k|. \quad (94)$$

We are sometimes also interested in a ‘‘conservative’’ error estimate given by the sum of the individual error indicators,

$$\text{sum of error indicators} = \epsilon = \sum_{k=1}^{N_k} \sum_{e=1}^{N_e} \epsilon_e^k. \quad (95)$$

6.3.2 Incorporating Space-Time Anisotropy

The indicator in (94) is sufficient for isotropic space-time refinement in which elements are targeted for refinement in both space and time. However, isotropic refinement can produce inefficient meshes. For example, if the discretization is under-resolved in the time domain, spatial elements would needlessly be refined on account of the temporal error. Therefore, important for an efficient adaptation strategy is a measure of the space-time anisotropy of the error; in other words, is the error on a given space-time element due primarily to the *spatial* or *temporal* discretization?

As a heuristic measure of space-time anisotropy, one could consider inter-element jumps in the solution. For each space-time element, the average jump in the state computed across the spatial interfaces and across the time slab interfaces could serve as an estimate of which direction (space or time) is better resolved. Although cheap to evaluate, such a heuristic requires normalization and some arbitrary decisions, especially for systems of equations.

We can eliminate some of the heuristics by basing the anisotropy measure on the output error estimate [33, 34]. Specifically, we calculate the error anisotropy using separate projections of the fine-space adjoint onto semi-coarsened spatial and temporal spaces, as illustrated in Figure 73. The spatial and temporal error estimates for space-time element (e, k) are obtained by using these projected adjoints in (92), resulting in separate $\varepsilon_e^{k,\text{space}}$ and $\varepsilon_e^{k,\text{time}}$ estimates. We then use the ratio of these values to estimate the fractions (β) of spatial and temporal error on element (e, k) as

$$\beta_e^{k,\text{space}} = \frac{|\varepsilon_e^{k,\text{space}}|}{|\varepsilon_e^{k,\text{space}}| + |\varepsilon_e^{k,\text{time}}|}, \quad \beta_e^{k,\text{time}} = 1 - \beta_e^{k,\text{space}}. \quad (96)$$

6.3.3 Space-Time Mesh Adaptation

With the above estimates of spatial and temporal error, we have the fundamental information needed for adaptation. We adapt by resolving the unsteady problem several times on successively refined space-time meshes, as illustrated in Figure 74. After each primal solve, the adjoint equations are marched backward in time, new error indicators are obtained, and a new adapted mesh is generated. This process is repeated until the output error drops below a specified tolerance. A key choice in adaptation is the decision of *what* to adapt. First, to address temporal errors, we could refine the time-step on each



73.1: Estimation of ϵ^{time} 73.2: Estimation of ϵ^{space}

Figure 73: Measure of space-time anisotropy via projections of the fine-space adjoint into semi-coarsened spaces. In this work, the coarse spatial and temporal resolutions are obtained by order decrement, and the projection is least-squares.




Figure 74: Adaptive solution process for unsteady problems. Adjoint vectors are written during each adaptive iteration, which consists of a forward and adjoint solve.

space-time element separately, and this would require a type of hanging-node treatment in time. However, to enable efficient solution of the DG-in-time discrete system by our approximate Newton solver, we limit the temporal refinement to entire time slabs, as described below. Second, to address spatial errors, we can adapt a mesh and leave it adapted for the course of the simulation (static spatial refinement) or we can devise an adaptive schedule in which the refinement changes during the course of the simulation (dynamic refinement). We consider both options, as described below.

Temporal Refinement: To address temporal errors, we adapt the DG-in-time discretization by adjusting the size of each time slab. When adaptation consists of only refinement (no coarsening), then this adjustment consists of simple bisection of those time slabs targeted for refinement, as shown in Figure 75.1. However, when coarsening also occurs, the boundaries of a coarsened slab will typically encroach on those of the neighboring slabs, and the entire temporal grid must be shuffled to make room for the coarsened slab, as illustrated in Figure 75.2. To perform this shuffling, time slabs are redistributed using one-dimensional metric-based meshing, the details of which are given in Appendix B.




Figure 75: Temporal mesh refinement using bisection and shuffling. Time slabs shaded in blue are flagged for refinement, and those in gold are flagged for coarsening.

Static Spatial Refinement: One simplification that minimizes storage and complexity of the data structures is to have refinement of the spatial mesh remain fixed throughout the unsteady simulation. That is, the number/location of elements and approximation order on each element are both constant in time, as illustrated in Figure 76. Although this can limit efficiency of the adaptation, especially for problems that exhibit spatially-localized sources of error that move in time, there are many problems for which such a static refinement strategy works well.

In static spatial refinement, the “objects” to be adapted are (i) spatial elements and (ii) time slabs. Adaptive indicators identifying the spatial error on each spatial element

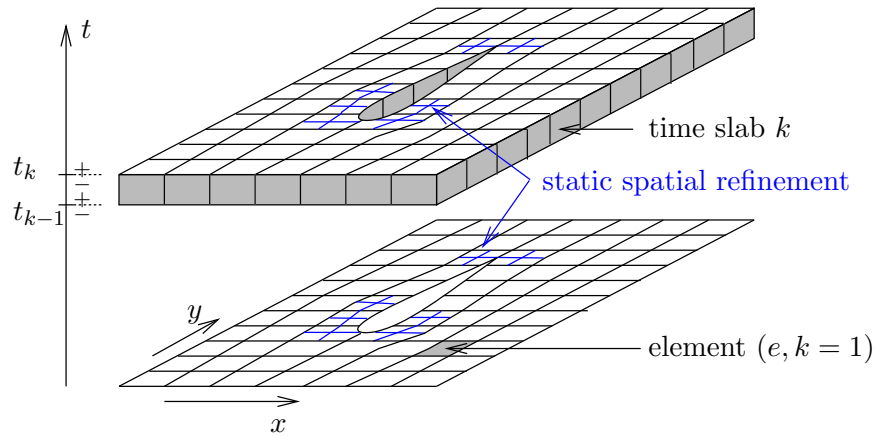


Figure 76: Illustration of static spatial refinement for a DG-in-time discretization, using hanging-node spatial mesh refinement. Static order refinement follows a similar approach.

and the temporal error on each slab are given by

$$\text{aggregate spatial indicator on element } e = \epsilon_e^{\text{space}} = \sum_{k=1}^{N_k} \epsilon_e^k \beta_e^{k,\text{space}}, \quad (97)$$

$$\text{aggregate temporal indicator on time slab } k = \epsilon_e^{k,\text{time}} = \sum_{e=1}^{N_e} \epsilon_e^k \beta_e^{k,\text{time}}, \quad (98)$$

With these indicators, we could then conceivably lump all time slabs and space-time elements into the same “bin,” rank them according to their error indicators, and determine which to adapt based on their relative positions in that ranking. However, this approach would neglect the *cost* of each refinement option. For example, if a time slab and a spatial element had the same error but there were many more time slabs than spatial elements, one would not want to weigh refining the time slab equally compared to refining the element. Rather, refining the time slab would be cheaper because it would add fewer degrees of freedom compared to refining a spatial element. Thus, rather than adapting directly on errors, we adapt on a slightly different figure of merit – the amount of error on a given element or slab (Eqns. 97 and 98) *divided by* the additional degrees of freedom associated with adapting that element or slab. This figure of merit then ensures that we eliminate the most output error for the least additional cost.

To first approximation, the number of degrees of freedom introduced in a time slab division is estimated as the number of spatial degrees of freedom in the current mesh, while the number of degrees of freedom introduced in a spatial refinement is estimated as N_k times the number of new spatial degrees of freedom obtained from a refinement of that element. The respective adaptive indicators $\epsilon_e^{k,\text{time}}$ and $\epsilon_e^{\text{space}}$ are divided by these quantities and then sorted highest to lowest. The element or time slab with the highest error indicator per proposed additional number of degrees of freedom is chosen for refinement first, and the process continues until a growth budget is reached or surpassed. This is effectively a fixed-growth adaptation strategy. We note that in some cases the refinement could target only spatial elements or only time slabs, depending on the relative resolution in time and space.

Dynamic Spatial Refinement: In dynamic spatial refinement, the spatial resolution of an element can change during the course of the simulation. This change could arise from hanging node refinement, mesh motion, or order refinement. We focus on order refinement, illustrated in Figure 68, as it is one of the simplest options for a DG discretization in space and time.

Recall that p_e^k is the order of spatial approximation on spatial element e at time slab k . So the “objects” to be adapted are both (i) individual space-time elements in order and (ii) time slabs. Adaptive indicators identifying the spatial error on each space-time element and the temporal error on each slab are given by

$$\text{spatial indicator on space-time element } e, k = \epsilon_{ek}^{\text{space}} = \epsilon_{ek} \beta_{e,k}^{\text{space}}, \quad (99)$$

$$\text{aggregate temporal indicator on time slab } k = \epsilon^{k,\text{time}} = \sum_{e=1}^{N_e} \epsilon_{e,k} \beta_{e,k}^{\text{time}}, \quad (100)$$

These definitions are similar to the static refinement case, except that now the spatial indicator exists for every space-time element, instead of just as a sum over time slabs for each element. As in static spatial refinement, we make a decision of what to adapt based on a figure of merit that incorporates both the localized error estimate and the cost (in terms of additional degrees of freedom) of the refinement operation. We now also allow for order decrease, i.e. coarsening of the spatial mesh, in addition to refinement. More details on this algorithm are given in Appendix B.

6.3.4 Implementation Notes

The adjoint equations require several derivative terms, including residual Jacobians and output linearizations. In the current work, we perform all differentiation analytically, with the exception of some terms required when using the geometric conservation law, which we evaluate using finite differences for ease of implementation.

When solving the adjoint equations, we use the entire time history of the primal state and GCL variable, which we store to disk during the primal solve. While for the present work this storage has not been prohibitive, for larger problems solution checkpointing [49] or local-in-time adjoint solvers [124] may be considered.

The space-time error indicators are computed on the fly during the backwards time-marching solution of the unsteady adjoint problem. Following each unsteady adaptive iteration, the error indicators are sorted and space-time elements are identified for refinement and coarsening. A schedule of orders p_e^k and time slab sizes is then written to disk for the next adaptive iteration. This schedule takes the form of individual files that are read in by the forward solver during the next forward solve.

Error estimation requires that the adjoint be solved in an enriched space, and for our work we use order refinement in both space and time: $p \rightarrow p+1$, and $r \rightarrow r+1$. For large (e.g. 3D) simulations, the cost of an adjoint solution on this space may become prohibitive due to memory limitations associated with the linear solve (we store the full residual Jacobian matrix). To address this issue, in these cases we employ a reconstruction-based strategy in lieu of computing the fine-space adjoint directly. That is, we compute the adjoint in the same space as the primal problem, and then reconstruct it in both space and time to obtain an approximation to the fine-space adjoint. This reconstruction is

performed locally, with patches of neighboring elements used to generate a least-squares spatial reconstruction, and pairs of neighboring time slabs used to obtain a high order temporal interpolant [33], as illustrated in Figure 77.

