An *a posteriori* Error Control Framework for Adaptive Precision Optimization using Discontinuous Galerkin Finite Element Method

by

James Ching-Chieh Lu

Submitted to the Department of Aeronautics and Astronautics
in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in Aeronautics and Astronautics
at the
MASSACHUSETTS INSTITUTE OF TECHNOLOGY
September 2005

Author ....
Department of Aeronautics and Astronautics
June 2, 2005

Certified by......
David L. Darmofal
Associate Professor
Thesis Supervisor

Certified by..
Jaime Peraire
Professor
Thesis Committee

Certified by.........
Karen E. Willcox
Assistant Professor
Thesis Committee

Accepted by ....
Jaime Peraire
Professor of Aeronautics and Astronautics
Chair, Committee on Graduate Students

AERO

Acknowledgments

Contents

List of Figures

10

11

12

List of Tables

Chapter 1

Introduction

Aerodynamic design optimization has seen significant development over the past decade. Adjoint-based shape design for elliptic systems was first proposed by Pironneau [69] and applied to transonic flow by Jameson [49]. A review of the aerodynamic shape optimization literature and a large list of references is given in [46]. Over the years much technology has been developed, allowing engineers to contemplate applying optimization methods to a wide variety of problems. In the context of structured grids, adjoint-based applications include multipoint, multi-objective airfoil design using compressible Navier-Stokes equations [64] and 3D multipoint design of aircraft configurations using inviscid Euler equations [75, 76]. There have also been significant effort in applying adjoint methods to the unstructured grid setting. In this context, Newman *et al.* [47, 45], Elliot and Peraire [21, 22] were among the first to develop discrete adjoint approaches for the inviscid Euler equations. The work of Elliot and Peraire was also extended to include laminar viscous effects [23]. For 2- and 3D turbulent flows respectively, Anderson and Bonhaus [4], Nielsen and Anderson [66] have developed discrete adjoint implementations for the one-equation turbulence model of Spalart-Allmaras. In [5], Anderson and Venkatakrishnan developed a continuous adjoint approach using unstructured grids. The reverse mode of automatic differentiation has also been applied to both inviscid and Navier-Stokes equations with a two-equation $k - \epsilon$ turbulence model [60].

Despite these significant advances in the development of computational implemen-

tations, there are still obstacles that stand in the way of automatic design methods being widely accepted and applied in the engineering community. In particular, an outstanding issue is the question of reliability of the discrete computational models and its impact on the resulting designs. Lack of trust in the computed results may lead to decreased acceptance and utility of automatic design tools in the engineering community, or in an attempt to minimize uncertainty in the computational results, the designer may use unnecessarily refined large-scale models for which the computational costs quickly become prohibitive.

Currently, the predominant approach towards the optimization of continuous systems is to apply a general nonlinear programming algorithm to discrete models that are at a precision that is fixed prior to optimization. Thus, the chosen algorithm attempts to attain the best possible performance for the discrete model but is not aware of the underlying continuous system of interest. An alternative approach described in this thesis is to adaptively control the precision of the discrete model during the optimization. In particular, we propose a method which can ensure that: (1) at each step of the optimization, the objective function for the underlying continuous system is improved; (2) stationary points of the continuous system can be approached to arbitrary accuracy given enough iterations of the optimization algorithm.

1.1 Objective

The objective of this work is, firstly to develop a framework to increase designer confidence in simulation-based design and, secondly to demonstrate the feasibility of the framework in the context of aerodynamic design.

1.2 Review of related prior work

A major source of model uncertainty arises from the use of coarse discretizations and incomplete solution iteration. To ensure the reliability of the design changes obtained by optimization algorithms based on these approximations, it is necessary

18

to accurately estimate the error contributions and effect a mechanism for control. The use of model precision adjustments (or variable fidelity) in optimization has been previously proposed. However, in contrast to our objective these prior efforts were largely driven by the desire to decrease the significant computational effort required to perform optimization, typically for a given high-fidelity model. Below, a review of these work is given in particular examining whether these approaches ensure reliable convergence towards a true optimum for the underlying continuous system, in the sense defined previously.

1.2.1 First-order approximation and model management

The approximation and model management (AMMO) approach proposed by Alexandrov *et al.* (see [2] for an overview) is a methodology for utilizing a computationally cheap but low-fidelity model in combination with an expensive, high-fidelity model so that global convergence to a local optimum of the high-fidelity model is guaranteed. In this approach, gradient-based optimization is performed using a low-fidelity model with occasional use of the high-fidelity model to provide a performance measure of the low-fidelity model's predictive quality and recalibrate it via a multiplicative correction. The correction term is constructed so that the low-fidelity model satisfies first-order consistency with the high-fidelity model. Denoting F_{lo}, F_{hi} to be the objective function obtained from low- and high-fidelity models, the corrected low-fidelity model $\tilde{F}_{lo}$ around the design $\mathbf{d}_k$ satisfies

$$\tilde{F}_{lo}(\mathbf{d}_k) = F_{hi}(\mathbf{d}_k), \quad \nabla \tilde{F}_{lo}(\mathbf{d}_k) = \nabla F_{hi}(\mathbf{d}_k). \tag{1.1}$$

A way to enforce the above is to obtain $\tilde{F}_{lo}$ from F_{lo} via a multiplicative correction [2],

$$\tilde{F}_{lo}(\mathbf{d}) = \beta(\mathbf{d}) F_{lo}(\mathbf{d}), \tag{1.2}$$

19

where $\beta(\mathbf{d})$ is a linear function constructed using information from low/high-fidelity models at $\mathbf{d}_k$ so as to ensure $\tilde{F}_{\text{lo}}$ satisfies (1.1). This consistency condition is crucial both theoretically for the convergence proof as well as practically in ensuring a good match of trends between the two models. For 2D and 3D wing optimization in inviscid Euler flows and utilizing low/high-fidelity models of same physics but half the mesh spacing, AMMO results in a factor of 2 to 3 compute time saving [2]. When low/high-fidelity models have variable physics as well, the computational saving of AMMO is more significant. The use of AMMO for variable physics models was first demonstrated by Alexandrov *et al.* [3] and more recently applied by Le Moigne and Qin [62] as well.

AMMO provides a general framework to automatically manage the use of variable fidelity models (of arbitrary accuracies) provided by the user. An inherent assumption in AMMO is that the (computable) high-fidelity model is a sufficiently accurate representation of the underlying continuous system. Hence, in the present context where we would like to ensure convergence to the (uncomputable) continuous system, the assumptions made in the AMMO approach are violated. In particular, the gradient information for the continuous system is not available; however, the error in the objective values can usually be estimated. Hence, a different approach based on this assumption is needed.

1.2.2 Progressive optimization

In a series of papers [18, 19, 20], Dadone and Grossman proposed an approach for increasing the efficiency of aerodynamic optimization that relies on converging the analysis and design process simultaneously using progressively finer grids. To decrease the computational costs associated with obtaining the objective function gradient, for inviscid design problems the adjoint state is not solved on the current working grid but on the coarsest grid. For viscous optimization cases, a further approximation is made for the adjoint by ignoring the viscous contribution to the residual. On each given mesh, the flow and adjoint equations are not solved exactly but only converged 1 to 2 orders of magnitude. Once the objective function has decreased

by an order of magnitude on the current mesh, the mesh spacing is halved and the same steps are carried out on the new mesh. Among the factors contributing to the efficiency of strategy, it was estimated that the most significant contribution comes from progressively converging the flow solution [19].

The methodology demonstrates significant compute time savings for many aerodynamic design problems. For our purpose of developing a framework that is applicable to general optimization problems for PDE systems, it is not clear that the prescribed strategy is easily extendable without considerable user experience in fine tuning the parameters. In particular, both the desired level of optimization convergence prior to refining the mesh and the drop in the solution residuals in each design cycle may be highly problem-dependent. We would like to develop an approach that incorporates a procedure to automatically detect the need to refine mesh or continue iterative solution, in a manner that is generally applicable.

1.2.3 Simultaneous analysis and design

In the simultaneous analysis and design (SAND) or one-shot approach, design updates are not computed from fully converged solutions. Rather, the design and solution approximation are evolved at the same time. Thus, in contast to the reduced variable approach where the primal state is fully determined from the design via the residual equations, for SAND the solution is not required to be feasible until the design approaches optimality. In [54], Kuruvila *et al.* propose an implementation where the geometry is updated in a hierarchical manner such that high frequency changes are done separately from low frequency changes. Hence, the optimization procedure is broken into a sequence of problems each of its own length scale so as to minimize computational costs and improve the conditioning for the optimization problems. The approach is applied to airfoil optimization using the potential flow equations, where the multigrid one-shot strategy is demonstrated to bring the cost of optimization down to two or three times the effort required for one analysis. In [48] the one-shot approach without the use of multiple grids is further applied to inviscid channel and Ringleb flow designs with shape updates obtained via the steepest descent method.

21

In [79], interior-point trust-region sequential quadratic programming (SQP) is applied to drag constrained 2D airfoil design in Euler flow. More recently, Sung and Kwon extends approach of [48] to more complex and challenging 2/3D design cases [82, 81]. In the context of optimal control of incompressible Navier-Stokes flows, Ghattas and Bark applied the one-shot strategy using a quasi-Newton approximation for the equations governing the control updates [27]. To improve convergence rate, Biros and Ghattas [14] proposed the use of Krylov method to solve the Newton system for the Karusch-Kuhn-Tucker (KKT) condition, preconditioned by quasi-Newton SQP with inexact forward and adjoint solves.

Although significant progress has been made, there are a number of issues that remain to be addressed for the SAND approaches. Firstly, it has been observed that these approaches tend to suffer more convergence difficulties in comparison to the traditional reduced-gradient approach [79, 27], motivating Biros et al. to develop globalizing strategies [14]. Also, owing to the lack of theoretical criterion to determine the adequate amount of solution convergence carried out in each design step, numerical experience is needed to find the appropriate trade-off between convergence robustness and efficiency [82]. In addition to convergence instability, another potential drawback to SAND approaches is the uncertainty associated with incomplete solution convergence introduced into the design procedure. For instance, in [82] Sung and Kwon described an airfoil optimization test case where the optimized results obtained from reduced-gradient and one-shot algorithms are dissimilar. This could be a manifestation of invalid optimization steps in the sense of leading to an increased objective function that allowed the design to escape the basin of attraction and converge instead to a neighboring local optimum. While in this case the design converged to another acceptable solution, in other cases this effect could lead to detrimentally degraded designs. A procedure of balancing the degree of feasibility and optimality in the design path to result in added robustness of the algorithm is clearly desirable.

Another issue that remains to be addressed is the incorporation of discretization levels into SAND approaches. Most of the SAND strategies are implemented on a single, fine grid. Although in Kuruvila et al. [54] and Shenoy et al. [79] a sequence of

refined meshes are used in conjunction with the SAND strategy, the refinement criteria used are purely heuristic. Lacking in the current SAND approaches are automatic mesh refinement procedures based on the approximation properties provided by the current mesh and taking into account both the current level of solution convergence and design optimality.

1.2.4 Adaptive precision

Recently, algorithm models for controlling the degree of discretization fidelity and iterative convergence have been proposed by Pironneau and Polak [70]. These constitute extensions of previous work by Polak *et al.* where only the effect of discretization fidelity was of concern [71, 53, 78]. Based on *a priori* known convergence properties of the discretization formulation and solution procedure, a number of algorithm models are proposed such that precision parameters are controlled within the optimization process. The framework of *quasi-consistent approximations* ensures that using any optimizer which produces sufficient decrease in the objective function away from points of zero gradient, every accumulation point of the sequence of iterates constructed by the algorithm is a stationary point for the underlying continuous system. The approach has been successfully applied to distributed control problems governed by elliptic equations.

For more complex problems, the approach of Pironneau and Polak may not be applicable since the discretization and iterative convergence properties are not known. A related issue is that although the algorithm would eventually converge to a stationary point of the continuous model, it does not guarantee that all design updates computed on intermediate models are valid improvements. Thus, upon termination at a finite optimization index the designer is left unsure whether the obtained design constitutes an improvement over the initial or the observed changes in the computed objective values arise from the use of numerical approximations. Hence, for the given goal of increasing reliability, a further extension is necessary.

1.3 Approach

The approach taken here is to develop a framework to increase designer confidence that is applicable to general contexts, including aerodynamic optimization. In particular, the approach under consideration is that of successive model refinement which is necessary in order to obtain converging approximations to the underlying continuous system. Furthermore, this approach is arguably more applicable to situations where certain solution features (in the primal and dual variables) may develop during the optimization procedure and hence (local) refinements in the model may be necessary.

The framework proposed in this thesis replaces the *a priori* error estimates utilized in Pironneau and Polak's work with *a posteriori* output error estimates. This approach reduces the uncertainty inherent in *a priori* error estimates while simultaneously targeting the outputs for which the optimization is focused on. This *a posteriori* framework is discussed in Section 1.3.1. The method is then applied to aerodynamic optimization using higher-order discontinuous Galerkin discretization of the compressible Euler and Navier-Stokes equations. In Section 1.3.2, the potential benefits of higher-order DGFEM in this context are discussed.

1.3.1 A posteriori error estimation and control in optimization

As reviewed in Section 1.2, in all the variable-fidelity techniques other than the adaptive precision method proposed by Pironneau and Polak [70] a fixed set of high and low fidelity models are chosen *a priori*, often simply constructed for instance by uniform, global mesh refinements. The computed sequence of designs are only guaranteed to converge to an optimal solution of a fixed finite dimensional model rather than to that of the underlying continuous system. Given the lack of feedback on the model accuracy, there exists no automatic precedure for increasing the refinement of the highest-fidelity model when it is in fact not sufficiently refined for the purpose of optimization or alternatively stopping the procedure when the design changes given by the optimizer may no longer improvements for the underlying continous system.

24

There is clearly a need for a precision adjustment framework to ensure the information provided by the optimization procedure (including the history of designs and computed values of the objective) can be relied upon by the designer.

Recent years has seen the development of *a posterior* error estimators and bounds in the context of computational simulation. Error estimates for functionals allow one to gain confidence in the computed accuracy for outputs of engineering interest and localization of these estimators give one an ability to perform local refinements where necessary [1, 13, 83, 84, 74, 33, 7]. For exact weak solutions of linear coercive PDEs, the existence of functional error bounds in fact allows one to certify the result of the simulation [77]. Application of duality-based analysis technique to the iterative solution of algebraic systems also results in output error estimates due to incomplete solution convergence [68, 37, 57]. In the context of optimal control, there has also been recent effort in using duality-based local error indicators to obtain a sequence of approximating meshes. For drag reduction in incompressible Navier-Stokes flow via Neumann and Dirichlet boundary control, by applying the general approach proposed for functional outputs [13] Becker used the Lagrangian for the discretized control problem computed with the converged primal and dual states to obtain the subsequent mesh via local mesh refinement [11, 12]. Other potential alternatives exist to obtain meshes that approximate the continuous problem. In the context of Neumann boundary control for elliptic systems, Liu *et al.* [56] perform error analysis for the sum of the norms of the state, adjoint and control errors and adapts the mesh to effect control on these quantities [55].

Although the use of error estimates is becoming prevalent in simulations and local error indicators have been applied to construct sequence of approximating meshes in the optimal control context, the quantitative estimate on the magnitude of the uncertainty in the objective function computed with the approximation models has yet to be incorporated within optimization procedures as a basis for controlling the level of model fidelity for reliability. In this thesis, *a posterior* error estimates are incorporated within the general adaptive precision framework of Pironneau and Polak [70]. By using discretization and iteration error estimates rather than *a priori* bound

functions containing unknown constants, the accuracy provided by the discretization in relation to optimization steps as well as the accuracy of solution iteration in relation to discretization level can be appropriately controlled so as to significantly increase the reliability of simulation-based design.

Several connections can be made between the proposed and existing approaches. For instance, the proposed approach of sequencing the grids within optimization can be viewed as an adaptive extension of progressive optimization described in Section 1.2.2 based on rigorous error estimates. Whereas the latter converges the flow solution by a fixed number of iterations, the proposed approach ensures the iterative error reaches an adaptively chosen tolerance. Also, instead of using a predetermined fine grid, the proposed approach successively refines grids via the current error estimator when and where necessary.

1.3.2 High-order DGFEM implementation

To demonstrate the practicality of the proposed methodology, the necessary analysis and computational tools are developed in the context of discontinuous Galerkin finite element method (DGFEM). DG schemes have recently become popular for convection-dominated flow problems with the potential of resulting in orders of magnitude decrease in simulation time compared to traditional low-order finite volume methods. At least for shock-free flows, it has been demonstrated that given a desired error tolerance on outputs of engineering interest, high-order interpolations can obtain estimates with orders of magnitude fewer degrees of freedom in comparison to the use of linear interpolation [25, 67]. In the context of optimal control and shape optimization, the use of high-order solution could similarly result in significant efficiency benefits.

Although DGFEM has seen significant development as an analysis tool, it has only recently been applied to the context of optimal control [16] and has yet to be demonstrated in an aerodynamic optimization setting. Owing to the variational properties inherent in its formulation, DGFEM is arguably more amenable to duality analysis than finite volume methods and hence more suitable to the setting of optimal con-

26

trol and error estimation. To enable the use of these adjoint-based techniques it is important to examine the dual-consistency property of DG schemes. In particular, the form of boundary treatment has significant effects on adjoint regularity and care needs to be taken to ensure that the primal boundary conditions and functionals are formulated in a dual-consistent manner. This variational property of the numerical scheme turns out also to be crucial for duality-based techniques to fully benefit from the use of high-order solution. Implications of dual-consistency demonstrated in this thesis include the convergence rate in certain error measures, as well as potentially benefitting both the effectivity of error estimates and the accuracy of shape sensitivities.

Another essential ingredient for the proposed adaptive precision methodology is the ability to efficiently estimate output error due to incomplete solution. The proposed approach is based on the development of a concurrent flow-adjoint solver. Although the feasibility of iterative error estimation via concurrent primal-dual iterations has been shown in a number of settings [68, 37, 57], it remains to demonstrate that this solution approach can be performed in an efficient manner. It turns out that by making use of the DG properties of nearest neighbor stencil as well as the algebraic construction of adjoint preconditioner and residual, the concurrent solver can obtain the adjoint solution at little additional cost over the flow algorithm. To summarize, a unified adjoint approach is developed in the present DGFEM context for all of discretization and iteration error estimation as well as the computation of shape sensitivities.

1.4 Contributions

The main contributions of this thesis are in two general areas: firstly, a strategy is proposed for the incorporation of *a posteriori* estimates into optimization for PDE systems; secondly, the feasibility of the proposed strategy is demonstrated via an application to aerodynamic design. In the latter area, DGFEM is demonstrated as an effective way to realize the proposed methodology. Summing up, the advances

made in this work include:

- Development of an *a posteriori* error estimation and control framework for adaptive precision optimization.

- Demonstration of the feasibility of the proposed framework to aerodynamic optimization.

- Development of duality techniques for high-order DG, including:

  - Dual-consistent boundary treatment;

  - Efficient concurrent flow-adjoint solution algorithm;

  - Accurate adjoint-based estimation of geometric design sensitivities.

1.5 Overview of thesis

In Chapter 2 the setting of *consistent approximations* is introduced and an adaptive precision framework based on *a posteriori* error estimates is developed. In Chapters 3 to 6 the overarching goals are to build the necessary computational tools for adjoint-based methods within the DG context and verify certain assumptions on the finite dimensional approximations made in the adaptive precision framework. In particular, the numerical examples given at the end of each chapter demonstrate the particular capability required for the adaptive precision computation carried out in Chapter 7. A number of contributions of independent interest are also made in each chapter. In Chapter 3, a dual-consistent boundary treatment for DG is proposed and implications are illustrated. In Chapter 4, error analysis and control for functional outputs is carried out within a general, optimal control framework applicable to DG schemes. This represents the first treatment of output-based error analysis and adaptation using the second form of Bassi-Rebay (BR2) discretization. Expressions for local error indicators are derived. Due in part to the dual-consistent property of the chosen DG scheme, the error indicators do capture the local error contribution and the output error is effectively controlled via p-adaptation. In Chapter 5, a concurrent

28

flow-adjoint solution algorithm is developed to enable adjoint-based estimation of iteration error. In particular, it is shown that the nearest-neighbor stencil property of BR2 discretization allows for an efficient adjoint solution algorithm. Furthermore, in the case that the full linearization cannot be stored in memory, the concurrent approach is shown to provide an attractive alternative to the sequential adjoint solution approach in regard to the computational cost. In Chapter 6, the use of an incomplete shape-sensitivity based on discrete adjoint solution is proposed. Instead of fully differentiating the location of all mesh nodes with respect to the design variables, only surface elements are perturbed while the interior mesh motion is ignored. It is demonstrated that accurate gradient approximations can be obtained from high-order interpolations without including interior mesh motions. To verify the gradient convergence assumption made in the adaptive precision framework, the convergence rate of the incomplete to full discrete adjoint sensitivities is studied for various solution orders. In Chapter 7 computational results of applying the adaptive precision framework to aerodynamic design cases are presented. Finally, conclusions and potential areas of future work are discussed in Chapter 8.