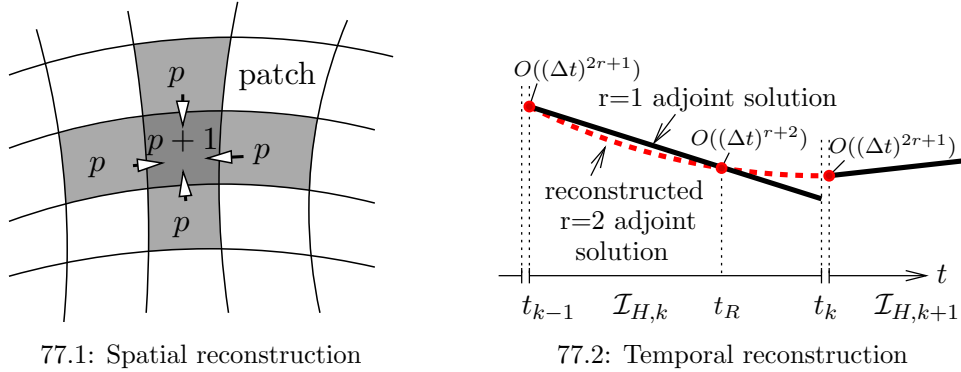


Figure 77: Illustration of spatial and temporal adjoint reconstructions. On the left, a patch of nearest-neighbor elements used for spatial high order reconstruction via least-squares interpolation. On the right, reconstruction of an $r = 1$ adjoint to $r = 2$ using the left node from the adjacent future time slab and superconvergent nodes on the current time slab. t_R indicates the root of the left Radau polynomial for $r = 1$.

6.4 Examples

6.4.1 Static h -Refinement for an Impulsively-Started Airfoil

In this example we demonstrate the performance of static hanging-node spatial mesh refinement together with time slab bisection for an unsteady computation on a static domain. The adaptive mechanics consist of only refinement (no coarsening) in space and time, driven by a prescribed fixed-growth of degrees of freedom, $f^{\text{growth}} = 2$, per adaptive iteration. Adaptation using the output-based error indicator is compared to several other strategies: adaptation using an indicator obtained from an unweighted residual; adaptation driven by a measure of inter-element jumps; and uniform space and time refinement. Details on these adaptive strategies are given in [33].

We consider an impulsively-started NACA 0012 airfoil in viscous flow, where for $t \geq 0$, the freestream conditions are $M_\infty = 0.25$, $\alpha = 8^\circ$, $Re = 5000$. To prevent a non-physical step change in the velocity of the fluid at the airfoil surface, the initial condition at $t = 0$ consists of the freestream with the velocity blended smoothly to zero in a circular disk around the airfoil. Specifically, the velocity in the blended region, $r_1 \leq r \leq r_2$ is $\mathbf{v} = \mathbf{V}_\infty(1 - \cos(\pi(r - r_1)/(r_2 - r_1)))/2$ where r_1 and r_2 are radial distances from the airfoil mid-chord (set to one and three chord lengths respectively) and \mathbf{V}_∞ is the freestream velocity. No steady solve is performed prior to the unsteady simulation.

Figure 78a shows the entropy contours at $t = 10$ units, the final time in the simulation. By this time an alternating pattern of shed vortices has developed and is clearly visible. The output of interest is the lift coefficient integral from $t = 9$ to $t = 10$, as illustrated in Figure 79a. A snapshot adjoint solution for the y -momentum equation at $t = 6$ is

illustrated in Figure 78b. A “reverse wake” is evident in the adjoint solution, signifying an oscillatory sensitivity of the output to y -momentum residual perturbations upstream.





Figure 78: Impulsively-started airfoil: primal state at the final time and the adjoint state at $t = 6$.


For the adaptive runs, an initial spatial mesh of 510 elements is used, and the initial temporal mesh contains 16 time slabs. The output convergence for the various indicators is shown in Figure 79b. The residual indicator does not perform well at all again: the output varies significantly from iteration to iteration. The other indicators converge, with the fastest being output-based adaptation, followed by approximation error and then uniform refinement. The advantage of the output-based refinement with degrees of freedom is a factor of 3-4 savings over the approximation error indicator. The error estimates under-predict the error in the middle stages of output-based refinement, while the conservative whiskers at $\pm\epsilon$ are more robust.

Figure 80a shows the time histories of the lift coefficient for adapted space-time meshes of similar size. The source of the error in the residual-adapted case is clear: it does not predict oscillatory vortex shedding, but rather an increasing lift coefficient. The other three adaptive indicators track the actual time history well. Figure 80b shows the L_2 time history error convergence for all of the methods, versus the cube root of the degrees of freedom. Under uniform refinement, the observed convergence rate is a suboptimal 1.5, which could be caused by high-order derivative discontinuities in the blended initial condition or the resolution not yet being in the asymptotic regime. Of primary interest, however, is the performance of output-based adaptation, which remains the fastest out of the methods tested. Even though the output is only measured on the final 10% of the simulation time, accurate resolution prior to this metric time is important as it affects the state at the start of the output measurement.

The corresponding adapted meshes are shown in Figure 81. Output-based adaptation targets the airfoil leading and trailing edges, the boundary-layer region above the front of the airfoil, and slightly the stagnation line in front and the wake behind the airfoil. The approximation-error indicator also targets the leading and trailing edges and puts




79.1: Output definition



79.2: Output convergence

Figure 79: Impulsively-started airfoil: time integral output definition and its convergence under the adaptive indicators. Error bars at $\pm\delta J$ and whiskers at $\pm\epsilon$ are shown for the output-based results. The “actual” output is computed on a uniformly-refined final output-adapted space-time mesh.



80.1: Sample output histories


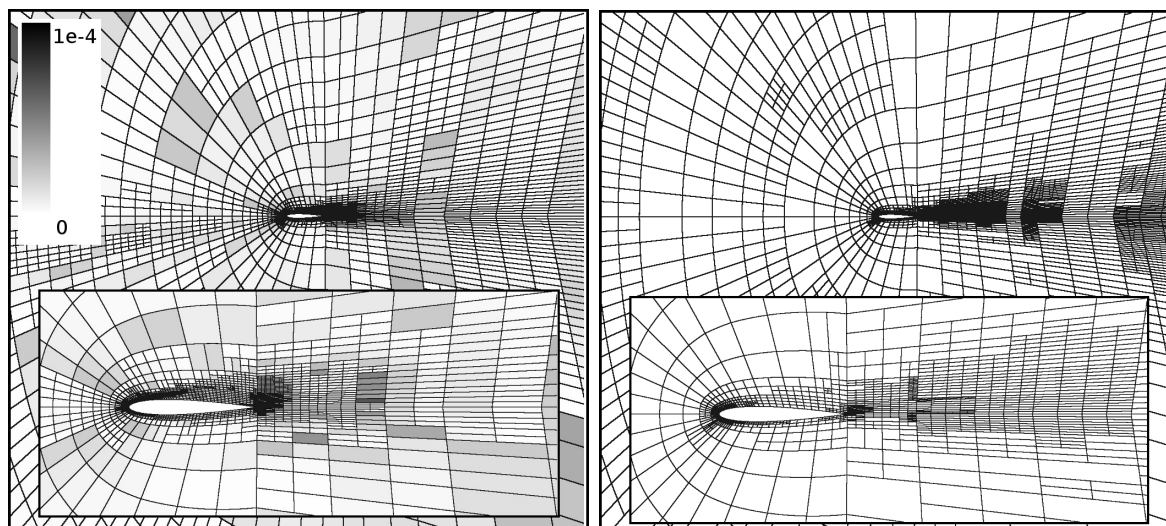
80.2: L_2 error convergence

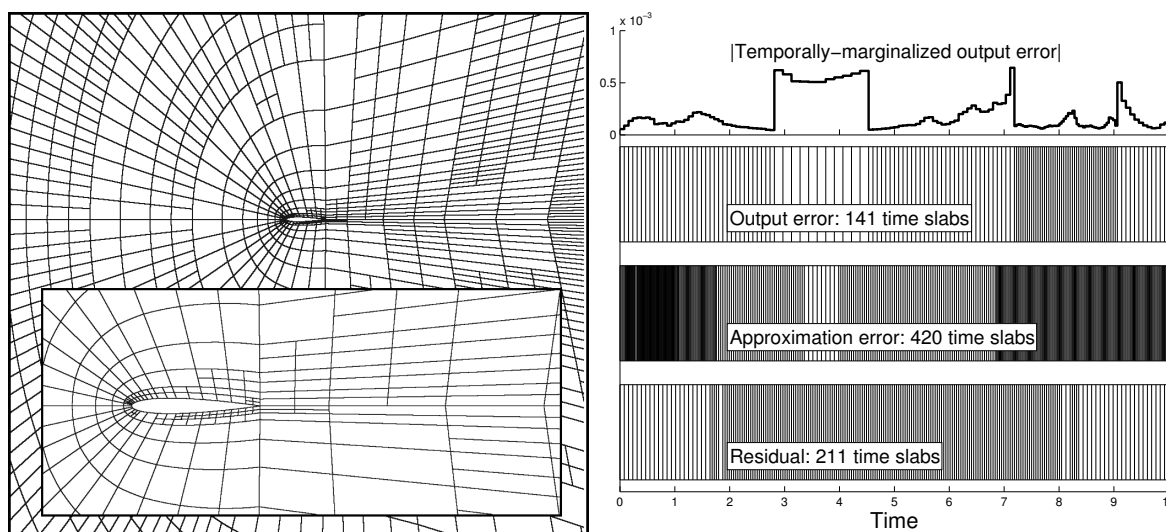
Figure 80: Impulsively-started airfoil: sample lift coefficient time histories and convergence of the L_2 time history error for various adaptive indicators. The “actual” time history is computed on a uniformly-refined final output-adapted space-time mesh.

more emphasis on the wake, where the shed vortices propagate. The residual indicator is distracted by effects of the initial condition: the velocity blending near the airfoil sends out acoustic waves that the residual indicator attempts to track as they propagate away from the airfoil.



81.1: Adapted on output error (5956 elements)

81.2: Adapted on approx. error (4585 elements)



81.3: Adapted on residual (7929 elements)

81.4: Temporal meshes

Figure 81: Impulsively-started airfoil: adapted spatial and temporal meshes for the seventh adaptive iteration. Localized output error estimates ϵ_e and ϵ^k are shown for the output-error adapted meshes.

The temporal meshes are shown in Figure 81.4. The output-based indicator creates a fairly uniform temporal refinement, with slightly higher resolution prior to the metric time. The approximation error focuses on the initial time, as it tracks the evolution of the blended velocity field, and the latter 1/3 time when the shed vortices develop. The residual indicator again creates a mostly-uniform temporal mesh as it tracks the acoustic waves.

6.4.2 Dynamic-Order Adaptation for Pitching and Plunging Airfoils

In this example we demonstrate the performance of dynamic-order spatial mesh refinement/coarsening together with time slab refinement/coarsening for an unsteady computation on a deformable domain. The adaptive mechanics consist of refinement driven by a prescribed fixed-growth factor for degrees of freedom of $f^{\text{growth}} = 1.35$ per adaptive iteration, together with a coarsening fraction of $f^{\text{coarsen}} = .05$. Adaptation using the output-based error indicator is compared to adaptation using an indicator obtained from an unweighted residual, and to uniform h and p refinement of the spatial mesh. The temporal adaptation is always time slab refinement/coarsening.

This case involves two airfoils pitching and plunging in series at low Reynolds number [67]. The airfoils start from an impulsive free-stream condition and undergo three periods of motion. The plunge amplitude is 0.25 chords, the pitch amplitude is 30° , and the period of both motions is $T = 2.5$. The Strouhal, Mach, and Reynolds numbers are $2/3$, 0.3, and 1200, respectively. The airfoils are offset 4.5 chords horizontally and 1 chord vertically, and are situated in a 60×60 chord-length mesh with 3,534 triangular elements.

Entropy contours at various phases of the motion are shown in Figure 82. A reverse Kármán vortex street develops behind each airfoil, and the second airfoil interacts with the wake from the first airfoil near the end of the simulation. Our output of interest is the lift on the second airfoil integrated from time $t = 7.25$ to $t = 7.5$ (the final time).

Adaptations were performed starting from an initial $p = 1$, 90 time step solution, with a 35% growth factor and 5% coarsening factor used for the output- and residual-based methods. The spatial order p was constrained to lie between 0 and 5, and a DG1 scheme was used in time.

The output convergence for each adaptive method as a function of total space-time degrees of freedom is shown in Figure 83. We see that the output-based adaptation converges much faster than uniform refinement, requiring roughly two orders of magnitude fewer degrees of freedom. These gains relative to uniform refinement are impressive, though not entirely unexpected. Equally interesting is the difference between the output-based and residual adaptation. The residual indicator targets regions of the domain where the governing equations are not well-satisfied, and hence usually performs well for static problems. However, in this case, its performance is erratic and no better than uniform refinement. This erratic behavior is a consequence of the acoustic waves that emanate from the airfoils as they pitch back and forth. The residual indicator becomes distracted by these waves and exhausts degrees of freedom trying to resolve them as they propagate throughout the domain. The output-based method, on the other hand, dismisses the majority of these waves as irrelevant and simply ignores them.

The spatial and temporal meshes from the final output-based adaptation are shown in Figures 84 and 85, respectively. We see that the near-airfoil and vortex shedding regions are targeted for adaptation, as well as the group of large elements surrounding the mesh motion regions. While somewhat difficult to observe in the still-frames, the initial vortex shed from the first airfoil is heavily targeted throughout the simulation, since this vortex later collides with the second airfoil near the final time.

To highlight one of the factors driving the adaptation, contours of the GCL adjoint are shown alongside the entropy contours in Figure 82. The time $t = 0.70T$ is the instant before the initial vortex is shed, and the large sensitivity of the output to this event can be




82.1: $t = 0.70T$ 82.2: $t = 1.25T$ 82.3: $t = 2.75T$

Figure 82: Two-airfoil case: Entropy (left) and GCL adjoint (right) contours at various stages of the motion on a fine mesh. The GCL adjoint contours have been re-scaled to more clearly show the features (black is -2, white is 1). Both acoustic and convective modes of error propagation can be seen in the first two contours, while at the final time, the adjoint field collapses on the second airfoil.




Figure 83: Two-airfoil case: Output convergence for various adaptive methods. The output-based method performs the best. Note that for all plots, the “actual” value is taken from the final uniform p -refinement.

seen in the adjoint contours. As the simulation proceeds, the output sensitivity gradually shifts from the first airfoil to the second, before collapsing upon the second airfoil at the final time.

Some other aspects of the GCL adjoint are worth pointing out. In the first two contours, the near-circular rings represent inward-moving (adjoint) acoustic waves, which converge upon a particular region as the simulation proceeds. The existence of a ring implies that an important event in space-time is about to occur, and any errors made within the circumference of the ring have the ability to influence this event. In this simulation, the important events tend to be instances of vortex shedding, and the rings converge on the trailing edge regions. Lastly, between the two airfoils, a path can be seen tethering them together. This path appears because any errors within it ultimately reach the second airfoil via convection, and can therefore directly affect the output.

Above, we solved the fine-space adjoint to machine precision to ensure accurate error estimates and efficient allocation of mesh degrees of freedom. In practice, CPU time is another important factor, and solving the adjoint to machine precision is generally not the most efficient strategy. In Figure 86, we show a wall time comparison for the various adaptive strategies, with the adjoint smoothed to a residual tolerance of 1×10^{-3} . This tolerance is not necessarily optimal, and we note that the code itself has not been optimized for CPU time. However, our aim is simply to give an idea of the relative timings. While the performance gains for the output-based method are not as substantial in this context, it converges roughly 5 – 10 times faster than the uniform refinement strategies. It also significantly outperforms the residual-based adaptation, which fails to converge in any reasonable amount of time.




84.1: $t = 0.70T$ 84.2: $t = 1.25T$ 84.3: $t = 2.75T$

Figure 84: Two-airfoil case: Output-adapted meshes at various stages of the motion. Blue is $p = 0$, red is $p = 5$.




Figure 85: Two-airfoil case: Temporal grids from the seventh adaptation of both output-based and residual runs. For clarity, only every other time slab is plotted.




Figure 86: Two-airfoil case: Wall-time comparison for the adaptive methods. The output-based method converges the fastest.

6.4.3 Dynamic-Order Adaptation for a Three-Dimensional Wing

In this example we demonstrate the performance of dynamic-order spatial mesh refinement/coarsening together with time slab refinement/coarsening for an unsteady computation on a deformable domain in three dimensions. The adaptive mechanics consist of refinement driven by a prescribed fixed-growth factor for degrees of freedom of $f^{\text{growth}} = 1.3$ per adaptive iteration, together with a coarsening fraction of $f^{\text{coarsen}} = .05$. Adaptation using the output-based error indicator is compared to adaptation using an indicator obtained from an unweighted residual, as well as to uniform h and p refinement of the spatial mesh – the temporal adaptation is always time slab refinement/coarsening.

In three dimensions, the underlying ideas for mesh motion and adaptation remain the same, but the increased algorithmic complexity and new dimensional scaling do not allow for a simple extrapolation of the 2D results. In particular, the extra dimension makes the scaling for the adjoint problem less favorable. For a $(p=1, r=1)$ primal solve in 2D, the dimension of the fine-space adjoint is approximately 3 times that of the primal, while in 3D this factor increases to over 5. When combined with additional CPU costs, computing (or even smoothing) the adjoint on the fine space becomes an expensive proposition. We therefore employ spatial and temporal reconstruction of the adjoint in lieu of a solution on the fine space, as discussed in Section 6.3.4.




Figure 87: 3D wing: Schematic of the flapping motion. The flow regime is approximately that of a housefly.

To test this strategy, we proceed in a similar manner as in 2D, though now simulating a full wing rather than just an airfoil. The wing is shown in Figure 87, along with a schematic of the flapping motion. The wing moves in all dimensions, with a 30° stroke angle, a slight vertical plunge simulating movement of the “shoulder joint,” and an angle of attack variation of $\pm 10^\circ$. The Mach number is 0.3, while the Reynolds number of 500, Strouhal number of 0.4, and aspect ratio of 1.5 place the wing in the flight regime of a small housefly.

Three periods of motion were simulated, and the solution at various times is shown in Figure 88. Strong vortex cores develop near the leading edge and wingtip regions before detaching and shedding into the wake. For adaptation purposes, the output of interest is taken to be the lift integrated over the final 5% of the simulation time (from $t = 7.1$

to $t = 7.5$). To solve the problem, output-based, residual, and uniform p -refinement strategies were employed. Adaptations were performed starting from an initial $p = 0$, 75 time step solution, the spatial order p was constrained to lie between 0 and 5, and a DG1 (i.e. $r = 1$) scheme was used in time. The convergence for each adaptive method is






Figure 88: 3D wing: Mach contours projected onto entropy isosurfaces, shown for several stages of the flapping motion. The maximum Mach number (in red) is approximately 0.5. The images are taken from the final ($p=3$) uniform refinement.

shown in Figure 89, and the results are encouraging. We see that although reconstructed adjoints were used, the output-based adaptation still performs well, and converges with about two orders of magnitude fewer degrees of freedom than uniform refinement. As expected, the accuracy of the error estimate itself is not great (judging by the fact that the corrected output converges no faster than the uncorrected one), but is enough to guide the adaptation in the correct direction. Convergence as a function of wall time is shown in Figure 89, and we see that the method also performs well in this context.



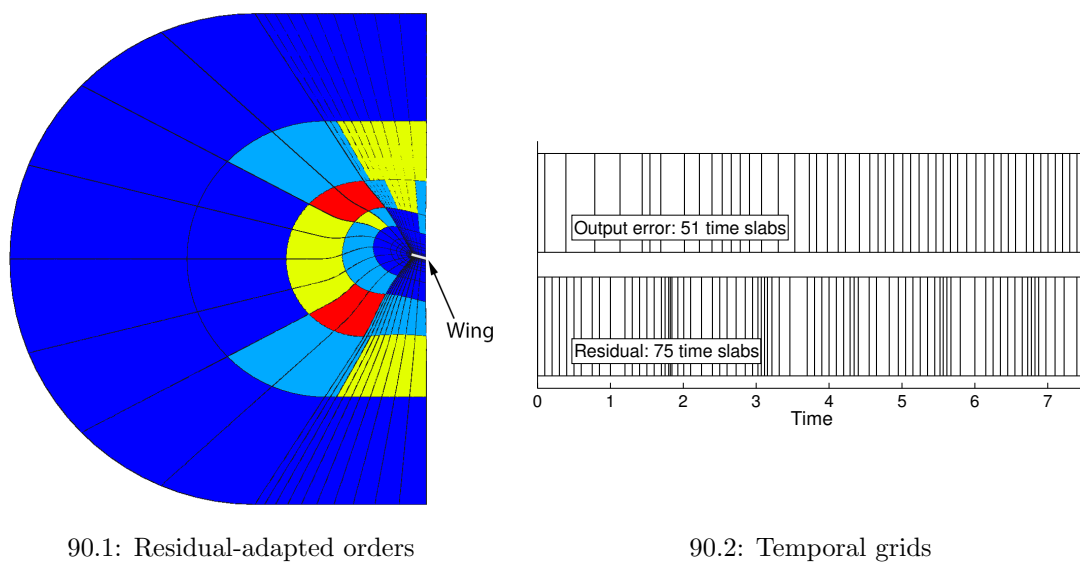
89.1: Convergence versus DOF



89.2: Convergence versus wall time

Figure 89: 3D wing: Output convergence as a function of degrees of freedom and wall time.