Chapter 2

Adaptive precision methodology

2.1 Introduction

This chapter is concerned with the development of an adaptive precision methodology based on the error estimates that will be developed in the subsequent chapters. Given the underlying concern for reliability and correctness, the main issues addressed here are conditions on the precision adjustments so that design changes computed on the approximation models are valid improvements as well as ensuring the convergence of a sequence of discrete solutions to local optima of the underlying continuous problem. The latter issue of convergence has been examined by Polak *et al.* in the context of computational optimal control of differential equations via discretized approximations [53, 78, 71] and has more recently been extended to include the use of iterative methods to solve the discrete approximations [70]. This general setting is introduced in Section 2.2. However, the issue of reliability in the algorithm is not addressed by the use of *a priori* bound functions with its unknown, multiplicative constants that have to be properly tuned in an implementation of the algorithm. The approach to improve the reliability proposed here is to incorporate *a posteriori* error estimates and is discussed in Section 2.3. By an appropriate choice of parameters in the algorithm, optimization steps on the approximation models are required to satisfy a descent condition for the underlying continuous problem.

31

2.2 Consistent approximations

Consider the optimization problem $\mathbf{P}$ of minimizing an objective function $\mathcal{J}(\cdot)$ over a normed space $(\mathbf{D}, \|\cdot\|_{\mathbf{D}})$:

$$(\mathbf{P}): \quad \min_{\mathbf{d} \in \mathbf{D}} \mathcal{J}(\mathbf{d}). \tag{2.1}$$

For finite dimensional implementations, consider $\mathbf{D}_h$ a sequence of dense finite dimensional subspaces of $\mathbf{D}$ and $\mathbf{P}_h$ a sequence of optimization problems for $\mathcal{J}_h(\cdot)$:

$$(\mathbf{P}_h): \quad \min_{\mathbf{d}_h \in \mathbf{D}_h} \mathcal{J}_h(\mathbf{d}_h). \tag{2.2}$$

The above setting includes the situation where an optimization problem on a PDE model with an infinite dimensional control space is approximated by a sequence of optimization problems consisting of increasingly finer discretizations over control spaces of expanding dimension. The problems $\mathbf{P}_h$ are assumed to provide approximations to $\mathbf{P}$, mathematically described as the convergence of the epigraphs of $\mathbf{P}_h$ to that of $\mathbf{P}$ as defined by Polak [71]:

Definition 1 *The problems* epi-converge $(\mathbf{P}_h \overset{epi}{\rightarrow} \mathbf{P})$ *if:*

1. For every $\mathbf{d} \in \mathbf{D}$, $\exists \mathbf{d}_h \in \mathbf{D}_h$ *such that* $\mathbf{d}_h \rightarrow \mathbf{d}$ *and* $\limsup \mathcal{J}_h(\mathbf{d}_h) \leq \mathcal{J}(\mathbf{d})$;

2. For every sequence $\mathbf{d}_h \in \mathbf{D}_h, \mathbf{d}_h \rightarrow \mathbf{d} \in \mathbf{D}$, $\liminf \mathcal{J}_h(\mathbf{d}_h) \geq \mathcal{J}(\mathbf{d})$.

Although epi-convergence ensures that global optimal solutions of $\mathbf{P}_h$ converge to that of $\mathbf{P}$, it does not ensure that local optima of $\mathbf{P}_h$ converge to stationary points of $\mathbf{P}$. As shown in [71], this may happen if the radius of attraction of the local minimizer for $\mathbf{P}_h$ is not bounded away from zero. This is due to the fact that epi-convergence prescribes only zeroth-order characterization of the approximation problems. To preclude this situation, optimality functions $\theta_h(\cdot), \theta(\cdot)$ are introduced to characterize the first-order (gradient) convergence of the approximations.

Definition 2 $\theta_h(\cdot), \theta(\cdot)$ *are* optimality functions *for* $\mathbf{P}_h, \mathbf{P}$ *if they are upper semicontinuous, non-positive functions which vanish at the local minimizers for* $\mathbf{P}_h, \mathbf{P}$

respectively.

For the case of continuously differentiable objectives, an example of optimality function is some chosen norm of the gradient. The *consistent approximation* qualification given below on the problem-optimality function pair $\{\mathbf{P}, \theta(\cdot)\}$ provides a sufficient condition for the local minima of the finite dimensional approximations to converge to stationary points for the original problem.

Definition 3 *The pair $\{\mathbf{P}_h, \theta_h(\cdot)\}$ form consistent approximations to $\{\mathbf{P}, \theta(\cdot)\}$ if $\mathbf{P}_h \overset{epi}{\to} \mathbf{P}$ and for every sequence $\mathbf{d}_h \in \mathbf{D}_h$ such that $\mathbf{d}_h \to \mathbf{d} \in \mathbf{D}$, $\limsup \theta_h(\mathbf{d}_h) \leq \theta(\mathbf{d})$.*

In particular, the above is satisfied if it can be shown that:

$$\lim_{h \to 0} \mathbf{d}_h \to \mathbf{d} \Rightarrow \lim_{h \to 0} \theta_h(\mathbf{d}_h) \to \theta(\mathbf{d}). \tag{2.3}$$

The above condition has been shown in a number of simple settings. For an inverse design problem on an elliptic PDE via Neumann boundary control, it has been shown that both the objective and optimality function are continuous with respect to boundary control in L_2 [61]. By discretizing the continuous system using standard conforming finite element method (FEM), the sequence of finite-dimensional problems obtained as the mesh diameter goes to zero are consistent approximations in the sense of Definition 3. In the setting of shape optimization, for an inverse design problem of nozzle flow modelled by Laplace's equation with homogeneous Neumann boundary condition enforced on the design surface, it has been shown that both the objective and the optimality function are continuous with respect to shape perturbations in H_0^2. Using standard conforming FEM to approximate the continuous problem, in the limit $h \to 0$ both the discrete objective and the optimality functions converge to the corresponding continuous functions as well.

In order to obtain an approximating sequence to some stationary point of problem $\mathbf{P}$ via nonlinear programming iterations on the finite-dimensional problems $\mathbf{P}_h$, it is necessary to dynamically adjust the precision parameter h at certain points of the

computation. In [71], algorithm models are proposed where refinements are based on tests involving either the comparison of optimality function or function changes produced by the underlying nonlinear programming algorithm, with bound functions on the precision of the model $\mathbf{P}_h$. In [70], the algorithm models are further extended to handle the situation where a significant number N of solver iterations are necessary to obtain approximation to the functional output. To decrease the computational time required for optimal control, N is dynamically set with respect to h in a manner so as to ensure convergence. In the situation involving both discretization and iteration parameters, the following assumptions are made on the behavior of the discretization and iteration error [70]. Consider optimization problems where the objective function is a functional of the state $\mathbf{u}(\mathbf{d}_h)$. Let $\mathbf{u}_h(\mathbf{d}_h)$ denote a finite dimensional approximation of the state. Also, let $\mathbf{u}_{h,N}(\mathbf{d}_h)$ denote an approximation to $\mathbf{u}_h(\mathbf{d}_h)$ obtained by N steps of iterative solution.

Assumption 1 *For every bounded set* $\mathbf{B} \in \mathbf{D}$, *there exists* $h_{max} \in \mathbb{R}_+$, $k < \infty$, $\Delta : \mathbb{R}_+ \to \mathbb{R}_+$ *such that* $\forall h \in (0, h_{max}]$, $\mathbf{d}_h \in \mathbf{D}_h \cap \mathbf{B}$:

$$|\mathcal{J}_h(\mathbf{u}_h(\mathbf{d}_h), \mathbf{d}_h) - \mathcal{J}(\mathbf{u}(\mathbf{d}_h), \mathbf{d}_h)| \leq k\Delta(h), \tag{2.4}$$

and for $N \in \mathbb{N}$ *there exists* $\varphi : \mathbb{R}_+ \times \mathbb{N} \to \mathbb{R}_+$:

$$|\mathcal{J}_h(\mathbf{u}_{h,N}(\mathbf{d}_h), \mathbf{d}_h) - \mathcal{J}_h(\mathbf{u}_h(\mathbf{d}_h), \mathbf{d}_h)| \leq k\varphi(h, N), \tag{2.5}$$

where the discretization and iterative bound functions $\Delta(\cdot), \varphi(\cdot, \cdot)$ *are naturally assumed to satisfy the limiting properties,*

$$
\begin{aligned}
\lim_{h \to 0} \Delta(h) &= 0, \\
\lim_{N \to \infty} \varphi(h, N) &= 0, \\
\exists N^*(h) : \quad \lim_{h \to 0} \varphi(h, N^*(h)) &= 0.
\end{aligned}
\tag{2.6}
$$

For a given h, equation (2.4) assumes the existence of a bound function $k\Delta(h)$ that holds uniformly on the set $\mathbf{D}_h \cap \mathbf{B}$ with k being a constant that absorbs the depen-

34

dence of the bound on the size of the bounded set **B**. To prove convergence results, the underlying nonlinear programming algorithm is required to satisfy a monotone, uniform descent condition [70]. The gradient-based nonlinear programming algorithm is denoted by the map of controls $\mathbf{d}_h^i \rightarrow C(\mathbf{d}_h^i, \mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N})$, where the gradient is computed using an adjoint approximation, $\boldsymbol{\psi}_{h,N}$, as discussed in Chapter 6.

Assumption 2 *For every* $\mathbf{d}^*$ *where* $\frac{d\mathcal{J}(\mathbf{d}^*)}{d\mathbf{d}} \neq 0$, $\exists \rho^*, \delta^*, h^* > 0, N^{**}(\cdot) < \infty$ *such that* $\forall h \leq h^*, N \geq N^{**}(h)$,

$$\mathcal{J}_h(\mathbf{u}_{h,N}(C(\mathbf{d}_h)), C(\mathbf{d}_h)) - \mathcal{J}_h(\mathbf{u}_{h,N}(\mathbf{d}_h), \mathbf{d}_h) \leq -\delta^*, \quad \forall \mathbf{d}_h \in \mathbf{D}_h \cap \mathbf{B}(\mathbf{d}^*, \rho^*). \quad (2.7)$$

The above condition stipulates that around every non-stationary point $\mathbf{d}^* \in \mathbf{D}$, there exists some ball **B** of radius ρ^* such that applying the nonlinear programming algorithm on all $\mathbf{d}_h \in \mathbf{D}_h \cap \mathbf{B}$ using gradient information obtained from primal and dual approximations with sufficiently fine h, N would produce an improvement in the computed objective function that is bounded away from zero. With the conditions set out in Assumption 1, the Algorithm Model 2 of [70] based on a nonlinear programming algorithm satisfying Assumption 2 has the property that if the constructed sequence has any accumulation point then the discretization parameter has to tend zero ($h \rightarrow 0$). Furthermore, every accumulation point of the constructed sequence are stationary points for $\{\mathbf{P}, \theta(\cdot)\}$. Hence, if $\mathcal{J}(\cdot)$ is strictly convex with bounded level sets, the algorithm converges to the unique optimum.

2.3   Algorithm based on error estimates

In this work, the bound functions are determined by *a posteriori* estimates rather than chosen *a priori* by the user. The former is preferred since the latter is often difficult to realize in many practical situations. An instance of this is the case where the mesh is updated by local p-refinements rather than global h-refinements where both the order of convergence and the required number of solver iterations are difficult to estimate *a priori*. Standard results in *a posteriori* error estimates are of the form (2.4) with $k = 1$

35

for a particular given $\mathbf{d}_h$ rather than over some bounded neighborhood. However, if the discretization error for the underlying problem has a smooth dependence on the design with respect to the norm $\| \cdot \|_{\mathbf{D}}$, within some small enough set $\Delta(h)$ can be approximated by an *a posteriori* estimate at the given single design point or obtained by interpolation at a certain set of neighboring design points. Similarly, if for each given discretization the iteration error is assumed to have a smooth dependence on the design then $\varphi(h, N)$ of (2.5) can be obtained via an iteration error estimate for the given design $\mathbf{d}_h$.

In the current work, the discretization bound function is simply set to the value of the discretization error estimate computed from partially converged primal and dual state approximations so that $\Delta = \Delta(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N})$. The specific form of the estimate will be discussed in Chapter 4. Similarly, the iterative bound function is simply set to the value of the iteration error estimate computed from primal and dual state approximations, $\varphi = \varphi(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N})$. The procedure for obtaining $\varphi(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N})$ is discussed in Chapter 5. Also necessary in the algorithm is a function $N^*(\cdot)$ satisfying (2.6). Given that asymptotically, $\lim_{N \to \infty} \Delta(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N}) = \Delta(\mathbf{u}_h, \boldsymbol{\psi}_h)$ and since $\Delta(\mathbf{u}_h, \boldsymbol{\psi}_h)$ vanishes in the limit as $h \to 0$, a choice for $N^*(\cdot)$ satisfying the condition (2.6) is to take it to be the smallest N such that the iterative error is less than a certain ζ multiple of $\Delta(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N})$:

$$N^*(h) \equiv \arg\min_N \left\{ \varphi(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N}) \leq \zeta \times \Delta(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N}) \right\}.$$

The above choice of $N^*(h)$ has the property that a correspondence is maintained between the tolerance level of iterative to discretization error. Using these ingredients, the proposed adaptive precision algorithm taking parameters within the range $\gamma, \zeta > 0$, $\tau, E \in (0, 1)$, $\omega \in (0, 1]$, $j_{max} \in \mathbb{Z}^+$ is shown below.

Adaptive Precision Algorithm $(\gamma, \zeta, \tau, \omega, E, j_{max})$

Initial control: $\mathbf{d}_h \in \mathbf{D}_h$.
Initial converged solution: $\mathbf{u}_h, \boldsymbol{\psi}_h$.
Set $\mathbf{u}_{h,N} := \mathbf{u}_h, \boldsymbol{\psi}_{h,N} := \boldsymbol{\psi}_h$.

Begin Outer Loop($i := 0$; $i \leq i_{max}$)

- Set $\Delta := \Delta(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N})$.

  While Inner Loop($j := 0$; $j \leq j_{max}$)

  1. **While Line-search**

     - Compute control update

     $$\tilde{\mathbf{d}}_h := C(\mathbf{d}_h, \mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N}). \tag{2.8}$$

     - Concurrently iterate state updates $\tilde{\mathbf{u}}_{h,N}(\tilde{\mathbf{d}}_h)$, $\tilde{\boldsymbol{\psi}}_{h,N}(\tilde{\mathbf{d}}_h)$ until:

     $$\varphi(\tilde{\mathbf{u}}_{h,N}, \tilde{\boldsymbol{\psi}}_{h,N}) \leq \zeta \tau^j \times \Delta(\tilde{\mathbf{u}}_{h,N}, \tilde{\boldsymbol{\psi}}_{h,N}). \tag{2.9}$$

     End Line-search

  2. Set $\varphi := \varphi(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N})$.

  3. **If $\mathcal{J}_h(\tilde{\mathbf{u}}_{h,N}, \tilde{\mathbf{d}}_h) - \mathcal{J}_h(\mathbf{u}_{h,N}, \mathbf{d}_h) < -\varphi^\omega$**

     - Exit **Inner Loop.**

     Else

     - $j := j + 1$.

     - Concurrently iterate $\mathbf{u}_{h,N}(\mathbf{d}_h)$, $\boldsymbol{\psi}_{h,N}(\mathbf{d}_h)$ until:

     $$\varphi(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N}) \leq \zeta \tau^j \times \Delta(\mathbf{u}_{h,N}, \boldsymbol{\psi}_{h,N}). \tag{2.10}$$

  End Inner Loop

- Set

$$\hat{\Delta}(h, N) := \Delta + \varphi.$$

- **If**

$$\mathcal{J}_h(\tilde{\mathbf{u}}_{h,N}, \tilde{\mathbf{d}}_h) - \mathcal{J}_h(\mathbf{u}_{h,N}, \mathbf{d}_h) > -\gamma \hat{\Delta}(h, N)^\omega \qquad (2.11)$$

- Call AdaptGrid$(E, \mathbf{d}_h, \mathbf{u}_{h,N}, \psi_{h,N})$.

Else

- Update valid control and states :

$$\left\{ \mathbf{d}_h; \ \mathbf{u}_{h,N}; \ \psi_{h,N} \right\} := \left\{ \tilde{\mathbf{d}}_h; \ \tilde{\mathbf{u}}_{h,N}; \ \tilde{\psi}_{h,N} \right\}. \qquad (2.12)$$

End Outer Loop

The algorithm controls the error in the objective function in a two-tiered manner. In the inner loop, at the trial update $\tilde{\mathbf{d}}_h$ the iterative error is initially made to be less than a ζ multiple of the discretization error term Δ^ω at the current design point $\mathbf{d}_h$. If the change in the approximate objective function is not sufficiently negative, as may happen if the approximate gradient does not result in a descent direction, additional solution iterations are performed to tighten the value of iterative error by the factor τ. If the iterative error test is satisfied, the computed change in the objective function is tested against γ multiple of the sum of discretization and iteration error contributions, $\hat{\Delta}(h, N)^\omega$. If the change is not sufficiently negative, the procedure denoted by AdaptGrid$(E, \mathbf{d}_h, \mathbf{u}_{h,N}, \psi_{h,N})$ refines the grid according to the local error indicator (as discussed in Chapter 4) to reduce the error bound in the objective function by the fraction E.

Given parameters in the valid range as described, if the algorithm produces an infinite sequence of iterates $\mathbf{d}_h^i$ that has at least one accumulation point, the model precision as governed by h, N can be proved to increase indefinitely. An additional criterion is needed to obtain convergence statements for subsequences of $\{\mathbf{d}_h^i\}$. A sufficient condition for every accumulation point $\mathbf{d}^*$ of the constructed sequence $\{\mathbf{d}_h^i\}$ to be stationary points for problem **P** is that for all large enough i, the change in the

"exact" objective function evaluated with $\{\mathbf{d}_h^i\}$ is negative (Theorem 5, [70]):

$$\mathcal{J}(\mathbf{u}(\mathbf{d}_h^{i+1}), \mathbf{d}_h^{i+1}) - \mathcal{J}(\mathbf{u}(\mathbf{d}_h^i), \mathbf{d}_h^i) \leq 0. \tag{2.13}$$

In the Algorithm Model 2 presented in [70], ω is chosen to be strictly less than 1. In this case, for all choice of precision parameters γ, ζ, the condition (2.13) would eventually be satisfied for sufficiently large i. In the case that the constant k of (2.4, 2.5) can be estimated effectively, the choice $\omega = 1$ can also be made to satisfy the improvement condition (2.13) provided γ is chosen appropriately. In view of the test (2.11) as well as error bounds (2.4) and (2.5), the change in the objective function given by iterates produced by the adaptive precision algorithm is bounded by,

$$
\begin{aligned}
\mathcal{J}(\mathbf{u}(\mathbf{d}_h^{i+1}), \mathbf{d}_h^{i+1}) - \mathcal{J}(\mathbf{u}(\mathbf{d}_h^i), \mathbf{d}_h^i) &\leq -\gamma \hat{\Delta}(h, N) + 2k\hat{\Delta}(h, N) \\
&= \hat{\Delta}(h, N)(-\gamma + 2k),
\end{aligned}
\tag{2.14}
$$

which is negative provided $\gamma \geq 2k$. Therefore, in the case that the *a posteriori* error estimate is tight ($k \approx 1$) for some appropriate range of designs, setting $\gamma \geq 2$ would ensure the inequality (2.13) which has the interpretation that the design updates are always valid in the sense of leading to improvements for the underlying problem $\mathbf{P}$. Given the goal of providing the user confidence in the design updates, the parameter values $\omega = 1, \gamma = 2$ is adopted in this thesis. For the algorithm to be stable and efficient, the value of the iterative precision parameter ζ should also be appropriately chosen. Since the discretization error is estimated using partially converged primal and dual states, a natural requirement is that the iterative error contribution should at least be small in comparison. However, for reasons of efficiency, it should not be chosen unnecessarily small. For computational results shown in Chapter 7, $\zeta = 0.2$ is chosen.

Chapter 3

Dual-consistent discretization

In this chapter, the dual-consistency of DGFEM discretizations are discussed. In Section 3.1, the concept of dual-consistency is defined. Past work related to dual-consistency for FEM are discussed in Section 3.2. In Sections 3.3 and 3.4, the dual-consistency of DGFEM discretizations of first- and second-order equations is analyzed. Finally, Section 3.5 demonstrates some of the implications of dual-consistency via application to the Euler and Navier-Stokes equations.