Of additional interest is the performance of the residual-based adaptation, which is so poor that the curve is easy to miss in the plots. The reason for this behavior is evident in Figure 90, which shows the orders midway through the final residual adaptation. The residual indicator flags only large elements in the mesh blending region for refinement, and leaves all elements near the wing at their original low order. The output-based adaptation, on the other hand, refines elements near the wing as needed. The adapted spatial meshes from the output-based runs are shown in Figure 91, while adapted temporal grids are shown in Figure 90. We see that, at early times, the output-based indicator leaves the mesh relatively coarse, but progressively increases the resolution in both space and time as the simulation progresses.



90.1: Residual-adapted orders

90.2: Temporal grids

Figure 90: : On the left, spatial orders from the residual adaptation midway through the simulation (blue is $p = 0$, red is $p = 3$). The adaptation targets only elements far from the wing, leading to poor output convergence. On the right, temporal grids from the final output-based and residual adaptations. The residual adaptation refines periodically with the motion, while the output-based adaptation increases the resolution near the final times.

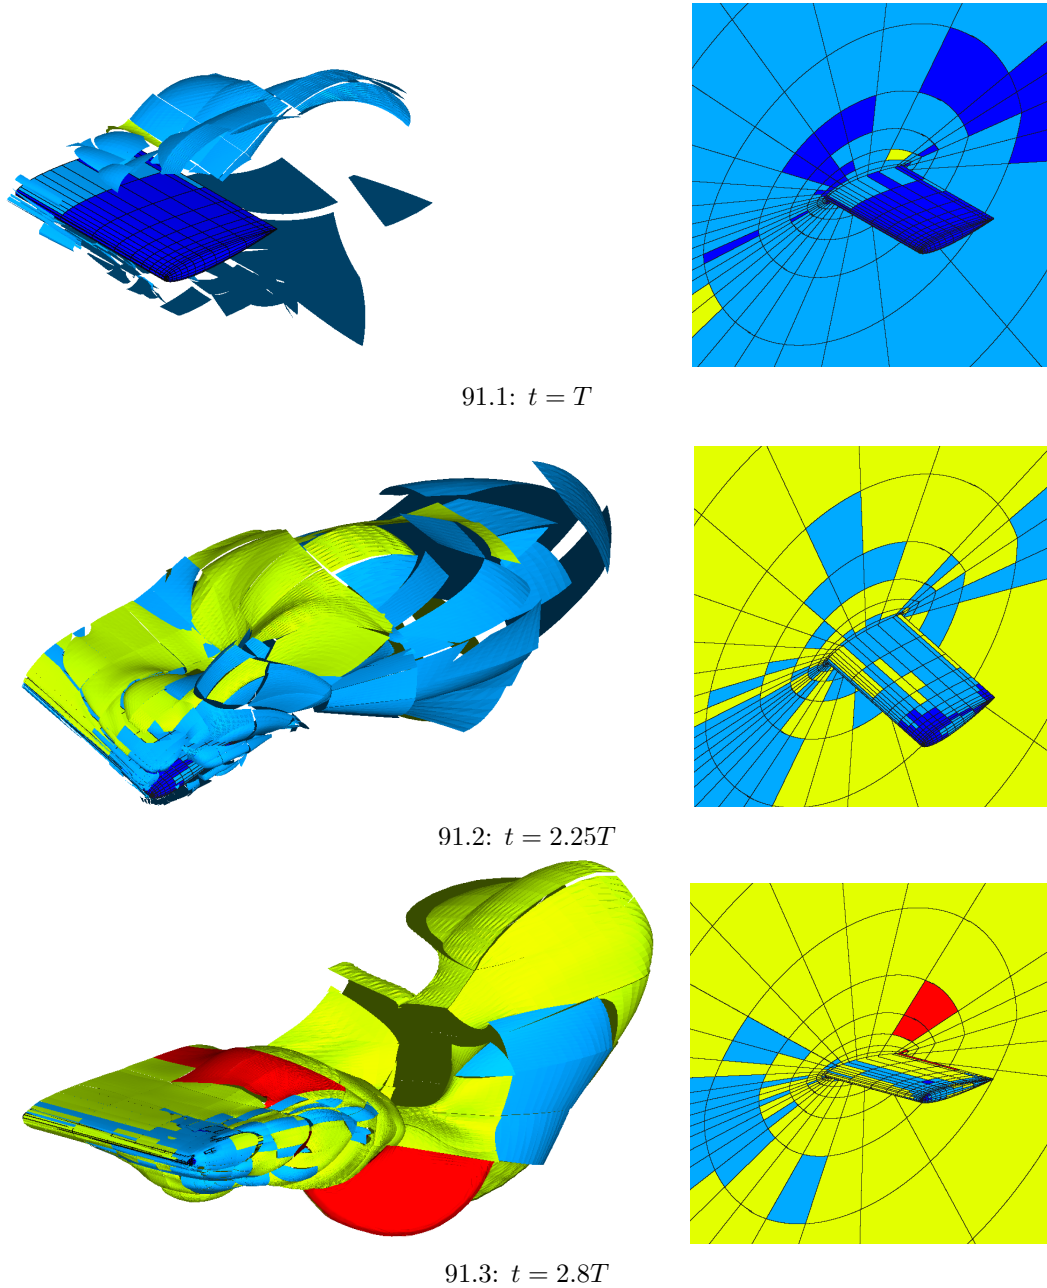


Figure 91: 3D wing: Spatial orders from the final output-based adaptation, shown at several stages of the flapping motion. Dark blue is $p = 0$, red is $p = 3$. The images on the left show the interpolation orders projected onto entropy isosurfaces.

7 Conclusions

In this set of lecture notes we have presented some of the core ideas, methods, and results for high-order output-based adaptive simulations of aerodynamic flows. Many of the presented methods can be generalized to more than one discretization; we chose the discontinuous Galerkin (DG) finite element method for various reasons, including ease of high-order implementation, adaptive flexibility including natural support for hp -adaptation, parallel scalability, and of course code availability.

Numerical error estimates that use an adjoint solution to target a specific scalar output are particularly useful for convection-dominated aerodynamic simulations, in which even seemingly-small errors in one area of the computational domain can affect a large portion of the flow field and significantly impact an output of interest. Weighting residuals by the adjoint solution produces an error estimate that takes such propagation effects into account. The result is a type of “error bar” on the chosen output of the numerical simulation. The cost of this error estimate lies in the computation of the adjoint solution via a linear solve, possibly in an enriched “fine” approximation space. For the complex nonlinear systems often encountered in aerodynamics, even a fine-space linear solve is generally cheaper than the solution of the original primal system on the coarse space.

Moreover, because an output-error estimate is residual-based, it can be localized to cells/elements in a finite volume/element calculation for the purpose of mesh adaptation. Targeted local adaptation, such as refinement of a fraction of the elements with the highest indicator, reduces local residuals, which then reduces the local contributions of the targeted elements to the error. Furthermore, global unstructured remeshing using models for the cost and error yields optimal meshes with equidistributed marginal error to cost ratios. As we have shown, mesh adaptation using output based methods extends to a wide variety of complex problems, including three-dimensional, transonic, Reynolds-averaged turbulent flow over a wing, and unsteady simulations on deforming domains. Below we highlight some key conclusions garnered from the presented examples.

- Output-based adaptation can quarantine singularities, as long as they do not span too much of the domain, so that these singularities do not affect the results and so that optimal convergence rates, such as $2p + 1$ or $2p$ for inviscid or viscous problems, are recovered. Uniform refinement, on the other hand, is limited by the lowest-order singularity.
- For many steady problems, output-based adaptation shows benefits compared to uniform refinement or heuristic indicators in both degrees of freedom and computational time. The latter is especially true for three-dimensional problems with complex physics, e.g. RANS, in which the cost of an adjoint solve is low relative to that of the primal solve.
- The mechanics of adaptation can be difficult in situations when curved anisotropic elements are required, in particular in three dimensions. We have presented adaptation results using hanging-node refinement of initially-structured quadrilateral and hexahedral meshes, and this works quite well in the near-body region, where the flow follows the body. However, such an approach requires a reasonable initial mesh

and exhibits much more limited resolution capabilities compared to anisotropic unstructured meshes. Improving the robustness of the latter is still an active research area.

- Unsteady problems add another dimension – time – as a candidate for refinement, and this generally bodes well for adaptation, as additional dimensions often translate to a larger portion of the domain being “unimportant” for the prediction of an output. With an ALE formulation, it is possible to handle complicated problems involving geometry deformation, and we have demonstrated adaptation for such problems with dynamic p refinement and coarsening. The adjoint becomes more expensive in the unsteady case due to storage requirements for nonlinear problems. Yet, for sufficiently strict error tolerances, we still observe benefits in computational time relative to uniform refinement and heuristic indicators. Adaptation becomes even trickier in unsteady problems with the added complexity of temporal refinement, and we have presented a projection-based measure of space-time anisotropy that helps guide the space versus time refinement decision.
- A sampling-based approach to metric-based anisotropic mesh regeneration can yield meshes optimized for the prediction of an output with minimal computational resources. The marginal error to cost ratio can be equidistributed over the elements in the domain using a simple iterative procedure. Anisotropy arises naturally in this approach from sampling the error and cost using a set of hypothetical cuts of the elements, and this anisotropy reflects efficiency gains due to features in both the primal and adjoint solutions, via the adjoint-weighted residual combination. Success of this method relies on the availability of a metric-based global re-mesher, which is a challenging task in three spatial dimensions.

Finally, adaptive high-order methods likely are not yet being used to their full potential. Additional challenges that need to be addressed include: reducing the cost of adjoint solves, in particular for unsteady problems; applying adjoint methods to chaotic simulations; calculating bounds instead of just estimates of the output error; investigating the possibility of cheaper surrogates such as the entropy adjoint; developing better adaptive mechanics for complex spatial meshes; and generalization of the unsteady adaptive methods to other discretizations, such as multi-step or multi-stage time integration.

8 Acknowledgments

These notes were made possible by contributions of numerous undergraduate students, graduate research assistants, and post-doctoral scholars working in the CFD Group at the University of Michigan. The research was supported by the University of Michigan, the U.S. Air Force Office of Scientific Research (grant FA9550-11-1-0081), the U.S. Department of Energy (grant DE-SC0010341), and student fellowships from the U.S. National Science Foundation and the U.S. Department of Defense.