3.1 Dual-consistency

Let $\mathcal{V}$ and $\mathcal{W}$ be appropriate function spaces. Let $\mathbf{u} \in \mathcal{V}$ be a weak solution to a partial differential equation (PDE) together with a certain set of boundary conditions (BCs), satisfying

$$F(\mathbf{u}) = 0, \tag{3.1}$$

where F is an operator mapping $\mathcal{V} \to \mathcal{W}'$, with $\mathcal{W}'$ being the dual space of $\mathcal{W}$. Let $\mathcal{J}(\cdot) : \mathcal{V} \to \mathbb{R}$ be the functional of interest. Introduce the Lagrangian $\mathcal{L} : \mathcal{V} \times \mathcal{W} \to \mathbb{R}$,

$$\mathcal{L}(\mathbf{u}, \boldsymbol{\psi}) \equiv \mathcal{J}(\mathbf{u}) - (F(\mathbf{u}), \boldsymbol{\psi})_{\mathcal{W}, \mathcal{W}'}, \tag{3.2}$$

41

where $(\cdot, \cdot)_{\mathcal{W},\mathcal{W}'}$ denotes duality pairing of $\mathcal{W}$ with $\mathcal{W}'$. Taking variations $\mathbf{u} \rightarrow \mathbf{u}+\delta\mathbf{v} \in \mathcal{V}$ permitted by the primal BCs and requiring the Lagrangian (3.2) be stationary with respect to permissible $\delta\mathbf{v}$, the equation for the adjoint variable ψ is obtained,

$$F'[\mathbf{u}]^*\psi = \mathcal{J}'[\mathbf{u}], \tag{3.3}$$

where the primed notation on an operator is used to denote the Frechét derivative, the square bracket $[\cdot]$ denotes the state about which linearization is performed and superscript $*$ denotes the adjoint operation. In applications such as *a posteriori* error estimation and optimal control of PDEs, it is important to obtain convergent, finite dimensional approximations to the PDE (continuous) adjoint. This could be done by directly discretizing (3.3). An alternative to discretizing the adjoint PDE directly would be to solve the discete adjoint problem corresponding to the Lagrangian functional for the discretized equation and output. With the algebraic system for the primal unknowns $\mathbf{U} \in \mathbb{R}^n$ denoted by,

$$\mathbf{A}(\mathbf{U}) = 0, \tag{3.4}$$

and $J(\cdot)$ the discrete functional, the same procedure as carried out in (3.3) can be performed at a discrete level. In this case, the discrete Lagrangian $\mathcal{L}: \mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ is introduced

$$\mathcal{L}(\mathbf{U}, \mathbf{\Psi}) \equiv J(\mathbf{U}) - \langle A(\mathbf{U}), \mathbf{\Psi} \rangle, \tag{3.5}$$

where $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product. Then the discrete adjoint $\mathbf{\Psi}$ is the solution to the linear system

$$\mathbf{A}'[\mathbf{U}]^T\mathbf{\Psi} = J'[\mathbf{U}]. \tag{3.6}$$

For finite element formulations, the transpose Jacobian $\mathbf{A}'[\mathbf{U}]^T$ in (3.6) is equivalent to permuting the role of trial and test functions. In the following discussion of general

42

finite element methods let $\mathcal{V}_h$ be a N-dimensional vector space and let the primal state $\mathbf{u}_h$ be sought in the affine space $\mathbf{u}_h^D + \mathcal{V}_h$ where $\mathbf{u}_h^D$ arises for instance in Dirichlet BC imposition. Let $\mathcal{W}_h$ be a possibly different trial space of the same dimension. With $R_h(\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ denoting a chosen finite element semilinear form (linear in the second argument), the primal solution $\mathbf{u}_h \in \mathbf{u}_h^D + \mathcal{V}_h$ satisfies:

$$R_h(\mathbf{u}_h, \mathbf{v}_h) = 0, \quad \forall \mathbf{v}_h \in \mathcal{W}_h. \tag{3.7}$$

Denote $\mathcal{J}_h(\cdot)$ to be the discrete functional of interest. Let $\{\phi_k\}$ and $\{\varphi_k\}$ be a set of bases for $\mathcal{V}_h$ and $\mathcal{W}_h$ respectively and let $\mathbf{U}_k, \mathbf{V}_k$ denote the coefficients of given $\mathbf{u}_h \in \mathcal{V}_h, \mathbf{v}_h \in \mathcal{W}_h$ in the bases:

$$
\begin{aligned}
\mathbf{u}_h &= \mathbf{u}_h^D + \sum_{i=1}^{N} \mathbf{U}_i \phi_i, \\
\mathbf{v}_h &= \sum_{i=1}^{N} \mathbf{V}_i \varphi_i.
\end{aligned}
\tag{3.8}
$$

To form the primal system (3.4), let the i-th component of the nonlinear system of equations for the unknown coefficients $\mathbf{U} \equiv \{\mathbf{U}_k\}$ be

$$R_h(\mathbf{u}_h^D + \sum_{k=1}^{N} \mathbf{U}_k \phi_k, \varphi_i) = 0, \quad i = 0, \ldots, N. \tag{3.9}$$

With the coefficients $\mathbf{U}_k$ as unknowns, the (i, j)-th entry of the transposed Jacobian matrix and i-th entry of the discrete functional derivative are given by,

$$
\begin{aligned}
\left[\mathbf{A}'[\mathbf{U}]^T \right]_{i,j} &\equiv \left[\mathbf{A}'[\mathbf{U}] \right]_{j,i} \\
&= R_h'[\mathbf{u}_h](\phi_i, \varphi_j),
\end{aligned}
$$

$$
\left[J'[\mathbf{U}] \right]_i = \mathcal{J}_h'[\mathbf{u}_h](\phi_i). \tag{3.10}
$$

where $R_h'[\mathbf{u}_h](\cdot, \cdot) : \mathbb{R}^n \times \mathbb{R}^n \to \mathbb{R}$ is the bilinear form obtained from the corresponding semilinear form by linearising about $\mathbf{u}_h$. Hence for the finite element formulation (3.7)

43

the discrete adjoint system (3.6) is

$$\sum_{j=1}^{N} R'_h[\mathbf{u}_h](\phi_i, \varphi_j)\Psi_j = \mathcal{J}'_h[\mathbf{u}_h](\phi_i), \quad \forall i = 0, \dots, N. \tag{3.11}$$

Since $R'_h[\mathbf{u}_h](\cdot, \cdot)$ is bilinear in the arguments, the above trivially implies the following variational characterization for the discrete adjoint solution $\psi_h \equiv \sum_{j=1}^{N} \Psi_j \varphi_j \in \mathcal{W}_h$,

$$R'_h[\mathbf{u}_h](\mathbf{v}_h, \psi_h) = \mathcal{J}'_h[\mathbf{u}_h](\mathbf{v}_h), \quad \forall \mathbf{v}_h \in \mathcal{V}_h. \tag{3.12}$$

To make the connection between discrete and continuous adjoint equations (3.3) and (3.12), the concept of dual-consistency is introduced. In addition to the property of consistency (as defined for example in [80, 24]), this variational characterization of finite element methods is of importance for duality-based analysis and techniques.

Definition 4 *The finite element formulation (3.7) together with the discrete functional $\mathcal{J}_h(\cdot)$ is* dual-consistent *if given* $\mathbf{u}$ *and* ψ *solutions to the primal (3.1) and adjoint PDE (3.3) respectively,*

$$R'_h[\mathbf{u}](\mathbf{v}_h, \psi) = \mathcal{J}'_h[\mathbf{u}](\mathbf{v}_h), \quad \forall \mathbf{v}_h \in \mathcal{V}_h. \tag{3.13}$$

The formulation is asymptotically dual-consistent *if the following holds:*

$$\lim_{h \to 0} \left(\sup_{\mathbf{v}_h \in \mathcal{V}_h} \frac{|R'_h[\mathbf{u}](\mathbf{v}_h, \psi) - \mathcal{J}'_h[\mathbf{u}](\mathbf{v}_h)|}{\|\mathbf{v}_h\|_{\mathcal{V}_h}} \right) = 0. \tag{3.14}$$

In Section 3.2, we give a review of factors contributing to dual-inconsistency for both the conforming finite element and DG methods as well as the implications for the numerical behavior in applications. These examples serve to illustrate some of the differences between the conforming and DG finite element methods as well as point out the potential pitfalls leading to dual-inconsistency. There are two parts to consider in showing dual-consistency: the semilinear form of the finite element method and the boundary treatment used. Whereas the choice of DG formulation on dual-consistency is widely appreciated [6], the latter issue of the effect of boundary treatment appears

to be much less understood. In fact, there is currently no general prescription for handling boundary conditions and functionals in a dual-consistent manner. Thus the main contribution of this chapter is the construction of state and flux mappings which are dual-consistent.

3.2 Review of dual-consistency and implications

3.2.1 Interior treatment

One source of dual-inconsistency arises from the use of stabilization terms present in some finite element methods. An illustrative example of a conforming, but dual-inconsistent method is the streamline upwind/Petrov Galerkin (SUPG) discretization [15]. This method was considered by Collis and Heinkenschloss [17] for a distributed optimal control problem governed by the linear advection-diffusion equation. It is well known [72] that SUPG has the consistency property that the stabilization term vanishes when evaluated with a solution to the underlying differential operator. However, as shown in [17], SUPG is not dual-consistent since the discrete adjoint of the stabilization operator does not vanish on the continuous adjoint solution [17]. By estimating the dual-consistency error, *a priori* analysis on the solution of the optimality system shows that the convergence rates of the norm of error in the computed control and the corresponding primal, dual states for the discretize-then-optimize approach are lower than the optimize-then-discretize approach. That is, using the discrete adjoint approximation rather than directly discretizing the adjoint PDE leads to a poorer approximation of the underlying control problem on the continuous system. This finding is verified numerically, demonstrating that whereas the convergence discrepancy is small when linear elements are used, for higher-order interpolations the degradation becomes significant. As an additional note, a closely related SUPG exists where the stabilization term is in fact dual-consistent [26]. In the context of *a posteriori* error analysis and adaptation for functionals using stabilized finite element methods, the issue of whether to use some approximate solution to the dual PDE or

45

the stabilization-dependent, discrete adjoint state obtained from the stabilized finite element formulation also arises and is discussed in detail in [44, 33].

The property of dual-consistency is also important for the analysis of solution convergence in various norms. For instance, for the purpose of analyzing L_2 convergence of DG formulations for second order elliptic operators, Arnold *et al.* [6] applied the standard Aubin-Nitsche duality technique [80, 24] via an adjoint state ψ for the primal state error, solving:

$$-\Delta\psi = u - u_h \text{ on } \Omega, \quad \psi = 0 \text{ on } \partial\Omega. \tag{3.15}$$

In this context, dual-consistent schemes are those where ψ satisfy the weak statement associated with the bilinear form,

$$R_h(v_h, \psi) = (u - u_h, v_h)_{L_2}, \quad \forall v \in \mathcal{V}_h^p. \tag{3.16}$$

Using regularity property of the elliptic operator, optimal L_2 convergence estimates of the following form can be obtained for dual-consistent DG schemes:

$$\|u - u_h\|_{0,\Omega} \le Ch^{p+1}|u|_{p+1,\Omega}. \tag{3.17}$$

Not all DG schemes allow for duality-based derivation of optimal estimates. For instance, the nonsymmetric interior penalty (NIP) method lacks the necessary dual-consistency property and in fact suboptimal L_2 convergence rates have been shown numerically. Similar sub-optimal convergence is observed when the error is measured with respect to certain functional outputs [39, 38]. In this case, the dual-inconsistency of the scheme is manifested as non-convergent, mesh-dependent discontinuities in the discrete adjoint solution [39, 38]. Since this violates the smoothness estimates of the dual problem used in the construction of superconvergent *a priori* bounds, NIP also suffers from degraded output convergence as compared to dual-consistent schemes. A general conclusion that may be drawn is that although consistency and stability of the method together implies the solutions converge at optimal order with respect

to the DG (energy) norm, the lack of dual-inconsistency results in suboptimal error convergence in both the L_2 norm and functional outputs. By appropriately controlling the consistency error such that the discretization is asymptotically dual-consistent it is possible to retain properties obtained via these duality-based techniques. For instance, by choosing the penalty parameter of NIP proportional to some negative power of the mesh size h, the resulting superpenalized version of NIP is asymptotically dual-consistent and in fact the optimal L_2 convergence rate of the form (3.17) is recovered [6].

3.2.2   Boundary treatment and functional

Another source of dual-inconsistency arises from the incompatibility of discrete boundary treatment and the functional with the dual variational property of the underlying continuous problem. Again, the lack of dual-consistency results in degraded convergence rate for functional outputs which can be connected with the resulting irregular discrete adjoint solution behavior at the boundary. Here we illustrate via an example that, even in the case of conforming finite element method, analysis is needed to examine the form of discrete functional so as to ensure Galerkin superconvergence. The concept of dual-consistency can be used to further provide an understanding into the well-known fact that for second-order elliptic systems, weighted surfrace flux integrals can be approximated more accurately using an equivalent, volume form.

Consider the following boundary value problem for u as described in Giles and Süli [33]:

$$
\begin{aligned}
-\nabla \cdot [\mathbf{A(x)}\nabla \mathbf{u}] &= \mathbf{f(x)}, \quad \text{in } \Omega, \\
\mathbf{u} &= 0, \quad \text{on } \partial\Omega,
\end{aligned}
\tag{3.18}
$$

where $\mathbf{A(x)}$ is an $m \times m$ smooth, positive definite matrix-valued function. The functional of interest is

$$
\mathcal{J}(\cdot) = \int_{\partial\Omega} \mathbf{g(x)}^T \hat{\mathbf{n}} \cdot [\mathbf{A(x)}\nabla(\cdot)] ds,
\tag{3.19}
$$

where $\mathbf{g}(\mathbf{x})$ is a smooth function on $\partial\Omega$. The standard conforming finite element formulation for (3.18) consists of: find $\mathbf{u}_h \in \mathcal{V}_h^0$ such that

$$R(\mathbf{u}_h, \mathbf{v}_h) = (\mathbf{f}, \mathbf{v}_h)_{L_2(\Omega)}, \quad \forall \mathbf{v}_h \in \mathcal{V}_h^0. \qquad (3.20)$$

where $(\cdot, \cdot)_{L_2}$ denotes the L_2 inner product on Ω, the bilinear form $R(\cdot, \cdot)$ is defined as,

$$R(\mathbf{u}_h, \mathbf{v}_h) \equiv \int_\Omega \nabla \mathbf{v}_h^T \mathbf{A}(\mathbf{x}) \nabla \mathbf{u}_h d\mathbf{x}, \qquad (3.21)$$

and $\mathcal{V}_h^0$ is the vector space of continuous, piecewise polynomial functions with zero trace on $\partial\Omega$. The straightforward way to compute an output estimate is to evaluate the discrete solution $\mathbf{u}_h$ in the functional (3.19). However, as shown by Giles and Süli [33] this choice is not optimal with regard to the convergence rate. For better convergence, an alternative volume form of the functional is proposed [33]:

$$\mathcal{J}^{\mathrm{vol}}(\cdot) = (\mathbf{f}, \mathbf{w}_h^{-\mathbf{g}})_{L_2} - R(\cdot, \mathbf{w}_h^{-\mathbf{g}}), \qquad (3.22)$$

where $\mathbf{w}_h^{-\mathbf{g}}$ is an arbitrary element in the affine space of functions $\mathcal{V}_h^{-\mathbf{g}}$ with trace $-\mathbf{g}$ on $\partial\Omega$. Using this functional, the output error can be precisely represented as the residual of $\mathbf{u}_h$ weighted by the interpolation error of the PDE adjoint solution. Using interpolation estimates for the adjoint solution a priori analysis has been carried out to derive a superconvergent rate for output convergence [33]. The enhanced convergence has been numerically demonstrated for conforming discretization of the Stokes and incompressible Navier-Stokes equations [28, 33]. In the case of drag estimation, whereas the standard surface functional of the form (3.19) produce only first-order convergent output estimates on linear finite element solutions, the volume functional (3.22) extracts second-order convergent drag estimates on quasi-uniform meshes. The volume form of the flux functional has also been applied to the context of control for incompressible Navier-Stokes equations [11] and goal-oriented error estimators for Stokes equations [59].

48

The lower convergence rate provided by functional (3.19) results from the lack of a well-posed dual statement. This in turn is a consequence of the lack of dual-consistency property in the numerical approximation used. Examine the discrete adjoint problem for the surface functional (3.19) : find $\psi_h \in \mathcal{V}_h^0$ such that

$$R(\mathbf{v}_h, \psi_h) = (\mathbf{g}, \hat{\mathbf{n}} \cdot [\mathbf{A}(\mathbf{x})\nabla \mathbf{v}_h])_{\partial\Omega}, \quad \forall \mathbf{v}_h \in \mathcal{V}_h^0. \tag{3.23}$$

However, the PDE adjoint ψ solves the following boundary value problem:

$$
\begin{aligned}
-\nabla \cdot [\mathbf{A}(\mathbf{x})^T \nabla \psi] &= 0, \quad \text{in } \Omega, \\
\psi &= -\mathbf{g}(\mathbf{x}), \quad \text{on } \partial\Omega.
\end{aligned}
\tag{3.24}
$$

Hence, solution ψ of (3.24) does not satisfy the discrete adjoint equations (3.23) since for a general $\mathbf{v}_h \in \mathcal{V}_h^0$,

$$
\begin{aligned}
R(\mathbf{v}_h, \psi) &= (\mathbf{v}_h, \hat{\mathbf{n}} \cdot [\mathbf{A}(\mathbf{x})^T \nabla \psi])_{\partial\Omega} && \text{(integration by parts)} \\
&= 0 && \text{(zero trace of } \mathbf{v}_h) \\
&\neq (\mathbf{g}, \hat{\mathbf{n}} \cdot [\mathbf{A}(\mathbf{x})\nabla \mathbf{v}_h])_{\partial\Omega}.
\end{aligned}
$$

The above is clearly non-vanishing even as the mesh is refined. This dual-inconsistent term forms the underlying obstruction to the convergence of discrete adjoint solution towards its continuous counterpart. In contrast to (3.23), for the choice of volume functional (3.22) the discrete adjoint ψ_h^{vol} satisfies the variational statement

$$R(\mathbf{v}_h, \psi_h^{\text{vol}} + \mathbf{w}_h^{-\mathbf{g}}) = 0, \quad \forall \mathbf{v}_h \in \mathcal{V}_h^0, \tag{3.25}$$

Hence, the PDE adjoint ψ satisfies the variational statement for $\psi_h^{\text{vol}} + \mathbf{w}_h^{-\mathbf{g}}$ as given above. Thus, this functional results in a dual-consistent formulation allowing the discrete adjoint to converge towards the continuous solution. To summarize, in this example the property of dual-consistency determines whether or not the discrete adjoint solution can approximate the continuous solution as the mesh is refined.

3.3 First-order conservation laws

Let $\Omega \subset \mathbb{R}^n$ be the computational domain of interest, $\mathcal{V}$, $\partial\mathcal{V}$ be function spaces on Ω and $\partial\Omega$ respectively such that there exists a bounded, surjective trace operator from $\mathcal{V}$ to $\partial\mathcal{V}$. Let $\mathbf{u} \in \mathcal{V}$ be a weak solution to the following conservation law,

$$\begin{cases} \nabla \cdot \mathcal{F}(\mathbf{u}) & = 0, \ \mathbf{x} \in \Omega, \\ \mathbf{D}(\mathbf{u}|_{\partial\Omega}, \mathrm{BC}) & = 0, \ \mathbf{x} \in \partial\Omega, \end{cases} \tag{3.26}$$

where the boundary operator $\mathbf{D}(\cdot, \mathrm{BC}) : \partial\mathcal{V} \to \partial\mathcal{V}$ imposes Dirichlet conditions with the associated input BC data. Following the prescription given in Section 3.1 and using integration by parts, it can be shown that the adjoint state ψ associated to the output $\mathcal{J}(\cdot) \equiv \int_{\Gamma_{\mathrm{output}}} J(\cdot) ds$, $\Gamma_{\mathrm{output}} \subset \partial\Omega$, satisfies the following PDE [31],

$$-\mathcal{F}'[\mathbf{u}]^T \cdot \nabla\psi = 0, \ \mathbf{x} \in \Omega, \tag{3.27}$$

subject to the adjoint boundary conditions,

$$\int_{\partial\Omega} \psi^T \hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}](\tilde{\mathbf{u}}) ds = \int_{\Gamma_{\mathrm{output}}} J'[\mathbf{u}](\tilde{\mathbf{u}}) ds, \ \forall \tilde{\mathbf{u}} \in \partial\mathcal{V}^0, \tag{3.28}$$

where $\partial\mathcal{V}^0$ denotes the trace space of $\mathcal{V}$ satisfying homogeneous Dirichlet condition,

$$\partial\mathcal{V}^0 \equiv \left\{ \tilde{\mathbf{u}} \in \partial\mathcal{V} : \mathbf{D}'[\mathbf{u}|_{\partial\Omega}](\tilde{\mathbf{u}}) = 0 \right\}. \tag{3.29}$$

The adjoint PDE (3.27) differs in character from the primal PDE (3.26) in that the former is not in conservative form, in contrast to the latter. However, the adjoint and primal BCs do exhibit a duality structure: components of variations $\tilde{\mathbf{u}}|_{\partial\Omega}$ allowed by the flow Dirichlet data give rise to constraints on $\psi|_{\partial\Omega}$, whereas components of $\tilde{\mathbf{u}}|_{\partial\Omega}$ fixed by the Dirichlet data do not constrain $\psi|_{\partial\Omega}$.