References

- [1] M.J. Aftosmis and M.J. Berger. Multilevel error estimation and adaptive h-refinement for Cartesian meshes with embedded boundaries. AIAA Paper 2002-14322, 2002.
- [2] Mark Ainsworth and Bill Senior. An adaptive refinement strategy for *hp*-finite element computations. *Applied Numerical Mathematics*, 26:165–178, 1998.
- [3] S.R. Allmaras, F.T. Johnson, and P.R. Spalart. Modifications and clarifications for the implementation of the Spalart-Allmaras turbulence model. Seventh International Conference on Computational Fluid Dynamics (ICCFD7) 1902, 2012.
- [4] Douglas N. Arnold, Franco Brezzi, Bernardo Cockburn, and L. Donatella Marini. Unified analysis of discontinuous Galerkin methods for elliptic problems. *SIAM Journal on Numerical Analysis*, 39(5):1749–1779, 2002.
- [5] Timothy J. Baker. Mesh adaptation strategies for problems in fluid dynamics. *Finite Elements in Analysis and Design*, 25:243–273, 1997.
- [6] Pinhas Bar-Yoseph and David Elata. An efficient L2 Galerkin finite element method for multi-dimensional non-linear hyperbolic systems. *International Journal for Numerical Methods in Engineering*, 29:1229–1245, 1990.
- [7] Garrett E. Barter. *Shock Capturing with PDE-Based Artificial Viscosity for an Adaptive Higher-Order Discontinuous Galerkin Finite Element Method*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2008.
- [8] T. J. Barth. Numerical methods for gasdynamic systems on unstructured meshes. In D. Kröner, M. Ohlberger, and C. Rhode, editors, *An Introduction to Recent Developments in Theory and Numerics for Conservation Laws, Proceedings of the International School on Theory and Numerics for Conservation Laws, Berlin, Lecture Notes in Computational Science and Engineering*. Springer-Verlag, 1999.
- [9] Timothy Barth and Mats Larson. A posteriori error estimates for higher order Godunov finite volume methods on unstructured meshes. In R. Herban and D. Kröner, editors, *Finite Volumes for Complex Applications III*, pages 41–63, London, 2002. Hermes Penton.
- [10] Timothy J. Barth. Space-time error representation and estimation in Navier-Stokes calculations. In Stavros C. Kassinos, Carlos A. Langer, Gianluca Iaccarino, and Parviz Moin, editors, *Complex Effects in Large Eddy Simulations*, pages 29–48. Springer Berlin Heidelberg, Lecture Notes in Computational Science and Engineering Vol 26, 2007.
- [11] F. Bassi and S. Rebay. GMRES discontinuous Galerkin solution of the compressible Navier-Stokes equations. In Karniadakis Cockburn and Shu, editors, *Discontinuous Galerkin Methods: Theory, Computation and Applications*, pages 197–208. Springer, Berlin, 2000.

-
- [12] F. Bassi and S. Rebay. Numerical evaluation of two discontinuous Galerkin methods for the compressible Navier-Stokes equations. *International Journal for Numerical Methods in Fluids*, 40:197–207, 2002.
- [13] R. Becker and R. Rannacher. A feed-back approach to error control in finite element methods: Basic analysis and examples. *East-West Journal of Numerical Mathematics*, 4(4):237–264, 1996.
- [14] R. Becker and R. Rannacher. An optimal control approach to a posteriori error estimation in finite element methods. In A. Iserles, editor, *Acta Numerica*, pages 1–102. Cambridge University Press, 2001.
- [15] M.J. Berger and A. Jameson. Automatic adaptive grid refinement for the Euler equations. *AIAA Journal*, 23:561–568, 1985.
- [16] Kim S. Bey and J. Tinsley Oden. *hp*-version discontinuous Galerkin methods for hyperbolic conservation laws. *Computer Methods in Applied Mechanics and Engineering*, 133:259–286, 1996.
- [17] H. Borouchaki, P. George, F. Hecht, P. Laug, and E. Saltel. Maillageur bidimensionnel de Delaunay gouverné par une carte de métriques. Partie I: Algorithmes. INRIA-Rocquencourt, France. Tech Report No. 2741, 1995.
- [18] F. Brezzi, L.D. Marini, and E. Süli. Discontinuous Galerkin methods for first-order hyperbolic problems. *Mathematical Models and Methods in Applied Sciences*, 14:1893–1903, 2004.
- [19] Nicholas K. Burgess and Dimitri J. Mavriplis. An *hp*-adaptive discontinuous Galerkin solver for aerodynamic flows on mixed-element meshes. AIAA Paper 2011-490, 2011.
- [20] Gustavo C. Buscaglia and Enzo A. Dari. Anisotropic mesh optimization and its application in adaptivity. *International Journal for Numerical Methods in Engineering*, 40(22):4119–4136, November 1997.
- [21] P. J. Capon and P. K. Jimack. On the adaptive finite element solution of partial differential equations using h-r refinement. Technical Report 96.03, University of Leeds, School of Computing, 1996.
- [22] M. J. Castro-Diaz, F. Hecht, B. Mohammadi, and O. Pironneau. Anisotropic unstructured mesh adaptation for flow simulations. *International Journal for Numerical Methods in Fluids*, 25:475–491, 1997.
- [23] Marco A. Ceze and Krzysztof J. Fidkowski. Output-driven anisotropic mesh adaptation for viscous flows using discrete choice optimization. AIAA Paper 2010-0170, 2010.
- [24] Marco A. Ceze and Krzysztof J. Fidkowski. A robust adaptive solution strategy for high-order implicit CFD solvers. AIAA Paper 2011-3696, 2011.

- [25] Marco A. Ceze and Krzysztof J. Fidkowski. An anisotropic hp-adaptation framework for functional prediction. *American Institute of Aeronautics and Astronautics Journal*, 51:492–509, 2013.
- [26] Bernardo Cockburn and Chi-Wang Shu. Runge-Kutta discontinuous Galerkin methods for convection-dominated problems. *Journal of Scientific Computing*, 16(3):173–261, 2001.
- [27] L. Demkowicz, W. Rachowicz, and Ph. Devloo. A fully automatic hp-adaptivity. *Journal of Scientific Computing*, 17:117–142, 2002.
- [28] Bruno Despres. Lax theorem and finite volume schemes. *Mathematics of Computation*, 73(247):1203–1234, 2003.
- [29] Julien Dompierre, Marie-Gabrielle Vallet, Yves Bourgault, Michel Fortin, and Wagdi G. Habashi. Anisotropic mesh adaptation: towards user-independent, mesh-independent and solver-independent CFD. Part III: Unstructured meshes. *International Journal for Numerical Methods in Fluids*, 39:675–702, 2002.
- [30] K. J. Fidkowski, M. A. Ceze, and P. L. Roe. Entropy-based drag error estimation and mesh adaptation in two dimensions. *AIAA Journal of Aircraft*, 49(5):1485–1496, September-October 2012.
- [31] K. J. Fidkowski and D. L. Darmofal. A triangular cut-cell adaptive method for high-order discretizations of the compressible Navier-Stokes equations. *Journal of Computational Physics*, 225:1653–1672, 2007.
- [32] Krzysztof J. Fidkowski. *A Simplex Cut-Cell Adaptive Method for High-order Discretizations of the Compressible Navier-Stokes Equations*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2007.
- [33] Krzysztof J. Fidkowski. Output error estimation strategies for discontinuous Galerkin discretizations of unsteady convection-dominated flows. *International Journal for Numerical Methods in Engineering*, 88(12):1297–1322, 2011.
- [34] Krzysztof J. Fidkowski. An output-based dynamic order refinement strategy for unsteady aerodynamics. AIAA Paper 2012-77, 2012.
- [35] Krzysztof J. Fidkowski and David L. Darmofal. An adaptive simplex cut-cell method for discontinuous Galerkin discretizations of the Navier-Stokes equations. AIAA Paper 2007-3941, 2007.
- [36] Krzysztof J. Fidkowski and David L. Darmofal. Review of output-based error estimation and mesh adaptation in computational fluid dynamics. *American Institute of Aeronautics and Astronautics Journal*, 49(4):673–694, 2011.
- [37] Krzysztof J. Fidkowski and Yuxing Luo. Output-based space-time mesh adaptation for the compressible Navier-Stokes equations. *Journal of Computational Physics*, 230:5753–5773, 2011.

- [38] Krzysztof J. Fidkowski, Todd A. Oliver, James Lu, and David L. Darmofal. p -Multigrid solution of high-order discontinuous Galerkin discretizations of the compressible Navier-Stokes equations. *Journal of Computational Physics*, 207:92–113, 2005.
- [39] Krzysztof J. Fidkowski and Philip L. Roe. An entropy adjoint approach to mesh refinement. *SIAM Journal on Scientific Computing*, 32(3):1261–1287, 2010.
- [40] L. Formaggia, S. Micheletti, and S. Perotto. Anisotropic mesh adaptation with applications to CFD problems. In H. A. Mang, F. G. Rammerstorfer, and J. Eberhardsteiner, editors, *Fifth World Congress on Computational Mechanics*, Vienna, Austria, July 7-12 2002.
- [41] Luca Formaggia and Simona Perotto. New anisotropic a priori error estimates. *Numerische Mathematik*, 89(4):641–667, 2001.
- [42] Luca Formaggia, Simona Perotto, and Paolo Zunino. An anisotropic a posteriori error estimate for a convection-diffusion problem. *Computing and Visualization in Science*, 4:99–104, 2001.
- [43] Lori A. Freitag and Carl Ollivier-Gooch. Tetrahedral mesh improvement using swapping and smoothing. *International Journal for Numerical Methods in Engineering*, 40:3979–4002, 1997.
- [44] Neal T. Frink. 3rd AIAA CFD drag prediction workshop gridding guidelines. NASA Langley, 2007. http://aaac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw/Workshop3/gridding_guidelines.html.
- [45] Neal T. Frink. Test case results from the 3rd AIAA drag prediction workshop. NASA Langley, 2007. http://aaac.larc.nasa.gov/tsab/cfdlarc/aiaa-dpw/Workshop3/final_results_jm.tar.gz.
- [46] M. B. Giles and N. A. Pierce. Adjoint equations in CFD: duality, boundary conditions and solution behavior. AIAA Paper 97-1850, 1997.
- [47] M. B. Giles and E. Süli. Adjoint methods for PDEs: a posteriori error analysis and postprocessing by duality. In *Acta Numerica*, volume 11, pages 145–236, 2002.
- [48] M.B. Giles and N.A. Pierce. Analytic adjoint solutions for the quasi-one-dimensional Euler equations. *Journal of Fluid Mechanics*, 426:327–345, 2001.
- [49] Andreas Griewank and Andrea Walther. Revolve: An implementation of checkpointing for the reverse or adjoint mode of computational differentiation. *ACM Transactions on Mathematical Software*, 26(1):19–45, 2000.
- [50] W. Gui and I Babuska. The h , p , and $h - p$ versions of the finite element method in 1 dimension. Part III. the adaptive $h - p$ version. *Numerische Mathematik*, 49(6):659–683, 1986.