3.3.1 DG discretization

Denote $\mathcal{V}_h^p$ to be the space of discontinuous, vector-valued polynomials of degree p on a subdivision T_h of domain Ω into elements such that $\bar{\Omega} = \bigcup_{\kappa \in T_h} \bar{\kappa}$. The discontinuous Galerkin discretization for systems of conservation law (3.26) is given as: find $\mathbf{u}_h \in \mathcal{V}_h^p$ such that $\forall \mathbf{v}_h \in \mathcal{V}_h^p$,

$$R_h(\mathbf{u}_h, \mathbf{v}_h) \equiv \sum_{\kappa \in T_h} \left\{ - \int_\kappa \nabla \mathbf{v}_h^T \cdot \mathcal{F}(\mathbf{u}_h) d\mathbf{x} + \int_{\partial\kappa \backslash \partial\Omega} \mathbf{v}_h^{+T} \mathcal{H}(\mathbf{u}_h^+, \mathbf{u}_h^-, \hat{\mathbf{n}}) ds \right.$$
$$\left. + \int_{\partial\kappa \cap \partial\Omega} \mathbf{v}_h^{+T} \mathcal{H}^b(\mathbf{u}_h^+, \mathbf{u}_h^b, \hat{\mathbf{n}}) ds \right\} = 0, \tag{3.30}$$

where on $\partial\kappa$, $\hat{\mathbf{n}}$ is the outward pointing normal, $(\cdot)^+$ and $(\cdot)^-$ are the interior and exterior traces with respect to element κ. $\mathcal{H}(\mathbf{u}_h^+, \mathbf{u}_h^-, \hat{\mathbf{n}})$ and $\mathcal{H}^b(\mathbf{u}_h^+, \mathbf{u}_h^-, \hat{\mathbf{n}})$ denote numerical flux functions used on interior edges and $\partial\Omega$, respectively. The boundary conditions on $\partial\Omega$ are imposed weakly through constructing a boundary state trace that is a function of the inner state and BC data, $\mathbf{u}_h^b(\mathbf{u}_h^+, \text{BCData})$. The associated mapping on the space of discrete boundary traces, $\mathbf{u}_h^b(\cdot, \text{BCData}) : \partial\mathcal{V}_h^p \to \partial\mathcal{V}_h^p$, is assumed to be projective: $\mathbf{u}_h^b(\mathbf{u}_h^b(\cdot)) = \mathbf{u}_h^b(\cdot)$. Differentiation of the mapping shows that $\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \bigg|_{\mathbf{u}_h^b}$ is a linear projection operator on $\partial\mathcal{V}_h^p$:

$$\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \bigg|_{\mathbf{u}_h^b} \left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \bigg|_{\mathbf{u}_h^b} = \left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \bigg|_{\mathbf{u}_h^b}. \tag{3.31}$$

Hence, the eigenvalues of the map $\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \bigg|_{\mathbf{u}_h^b}$ are either 0 or 1.

Below, the dual-consistency of the DG discretization (3.30) together with functional implementations are examined. This is done for the two components of the discrete adjoint equations: firstly, those terms arising from the trace of test functions on the boundary domain are examined in Section 3.3.2; subsequently, those terms arising from the domain interior are examined in Section 3.3.3. For the scheme to be dual-consistent, both contributions have to vanish.

51

3.3.2 Boundary treatment

To show dual-consistency for appropriately chosen boundary flux and functional, a number of assumptions are made regarding the problem setting.

Assumption 3 *The discrete and continuous boundary traces satisfy the inclusion property:*

$$\partial \mathcal{V}_h^p \subset \partial \mathcal{V}. \tag{3.32}$$

The above is satisfied for instance if $\partial \mathcal{V} \subset L^2(\partial \Omega)$. Also assumed is a correspondence between the boundary state map $\mathbf{u}_h^b(\cdot, \mathrm{BCData})$ and the boundary operator $\mathbf{D}(\cdot, \mathrm{BC})$.

Assumption 4 *The domain of the map $\mathbf{u}_h^b$ is assumed to be extendable to $\partial \mathcal{V}$ so that the set of boundary traces satisfying Dirichlet BC are fixed points of the map:*

$$\mathbf{D}(\mathbf{u}|_{\partial\Omega}, \mathrm{BC}) = 0 \Rightarrow \mathbf{u}_h^b(\mathbf{u}|_{\partial\Omega}, \mathrm{BCData}) = \mathbf{u}|_{\partial\Omega}. \tag{3.33}$$

Furthermore, the following inclusion relation holds for the homogeneous trace spaces:

$$\partial \mathcal{V}_h^{p,0} \;\equiv\; \left\{ \mathbf{v}_h^+ \in \partial \mathcal{V}_h^p : \left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \mathbf{v}_h^+ = \mathbf{v}_h^+ \right\} \subset \partial \mathcal{V}^0. \tag{3.34}$$

Now we consider the following choice of boundary flux and discrete functional $\mathcal{H}^b(\cdot), \mathcal{J}_h(\cdot)$: $\partial \mathcal{V}_h^p \to \mathbb{R}$,

$$\begin{aligned}
\mathcal{H}^b(\cdot) &= \hat{\mathbf{n}} \cdot \mathcal{F}(\mathbf{u}_h^b(\cdot, \mathrm{BCData})) \\
\mathcal{J}_h(\cdot) &= \int_{\Gamma_{\mathrm{output}}} J(\mathbf{u}_h^b(\cdot, \mathrm{BCData})) ds.
\end{aligned} \tag{3.35}$$

Examine the discrete adjoint equations (3.12) arising from the boundary trace of test functions, $\mathbf{v}_h^+ \in \partial \mathcal{V}_h^p$. From the expression for the DG semilinear form (3.30), the boundary contribution to the bilinear form evaluated with solutions $\mathbf{u}$ and ψ is

$$R_h'[\mathbf{u}](\mathbf{v}_h|_{\partial\Omega}, \psi) = \int_{\partial\Omega} \psi^T \hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}]\left(\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \mathbf{v}_h^+ \right) ds. \tag{3.36}$$

Also, the functional contribution is simply

$$\mathcal{J}_h'[\mathbf{u}](\mathbf{v}_h|_{\partial\Omega}) = \int_{\Gamma_{\text{output}}} J'[\mathbf{u}](\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+}\right] \mathbf{v}_h^+) ds. \tag{3.37}$$

Consider first test functions $\mathbf{v}_h|_{\partial\Omega}$ in the quotient space $\partial\mathcal{V}_h^p \backslash \partial\mathcal{V}_h^{p,0}$. From the projection property (3.31) it follows that this space can be characterized by:

$$\partial\mathcal{V}_h^p \backslash \partial\mathcal{V}_h^{p,0} = \left\{ \mathbf{v}_h^+ \in \partial\mathcal{V}_h^p : \left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+}\right] \mathbf{v}_h^+ = 0 \right\}. \tag{3.38}$$

Since all discrete adjoint terms on the boundary are multiplied by $[\partial\mathbf{u}_h^b/\partial\mathbf{u}_h^+]$, these all vanish for $\mathbf{v}_h^+ \in \partial\mathcal{V}_h^p \backslash \partial\mathcal{V}_h^{p,0}$ and therefore ψ trivially satisfies the discrete adjoint equations on this subspace. Now consider $\mathbf{v}_h^+ \in \partial\mathcal{V}_h^{p,0}$. Since the adjoint solution ψ satisfies the boundary condition (3.28), it follows from the inclusion property (3.34) of Assumption 4 that the discrete dual variational statement is satisfied identically:

$$\int_{\partial\Omega} \psi^T \mathcal{H}^{b'}[\mathbf{u}](\mathbf{v}_h^+) ds = \int_{\Gamma_{\text{output}}} J'[\mathbf{u}](\mathbf{v}_h^+) ds, \quad \forall \mathbf{v}_h^+ \in \partial\mathcal{V}_h^{p,0},$$

thus verifying dual-consistency for the choice (3.35).

In the above demonstration, it is crucial for the boundary flux and functional to depend purely on $\mathbf{u}_h^b$. Suppose now the functional $\mathcal{J}_h(\cdot)$ is fixed as before but the same numerical flux function that is used on interior edges is also used on the boundary, so that there is explicit dependence on $\mathbf{u}_h^+$:

$$\mathcal{H}^b(\cdot) = \mathcal{H}(\cdot, \mathbf{u}_h^b(\cdot, \text{BCData}), \hat{\mathbf{n}}). \tag{3.39}$$

Whereas previously all terms vanish for $\mathbf{v}_h^+ \in \partial\mathcal{V}_h^p \backslash \partial\mathcal{V}_h^{p,0}$, now there exists the term

$$\int_{\partial\Omega} \psi^T \frac{\partial\mathcal{H}}{\partial\mathbf{u}_h^+}(\mathbf{v}_h^+) ds,$$

which is not required to vanish as part of the adjoint BC for ψ. Note that in particular, the above error term remains constant even as the mesh size $h \to 0$ and hence the

formulation is not asymptotically dual-consistent. This conclusion holds in general if proper care is not taken in the boundary treatment. In output-based adaptation, this could result in detrimental situations where the discrete adjoint solution error does not vanish even as the boundary elements are targeted for mesh refinement.

As another illustration, consider a more specialized case where the output of interest is in fact a functional of the boundary flux, $\mathcal{J}(\mathbf{u}) \equiv \mathcal{J}_{flux}(\hat{\mathbf{n}} \cdot \mathcal{F}(\mathbf{u}))$. To satisfy dual-consistency, it is not sufficient to evaluate the output based on the same boundary flux function as that used in the underlying numerical scheme. Assume that a generic numerical flux function is used on the domain boundary, from which the discrete output is computed. That is, the boundary flux and functional are given by:

$$
\begin{aligned}
\mathcal{H}^b(\cdot) &= \mathcal{H}(\cdot, \mathbf{u}_h^b(\cdot, \mathrm{BCData}), \hat{\mathbf{n}}), \\
\mathcal{J}_h(\cdot) &= \int_{\Gamma_{\mathrm{output}}} J_{flux}(\mathcal{H}^b(\cdot)) ds.
\end{aligned}
\tag{3.40}
$$

Then, the discrete adjoint equations on the boundary are:

$$
\begin{aligned}
&\int_{\partial\Omega} \psi^T \left(\frac{\partial \mathcal{H}}{\partial \mathbf{u}_h^+} \mathbf{v}_h^+ + \frac{\partial \mathcal{H}}{\partial \mathbf{u}_h^-} \left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \mathbf{v}_h^+ \right) ds \\
&= \int_{\Gamma_{\mathrm{output}}} J'_{flux}[\mathcal{H}](\frac{\partial \mathcal{H}}{\partial \mathbf{u}_h^+} \mathbf{v}_h^+ + \frac{\partial \mathcal{H}}{\partial \mathbf{u}_h^-} \left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \mathbf{v}_h^+) ds, \quad \forall \mathbf{v}_h^+ \in \partial\mathcal{V}_h^p.
\end{aligned}
\tag{3.41}
$$

Hence, for $\mathbf{v}_h^+ \in \partial\mathcal{V}_h^p \backslash \partial\mathcal{V}_h^{p,0}$ in which case the terms multiplying $\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right]$ vanish, the PDE adjoint ψ does not satisfy the above discrete adjoint equation if $\frac{\partial \mathcal{H}}{\partial \mathbf{u}_h^+} \neq 0$. Thus, it is still necessary that both the output and boundary flux are based purely on a boundary state map that has the necessary projection properties. Therefore, requiring dual-consistency is a stronger condition than that obtained by conservation arguments that are typically used to motivate the appropriate form of functional for good output estimates.

54

3.3.3 Interior treatment

A review of previous work on the discretization of the adjoint PDE is discussed in Appendix A. In particular, a non-conservative finite-volume formulation is described that was studied by Anderson and Venkatakrishnan [5]. In this section, terms of the DG discrete adjoint formulation are shown to correspond to this non-conservative formulation.

Making use of the conservative property $\mathcal{H}(\mathbf{u}, \mathbf{v}, \hat{\mathbf{n}}) = -\mathcal{H}(\mathbf{v}, \mathbf{u}, -\hat{\mathbf{n}})$, the discrete adjoint equations (3.13) for DG discretization (3.30) on interior elements ($\bar{\kappa} \cap \partial\Omega = \emptyset$) are: $\forall \mathbf{v}_h \in \mathcal{V}_h^p|_\kappa$,

$$-\int_\kappa \nabla \psi_h^T \cdot \mathcal{F}'[\mathbf{u}_h](\mathbf{v}_h) d\mathbf{x} + \int_{\partial\kappa} (\psi_h^+ - \psi_h^-)^T \frac{\partial\mathcal{H}}{\partial\mathbf{u}_h^+}(\mathbf{v}_h^+) ds = 0.$$

For the purpose of comparing the DG discrete adjoint formulation with the finite volume discretizations (A.1) of the adjoint PDE, rewrite the above using integration by parts:

$$\int_\kappa \psi_h^T \left(\nabla \cdot \mathcal{F}'[\mathbf{u}_h](\mathbf{v}_h) \right) d\mathbf{x} - \int_{\partial\kappa} \psi_h^{+T} \left(\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}_h^+](\mathbf{v}_h^+) \right) ds$$
$$+ \int_{\partial\kappa} (\psi_h^+ - \psi_h^-)^T \frac{\partial\mathcal{H}}{\partial\mathbf{u}_h^+}(\mathbf{v}_h^+) ds = 0. \tag{3.42}$$

The first term of the above has no correspondence in the finite volume schemes of Anderson *et al.* [5] and Baysal *et al.*[10]. However, this discrepancy only appears at higher-orders as it vanishes for piecewise-constant solutions. The remaining terms do have correspondence to these finite volume schemes: if the Roe flux is used for $\mathcal{H}$ with $\tilde{\mathbf{A}}$ denoting the flux Jacobian evaluated at the Roe-averaged state, the coefficient of $\mathbf{v}_h^+$ for the sum of second and third terms above is

$$\mathbf{G}_{\mathbf{v}_h^+} \equiv -\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}_h^+]^T \psi_h^+ + \left[\frac{\partial\mathcal{H}}{\partial\mathbf{u}_h^+}\right]^T (\psi_h^+ - \psi_h^-)$$

$$= -\frac{1}{2}\left[\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}_h^+]^T(\psi_h^+ + \psi_h^-)\right.$$

$$+\frac{\partial}{\partial \mathbf{u}_h^+}\left(|\hat{\mathbf{n}}\cdot\tilde{\mathbf{A}}(\mathbf{u}_h^+,\mathbf{u}_h^-)|(\mathbf{u}_h^+ - \mathbf{u}_h^-)\right)^T(\boldsymbol{\psi}_h^+ - \boldsymbol{\psi}_h^-)\Big], \qquad (3.43)$$

which is analogous to the adjoint flux expression given in (A.2).

3.4 Second-order elliptic systems

In Section 3.4.1 the DG treatment of second-order operators proposed by Bassi-Rebay is given. Then, in Section 3.4.2 dual-consistency analysis is carried out for boundary terms arising from the formulation and functional. Subsequently, the interior terms are examined in Section 3.4.3.

Let Ω denote the domain of interest and consider $\mathbf{u} \in [H^1(\Omega)]^m$ a weak solution to the following m-component second-order system,

$$\nabla \cdot (\boldsymbol{\mathcal{A}}_v(\mathbf{u})\nabla\mathbf{u}) = 0, \quad \mathbf{x} \in \Omega, \qquad (3.44)$$

with $\boldsymbol{\mathcal{A}}_v(\mathbf{u})$ the coefficient matrix, subject to Dirichlet and Neumann boundary conditions:

$$\begin{cases} \mathbf{D}(\mathbf{u}|_{\partial\Omega}, \mathrm{BC}) & = 0, \quad \mathbf{x} \in \partial\Omega, \\ \mathbf{N}(\hat{\mathbf{n}}\cdot\boldsymbol{\mathcal{A}}_v(\mathbf{u})\nabla\mathbf{u}|_{\partial\Omega}, \mathrm{BC}) & = 0, \quad \mathbf{x} \in \partial\Omega, \end{cases}$$

where $\mathbf{D}(\cdot, \mathrm{BC}) : \left[H^{1/2}(\partial\Omega)\right]^m \rightarrow \left[H^{1/2}(\partial\Omega)\right]^m$ and $\mathbf{N}(\cdot, \mathrm{BC}) : \left[H^{-1/2}(\partial\Omega)\right]^m \rightarrow \left[L^2(\partial\Omega)\right]^m$ are boundary operators. Consider outputs of interest of the following general form,

$$\mathcal{J}(\mathbf{u}, \nabla\mathbf{u}) \equiv \int_{\Gamma_{\text{output}}} [J_I(\mathbf{u}) + J_{II}(\hat{\mathbf{n}}\cdot\boldsymbol{\mathcal{A}}_v\nabla\mathbf{u})]\, ds, \qquad (3.45)$$

where $J_I(\cdot) : [L^2(\partial\Omega)]^m \rightarrow \mathbb{R}$, $J_{II}(\cdot) : \left[H^{-1/2}(\partial\Omega)\right]^m \rightarrow \mathbb{R}$ are bounded functionals. Via the prescription given in Section 3.1 and integration by parts, it can be shown that the corresponding adjoint state $\boldsymbol{\psi} \in [H^1(\Omega)]^m$ solves the following PDE in the

weak sense,

$$\nabla \cdot \left(\boldsymbol{\mathcal{A}}_v^T \nabla \psi \right) - \left(\boldsymbol{\mathcal{A}}_v' \nabla \mathbf{u} \right)^T \cdot \nabla \psi \;\; = 0, \;\; \mathbf{x} \in \Omega. \tag{3.46}$$

Adjoint boundary condition are obtained by examining the permissible boundary variations. Dirichlet BCs on $\psi|_{\partial\Omega} \in \left[H^{1/2}(\partial\Omega) \right]^m$ are obtained from the primal boundary flux variations: $\forall \tilde{\mathbf{u}} \in \left[H_0^{-1/2}(\partial\Omega) \right]^m$,

$$\int_{\partial\Omega} \tilde{\mathbf{u}}^T \psi ds = \int_{\Gamma_{\text{output}}} J_{II}'[\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}](\tilde{\mathbf{u}}) ds, \tag{3.47}$$

and Robin BCs are obtained from the permissible state variations: $\forall \tilde{\mathbf{u}} \in \left[H_0^{1/2}(\partial\Omega) \right]^m$,

$$-\int_{\partial\Omega} \tilde{\mathbf{u}}^T \left(\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v^T \nabla \psi \right) ds + \int_{\partial\Omega} \left(\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v'[\mathbf{u}](\tilde{\mathbf{u}}) \nabla \mathbf{u} \right)^T \psi ds = \int_{\Gamma_{\text{output}}} J_I'[\mathbf{u}](\tilde{\mathbf{u}}) ds, \tag{3.48}$$

where the homogeneous trace spaces are defined as,

$$\left[H_0^{1/2}(\partial\Omega) \right]^m \equiv \left\{ \tilde{\mathbf{u}} \in \left[H^{1/2}(\partial\Omega) \right]^m : \; \mathbf{D}'[\mathbf{u}|_{\partial\Omega}, \text{BC}](\tilde{\mathbf{u}}) = 0 \right\},$$

$$\left[H_0^{-1/2}(\partial\Omega) \right]^m \equiv \left\{ \tilde{\mathbf{u}} \in \left[H^{-1/2}(\partial\Omega) \right]^m : \; \mathbf{N}'[\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}|_{\partial\Omega}, \text{BC}](\tilde{\mathbf{u}}) = 0 \right\}.$$

The variational characterizations (3.47) and (3.48) imply that the corresponding adjoint BCs are satisfied almost everywhere on the boundary.