- [51] Wagdi G. Habashi, Julien Dompierre, Yves Bourgault, Djaffar Ait-Ali-Yahia, Michel Fortin, and Marie-Gabrielle Vallet. Anisotropic mesh adaptation: towards user-independent, mesh-independent and solver-independent CFD. Part I: general principles. *International Journal for Numerical Methods in Fluids*, 32:725–744, 2000.
- [52] K. Harriman, D. Gavaghan, and E. Süli. The importance of adjoint consistency in the approximation of linear functionals using the discontinuous Galerkin finite element method. Technical Report Technical Report NA 04/18, Oxford University Computing Lab Numerical Analysis Group, 2004.
- [53] K. Harriman, P. Houston, B. Senior, and Endre Süli. hp-version discontinuous Galerkin methods with interior penalty for partial differential equations with non-negative characteristic form. Technical Report Technical Report NA 02/21, Oxford University Computing Lab Numerical Analysis Group, 2002.
- [54] R. Hartmann and P. Houston. Goal-oriented a posteriori error estimation for multiple target functionals. In T.Y. Hou and E. Tadmor, editors, *Hyperbolic Problems: Theory, Numerics, Applications*, pages 579–588. Springer-Verlag, 2003.
- [55] R. Hartmann and P. Houston. Error estimation and adaptive mesh refinement for aerodynamic flows. In H. Deconinck, editor, *36th CFD/ADIGMA course on hp-adaptive and hp-multigrid methods: VKI Lecture Series 2010-01 (Oct. 26-30, 2009)*. von Karman Institute for Fluid Dynamics, 2010.
- [56] Ralf Hartmann. Adjoint consistency analysis of discontinuous Galerkin discretizations. *SIAM Journal on Numerical Analysis*, 45(6):2671–2696, 2007.
- [57] Ralf Hartmann and Paul Houston. Adaptive discontinuous Galerkin finite element methods for the compressible Euler equations. *Journal of Computational Physics*, 183(2):508–532, 2002.
- [58] F. Hecht. BAMG: Bidimensional anisotropic mesh generator. INRIA–Rocquencourt, France, 1998. <http://pauillac.inria.fr/cdrom/www/bamg/bamg/eng.htm>.
- [59] V. Heuveline and R. Rannacher. Duality-based adaptivity in the hp-finite element method. *Journal of Numerical Mathematics*, 11(2):95–113, 2003.
- [60] P. Houston, R. Hartmann, and E. Süli. Adaptive discontinuous Galerkin finite element methods for compressible fluid flows. In M. Baines, editor, *Numerical Methods for Fluid Dynamics VII, ICFD*, pages 347–353, 2001.
- [61] P. Houston, B. Senior, and E. Süli. Sobolev regularity estimation for hp-adaptive finite element methods. In F. Brezzi, A. Buffa, S. Corsaro, and A. Murli, editors, *Numerical Mathematics and Advanced Applications*, pages 619–644. Springer-Verlag, 2003.
- [62] P. Houston and E. Süli. hp-adaptive discontinuous Galerkin finite element methods for first-order hyperbolic problems. *SIAM Journal on Scientific Computing*, 23(4):1226–1252, 2001.

- [63] Paul Houston, Emmanuil H. Georgoulis, and Edward Hall. Adaptivity and a posteriori error estimation for DG methods on anisotropic meshes. In G. Lube and G. Rapin, editors, *Proceedings of the International Conference on Boundary and Interior Layers (BAIL)*. University of Göttingen, 2006.
- [64] Paul Houston, Bill Senior, and Endre Süli. hp-Discontinuous Galerkin finite element methods for hyperbolic problems: Error analysis and adaptivity. *International Journal for Numerical Methods in Fluids*, 40:153–169, 2002.
- [65] Paul Houston and Endre Süli. A note on the design of hp-adaptive finite element methods for elliptic partial differential equations. *Computer Methods in Applied Mechanics and Engineering*, 194:229–243, 2005.
- [66] Hans Johansen and Philip Colella. A Cartesian grid embedded boundary method for Poisson’s equation on irregular domains. *Journal of Computational Physics*, 147:60–85, 1998.
- [67] Steven M. Kast and Krzysztof J. Fidkowski. Output-based mesh adaptation for high order Navier-Stokes simulations on deformable domains. *Journal of Computational Physics*, 252(1):468–494, 2013.
- [68] C.M. Klaij, J.J.W. van der Vegt, and H. van der Ven. Space-time discontinuous Galerkin method for the compressible Navier-Stokes equations. *Journal of Computational Physics*, 217:589–611, 2006.
- [69] Tobias Leicht and Ralf Hartmann. Multitarget error estimation and adaptivity in aerodynamic flow simulations. *International Journal for Numerical Methods in Fluids*, 56:2111–2138, 2008.
- [70] R.W. Logan and C.K. Nitta. Comparing 10 methods for solution verification and linking to model validation. *Journal of Aerospace Computing, Information, and Communication*, 3:354–373, 2006.
- [71] Robert B. Lowrie, Philip L. Roe, and Bram van Leer. Properties of space-time discontinuous Galerkin. Los Alamos Technical Report LA-UR-98-5561, 1998.
- [72] James Lu. *An a Posteriori Error Control Framework for Adaptive Precision Optimization Using Discontinuous Galerkin Finite Element Method*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2005.
- [73] Karthik Mani and Dimitri J. Mavriplis. Discrete adjoint based time-step adaptation and error reduction in unsteady flow problems. AIAA Paper 2007-3944, 2007.
- [74] Karthik Mani and Dimitri J. Mavriplis. Error estimation and adaptation for functional outputs in time-dependent flow problems. *Journal of Computational Physics*, 229:415–440, 2010.
- [75] D. J. Mavriplis. Adaptive mesh generation for viscous flows using Delaunay triangulation. *Journal of Computational Physics*, 90:271–291, 1990.

- [76] D. J. Mavriplis. Results from the 3rd drag prediction workshop using the NSU3D unstructured mesh solver. AIAA Paper 2007-256, 2007.
- [77] Dimitri J. Mavriplis, John C. Vassberg, Edward N. Tinoco, Mori Mani, Olaf P. Brodersen, Bernhard Eisfeld, Richard A. Wahls, Joseph H. Morrison, Tom Zickuhr, David Levy, and Mitsuhiro Murayama. Grid quality and resolution issues from the drag prediction workshop series. AIAA Paper 2008-930, 2008.
- [78] D. Scott McRae. r-Refinement grid adaptation algorithms and issues. *Computer Methods in Applied Mechanics and Engineering*, 2000(4):1161–1182, 189.
- [79] D. Meidner and B. Vexler. Adaptive space-time finite element methods for parabolic optimization problems. *SIAM Journal on Control Optimization*, 46(1):116–142, 2007.
- [80] Peter K. Moore. Applications of lobatto polynomials to an adaptive finite element method: A posteriori error estimates for hp-adaptivity and grid-to-grid interpolation. *Numerische Mathematik*, 94:367–401, 2003.
- [81] Joseph H. Morrison and Michael J. Hemsch. Statistical analysis of CFD solutions from the third AIAA drag prediction workshop. AIAA Paper 2007-254, 2007.
- [82] Marian Nemec and Michael J. Aftosmis. Error estimation and adaptive refinement for embedded-boundary Cartesian meshes. AIAA Paper 2007-4187, 2007.
- [83] Marian Nemec, Michael J. Aftosmis, and Mathias Wintzer. Adjoint-based adaptive mesh refinement for complex geometries. AIAA Paper 2008-0725, 2008.
- [84] William L. Oberkampf and Matthew F. Barone. Measures of agreement between computation and experiment: Validation metrics. *Journal of Computational Physics*, 217:5–36, 2006.
- [85] Todd A. Oliver. *A High-order, Adaptive, Discontinuous Galerkin Finite Element Method for the Reynolds-Averaged Navier-Stokes Equations*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2008.
- [86] Todd A. Oliver and David L. Darmofal. Impact of turbulence model irregularity on high-order discretizations. AIAA Paper 2009-953, 2009.
- [87] Doug Pagnutti and Carl Ollivier-Gooch. A generalized framework for high order anisotropic mesh adaptation. *Computers and Structures*, 87(11-12):670 – 679, 2009.
- [88] M. A. Park. Adjoint-based, three-dimensional error prediction and grid adaptation. AIAA Paper 2002-3286, 2002.
- [89] M. A. Park. Three-dimensional turbulent RANS adjoint-based error correction. AIAA Paper 2003-3849, 2003.
- [90] Michael A. Park. *Anisotropic Output-Based Adaptation with Tetrahedral Cut Cells for Compressible Flows*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2008.

-
- [91] Xavier Pennec, Pierre Fillard, and Nicholas Ayache. A riemannian framework for tensor computing. *International Journal of Computer Vision*, 66(1):41–66, 2006.
- [92] J. Peraire and P.-O. Persson. The compact discontinuous Galerkin (CDG) method for elliptic problems. *SIAM Journal on Scientific Computing*, 2007.
- [93] J. Peraire, M. Vahdati, K. Morgan, and O. C. Zienkiewicz. Adaptive remeshing for compressible flow computations. *Journal of Computational Physics*, 72:449–466, 1987.
- [94] P.-O. Persson, J. Bonet, and J. Peraire. Discontinuous Galerkin solution of the Navier-Stokes equations on deformable domains. *Computer Methods in Applied Mechanics and Engineering*, 198:1585–1595, 2009.
- [95] P.-O. Persson and J. Peraire. Sub-cell shock capturing for discontinuous Galerkin methods. AIAA Paper 2006-112, 2006.
- [96] Niles A. Pierce and Michael B. Giles. Adjoint recovery of superconvergent functionals from PDE approximations. *SIAM Review*, 42(2):247–264, 2000.
- [97] W. Rachowicz, L. Demkowicz, and J.T. Oden. Toward a universal $h - p$ adaptive finite element strategy, part 3. design of $h - p$ meshes. *Computer Methods in Applied Mechanics and Engineering*, 77:181–212, 1989.
- [98] R. Rannacher. Adaptive Galerkin finite element methods for partial differential equations. *Journal of Computational and Applied Mathematics*, 128:205–233, 2001.
- [99] Thomas Richter. A posteriori error estimation and anisotropy detection with the dual-weighted residual method. *International Journal for Numerical Methods in Fluids*, 62:90–118, 2010.
- [100] Patrick J. Roache. Error bars for CFD. AIAA Paper 2003-408, 2003.
- [101] P. L. Roe. Approximate Riemann solvers, parameter vectors, and difference schemes. *Journal of Computational Physics*, 43:357–372, 1981.
- [102] Philip Roe and Hiroaki Nishikawa. Adaptive grid generation by minimizing residuals. *International Journal for Numerical Methods in Fluids*, 40:121–136, 2002.
- [103] Christopher J. Roy. Strategies for driving mesh adaptation in CFD. AIAA Paper 2009-1302, 2009.
- [104] Christopher J. Roy and Frederick G. Blottner. Review and assessment of turbulence models for hypersonic flows. *Progress in Aerospace Sciences*, 42:469–530, 2006.
- [105] Youcef Saad. A flexible inner-outer preconditioned GMRES algorithm. *SIAM Journal on Scientific Computing*, 14(2):461–469, 1993.
- [106] Youcef Saad and Martin H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific Computing*, 7(3):856–869, 1986.

-
- [107] E. Schall, D. Leservoisier, A. Dervieux, and B. Koobus. Mesh adaptation as a tool for certified computational aerodynamics. *International Journal for Numerical Methods in Fluids*, 45(2):179–196, 2004.
- [108] Michael Schmich and Boris Vexler. Adaptivity with dynamic meshes for space-time finite element discretizations of parabolic equations. *SIAM Journal on Scientific Computing*, 30(1):369–393, 2008.
- [109] R. Schneider and P. K. Jimack. Toward anisotropic mesh adaptation based upon sensitivity of a posteriori estimates. Technical Report 2005.03, University of Leeds, School of Computing, 2005.
- [110] J.R. Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry: Theory and Applications*, 22:21–74, 2002.
- [111] Kunibert G. Siebert. An a posteriori error estimator for anisotropic refinement. *Numerische Mathematik*, 73:373–398, 1996.
- [112] P. Solín and L. Demkowicz. Goal-oriented hp-adaptivity for elliptic problems. *Computer Methods in Applied Mechanics and Engineering*, 193:449–468, 2004.
- [113] Thomas Sonar. Strong and weak norm refinement indicators based on the finite element residual for compressible flow computation. *Impact of Computing in Science and Engineering*, 5:111–127, 1993.
- [114] Shuyu Sun and Mary Wheeler. Mesh adaptation strategies for discontinuous Galerkin methods applied to reactive transport problems. In H.W. Chu, M. Savoie, and B. Sanchez, editors, *International Conference on Computing, Communication and Control Technologies*, volume 1, pages 223–228, Austin, Texas, August 2004.
- [115] Shuyu Sun and Mary F. Wheeler. Anisotropic and dynamic mesh adaptation for discontinuous Galerkin methods applied to reactive transport. Technical Report 05-15, ICES, 2005.
- [116] Barna A. Szabo. Estimation and control of error based on p convergence. In I. Babuska, O. C. Zienkiewicz, J. Gago, and E. R. de Oliveira, editors, *Accuracy Estimates and Adaptive Refinements in Finite Element Computations*, pages 61–78. John Wiley & Sons Ltd., 1986.
- [117] J.J.W. van der Vegt. Space-time discontinuous Galerkin finite element method with dynamic grid motion for inviscid compressible flows. National Aerospace Laboratory NLR-TP-98239, 1998.
- [118] H. van der Ven and J.J.W. van der Vegt. Space-time discontinuous Galerkin finite element method with dynamic grid motion for inviscid compressible flows II. Efficient flux quadrature. *Computer Methods in Applied Mechanics and Engineering*, 191:4747–4780, 2002.
- [119] D. A. Venditti and D. L. Darmofal. Anisotropic grid adaptation for functional outputs: application to two-dimensional viscous flows. *Journal of Computational Physics*, 187(1):22–46, 2003.

-
- [120] V. Venkatakrishnan, S. R. Allmaras, D. S. Kamenetskii, and F. T. Johnson. Higher order schemes for the compressible Navier-Stokes equations. AIAA Paper 2003-3987, 2003.
- [121] G. P. Warren, W. K. Anderson, J. L. Thomas, and S. L. Krist. Grid convergence for adaptive methods. AIAA Paper 1991-1592, 1991.
- [122] William A. Wood and William L. Kleb. On multi-dimensional unstructured mesh adaptation. AIAA Paper 99-3254, 1999.
- [123] Guoping Xia, Ding Li, and Charles L. Merkle. Anisotropic grid adaptation on unstructured meshes. AIAA Paper 2001-0443, 2001.
- [124] Nail K. Yamaleev, Boris Diskin, and Eric J. Nielsen. Local-in-time adjoint-based method for design optimization of unsteady flows. *Journal of Computational Physics*, 229:5394–5407, 2010.
- [125] M. Yano, J.M. Modisette, and D.L. Darmofal. The importance of mesh adaptation for higher-order discretizations of aerodynamics flows. AIAA Paper 2011-3852, 2011.
- [126] Masayuki Yano. *An Optimization Framework for Adaptive Higher-Order Discretizations of Partial Differential Equations on Anisotropic Simplex Meshes*. PhD thesis, Massachusetts Institute of Technology, Cambridge, Massachusetts, 2012.

A Compressible Navier-Stokes Equations

This appendix presents the Euler, compressible Navier-Stokes, and Reynolds-averaged compressible Navier-Stokes equations.

A.1 Euler Equations

The Euler equations of gas dynamics are

$$\begin{aligned} \partial_t \rho &+ \partial_i(\rho u_i) &= 0 \\ \partial_t(\rho u_j) &+ \partial_i(\rho u_i u_j + p \delta_{ij}) &= 0 \\ \partial_t(\rho E) &+ \partial_i(\rho u_i H) &= 0 \end{aligned}$$

where ρ is the density ρu_j is the j^{th} momentum component, and ρE is the total energy. The pressure and total enthalpy are given by

$$\begin{aligned} p &= (\gamma - 1) \left(\rho E - \frac{1}{2} \rho u_k u_k \right), \\ H &= E + \frac{p}{\rho}. \end{aligned}$$

Note, i, j, k index the spatial dimensions and summation is implied over repeated indices. These equations can be written in compact conservation form as

$$\partial_t \mathbf{u} + \partial_i \mathbf{F}_i = \mathbf{0},$$

where

$$\mathbf{u} = \begin{bmatrix} \rho \\ \rho u_j \\ \rho E \end{bmatrix}, \quad \mathbf{F}_i = \begin{bmatrix} \rho u_i \\ \rho u_i u_j + p \delta_{ij} \\ \rho u_i H \end{bmatrix}.$$

A.2 Compressible Navier-Stokes

The compressible Navier-Stokes equations are given by

$$\begin{aligned} \partial_t \rho &+ \partial_i(\rho u_i) &= 0 \\ \partial_t(\rho u_i) &+ \partial_i(\rho u_i u_j + p \delta_{ij}) &= \partial_i \tau_{ij} \\ \partial_t(\rho E) &+ \partial_i(\rho u_i H) &= \partial_i(\tau_{ij} u_j - q_i) \end{aligned}$$

where the viscous stress tensor and the heat flux vector are

$$\begin{aligned} \tau_{ij} &= \mu(\partial_i u_j + \partial_j u_i) + \delta_{ij} \lambda \partial_m u_m, \\ q_i &= -\kappa_T \partial_i T. \end{aligned}$$

Relevant physical quantities for air are,

$$\begin{aligned}
 \text{Dynamic viscosity: } \mu &= \mu_{\text{ref}} \left(\frac{T}{T_{\text{ref}}} \right)^{1.5} \left(\frac{T_{\text{ref}} + T_s}{T + T_s} \right), \\
 &\quad (\text{Sutherland's law: } T_{\text{ref}} = 288.15\text{K}, T_s = 110\text{K}) \\
 \text{Bulk viscosity coefficient: } \lambda &= -\frac{2}{3}\mu, \\
 \text{Kinematic viscosity: } \nu &= \frac{\mu}{\rho}, \\
 \text{Thermal conductivity: } \kappa_T &= \frac{\gamma\mu R}{(\gamma - 1)Pr}, \\
 \text{Specific-heat ratio: } \gamma &= 1.4, \\
 \text{Prandtl number: } Pr &= 0.71, \\
 \text{Gas constant: } R &.
 \end{aligned}$$

To clarify the viscous implementation, we make the dependence of the viscous flux on the gradient of the state vector explicit by writing the equations in the form

$$\partial_t \mathbf{u} + \partial_i \mathbf{F}_i - \partial_i (\mathbf{K}_{ij} \partial_j \mathbf{u}) = \mathbf{0},$$

where $\mathbf{u} = [\rho, \rho u_j, \rho E]^T$ is the state vector and $\vec{\mathbf{F}}$ is the inviscid flux vector. The viscous flux coefficient matrices can be written as

$$\mathbf{K}_{ij} = \begin{bmatrix} 0 & \mathbf{0} & 0 \\ K_{ij}^{\rho u_d, \rho} & K_{ij}^{\rho u_d, \rho u_c} & \mathbf{0} \\ K_{ij}^{\rho E, \rho} & K_{ij}^{\rho E, \rho u_c} & K^{\rho E, \rho E} \end{bmatrix},$$

where for a given pair i, j ,

$$\begin{aligned}
 K_{ij}^{\rho u_d, \rho} &= (\text{dim} \times 1) \text{ matrix} = \nu \left(u_d \delta_{ij} - u_i \delta_{jd} + \frac{2}{3} u_j \delta_{id} \right) \\
 K_{ij}^{\rho u_d, \rho u_c} &= (\text{dim} \times \text{dim}) \text{ matrix} = \nu \left(\delta_{dc} \delta_{ij} + \delta_{dj} \delta_{ci} - \frac{2}{3} \delta_{di} \delta_{cj} \right) \\
 K_{ij}^{\rho E, \rho} &= (1 \times 1) \text{ matrix} = \nu \left(-u_k u_k \delta_{ij} - \frac{1}{3} u_i u_j + \kappa_T T_\rho \right) \\
 K_{ij}^{\rho E, \rho u_c} &= (1 \times \text{dim}) \text{ matrix} = \nu \left(-\frac{2}{3} u_i \delta_{cj} + u_j \delta_{ci} + u_c \delta_{ij} + \kappa_T T_{\rho u_c} \right) \\
 K_{ij}^{\rho E, \rho E} &= (1 \times 1) \text{ matrix} = \nu (\kappa_T T_\rho)
 \end{aligned}$$

and the temperature derivative vector is

$$T_{\mathbf{u}} = [T_\rho, T_{\rho u_c}, T_{\rho E}] = \frac{\gamma - 1}{R\rho} [-E + u_k u_k, -u_c, 1].$$

Note that the indices d, c, k index the spatial dimension.