3.4.1 DG discretization

The viscous terms are treated using the second form of Bassi and Rebay (BR2) [9] discretization. The notations $[\![\cdot]\!]$, $\{ \cdot \}$ are used to denote jump and average operators on interior faces,

$$[\![s]\!] \equiv s^+ \hat{\mathbf{n}}^+ + s^- \hat{\mathbf{n}}^-, \qquad \{ \varphi \} \equiv \frac{1}{2}(\varphi^+ + \varphi^-), \tag{3.49}$$

where s is a scalar quantity and φ is a vector quantity. The semilinear form for the discretization of (3.44) is given by:

$$
\begin{aligned}
R_h(\mathbf{u}_h, \mathbf{v}_h) \equiv \;& -\sum_{\kappa \in T_h} \int_\kappa \nabla \mathbf{v}_h^T \cdot (\boldsymbol{\mathcal{A}}_v(\mathbf{u}_h) \nabla \mathbf{u}_h) dx \\
& + \int_{\Gamma_i} \Big([\![\mathbf{u}_h]\!]^T \cdot \{\boldsymbol{\mathcal{A}}_v(\mathbf{u}_h)^T \nabla \mathbf{v}_h\} + [\![\mathbf{v}_h]\!]^T \cdot \{\boldsymbol{\mathcal{A}}_v(\mathbf{u}_h) \nabla \mathbf{u}_h\} \Big) ds \\
& - \int_{\Gamma_i} \eta_f [\![\mathbf{v}_h]\!]^T \cdot \{\boldsymbol{\delta}_f\} ds - \int_{\partial\Omega} (\mathbf{u}_h^b - \mathbf{u}_h^+)^T (\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_h^b)^T \nabla \mathbf{v}_h^+) ds \\
& + \int_{\partial\Omega} \mathbf{v}_h^{+T} \mathbf{F}_h^b \left(\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_h^b) \nabla \mathbf{u}_h^+ - \eta_f \hat{\mathbf{n}} \cdot \boldsymbol{\delta}_f^b \right) ds,
\end{aligned}
\tag{3.50}
$$

where η_f is a stabilization parameter, $\boldsymbol{\delta}_f$ and $\boldsymbol{\delta}_f^b$ are auxiliary variables associated with interior and boundary faces, defined by the following weak statements: find $\boldsymbol{\delta}_f \in [\mathcal{V}_h^p]^2$ such that $\forall \boldsymbol{\tau}_h \in [\mathcal{V}_h^p]^2$,

$$
\int_\Omega \boldsymbol{\tau}_h^T \cdot \boldsymbol{\delta}_f \, dx = \int_{\sigma_f} [\![\mathbf{u}_h]\!]^T \cdot \{\boldsymbol{\mathcal{A}}_v(\mathbf{u}_h)^T \boldsymbol{\tau}_h\} \, ds,
\tag{3.51}
$$

and $\boldsymbol{\delta}_f^b \in [\mathcal{V}_h^p]^2$ such that

$$
\int_\Omega \boldsymbol{\tau}_h^T \cdot \boldsymbol{\delta}_f^b dx = \int_{\sigma_f} (\mathbf{u}_h^+ - \mathbf{u}_h^b)^T [\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_h^b)^T \boldsymbol{\tau}_h^+] \, ds,
\tag{3.52}
$$

where σ_f denotes the face indexed by f. Neumann BCs are set via a boundary flux map $\mathbf{F}_h^b(\cdot, \text{BCData}) : \partial(\nabla \mathcal{V}_h^p) \to \partial(\nabla \mathcal{V}_h^p)$ on the space of gradient traces.

3.4.2 Boundary treatment

In the present setting, because the DG trace space on the boundary is discontinuous and hence does not have the required regularity to lie in $\left[H^{1/2}(\partial\Omega) \right]^m$, (3.32) of Assumption 3 no longer holds. However, as before the boundary state map $\mathbf{u}_h^b(\cdot, \text{BCData})$ is assumed to satisfy the correspondence assumption of (4) with the boundary operator $\mathbf{D}(\cdot, \text{BC})$. Analogously, the boundary flux map $\mathbf{F}_h^b(\cdot, \text{BCData})$ is assumed to be extendable to $\left[H^{-1/2}(\partial\Omega) \right]^m$ so as to satisfy the following correspondence with $\mathbf{N}(\cdot, \text{BC})$.

Assumption 5 *The set of boundary fluxes satisfying Neumann BC are a fixed points of the map:*

$$\mathbf{N}(\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}|_{\partial\Omega}, \mathrm{BC}) = 0 \Rightarrow \mathbf{F}_h^b(\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}|_{\partial\Omega}, \mathrm{BCData}) = \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}|_{\partial\Omega}. \quad (3.53)$$

Furthermore,

$$\partial(\nabla \mathcal{V}_h^p)^0 \equiv \left\{ \nabla \mathbf{v}_h^+ \in \partial(\nabla \mathcal{V}_h^p) : \left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \nabla \mathbf{v}_h^+ = \nabla \mathbf{v}_h^+ \right\} \subset \left[H_0^{-1/2}(\partial\Omega) \right]^m. \quad (3.54)$$

Here, we consider the following form of discrete output:

$$\mathcal{J}_h(\mathbf{u}_h, \nabla \mathbf{u}_h) \equiv \int_{\Gamma_{\text{output}}} \left[J_I(\mathbf{u}_h^b(\mathbf{u}_h^+)) + J_{II}(\mathbf{F}_h^b(\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}_h^+ - \eta_f \hat{\mathbf{n}} \cdot \boldsymbol{\delta}_f^b)) \right] ds. \quad (3.55)$$

To obtain the boundary contribution to the discrete adjoint residual, the expression for the BR2 semilinear form is first rewritten via integration by parts. Evaluating $\mathbf{u}$, $\boldsymbol{\psi}$ in the discrete adjoint residual together with the fact that on a continuous primal state $\boldsymbol{\delta}_f^b = 0$, gives the following expressions involving $\mathbf{v}_h^+$ and $\nabla \mathbf{v}_h^+$ respectively:

$$
\begin{aligned}
\mathbf{v}_h^+ \quad : \quad & -\int_{\partial\Omega} \left(\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \mathbf{v}_h^+ \right)^T \boldsymbol{\mathcal{A}}_v^T \nabla \boldsymbol{\psi}\, ds \\
& + \int_{\partial\Omega} \left(\left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v'[\mathbf{u}] \left(\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \mathbf{v}_h^+ \right) \nabla \mathbf{u} \right)^T \boldsymbol{\psi}\, ds \\
& + \int_{\partial\Omega} \left(\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \mathbf{v}_h^+ - \mathbf{v}_h^+ \right)^T (\eta_f \hat{\mathbf{n}} \cdot \boldsymbol{\delta}_f^{b,du})\, ds - \int_{\Gamma_{\text{output}}} J_I'[\mathbf{u}] \left(\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] \mathbf{v}_h^+ \right) ds,
\end{aligned}
$$

$$
\begin{aligned}
\nabla \mathbf{v}_h^+ \quad : \quad & \int_{\partial\Omega} \boldsymbol{\psi}^T \left(\left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{v}_h^+ \right) ds \\
& - \int_{\Gamma_{\text{output}}} J_{II}'[\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}] \left(\left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{v}_h^+ \right) ds,
\end{aligned}
$$

where the auxiliary variables $\boldsymbol{\delta}_f^{b,du}$ are defined via: $\forall \boldsymbol{\tau}_h \in [\mathcal{V}_h^p]^2$,

$$\int_\Omega \boldsymbol{\tau}_h^T \cdot \boldsymbol{\delta}_f^{b,du} d\mathbf{x} = \int_{\sigma_f} (\boldsymbol{\psi}^T - J_{II}'[\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}]) \left(\left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \boldsymbol{\tau}_h^+ \right) ds. \quad (3.56)$$

It remains to check whether the above vanish for all traces of test functions and its gradients in the respective discrete spaces, $\mathbf{v}_h^+ \in \partial \mathcal{V}_h^p$ and $\nabla \mathbf{v}_h^+ \in \partial (\nabla \mathcal{V}_h^p)$. Firstly, note that $\delta_f^{b,du} = 0$ by the Dirichlet condition (3.47) on ψ and the inclusion relation (3.54), hence has a zero contribution to the residual expression involving $\mathbf{v}_h^+$. Also, the discrete adjoint equations are trivially satisfied for $\mathbf{v}_h^+ \in \partial \mathcal{V}_h^p \backslash \partial \mathcal{V}_h^{p,0}$ and $\nabla \mathbf{v}_h^+ \in \partial (\nabla \mathcal{V}_h^p) \backslash \partial (\nabla \mathcal{V}_h^p)^0$. Otherwise, using the adjoint BCs (3.47) and (3.48) on ψ it follows that all boundary terms vanish as well. Thus, the discrete formulation and functional (3.55) constitute a dual-consistent treatment. The above analysis also shows the importance of including δ_f^b in the output (3.55): without it, $\delta_f^{b,du}$ as defined in (3.56) is non-zero in general and therefore constitutes a non-vanishing error term.

3.4.3   Interior treatment

To study the dual-consistency of interior treatment, a stronger regularity assumption of $\mathbf{u}, \psi \in [H^2(\Omega)]^m$ is made, as is done in Arnold *et al.* [6]. For the case that the coefficient matrix has no $\mathbf{u}$ dependence, $\mathcal{A}_v = \mathcal{A}_v(\mathbf{x})$, dual-consistency of BR2 scheme follows from the analysis carried out in [6]. Hence, it remains to examine only the $\mathbf{u}$ dependent contribution of the coefficient matrix. Firstly, since $\mathbf{u}$ is continuous across element faces, the contribution arising from the auxiliary variable δ_f is zero. The remaining terms are,

$$
-\sum_{\kappa \in T_h} \int_\kappa (\mathcal{A}_v'[\mathbf{u}](\mathbf{v}_h) \nabla \mathbf{u})^T \cdot \nabla \psi dx
$$

$$
+ \int_{\Gamma_i} \left([\![\mathbf{u}]\!]^T \cdot \{\mathcal{A}_v'[\mathbf{u}](\mathbf{v}_h)^T \nabla \psi\} + [\![\psi]\!]^T \cdot \{\mathcal{A}_v'[\mathbf{u}](\mathbf{v}_h) \nabla \mathbf{u}\} \right).ds \qquad (3.57)
$$

where the first term is dual-consistent and the remaining terms involving $[\![\mathbf{u}]\!]$, $[\![\psi]\!]$ are both zero as a result of the regularity assumption. To conclude, the BR2 scheme is shown to be dual-consistent for the nonlinear second-order system of the form (3.44), irrespective of the the stabilization parameter η_f.

3.5 Computational demonstrations

In this section, the general analysis carried out in the previous sections is verified via numerical experiments. In particular, the effect of dual-consistency arising from different boundary treatments is demonstrated. This is first done in Section 3.5.1 for the case of inviscid Euler equations, followed by Section 3.5.2 with the inclusion of viscous terms in the compressible Navier-Stokes equations.

3.5.1 Inviscid Euler equations

For 2D Euler's flow equations the conservative state $\mathbf{u}$ is given by,

$$
\mathbf{u} = \begin{pmatrix} \rho \\ \rho u \\ \rho v \\ \rho E \end{pmatrix},
$$

and the expression for the fluxes $\boldsymbol{\mathcal{F}} = [\mathcal{F}^x, \mathcal{F}^y]$ are,

$$
\mathcal{F}^x = \begin{pmatrix} \rho u \\ \rho u^2 + p \\ \rho u v \\ \rho u H \end{pmatrix}, \quad \mathcal{F}^y = \begin{pmatrix} \rho v \\ \rho u v \\ \rho v^2 + p \\ \rho v H \end{pmatrix}.
$$

The total enthalpy is given by $H = E + p/\rho$ and the pressure given by

$$
p = (\gamma - 1) \left[\rho E - \frac{1}{2}\rho \left(u^2 + v^2 \right) \right].
$$

Discrete adjoint regularity

The implication of boundary treatment for discrete adjoint solution behavior is demonstrated here for a smooth, subsonic Euler flow test case. A standard DG formulation using Roe flux function is used on uniformly-spaced grids of 587, 2325 and 9169 el-

ements with $q = 4$ boundary representation [67] to discretize a 2D, $M_\infty = 0.5$ duct flow over a Gaussian-shaped bump. The flow domain is described as:

$$\Omega = \left\{ (x,y) \in (-6,6) \times (0,6) : y > \frac{1}{5}e^{-2x^2} \right\}. \tag{3.58}$$

As the output of interest, a weighted lift on the bump surface of the duct is used:

$$\mathcal{J}(\mathbf{u}) = \int_{x \in [-6,6],\ y = \frac{1}{5}e^{-2x^2}} n_y p(\mathbf{u}) e^{-\frac{1}{2}x^2} ds, \tag{3.59}$$

where the Gaussian weight is used to localize the output around the bump portion of the channel. On the inflow boundary the total pressure and temperature are prescribed and at the outflow boundary the pressure is prescribed. At the duct walls, the flow-tangency condition is set. To impose the boundary conditions and compute outputs in a dual-consistent manner, the functional form (3.35) is used. Figure 3-1 shows the x-momentum component of the discrete adjoint solution, $\psi(2)$, computed using $p = 3$ interpolation on the fine mesh. As an example of boundary treatments that are dual-inconsistent, consider evaluating $\mathbf{u}_h^+$ in the following:

$$\begin{aligned}
\mathcal{H}^b(\cdot) &= \mathcal{H}(\cdot, \mathbf{u}_h^b(\cdot, \mathrm{BCData}), \hat{\mathbf{n}}), \\
\mathcal{J}_h(\cdot) &= \int_{\Gamma_{\mathrm{output}}} J(\cdot) ds.
\end{aligned} \tag{3.60}$$

In this case, both the boundary flux and output has explicit dependence on $\mathbf{u}_h^+$ rather than purely through $\mathbf{u}_h^b$ as the analysis requires. As Figure 3-2 shows, the resulting discrete adjoint using the same mesh and solution order has significant irregularity near the domain boundary, as well as large pollution error in the whole domain. Consider also a "conservative" but dual-inconsistent treatment described previously whose functional form is given by expression (3.40). As Figure 3-3 shows, in this case the discrete adjoint solution similarly has spurious irregularity near the domain boundary.

62

Figure 3-1: Adjoint behavior: dual-consistent boundary treatment using only $\mathbf{u}_h^b(\mathbf{u}_h^+)$. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.

Figure 3-2: Adjoint behavior: dual-inconsistent boundary treatment using both $\mathbf{u}_h^+$ and $\mathbf{u}_h^b(\mathbf{u}_h^+)$. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.

Figure 3-3: Adjoint behavior: conservative but dual-inconsistent boundary treatment based on numerical flux function. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.

Output convergence

Here, output convergence for the previously described duct flow test case is examined. The "truth" solution is taken to be the $p = 5$ result on the fine, 9169 element mesh. For the computational result shown in Figure 3-4, the dual-consistent treatment given in (3.35) is used. It is shown that at least for $p = 1$ and $p = 2$ solutions, the boundary output converges at the superconvergent rate of $\mathcal{O}(h^{2p+1})$ which is significantly higher than the $\mathcal{O}(h^{p+1})$ rate achieved for the $L_2(\Omega)$ norm of the solution error for smooth cases. For the $p = 3$ case, the convergence rate does not attain the expected value of $\mathcal{O}(h^{2p+1})$. This may be due to a number of reasons. Firstly, the boundary is represented using $q = 4$ interpolation and it may be that higher order is necessary to attain the full convergence rate. Secondly, there may be insufficient smoothness in either the flow or adjoint solution (possibly arising from the four corners of the channel), thereby hindering the expected superconvergence. For comparison, the output convergence using the dual-inconsistent boundary treatment (3.60) is given in Figure 3-5. The "true" output is the same as that used in the dual-consistent convergence plot of Figure 3-4. However, the plot is virtually indistinguishable if instead, the "true" output is taken to be the dual-inconsistent result using $p = 5$ interpolation on the same 9169 element mesh. It turns out that the error arising from dual-inconsistency for $p = 1$ to $p = 3$ solutions are simply significantly larger than corresponding output discrepancies on the $p = 5$ solution. In the dual-inconsistent case the output convergence rate obtained from the $p = 1$ solution appears to be only second-order. Furthermore, within the range of mesh refinements examined degradation in convergence becomes more significant at higher-orders. In fact, a conclusion that might be drawn for this dual-inconsistent treatment is that, for the purpose of obtaining accurate functional outputs, the benefit of higher-order is dubious. Consider also the conservative functional treatment given by expression (3.40), whose convergence is given in Figure 3-6. Comparison with Figure 3-5 shows that with this choice of flux-based functional the convergence of outputs are significantly better. Nevertheless, the full superconvergent rate of $\mathcal{O}(h^{2p+1})$ is not achieved. Instead, it appears that the convergence rate is only
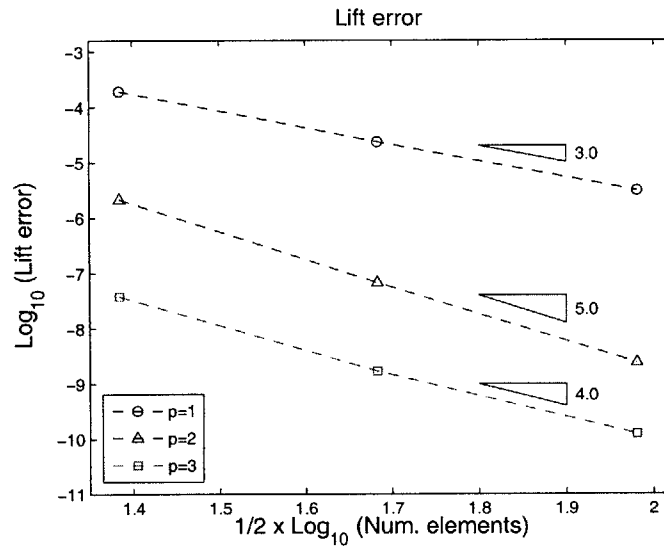
Figure 3-4: Output convergence: dual-consistent boundary treatment using only $\mathbf{u}_h^b(\mathbf{u}_h^+)$. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.

$\mathcal{O}(h^{p+1})$ (at least for $p = 1$ and $p = 2$ solutions). Thus, the overall conclusion is that to attain the best convergence rate the functional should be not only conservative but satisfy dual-consistency as well.

Figure 3-5: Output convergence: dual-inconsistent boundary treatment using both $\mathbf{u}_h^+$ and $\mathbf{u}_h^b(\mathbf{u}_h^+)$. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.
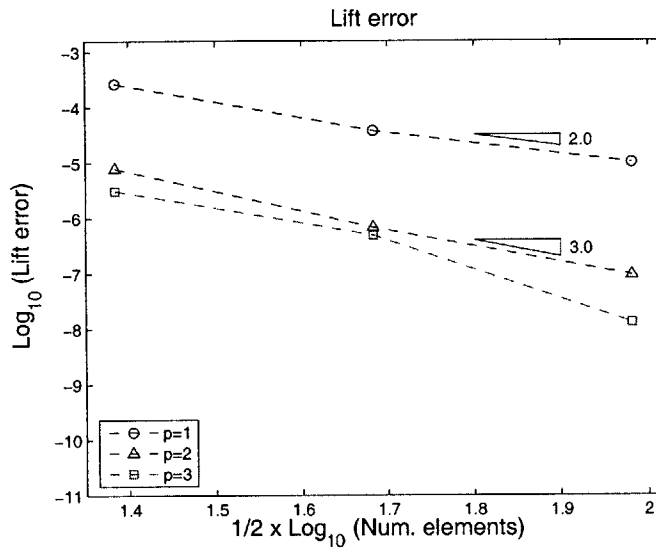
Figure 3-6: Output convergence: conservative but dual-inconsistent boundary treatment based on numerical flux function. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.

3.5.2 Compressible Navier-Stokes equations

Here, the x- and y-components of the viscous flux contribution are given respectively by,

$$\hat{\mathbf{i}} \cdot \mathcal{A}_v(\mathbf{u})\nabla\mathbf{u} = \begin{pmatrix} 0 \\ \frac{2}{3}\mu(\mathbf{u})(2\frac{\partial u}{\partial x} - \frac{\partial v}{\partial y}) \\ \mu(\mathbf{u})(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}) \\ \frac{2}{3}\mu(\mathbf{u})(2\frac{\partial u}{\partial x} - \frac{\partial v}{\partial y})u + \mu(\mathbf{u})(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x})v + \kappa(\mathbf{u})\frac{\partial T}{\partial x} \end{pmatrix}$$

$$\hat{\mathbf{j}} \cdot \mathcal{A}_v(\mathbf{u})\nabla\mathbf{u} = \begin{pmatrix} 0 \\ \mu(\mathbf{u})(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}) \\ \frac{2}{3}\mu(\mathbf{u})(2\frac{\partial v}{\partial y} - \frac{\partial u}{\partial x}) \\ \frac{2}{3}\mu(\mathbf{u})(2\frac{\partial v}{\partial y} - \frac{\partial u}{\partial x})v + \mu(\mathbf{u})(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x})u + \kappa(\mathbf{u})\frac{\partial T}{\partial y} \end{pmatrix}.$$

where $\kappa(\mathbf{u})$ is the thermal conductivity

$$\kappa(\mathbf{u}) = \frac{c_p}{Pr}\mu(\mathbf{u}), \tag{3.61}$$

Pr the Prandtl number, c_v the specific heat at constant volume and $\mu(\mathbf{u})$ is the dynamic viscosity, determined via Sutherland's law:

$$\mu(\mathbf{u}) = \mu_0 \frac{T_0 + S}{T + S}\left(\frac{T}{T_0}\right)^{3/2}, \tag{3.62}$$

where μ_0, T_0 are free-stream temperature and viscosity, S the Sutherland's constant.

Discrete adjoint regularity

The implication of boundary treatment in the discretization of second-order operators is demonstrated here for a subsonic, compressible Navier-Stokes test case. In addition to the standard treatment of inviscid terms, the BR2 discretization of viscous contributions as described in Section 3.4.1 is used to approximate a $M_\infty = 0.5$, $Re = 5000$ laminar flow over a NACA 0012 airfoil at angle of attack $\alpha = 2°$. The results shown here are obtained using structured grids containing 672, 2688, and 10752 elements

with $q = 3$ boundary representation generated by Oliver [67] from a baseline grid provided by Swanson [73]. The finest, 10752 element mesh used in the following results is shown in Figure 3-7. The airfoil surface is analytically defined as

$$y = \pm 0.6(0.2969\sqrt{x} - 0.1260x - 0.3516x^2 + 0.2843x^3 - 0.1036x^4), \quad x \in [0,1].$$

On the airfoil surface, the Dirichlet boundary condition of zero velocity together with the Neumann condition of adiabatic wall are set. On the inflow boundary, the Dirichlet conditions consist of prescribing the total temperature and pressure together with the flow angle. On the outflow boundary, the Dirichlet condition of prescribing the static pressure together with the Neumann condition of zero viscous momentum fluxes are set. The drag functional output is of the generic form (3.45) where $J_I(\mathbf{u})$, $J_{II}(\hat{\mathbf{n}} \cdot \mathcal{A}_v \nabla \mathbf{u})$ consist of the pressure and viscous stress contributions respectively. As the previous analysis shows, dual-consistent treatment of the discrete functional is of the form given in (3.55). Figure 3-8 shows the corresponding discrete adjoint solution behavior. As an example of dual-inconsistent boundary treatment, consider the following form of discrete functional where the auxiliary variable contributions are ignored:

$$\mathcal{J}_h(\mathbf{u}_h, \nabla \mathbf{u}_h) \equiv \int_{\Gamma_{\text{output}}} \left[J_I(\mathbf{u}_h^b(\mathbf{u}_h^+)) + J_{II}(\mathbf{F}_h^b(\hat{\mathbf{n}} \cdot \mathcal{A}_v \nabla \mathbf{u}_h^+)) \right] ds. \tag{3.63}$$

The top plot of Figure 3-9 shows the first component of the adjoint solution around the leading edge of the airfoil computed using $p = 3$ interpoluation on the fine, 10752 element mesh. As can be clearly seen, the resulting discrete adjoint has an irregular layer in the elements lying on the boundary. Furthermore, the error stemming from boundary treatment results in significant pollution error around the mid-section of the airfoil as can be seen by comparing Figure 3-8 and the bottom plot of Figure 3-9. The plots clearly demonstrate the need to include the stabilization contribution δ_f^b in the functional even when fine meshes and high-order interpolations are used.

Figure 3-7: Fine NACA 0012 grid, 10752 elements.

Figure 3-8: Adjoint behavior: dual-consistent treatment with the inclusion of δ_f^b in functional. Laminar flow over NACA 0012 airfoil, $M_\infty = 0.5$, $Re = 5000$, $\alpha = 2.0°$.

(a) Leading edge close-up

(b) Full view

Figure 3-9: Adjoint behavior: dual-inconsistent treatment without the inclusion of $\boldsymbol{\delta}_f^b$ in functional. Laminar flow over NACA 0012 airfoil, $M_\infty = 0.5$, $Re = 5000$, $\alpha = 2.0°$.

Output convergence

Here, output convergence for the previously described functional implementations are compared to study the effect of solution post-processing. Figure 3-10 shows the drag convergence on the NACA 0012 airfoil using the dual-consistent functional. A sequence of $p = 1$ to $p = 3$ solutions are compared to the "truth" solution taken to be the $p = 4$ result on the fine, 10752 element mesh. Here, the line-fit through $p = 1$ result matches the expected $\mathcal{O}(h^{2p})$ rate while the $p = 2$ result is somewhat better than expected. The $p = 3$ result does not attain the full expected rate of $\mathcal{O}(h^6)$ in going from the medium to fine mesh; this may possibly due to the lack of flow or adjoint solution regularity near the trailing edge. Figure 3-11 shows the drag convergence using the above described dual-inconsistent functional. Again, the "truth" value is the same as that used for Figure 3-10, although the plot is virtually indistinguishable if the dual-inconsistent output is used instead. In comparison, the $p = 1$ rate is significantly regraded especially as the mesh is refined. The $p = 2$ rate is less affected, but nevertheless the convergence line is shifted upwards. The $p = 3$ convergence is clearly quite erratic. In fact, the error on the dual-consistent drag for the $p = 2$ solution on the fine, 10752 element mesh is smaller by a factor of 5 than that of using an inconsistent functional to post-process the $p = 3$ solution on the same mesh. Hence, it is important to use dual-consistent solution post-processing to fully benefit from the use of high-order solution.

72

Figure 3-10: Drag convergence: dual-consistent treatment with the inclusion of δ_f^b in functional. Laminar flow over NACA 0012 airfoil, $M_\infty = 0.5$, $Re = 5000$, $\alpha = 2.0°$.

Figure 3-11: Drag convergence: dual-inconsistent treatment without δ_f^b in functional. Laminar flow over NACA 0012 airfoil, $M_\infty = 0.5$, $Re = 5000$, $\alpha = 2.0°$.

Chapter 4

Error estimation and adaptation

4.1 Background

Again consider the general setting introduced in Chapter 3: let $\mathbf{u} \in \mathcal{V}$ be a weak solution to

$$F(\mathbf{u}) = 0, \tag{4.1}$$

where $F : \mathcal{V} \to \mathcal{W}'$ is a differential operator mapping $\mathcal{V}$ to $\mathcal{W}'$. Let $\mathcal{J}(\cdot) : \mathbf{u} \in \mathcal{V} \to \mathbb{R}$ be a general nonlinear functional of interest. The goal of this chapter is to derive general error estimates for functional outputs using DG approximations of the underlying PDE (4.1), with specific results pertaining to the inviscid Euler and compressible Navier-Stokes equations. In the context of DG discretization for first-order hyperbolic systems, duality-based analysis has been carried out by Hartmann and Houston [40], Barth [8], Giles and Süli [33]. For second-order differential systems, the same analysis has been applied to the advection-diffusion-reaction [39] and compressible Navier-Stokes equations [41] using the symmetric interior penalty discretization of the diffusive operator. Common in the analysis is the use of consistency of the DG scheme, that is $\mathbf{u}$ of (4.1) satisfies the finite-dimensional weak statement:

$$R_H(\mathbf{u}, \mathbf{v}_H) = 0, \quad \forall \mathbf{v}_H \in \mathcal{V}_H^P, \tag{4.2}$$

where $R_H(\cdot,\cdot) : \mathcal{V}_H^P \times \mathcal{V}_H^P \to \mathbb{R}$ is the mesh-dependent DG semilinear form (linear in the second argument) and $\mathcal{V}_H^P$ is the element-wise polynomial function space on the (coarse) working mesh. Assume the functional $\mathcal{J}(\cdot)$ may be extended to be defined on $\mathcal{V}_H^P$ as well. The error in the value of the functional output may be written in the following manner:

$$
\begin{aligned}
\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H) &= \int_0^1 \mathcal{J}'[\theta\mathbf{u} + (1-\theta)\mathbf{u}_H](\mathbf{u} - \mathbf{u}_H)d\theta, \\
&\equiv \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H).
\end{aligned}
\tag{4.3}
$$

where the mean-value linearization $\bar{\mathcal{J}}$ is introduced, which is linear in the last argument. Similarly, the DG semilinear form may be written as:

$$
\begin{aligned}
R_H(\mathbf{u}, \mathbf{v}) - R_H(\mathbf{u}_H, \mathbf{v}) &= \int_0^1 R_H'[\theta\mathbf{u} + (1-\theta)\mathbf{u}_H](\mathbf{u} - \mathbf{u}_H, \mathbf{v})d\theta \\
&\equiv \bar{R}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{u} - \mathbf{u}_H, \mathbf{v}).
\end{aligned}
\tag{4.4}
$$

Let ψ denote the dual state for the mean-value linearized functional and DG semilinear form,

$$
\bar{R}_H(\mathbf{u}, \mathbf{u}_H; \mathbf{v}, \psi) = \bar{\mathcal{J}}(\mathbf{u}, \mathbf{u}_H; \mathbf{v}), \quad \forall \mathbf{v} \in \mathcal{V},
\tag{4.5}
$$

then it may be shown that the error in the output is given by:

$$
\mathcal{J}(\mathbf{u}) - \mathcal{J}(\mathbf{u}_H) = -R_H(\mathbf{u}_H, \psi - \psi_H).
\tag{4.6}
$$

Since the exact primal solution $\mathbf{u}$ in (4.5) is unknown, in practice some approximation has to be made for the mean-value linearization. Hartmann *et al.* employs the straightforward strategy of replacing the mean-value linearizations with linearizing at $\mathbf{u}_H$. Barth [8] employs both this method and the more involved technique of replacing path integration in the definition of $\bar{R}_H, \bar{\mathcal{J}}$ with numerical quadrature between the working approximation and a higher-order approximation obtained via interpolation. Although these approximation approaches have been demonstrated, there does

76

not appear to be conclusive evidence indicating the advantage of one over the other. Also needed in the expression (4.6) is an approximation to the adjoint solution ψ. Hartmann *et al.* computes a more refined approximation to the adjoint by using a solution order that is globally one polynomial degree higher but on the same underlying mesh. In the context of Godunov finite volume methods, Barth [8] examined the use of adjoint approximation via reconstructive interpolation on a test case using 2D linear advection equation for which analytical expressions for the adjoint solutions are available. Using various surface and volume integral functionals, the accuracy provided by the reconstructed adjoint closely matche that of the analytical adjoint solution and in most cases provide a better approximation than the original, working adjoint approximation with regard to the error estimate obtained.

4.2 Optimal control framework

The expression for the error estimator given in (4.6) involves only the residual of the primal problem but not the dual. However, it would appear that adapting on both the primal and dual residuals should be advantageous with respect to the robustness of the procedure [83]. The *optimal control* framework to error estimation for functional output introduced by Rannacher *et al.* is formulated by the introduction of a Lagrangian functional [13, 7]. Using this framework, an error estimator involving both the primal and dual residuals naturally arises. The framework has been applied to a wide variety of problems discretized with conforming finite element methods. Here, it is applied to DG discretization of inviscid Euler and compressible Navier-Stokes equations. The present analysis is carried out by considering a working approximation space $\mathcal{V}_H^P$ and a refined space $\mathcal{V}_h^p$, such that the inclusion relation $\mathcal{V}_H^P \subset \mathcal{V}_h^p$ holds. The corresponding approximate solutions $\mathbf{u}_H \in \mathcal{V}_H^P, \mathbf{u}_h \in \mathcal{V}_h^p$ satisfy respectively,

$$R_H(\mathbf{u}_H, \mathbf{v}_H) = 0, \quad \forall \mathbf{v}_H \in \mathcal{V}_H^P,$$
$$R_h(\mathbf{u}_h, \mathbf{v}_h) = 0, \quad \forall \mathbf{v}_h \in \mathcal{V}_h^p. \tag{4.7}$$

To apply the optimal-control framework of error estimation, we consider the situation where the coarse and fine semilinear forms are equivalent on the coarse space:

$$R_H(\mathbf{u}_H, \mathbf{v}_H) = R_h(\mathbf{u}_H, \mathbf{v}_H), \quad \forall \mathbf{u}_H, \mathbf{v}_H \in \mathcal{V}_H^P. \tag{4.8}$$

Similarly, the functionals of interest are also assumed to satisfy the coarse space equivalence:

$$\mathcal{J}_H(\mathbf{u}_H) = \mathcal{J}_h(\mathbf{u}_H), \quad \forall \mathbf{u}_H \in \mathcal{V}_H^P. \tag{4.9}$$

The conditions (4.8) and (4.9) are satisfied by almost all DG schemes when the fine mesh is obtained by increasing the solution order while keeping the underlying triangulation fixed. Introduce the Lagrangian functional $\mathcal{L}_h(\cdot, \cdot)$ on $\mathcal{V}_h^p \times \mathcal{V}_h^p$:

$$\mathcal{L}_h(\mathbf{v}_h, \varphi_h) \equiv \mathcal{J}_h(\mathbf{v}_h) - R_h(\mathbf{v}_h, \varphi_h), \quad \forall \mathbf{v}_h, \varphi_h \in \mathcal{V}_h^p. \tag{4.10}$$

Using properties (4.8) and (4.9) in the residual equations (4.7) defining $\mathbf{u}_H$ and $\mathbf{u}_h$, the outputs $\mathcal{J}_H(\mathbf{u}_H)$ and $\mathcal{J}_h(\mathbf{u}_h)$ are the values of $\mathcal{L}_h(\cdot, \cdot)$ at stationary points in the spaces $\mathcal{V}_H^P$ and $\mathcal{V}_h^p$ respectively. Using error representation of the following form,

$$\int_0^1 f(t)dt = \frac{1}{2}(f(0) + f(1)) + \frac{1}{2}\int_0^1 s(s-1) \times f''(s)ds, \tag{4.11}$$

it may be shown that the difference in the outputs obtained from $\mathbf{u}_h$ and $\mathbf{u}_H$ consists of weighted primal and dual residuals plus a remainder term $\mathcal{R}^{(3)}$ that is cubic in the primal and adjoint state errors [13, 7]:

$$
\begin{aligned}
\mathcal{J}_h(\mathbf{u}_h) - \mathcal{J}_h(\mathbf{u}_H) \;=\; \frac{1}{2}\Bigg\{ &- R_h(\mathbf{u}_H, \psi_h - \phi_H) \\
&+ \mathcal{J}_h'[\mathbf{u}_H](\mathbf{u}_h - \mathbf{v}_H) - R_h'[\mathbf{u}_H](\mathbf{u}_h - \mathbf{v}_H, \psi_H) \Bigg\} \\
&+ \mathcal{R}^{(3)}, \quad \forall \phi_H, \mathbf{v}_H \in \mathcal{V}_H^P,
\end{aligned}
\tag{4.12}
$$

where Galerkin orthogonality permits any choice of $\phi_H, \mathbf{v}_H \in \mathcal{V}_H^P$ in the estimator. Since the remainder term $\mathcal{R}^{(3)}$ is of higher order in comparison to terms of (4.12) in the curly braces, throughout this work only the latter is used in error estimation and adaptation procedures.

4.3   Higher-order reconstruction

The expression (4.12) involves unavailable fine mesh solutions, $\mathbf{u}_h, \psi_h$. One way to obtain estimates for these fine mesh solutions is to perform reconstructive post-processing on the coarse mesh solutions $\mathbf{u}_H, \psi_H$. This class of procedures can be motivated with the observation that since the output error estimator holds for all $\phi_H, \mathbf{v}_H \in \mathcal{V}_H^P$, an equivalent expression for the terms in the curly brackets is:

$$
\begin{aligned}
& \inf_{\phi_H \in \mathcal{V}_H^P} -\frac{1}{2} R_h(\mathbf{u}_H, \psi_h - \phi_H) \\
+ & \inf_{\mathbf{v}_H \in \mathcal{V}_H^P} -\frac{1}{2} \left(\mathcal{J}_h'[\mathbf{u}_H](\mathbf{u}_h - \mathbf{v}_H) - R_h'[\mathbf{u}_H](\mathbf{u}_h - \mathbf{v}_H, \psi_H) \right).
\end{aligned}
\tag{4.13}
$$

This shows that in the expression for the estimator, primal and dual residuals are weighted by quantities that estimate the lack of ability of the coarse mesh function space to interpolate fine mesh dual and primal states respectively. Thus, patch-wise least-squares reconstruction procedures could be used to estimate these approximation properties. Although there may be situations where reconstructive procedures do not capture interpolative error to a sufficient accuracy, there has been demonstrated success in applying these reconstruction techniques for the estimation of functional outputs. For practical cases using finite element methods with linear interpolation and quadratic reconstruction, Rannacher *et al.* have observed that the quality of the resulting estimates are satisfactory with the effectivity index, defined as the estimated to true error, typically in the range of $I_{eff} \approx 1.2 \sim 3$ [74].

In the current work, the fine mesh is obtained from the working mesh by increasing the solution order P but keeping the underlying triangulation fixed. Let κ be an element for which high-order reconstruction is desired, and let $\mathcal{N}(\kappa)$ be a patch of

neighboring elements surrounding κ. Given the coarse mesh solution $\mathbf{u}_H$, the increased order flow state estimate $\tilde{\mathbf{u}}_H$ on κ is given by:

$$\tilde{\mathbf{u}}_H|_\kappa = \arg|_\kappa \min_{\mathbf{v}_H \in \mathcal{V}_H^p(\mathcal{N}(\kappa))} \|\mathbf{v}_H - \mathbf{u}_H\|_{L_2(\mathcal{N}(\kappa))}. \qquad (4.14)$$

The same procedure is applied to obtain fine mesh adjoint estimate, $\tilde{\psi}_H$. Similar recovery techniques are widely used in output error estimation and control [74, 84, 8]. For each element κ, the above results in a local system of normal equations that may be solved by direct methods.

4.4 Localization

The error estimator (4.12) for DG discretizations consists of a sum over elements of area and surface integrals involving primal/dual state errors. Hence, it is natural to consider local error indicators that are simply inner products of the discrete residuals weighted by the corresponding elemental state errors. Here, results are given in the general setting assuming only the inclusion relation of the finite element spaces,

$$\mathcal{V}_H^P \subset \mathcal{V}_h^p, \qquad (4.15)$$

rather than restricting to the case of increasing only the order of the interpolation as done in the expression (4.14).

4.4.1 First-order conservation laws

For the simplicity of presentation, the analysis is presented only for the case where the functional is an integral over some portion of the domain boundary: $\mathcal{J}_h \equiv \int_{\Gamma_{\text{output}}} J(\mathbf{u})ds$, $\Gamma_{\text{output}} \subseteq \partial\Omega$. For the primal problem, the fine mesh DG semilinear form (3.3.1) can be written as a sum of elemental contributions with each term

involving the test function supported on the element of interest:

$$R_h(\mathbf{u}_h, \mathbf{v}_h) = \sum_{\kappa \in T_h} R_h(\mathbf{u}_h, \mathbf{v}_h|_\kappa), \tag{4.16}$$

where

$$R_h(\mathbf{u}_h, \mathbf{v}_h|_\kappa) \equiv -\int_\kappa \nabla \mathbf{v}_h^T \cdot \mathcal{F}(\mathbf{u}_h)d\mathbf{x} + \int_{\partial\kappa\backslash\partial\Omega} \mathbf{v}_h^{+T}\mathcal{H}(\mathbf{u}_h^+, \mathbf{u}_h^-, \hat{\mathbf{n}})ds$$
$$+ \int_{\partial\kappa\cap\partial\Omega} \mathbf{v}_h^{+T}\hat{\mathbf{n}} \cdot \mathcal{F}(\mathbf{u}_h^b(\mathbf{u}_h^+))ds. \tag{4.17}$$

For the dual problem, the linearized fine mesh DG bilinear form may be similarly written as the sum $R'_h[\mathbf{u}_h](\mathbf{v}_h, \boldsymbol{\psi}_h) = \sum_{\kappa \in T_h} R'_h[\mathbf{u}_h](\mathbf{v}_h|_\kappa, \boldsymbol{\psi}_h)$, where

$$R'_h[\mathbf{u}_h](\mathbf{v}_h|_\kappa, \boldsymbol{\psi}_h) \equiv -\int_\kappa \nabla \boldsymbol{\psi}_h^T \cdot \mathcal{F}'[\mathbf{u}_h](\mathbf{v}_h)d\mathbf{x} + \int_{\partial\kappa\backslash\partial\Omega} (\boldsymbol{\psi}_h^+ - \boldsymbol{\psi}_h^-)^T \frac{\partial\mathcal{H}}{\partial\mathbf{u}_h^+}(\mathbf{v}_h^+)ds$$
$$+ \int_{\partial\kappa\cap\partial\Omega} \boldsymbol{\psi}_h^{+T}\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}_h^b(\mathbf{u}_h^+)] \left(\left[\frac{\partial\mathbf{u}_h^b}{\partial\mathbf{u}_h^+}\right]\mathbf{v}_h^+ \right)ds. \tag{4.18}$$

The above expansions show that the estimator expression (4.12) may naturally be written as a sum over elements of primal and dual residual contributions. From the inclusion assumption (4.15), any element in the coarse space $\kappa \in T_H$ may be written as a sum over fine mesh elements:

$$\kappa = \bigcup_{\kappa' \subset \kappa} \kappa', \quad \kappa' \in T_h. \tag{4.19}$$

It follows that $\mathbf{u}_H$ and $\boldsymbol{\psi}_H$ are continuous across $\partial\kappa'\backslash\partial\kappa$. The local error indicator on a coarse element κ may be written as a sum of two contributions $\eta_\kappa = \eta_\kappa^{\text{pr}} + \eta_\kappa^{\text{du}}$, where:

$$\eta_\kappa^{\text{pr}} \equiv \left(\sum_{\kappa' \subset \kappa} \int_{\kappa'} \nabla(\tilde{\boldsymbol{\psi}}_h - \boldsymbol{\psi}_H)^T \cdot \mathcal{F}(\mathbf{u}_H)d\mathbf{x} - \int_{\partial\kappa'\backslash\partial\Omega} (\tilde{\boldsymbol{\psi}}_h^+ - \boldsymbol{\psi}_H^+)^T \mathcal{H}(\mathbf{u}_H^+, \mathbf{u}_H^-, \hat{\mathbf{n}})ds \right)$$
$$- \int_{\partial\kappa\cap\partial\Omega} (\tilde{\boldsymbol{\psi}}_h^+ - \boldsymbol{\psi}_H^+)^T \hat{\mathbf{n}} \cdot \mathcal{F}(\mathbf{u}_h^b(\mathbf{u}_H^+))ds, \tag{4.20}$$

and

$$
\begin{aligned}
\eta_\kappa^{\mathrm{du}} \equiv & \left(\sum_{\kappa' \subset \kappa} \int_{\kappa'} \nabla \psi_H^T \cdot \mathcal{F}'[\mathbf{u}_H](\tilde{\mathbf{u}}_h - \mathbf{u}_H) d\mathbf{x} \right) \\
& - \int_{\partial\kappa \backslash \partial\Omega} (\psi_H^+ - \psi_H^-)^T \frac{\partial \mathcal{H}}{\partial \mathbf{u}_h^+} \left(\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+ \right) ds \\
& - \int_{\partial\kappa \cap \partial\Omega} \psi_H^{+T} \hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}_h^b(\mathbf{u}_H)] \left(\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] (\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+) \right) ds \\
& + \int_{\partial\kappa \cap \partial\Gamma_{\text{output}}} J'[\mathbf{u}_H] \left(\left[\frac{\partial \mathbf{u}_h^b}{\partial \mathbf{u}_h^+} \right] (\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+) \right) ds. \tag{4.21}
\end{aligned}
$$

As a practical note, the elemental error indicators (4.20), (4.21) are not computed via approximating the respective volume and surface integrals. Instead, the expressions are equivalently evaluated in an algebraic manner by taking the Euclidean inner product of the residual vectors with the degrees of freedoms representing the state error as suggested by the decomposition (4.16). This is analogous to the form of the local error indicator used in the works of Venditti and Darmofal where the expression consists of vector inner products of the discrete residuals with the state error degrees of freedom [83, 84].