A.3 Reynolds-Averaged Compressible Navier-Stokes

We use the Spalart-Allmaras one-equation turbulence model to close the Reynolds-averaged Navier-Stokes (RANS) equations. The augmented equations are

$$\begin{aligned} \partial_t \rho + \partial_i(\rho u_i) &= 0 \\ \partial_t(\rho u_i) + \partial_i(\rho u_i u_j + p \delta_{ij}) &= \partial_i \tau_{ij} \\ \partial_t(\rho E) + \partial_i(\rho u_i H) &= \partial_i(\tau_{ij} u_j - q_i) \\ \partial_t(\rho \tilde{\nu}) + \partial_i(\rho u_i \tilde{\nu}) &= \partial_i \left(\frac{\eta}{\sigma} \partial_i \tilde{\nu} \right) + \frac{c_{b2} \rho}{\sigma} \partial_i \tilde{\nu} \partial_i \tilde{\nu} + P - D \end{aligned}$$

where the Reynolds stress, τ_{ij} , is

$$\tau_{ij} = 2(\mu + \mu_t) \epsilon_{ij}, \quad \epsilon_{ij} = \frac{1}{2}(\partial_i u_j + \partial_j u_i) - \frac{1}{3} \partial_k u_k \delta_{ij}.$$

μ is the laminar dynamic viscosity and the eddy viscosity, μ_t , is

$$\mu_t = \begin{cases} \rho \tilde{\nu} f_{v1} & \tilde{\nu} \geq 0 \\ 0 & \tilde{\nu} < 0 \end{cases}, \quad f_{v1} = \frac{\chi^3}{\chi^3 + c_{v1}^3}, \quad \chi \equiv \frac{\tilde{\nu}}{\nu}.$$

The heat flux, q_i , is given by

$$q_i = (k + k_t) \partial_i T, \quad k = C_p \mu / Pr, \quad k_t = C_p \mu_t / Pr_t.$$

The diffusivity for the SA working variable, $\tilde{\nu}$, is given by η/σ , where

$$\eta = \begin{cases} \mu(1 + \chi) & \chi \geq 0 \\ \mu(1 + \chi + \chi^2) & \chi < 0 \end{cases}$$

The production term, P , is

$$P = \begin{cases} c_{b1} \tilde{S} \rho \tilde{\nu} & \chi \geq 0 \\ c_{b1} S \rho \tilde{\nu} g_n & \chi < 0 \end{cases}, \quad \tilde{S} = \begin{cases} S + \bar{S} & \bar{S} \geq -c_{v2} S \\ S + \frac{S(c_{v2}^2 S + c_{v3} \bar{S})}{(c_{v3} - 2c_{v2})S - \bar{S}} & \bar{S} < -c_{v2} S \end{cases},$$

where $S = \sqrt{2\Omega_{ij}\Omega_{ij}}$ is the vorticity magnitude and

$$g_n = 1 - \frac{1000\chi^2}{1 + \chi^2}$$

Also,

$$\bar{S} = \frac{\tilde{\nu} f_{v2}}{\kappa^2 d^2}, \quad f_{v2} = 1 - \frac{\chi}{1 + \chi f_{v1}}.$$

The destruction term, D , is given by

$$D = \begin{cases} c_{w1} f_w \frac{\rho \tilde{\nu}^2}{d^2} & \chi \geq 0 \\ -c_{w1} \frac{\rho \tilde{\nu}^2}{d^2} & \chi < 0 \end{cases}$$

where

$$f_w = g \left(\frac{1 + c_{w3}^6}{g^6 + c_{w3}^6} \right)^{1/6}, \quad g = r + c_{w2}(r^6 - r), \quad r = \frac{\tilde{\nu}}{\tilde{S}\kappa^2 d^2}.$$

The closure coefficients are

$$\begin{aligned} c_{b1} &= 0.1355 & c_{w2} &= 0.3 \\ \sigma &= 2/3 & c_{w3} &= 2 \\ c_{b2} &= 0.622 & c_{v1} &= 7.1 \\ \kappa &= 0.41 & c_{v2} &= 0.7 \\ c_{w1} &= \frac{c_{b1}}{\kappa^2} + \frac{1 + c_{b2}}{\sigma} & c_{v3} &= 0.9 \\ Pr_t &= 0.9 \end{aligned}$$

In discretizing the $\tilde{\nu}$ viscous term, we have to be aware that we store $\rho\tilde{\nu}$, so that

$$F_{\tilde{\nu}i}^v = \frac{\eta}{\sigma} \partial_i \tilde{\nu} = -\frac{\eta\tilde{\nu}}{\rho\sigma} \partial_i \rho + \frac{\eta}{\rho\sigma} \partial_i (\rho\tilde{\nu}),$$

from which the contributions to the diffusion matrix can be read off directly.

The SA working variable, $\tilde{\nu}$, will generally be orders of magnitude smaller than the other state components. Scaling or “non-dimensionalization” of $\tilde{\nu}$ is used to make the stored values of $\tilde{\nu}$ closer to the other state components, which helps during the discrete linear solves and when using finite residual convergence tolerances. Instead of $\tilde{\nu}$ we store $\tilde{\nu}'$, given by

$$\tilde{\nu}' = \frac{\tilde{\nu}}{ND},$$

where ND is a scaling factor, usually on the order of $\nu\kappa_{SA}$, where ν is the laminar kinematic viscosity. κ_{SA} is a user-prescribed factor, typically 100 or 1000, that makes $\tilde{\nu}'$ on the order of unity. The associated changes in the above expressions are obtained by re-writing all of them in terms of $\tilde{\nu}'$. In addition, the SA equation is divided by ND ,

$$\partial_t(\rho\tilde{\nu}') + \partial_i(\rho u_i \tilde{\nu}') = \partial_i \left(\frac{\eta}{\sigma} \partial_i \tilde{\nu}' \right) + ND \frac{c_{b2}\rho}{\sigma} \partial_i \tilde{\nu}' \partial_i \tilde{\nu}' + \frac{P - D}{ND}.$$

This has the effect of bringing the SA equation residual to the same order as that of the other equations.

B Unsteady Adaptive Strategy

This appendix describes the merit-function-based adaptive strategy used to adapt the spatial and temporal discretizations in an unsteady simulation. This strategy is geared for a slab-based temporal discretization (e.g. DG-in-time), where the size of the time slabs can vary in time, and for dynamic hanging-node or order refinement in space.

B.1 Greedy Fixed Growth Selection Algorithm

The error-addressed versus cost-added figure of merit is used in a fixed-growth adaptive strategy in which some combination of time slabs and space-time elements are marked for coarsening or refinement. The user specifies the fraction of degrees of freedom to be coarsened (f^{coarsen}), as well as the growth factor f^{growth} , which dictates the total number of degrees of freedom in the next mesh as $D^{\text{next}} = f^{\text{growth}} D^{\text{current}}$ (where D^{current} is the number of current degrees of freedom). The coarsening and refinement budgets are then

$$\begin{aligned} B^{\text{coarsen}} &= f^{\text{coarsen}} D^{\text{current}}, \\ B^{\text{refine}} &= (f^{\text{growth}} - 1) D^{\text{current}} + B^{\text{coarsen}}. \end{aligned}$$

In practice, we typically use a coarsening fraction of $\sim 5\%$ ($f^{\text{coarsen}} = 0.05$) and a growth factor of 1.30-1.35. With these budgets defined, the following algorithm is then used to decide which space-time elements or time slabs to adapt:

1. Sort

Sort all space-time elements and time slabs based on the figure of merit: the amount of output error addressed divided by the degrees of freedom added if the element/slab were to be refined. For temporal refinement, the latter is approximated as the degrees of freedom in the targeted slab k , $\text{dof}_k \equiv \sum_e \text{dof}(p_{e,k})$, and for spatial refinement as the number of additional degrees of freedom $\text{dof}(p_{e,k} + 1) - \text{dof}(p_{e,k})$ associated with an order increase of element e, k .

2. Coarsen

Set coarsening degree-of-freedom tally to zero. Choose an unmarked space-time element or time slab with the lowest merit function. If a time slab was chosen, mark it for a factor of 2 coarsening and add 0.5dof_k to the coarsening tally. If a space-time element was chosen, mark it for an order decrement and add $\text{dof}(p_{e,k}) - \text{dof}(p_{e,k}-1)$ to the coarsening tally. If the tally meets or exceeds the coarsening budget, B^{coarsen} , stop. Otherwise choose the next space-time element or time slab and repeat.

3. Refine

Set refinement degree-of-freedom tally to zero. Choose an unmarked space-time element or time slab with the highest merit function. If a time slab was chosen, mark it for a factor of 0.5 refinement and add dof_k to the refinement tally. If a space-time element was chosen, mark it for an order increment and add $\text{dof}(p_{e,k}+1) - \text{dof}(p_{e,k})$ to the refinement tally. If the refinement budget, B^{refine} , is met or exceeded, stop. Else, choose the next element or time slab.

B.2 Temporal-Mesh Optimization

When we say ‘‘factor of 2 coarsening’’ and ‘‘factor of 0.5 refinement’’, this essentially means that the width of a slab marked for coarsening will be doubled, while the width of a slab marked for refinement will be halved. If only refinement occurs, this is exactly the case, as each marked slab will be perfectly bisected; however, when coarsening also occurs, the

boundaries of a coarsened slab will typically encroach on those of the neighboring slabs, and the entire temporal grid must be shuffled to make room for the coarsened slab. We shuffle time slabs using one-dimensional metric-based meshing, as described below.

1. For each time slab k , define $\widetilde{\Delta t}_k^{\text{desired}} = c\Delta t_k^{\text{current}}$ where c is 0.5 for refinement, 2 for coarsening, and 1 if the time slab is not marked. $\widetilde{\Delta t}_k^{\text{desired}}$ represents the new time step size that we desire on the current time slab k . The given choices of c are consistent with the choices made for the degree-of-freedom counts above.
2. Define $N^{\text{desired}} = \lceil \sum_k (\Delta t_k^{\text{current}} / \widetilde{\Delta t}_k^{\text{desired}}) \rceil$, where $\lceil \cdot \rceil$ is the greatest integer (ceiling) function. This will be the total number of time slabs on the new temporal mesh. Note that the ceiling function ensures that this number is an integer.
3. To ensure that the desired time step size is consistent with this total number of time steps, define the new desired time step size as

$$\Delta t_k^{\text{desired}} = \widetilde{\Delta t}_k^{\text{desired}} \frac{\sum_k (\Delta t_k^{\text{current}} / \widetilde{\Delta t}_k^{\text{desired}})}{N^{\text{desired}}}.$$

4. Finally, define a function $n(t)$ that is piecewise-constant over the current time slabs, and that takes on the value $1/\Delta t_k^{\text{desired}}$ on each current slab k . Define the new time slab breakpoints as times t_l where $\int_0^{t_l} n(t)dt$ is an integer.




Figure 92: Demonstration of the one-dimensional remeshing algorithm used to define new time slab breakpoints.