To further elucidate the role of adjoint state as the output sensitivity to the consistency error of the discrete solution $\mathbf{u}_H$ for the differential operator, expression (4.20) may be integrated by parts and use flux consistency $\mathcal{H}(\mathbf{u}_H, \mathbf{u}_H, \hat{\mathbf{n}}) = \hat{\mathbf{n}} \cdot \mathcal{F}(\mathbf{u}_H)$ to yield:

$$
\begin{aligned}
\eta_\kappa^{\mathrm{pr}} = & - \left(\sum_{\kappa' \subset \kappa} \int_{\kappa'} \left(\tilde{\psi}_h - \psi_H \right)^T \nabla \cdot \mathcal{F}(\mathbf{u}_H) d\mathbf{x} \right) \\
& - \int_{\partial\kappa \backslash \partial\Omega} (\tilde{\psi}_h^+ - \psi_H^+)^T (\mathcal{H} - \hat{\mathbf{n}} \cdot \mathcal{F}) ds \\
& - \int_{\partial\kappa \cap \partial\Omega} (\tilde{\psi}_h^+ - \psi_H^+)^T \hat{\mathbf{n}} \cdot \left[\mathcal{F}(\mathbf{u}_h^b(\mathbf{u}_H^+)) - \mathcal{F}(\mathbf{u}_H^+) \right] ds. \tag{4.22}
\end{aligned}
$$

The first line of (4.22) gives the adjoint solution its interpretation as a Lagrange multiplier to the primal differential operator. The second line may be similarly interpreted but now with generalized flux derivative arising from flow discontinuity across elements. The term on $\partial\kappa \cap \partial\Omega$ in the last line can be better understood via linear

82

analysis. Assuming $[\mathbf{B}] : (\cdot) \to \mathbf{u}^b(\cdot)$ to be a linear projective mapping of boundary states and making linear approximation of the flux, obtain:

$$\mathcal{F}(\mathbf{u}^b(\mathbf{u}_H^+)) - \mathcal{F}(\mathbf{u}_H^+) \approx \frac{d\mathcal{F}}{d\mathbf{u}}([\mathbf{B}]\mathbf{u}_H^+ - \mathbf{u}_H^+), \tag{4.23}$$

which is zero for the component of $\mathbf{u}_H^+$ that lies in the range of $[\mathbf{B}]$. This shows that the only non-zero contribution from the last line of (4.22) comes from the component of $\mathbf{u}_H^+$ which lies in the null space of $[\mathbf{B}]$; the component of $\mathbf{u}_H^+$ in range($[\mathbf{B}]$) results in no contribution. Thus the error contribution from the sides of element κ lying on the domain boundary $\partial\Omega$ consists only of flux consistency error of the interior state with the prescribed boundary data. The multiplication of the adjoint state with the flux error is in contrast to the resulting expression in the case of *strong* enforcement of boundary condition where the adjoint states takes on a different character on the boundary as the Lagrange multiplier to the boundary condition. Similar result holds for the adjoint error estimator (4.21) supported on $\partial\Omega$. In this case, only the component of $\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+$ in range($[\mathbf{B}]$) gives rise to a non-trivial contribution; the component of the state error in null($[\mathbf{B}]$) does not contribute. This is consistent with the fact that adjoint boundary condition is prescribed via the flux for the space of test functions in range($[\mathbf{B}]$) but allowed to be free otherwise.

4.4.2 Second-order systems

In a manner similar to that carried out for expression (4.22), integrating by parts on the BR2 semilinear form (3.50) together with manipulations detailed in Appendix B show that the primal residual contribution to the local error indicator takes the form,

$$
\begin{aligned}
\eta_\kappa^{pr} = &- \left(\sum_{\kappa' \subset \kappa} \int_{\kappa'} (\tilde{\psi}_h - \psi_H)^T \nabla \cdot (\mathcal{A}_v(\mathbf{u}_H)\nabla\mathbf{u}_H) dx \right) \\
&+ \frac{1}{2} \int_{\partial\kappa\backslash\partial\Omega} (\tilde{\psi}_h^+ - \psi_H^+)^T \left([\![\mathcal{A}_v(\mathbf{u}_H)\nabla\mathbf{u}_H]\!] + 2\eta_f \hat{\mathbf{n}} \cdot \{\boldsymbol{\delta}_f\} \right) ds \\
&- \frac{1}{2} \int_{\partial\kappa\backslash\partial\Omega} [\![\mathbf{u}_H]\!]^T \cdot \mathcal{A}_v(\mathbf{u}_H^+)^T \nabla(\tilde{\psi}_h^+ - \psi_H^+) ds
\end{aligned}
$$

$$
- \int_{\partial\kappa \cap \partial\Omega} (\tilde{\psi}_h^+ - \psi_H^+)^T \left[\mathbf{F}_h^b(\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_H^b)\nabla\mathbf{u}_H^+ - \eta_f\hat{\mathbf{n}} \cdot \boldsymbol{\delta}_f^b) - \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}(\mathbf{u}_H^+)\nabla\mathbf{u}_H^+ \right] ds
$$
$$
+ \int_{\partial\kappa \cap \partial\Omega} (\mathbf{u}_H^b - \mathbf{u}_H^+)^T \left[\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_H^b)^T \nabla(\tilde{\psi}_h^+ - \psi_H^+) \right] ds, \tag{4.24}
$$

where the abbreviation $\mathbf{u}_H^b \equiv \mathbf{u}_h^b(\mathbf{u}_H^+)$ is used. The volume term on the first line measures the consistency error with respect to the underlying differential operator weighted by an estimate of the adjoint error. Similarly, the second line measures error arising from the jump of viscous flux across element interfaces. The third line measures conformity error arising from discontinuities in the state which is weighted by the gradient in the adjoint error. The last two lines are analogous to the previous interior boundary terms, except that only the errors with respect to the imposed viscous flux and boundary state data are measured. The dual residual contribution to the local error indicator may be similarly derived, shown in detail in Section B.2 of Appendix B. For the simplicity of presentation, the resulting expression is shown here only for the case where the coefficient matrix is state-independent, $\boldsymbol{\mathcal{A}}_v = \boldsymbol{\mathcal{A}}_v(\mathbf{x})$. For terms arising from the state-dependency $\boldsymbol{\mathcal{A}}_v = \boldsymbol{\mathcal{A}}_v(\mathbf{u})$, refer to Section B.2.2 of Appendix B. Given a coarse mesh element $\kappa \in T_H$, the dual local error estimator is:

$$
\begin{aligned}
\eta_\kappa^{du} = & -\left(\sum_{\kappa' \subset \kappa} \int_{\kappa'} (\tilde{\mathbf{u}}_h - \mathbf{u}_H)^T \nabla \cdot (\boldsymbol{\mathcal{A}}_v^T \nabla\psi_H) d\mathbf{x} \right) \\
& + \frac{1}{2} \int_{\partial\kappa \backslash \partial\Omega} (\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+)^T (\llbracket \boldsymbol{\mathcal{A}}_v^T \nabla\psi_H \rrbracket + 2\eta_f\hat{\mathbf{n}} \cdot \{\boldsymbol{\delta}_f^{du}\}) ds \\
& - \frac{1}{2} \int_{\partial\kappa \backslash \partial\Omega} \llbracket \psi_H \rrbracket^T \cdot \boldsymbol{\mathcal{A}}_v \nabla(\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+) ds \\
& + \int_{\partial\kappa \cap \partial\Omega} \left(\left[\frac{\partial\mathbf{u}_H^b}{\partial\mathbf{u}_H^+} \right] (\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+) \right)^T (\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v^T \nabla\psi_H^+ - \eta_f\hat{\mathbf{n}} \cdot \boldsymbol{\delta}_f^{b,du}) ds \\
& + \int_{\partial\kappa \cap \Gamma_{\mathrm{output}}} J_I'[\mathbf{u}_H^b] \left(\left[\frac{\partial\mathbf{u}_H^b}{\partial\mathbf{u}_H^+} \right] (\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+) \right) ds \\
& + \int_{\partial\kappa \cap \partial\Omega} (\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+)^T \eta_f\hat{\mathbf{n}} \cdot \boldsymbol{\delta}_f^{b,du} ds \\
& - \int_{\partial\kappa \cap \partial\Omega} \psi_H^{+T} \left(\left[\frac{\partial\mathbf{F}_H^b}{\partial\mathbf{F}_H^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla(\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+) \right) ds \\
& + \int_{\partial\kappa \cap \Gamma_{\mathrm{output}}} J_{II}'[\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla\mathbf{u}_H] \left(\left[\frac{\partial\mathbf{F}_H^b}{\partial\mathbf{F}_H^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla(\tilde{\mathbf{u}}_h^+ - \mathbf{u}_H^+) \right) ds. \tag{4.25}
\end{aligned}
$$

84

Many of the terms of the dual error indicator have analogues to the primal expression (4.24). The first volume term measure the consistency error with respect to the adjoint differential operator (3.46), weighted by an estimate of the state error. The second line also measures consistency error but for the jump of adjoint viscous flux across element interfaces, where the auxiliary variable δ_f^{du} is defined in (B.6). The third line measures the conforming error of the adjoint solution where the interelement jumps are weighted by gradient of the primal solution. The boundary terms of lines 4-5 measure consistency error for the adjoint Neumann boundary conditions (3.48), measured only for the component of the primal state error in the range of the linearized boundary state mapping, $[\partial \mathbf{u}_H^b / \partial \mathbf{u}_H^+]$. Line 6 has no correspondence in the primal expression (4.24), which from the definition of $\delta_f^{b,du}$ given in (B.7) is only non-zero if the adjoint approximation does not satisfy the Dirichlet boundary condition (3.47). The last two lines measures the adjoint consistency error for the Dirichlet boundary condition (3.47), weighted by the component of the primal flux error in the range of $[\partial \mathbf{F}_H^b / \partial \mathbf{F}_H^+]$.

4.5 Adaptation strategy

For a given choice of local error indicator, there are several possibilities for mesh adaptation criterion. These include [33, 7]: error-balancing strategy, where elements are marked for refinement if the local error contribution is larger than some desired tolerance divided by the total number of elements; mesh-optimization strategy, where one seeks to minimize the number of elements subject to meeting some desired error tolerance; fixed-fraction strategy, where a fixed fraction of the total number of elements are refined or coarsened according to the ordered size of the error contribution.

Carrying out mesh adaptation within the context of optimization, it is reasonable to choose an adaptation criterion such that the precision estimate of the subsequent approximation model is a certain fraction of that for the current model. Hence, the fixed-error-reduction strategy as described in [7] is a natural choice and is adopted throughout the thesis. Let E denote the error reduction desired at each adaptation

iteration. The total number of elements n_{Tot} are ordered by the size of the local error contribution so that $|\eta_{\kappa_1}| \geq |\eta_{\kappa_2}| \geq \cdots \geq |\eta_{\kappa_n}|$ and the number n^* is determined such that the partial sum of the absolute value of the error contribution up to element number n^* is E fraction of the total:

$$\sum_{i=1}^{n^*} |\eta_{\kappa_i}| \approx E \times \sum_{i=1}^{n_{Tot}} |\eta_{\kappa_i}|. \tag{4.26}$$

Subsequently, the elements $\kappa_1, \cdots, \kappa_{n^*}$ are marked for refinement. In the current work, refinement is done only by increasing the order of interpolation by 1.

4.6 Results

Inviscid test case

In this section, error estimation for smooth solutions to the inviscid Euler equations is demonstrated. In particular, the test case involves a $M_\infty = 0.5$ subsonic flow over Gaussian-shaped bump as described in Section 3.5.1. The effectivity of the error estimator for the weighted lift output (3.59) is investigated for $p = 1$ to $p = 3$ solutions on a sequence of uniformly refined meshes containing 587, 2325 and 9169 elements. Figure 4-1 shows the result obtained using the dual-consistent treatment, for which the output convergence is given in Figure 3-4. It can be seen that the effectivity for the estimate is close to 1 for all solution orders and mesh sizes.

The results for the two forms of dual-inconsistent boundary treatments are shown in Figure 4-2 and 4-3 respectively. It can be seen that there is significantly more uncertainty in the error estimates. This is partly due to the fact that the discrete adjoint solution has strong, mesh depedent irregularities and thus rendering higher-order reconstruction for the adjoint ineffective.

Figure 4-1: Error estimate: dual-consistent boundary treatment using only $\mathbf{u}_h^b(\mathbf{u}_h^+)$. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.
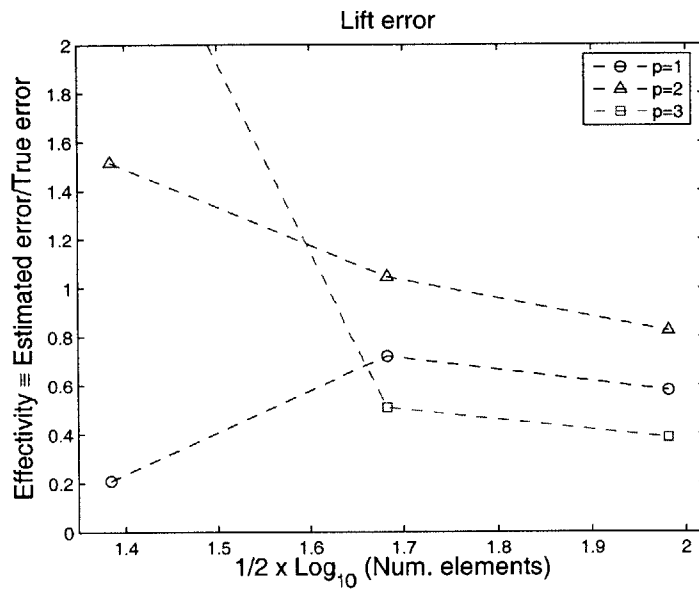

Figure 4-2: Error estimate: dual-inconsistent boundary treatment using both $\mathbf{u}_h^+$ and $\mathbf{u}_h^b(\mathbf{u}_h^+)$. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.

Figure 4-3: Error estimate: conservative but dual-inconsistent boundary treatment based on numerical flux function. Inviscid Euler flow over Gaussian bump, $M_\infty = 0.5$.

Viscous test case

The test case considered is a NACA 0012 airfoil in subsonic, laminar flow with free-stream Mach number $M_\infty = 0.5$, Reynolds number $Re = 5000$ and angle of attack $\alpha = 2.0°$ as used in Section 3.5.2. The p-adaptation is based on the drag output with error reduction parameter $E = 0.75$, performed on a mesh containing 2688 elements and initiated with a uniform $p = 1$ solution. The "true" drag is taken to be the $p = 4$ result on a finer, 10752 element mesh also used in Section 3.5.2. As Figure 4-4 shows, within the range of order refinement the drag converges exponentially in the degrees of freedom. On the sequence of output-adapted meshes, the effectivity index for the estimator (4.12) based on higher-order reconstruction is centered around 0.8 and is bounded below and above by 0.6 and 1.5 respectively. As Figure 4-7 shows, output-based adaptation gives rise to high order ($p = 4$) interpolation being used at the trailing edge, in the wake of the airfoil and upstream of the stagnation point. This is consistent with the presense of significant flow solution gradient in the wake, a localized region of sharp change in $\psi(3)$ at the airfoil trailing edge and an elongated layer in the adjoint solution emanating out of the stagnation point, as shown in Figure 4-5. To demonstrate the effect of p-adaptation on the flow solution, in Figure 4.6 the Mach contour is plotted for the original, uniform $p = 1$ solution as well as the last adaptive iteration.

The dual-consistency of the boundary treatment goes towards contributing to the monotonicity of error convergence shown in Figure 4-4. This can be demonstrated by comparison with adaptation result obtained using a dual-inconsistent form of functional without the penalty parameter δ_f^b, for which the lack of adjoint regularity is shown in Section 3.5.2. The resulting output convergence and the effectivity of error estimate is shown in Figure 4-8. In contrast to the result shown in Figure 4-4, the output error initially increases as the solution order is refined. Furthermore, due to the discrete adjoint irregularity introduced, adaptation results in significantly refined solution orders for elements lying on the airfoil boundary, shown in Figure 4-9 for the leading edge close-up. This result is in agreement with the computational re-
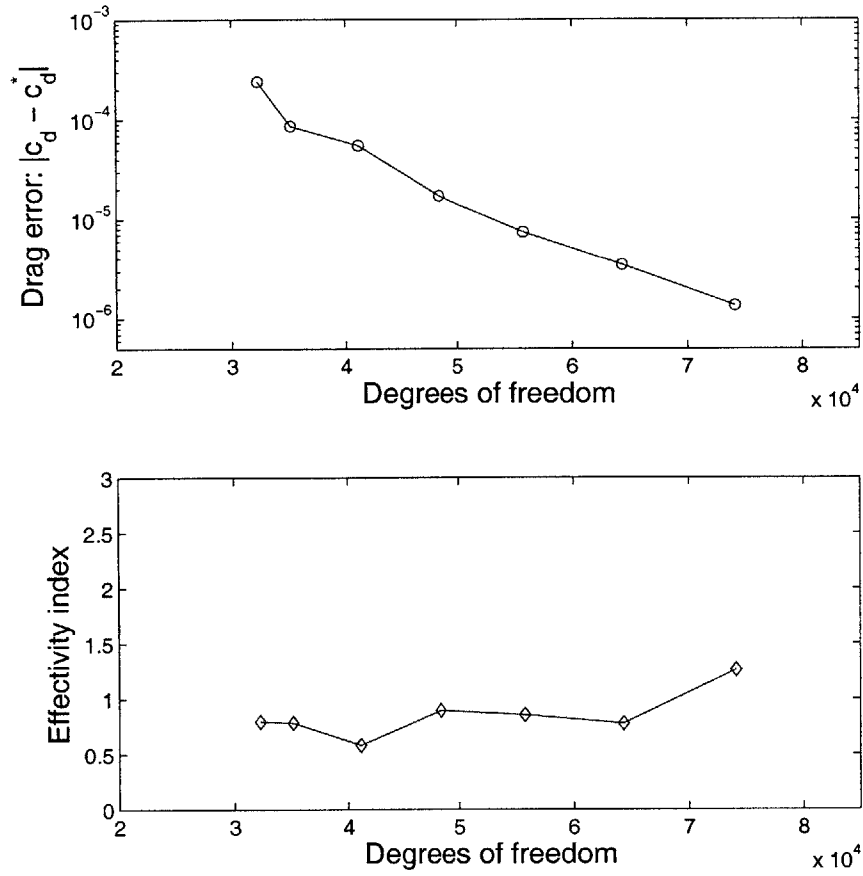
89

Figure 4-4: p-adaptation for drag: output convergence and error estimate. NACA 0012 airfoil. $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

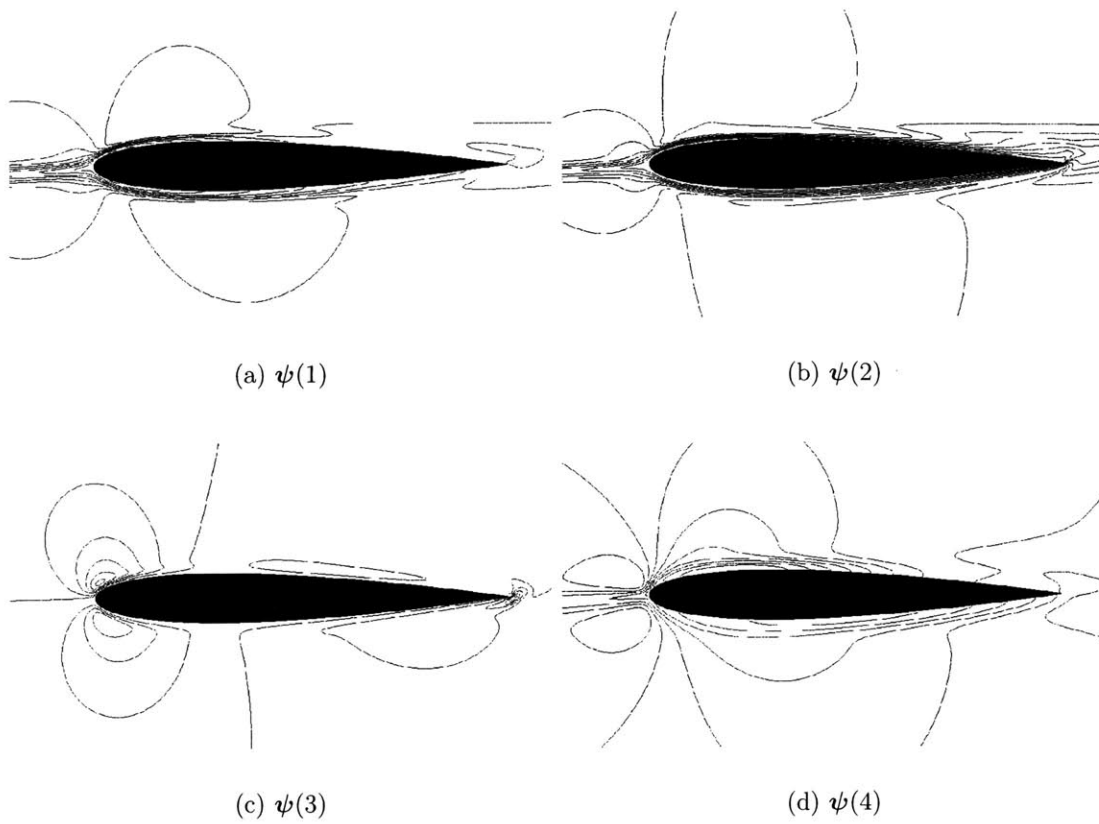(a) $\psi(1)$ (b) $\psi(2)$

(c) $\psi(3)$ (d) $\psi(4)$

Figure 4-5: Contour plots of drag adjoint solution. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

(a) Initial uniform $p = 1$ interpolation

(b) Final, 6th adaptive iteration

Figure 4-6: Mach contours. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

Figure 4-7: Adapted solution orders: p=1 to p=4. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.
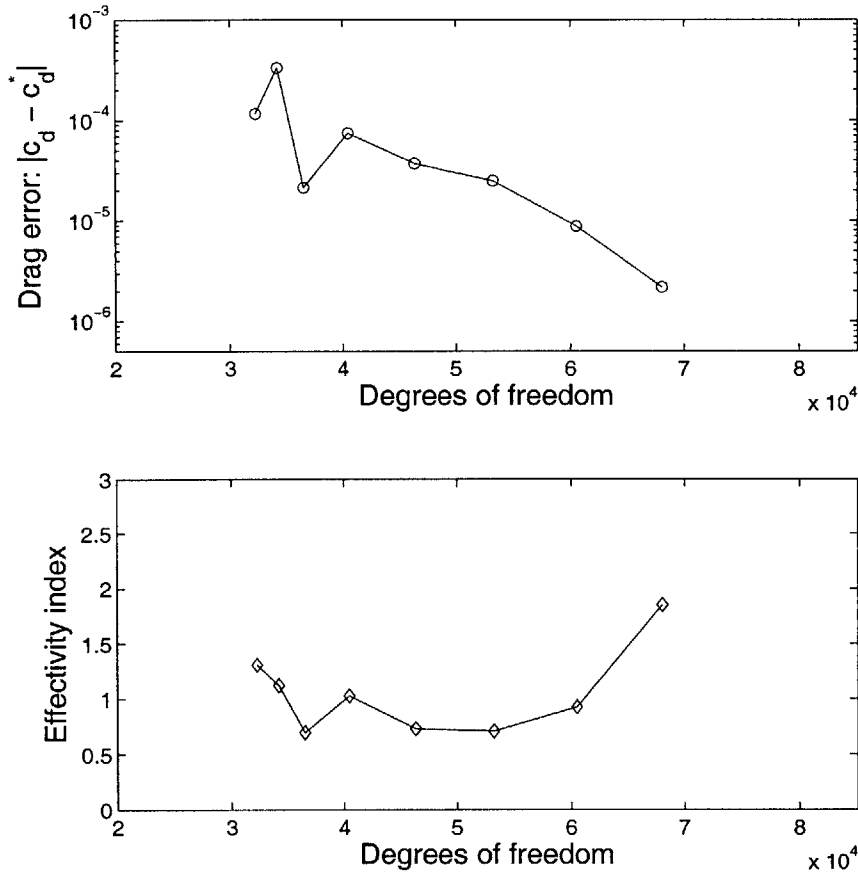
Figure 4-8: p-adaptation for drag using dual-inconsistent treatment: output convergence and error estimate. NACA 0012 airfoil. $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

sults of Hartmann and Houston [40] in the context of h-adaptation for inviscid Euler equations, where the irregularity in the discrete adjoint solution (introduced by the boundary treatment) leads to small mesh sizes near the airfoil surface. To conclude: by ensuring dual-consistency, the error in the output can be effectively controlled without requiring significant refinement at the boundary.

(a) Dual-consistent functional treatment

(b) Dual-inconsistent functional treatment

Figure 4-9: Comparison of adapted solution orders near leading edge, showing $p = 4$ being utilized for region near the airfoil in the dual-inconsistent case. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

Chapter 5

Concurrent flow-adjoint solution

5.1 Introduction

This chapter is concerned with the development of computational tools enabling the estimation of output error due to incomplete solution iteration, within the framework of adaptive precision methodology discussed in Chapter 2. In particular, an adjoint solution strategy is developed in which dual approximations are iterated at the same time as the nonlinear primal problem. For the proposed error estimation approach to be computationally viable in analysis and design settings, the resulting concurrent flow-adjoint solution algorithm has to be computationally efficient in comparison with the underlying flow solution algorithm, as well as being competitive with the traditional, sequential adjoint solution approach. It is shown here that the computational overhead associated with concurrently iterating the adjoint equations is a small fraction of the baseline flow algorithm. Furthermore, in practical situations where full linearization of the computational problem cannot be stored in memory, the concurrent approach can in fact be a better alternative to the sequential approach in terms of the computational cost. Using the adjoint approximations available from the concurrent algorithm, duality-based error analysis can be carried out to estimate the leading order error in the functional outputs of interest.

In Section 5.2 a duality-based technique is carried out to show how by an appropriate choice of preconditioner, the adjoint algorithm can be derived to match the

primal in the sense of outputs. In Section 5.3 the concurrent solution approach is introduced, with a discussion given on how the nearest neighbor stencil property of the chosen DG scheme can be used to construct an efficient adjoint solver. In Section 5.4 analysis is carried out to show the readily available adjoint iterates may be used to form a first order correction to the output estimates. In Section 5.5 computational results are given, in particular showing that adjoint estimates can be obtained concurrently at a fractional cost over the baseline (primal) algorithm and in fact its efficiency ratio can scale favorably with the solution order. The advantage of being able to compute the discrete adjoint residual in a purely algebraic manner that is independent of either the solution order or the underlying physics is particularly clear for the case of compressible Navier-Stokes equations.

5.2   Exact-dual solution method

Denote a general, nonlinear algebraic system of equations for the primal unknowns $\mathbf{U}$ as

$$\mathbf{A}(\mathbf{U}) = 0, \tag{5.1}$$

and let $J(\mathbf{U})$ be an output of interest. The corresponding discrete adjoint state $\mathbf{\Psi}$ is the solution to the linear system

$$\mathbf{A}'[\mathbf{U}]^T \mathbf{\Psi} = J'[\mathbf{U}]. \tag{5.2}$$

Now consider the special case of a linearized primal problem, $\mathbf{A}(\mathbf{U}) = \mathbf{A}\mathbf{U} - \mathbf{f}$. Denote a preconditioned iterative scheme for the linearized primal problem as,

$$\mathbf{U}^0 = 0$$
$$\mathbf{U}^{n+1} = \mathbf{U}^n - \mathbf{P}_n^{-1}(\mathbf{A}\mathbf{U}^n - \mathbf{f}), \quad n = 1, 2, \cdots, N, \tag{5.3}$$

where $\mathbf{P}_n$ is the preconditioner which may in general depend on the iteration index n. Denote the linearized primal functional as

$$J^{\mathrm{pr}}(\cdot) \equiv \langle \mathbf{g}, \cdot \rangle. \tag{5.4}$$

where $\langle \cdot, \cdot \rangle$ is the Euclidean inner product. Giles proposed an approach to derive iterative schemes for the adjoint equations via maintaining a duality property [29]. By introducing Lagrange multipliers $\boldsymbol{\lambda}^{n+1}$ to the primal solution scheme (5.3), the linear output at final iteration N can be trivially rewritten as:

$$
\begin{aligned}
J^{\mathrm{pr}}(\mathbf{U}^N) &\equiv \langle \mathbf{g}, \mathbf{U}^N \rangle \\
&= \langle \mathbf{g}, \mathbf{U}^N \rangle - \sum_{n=0}^{N-1} \langle \boldsymbol{\lambda}^{n+1}, \mathbf{U}^{n+1} - \mathbf{U}^n + \mathbf{P}_n^{-1}(\mathbf{A}\mathbf{U}^n - \mathbf{f}) \rangle.
\end{aligned}
\tag{5.5}
$$

Using the summation by parts formula for arbitrary sequences $\{\mathbf{a}^n\}$, $\{\mathbf{b}^n\}$:

$$\sum_{n=0}^{N-1} \langle \mathbf{a}^{n+1}, \mathbf{b}^{n+1} - \mathbf{b}^n \rangle = \langle \mathbf{a}^N, \mathbf{b}^N \rangle - \langle \mathbf{a}^0, \mathbf{b}^0 \rangle - \sum_{n=0}^{N-1} \langle \mathbf{a}^{n+1} - \mathbf{a}^n, \mathbf{b}^n \rangle, \tag{5.6}$$

expression (5.5) is

$$
\begin{aligned}
J^{\mathrm{pr}}(\mathbf{U}^N) &= \langle \mathbf{g} - \boldsymbol{\lambda}^N, \mathbf{U}^N \rangle + \langle \boldsymbol{\lambda}^0, \mathbf{U}^0 \rangle \\
&\quad + \sum_{n=0}^{N-1} \langle \mathbf{P}_n^{-T} \boldsymbol{\lambda}^{n+1}, \mathbf{f} \rangle + \langle \boldsymbol{\lambda}^{n+1} - \boldsymbol{\lambda}^n - (\mathbf{P}_n^{-1}\mathbf{A})^T \boldsymbol{\lambda}^{n+1}, \mathbf{U}^n \rangle.
\end{aligned}
\tag{5.7}
$$

Suppose now the multipliers are chosen to satisfy the following:

$$
\begin{aligned}
\boldsymbol{\lambda}^N &= \mathbf{g}, \\
\boldsymbol{\lambda}^n &= \boldsymbol{\lambda}^{n+1} + \mathbf{A}^T \mathbf{P}_n^{-T} \boldsymbol{\lambda}^{n+1}, \quad n = N-1, N-2, \cdots, 0,
\end{aligned}
\tag{5.8}
$$

then the value of primal output equals

$$J^{\mathrm{pr}}(\mathbf{U}^N) = \sum_{n=0}^{N-1} \langle \mathbf{P}_n^{-T} \boldsymbol{\lambda}^{n+1}, \mathbf{f} \rangle. \tag{5.9}$$

99

The multipliers $\boldsymbol{\lambda}^{n+1}$ can be regrouped so that (5.8) takes a form similar to the primal iterative scheme (5.3). Define

$$\boldsymbol{\Psi}^n \equiv \sum_{i=N-n}^{N-1} \mathbf{P}_i^{-T} \boldsymbol{\lambda}^{i+1}. \tag{5.10}$$

Inductively, it can be shown via (5.8) and (5.10) that

$$\boldsymbol{\lambda}^{N-n} = \mathbf{g} - \mathbf{A}^T \boldsymbol{\Psi}^n, \quad 0 \leq i < N. \tag{5.11}$$

Hence, the iterative scheme for the regrouped adjoint variable $\{\boldsymbol{\Psi}^n\}$ is:

$$\boldsymbol{\Psi}^0 = \mathbf{0},$$
$$\boldsymbol{\Psi}^{n+1} = \boldsymbol{\Psi}^n - \mathbf{P}_{N-n}^{-T} \left(\mathbf{A}^T \boldsymbol{\Psi}^n - \mathbf{g} \right), \tag{5.12}$$

If the dual functional is defined as:

$$J^{\mathrm{du}}(\cdot) = \langle \mathbf{f}, \cdot \rangle, \tag{5.13}$$

it follows from (5.9) that the following output equivalence relation holds:

$$J^{\mathrm{pr}}(\mathbf{U}^N) = J^{\mathrm{du}}(\boldsymbol{\Psi}^N). \tag{5.14}$$

The exact-dual derivation shown above may be generalized to a wide class of iterative methods, as has been carried out for multigrid solution by Giles *et al.* [30] or the inclusion of inner/outer iterations by Nielsen *et al.* [65].

5.3 Concurrent iteration

Since the adjoint $\boldsymbol{\Psi}$ satisfying (5.2) is defined in terms of the linearization of the algebriac system at its true solution $\mathbf{U}$, the traditional solution approach is to commence adjoint iterations only after the primal solution (or some accurate approximation

thereof) has been obtained, in a sequential manner. There have been recent work exploring the benefit of simultaneously iterating the primal and dual. In [68], Patera and Rønquist applied a general bound result for functional outputs to the situation of estimating the output error due to incomplete iteration of a linear system. In particular, a bound procedure is developed for the preconditioned conjugate gradient method for the solution of symmetric positive definite (SPD) linear systems in which the primal and dual states are solved simultaneosly via the respective Lanczos processes. At each iteration index n, the procedure involves the computation of the minimum eigenvalue of the tridiagonal $n \times n$ Lanczos matrix which is naturally available from the solution procedure. For sufficiently large n, the computable minimum eigenvalue estimates are assured to closely approximate that for the underlying coefficient matrix hence the bound expression produces true lower and upper bounds for the exact output.

In the context of nonlinear systems, Griewank described a "piggy-back" approach of simultaneously iterating the primal and dual states via fixed-point iterations [35] to obtain improved output estimates via duality corrections. This concurrent approach is also utilized by Lu and Darmofal [58] to estimate output error due to incomplete convergence. In particular, the adjoint estimate Ψ_n is iterated using a preconditioner which is the transpose of the primal, driven by a forcing term consisting of the linearization of the nonlinear system at the current primal iterate $\mathbf{U}_n$. That is, the primal and dual systems are iterated in the following manner:

$$
\begin{aligned}
\mathbf{U}_{n+1} &= \mathbf{U}_n - \mathbf{P}(\mathbf{U}_n)^{-1}\mathbf{A}(\mathbf{U}_n) \\
\Psi_{n+1} &= \Psi_n - \mathbf{P}(\mathbf{U}_n)^{-T}\left(\mathbf{A}'[\mathbf{U}_n]^T \Psi_n - J'[\mathbf{U}_n]^T\right),
\end{aligned}
\tag{5.15}
$$

where in general the preconditioners have state-dependence as well. Since the adjoint residual at iteration n depends on the primal state $\mathbf{U}_n$, analysis shows that the resulting adjoint convergence rate is limited in part by the primal residual convergence rate [35]. However, since the asymptotic convergence rate of the primal problem should be the same as that for the adjoint, the convergence rate of the concurrent

solution method is not expected to be adversely affected by the use of non-stationary linearization. This is verified in Griewank and Faure [36, 37] on 2D Euler flow airfoil test cases with lift and drag as functionals of interest, where it is shown that the concurrent approach has almost identical convergence behavior to that of the sequential approach.

Although the concurrent flow-adjoint solution methodology has been successfully demonstrated as a viable approach for giving estimates or bounds for the iterative error in the output, to establish its computational benefit in various analysis and design settings, important questions that need to be addressed include the amount of additional computational cost necessary over the baseline flow algorithm to iterate the adjoint equations as well as the efficiency of concurrent adjoint solution in comparison to the standard sequential approach. In the context of iterative solution for high-order discretization of PDEs, a related question to consider is the scalability of the efficiency for the concurrent solver with the interpolation order.

In the present context of DG formulation, the state degrees of freedom are grouped by elements and the computational stencil consists of nearest neighbor elements for all solution orders. Thus, the coefficient matrix for the adjoint problem has exactly the same block non-zero pattern as that for the primal problem. This implies that the adjoint residual can be computed simply using the available linearization as the primal residual and preconditioner are assembled, instead of needing to resort to different loops over the elements. In the latter case, linearizations required for primal and dual residual calculations may have to be duplicated if the full Jacobian matrix is not stored. By directly using the flow linearization, the computation required for the adjoint residual in the concurrent approach can be a small fraction of that needed for the baseline flow algorithm. In practical, high-dimensional applications one usually cannot afford to store in memory the full linearization for the whole computational domain. Instead, the residual and its linearization may only be performed over some sub-domain such as element blocks and lines, which are then used to compute an update for the degrees of freedom on the same sub-domain. Thus, in the case that full linearization cannot be stored, for the sequential adjoint approach the linearization

102

about the converged primal state has to be performed for each step of the adjoint solver. This can be significantly more expensive than the concurrent approach where linearization is performed once to compute both flow and adjoint residuals. The nearest stencil property is also advantageous when matrix-vector products of the adjoint preconditioner $\mathbf{P}^{-T}$ with the corresponding residual vector is needed to obtain the adjoint updates. For instance, in the case where the preconditioner $\mathbf{P}^{-1}$ consist of a set of block-tridiagonal pieces of the Jacobian obtained from line elements as described in [25], entries of the transposed matrix $\mathbf{P}^{-T}$ may be obtained from $\mathbf{P}^{-1}$ simply by reversing the role of upper and lower block diagonal as well as the row and column indices within the individual square blocks. The implications of compact stencil property on the computational cost of adjoint solver are illustrated by results shown in Section 5.5.

5.4 Superconvergent output estimates

Consider again the primal iterative scheme (5.3) and its dual (5.12) in the linear case, where now the primal preconditioner $\mathbf{P}$ is assumed fixed for all iterations. In this case the primal and dual iterates may be written explicitly as

$$
\begin{aligned}
\mathbf{U}^n &= -\sum_{i=0}^{n-1}(\mathbf{I} - \mathbf{P}^{-1}\mathbf{A})^i\mathbf{P}^{-1}\mathbf{f} \\
\mathbf{\Psi}^n &= -\sum_{i=0}^{n-1}(\mathbf{I} - \mathbf{P}^{-T}\mathbf{A}^T)^i\mathbf{P}^{-T}\mathbf{g}.
\end{aligned}
\tag{5.16}
$$

Define the primal and adjoint residual as:

$$
\begin{aligned}
\mathbf{R}_n^{\mathrm{pr}} &\equiv \mathbf{A}\mathbf{U}^n - \mathbf{f}, \\
\mathbf{R}_n^{\mathrm{du}} &\equiv \mathbf{A}^T\mathbf{\Psi}^n - \mathbf{g}.
\end{aligned}
\tag{5.17}
$$

103

Then, under the fixed preconditioner assumption, the residuals may be written explicitly as:

$$\begin{aligned}
\mathbf{R}_n^{\mathrm{pr}} &= (\mathbf{I} - \mathbf{A}\mathbf{P}^{-1})^n \mathbf{f}, \\
\mathbf{R}_n^{\mathrm{du}} &= (\mathbf{I} - \mathbf{A}^T \mathbf{P}^{-T})^n \mathbf{g}.
\end{aligned} \tag{5.18}$$

Now examine the output estimates obtained from the primal state approximation together with a correction term consisting of the inner product of primal residual and adjoint state approximations. The following exact result holds [36]:

$$\begin{aligned}
J^{\mathrm{pr}}(\mathbf{U}^n) - \langle \boldsymbol{\Psi}^n, \mathbf{R}_n^{\mathrm{pr}} \rangle &= \langle \mathbf{g}, -\sum_{i=0}^{n-1} (\mathbf{I} - \mathbf{P}^{-1}\mathbf{A})^i \mathbf{P}^{-1}\mathbf{f} \rangle + \langle \mathbf{g}, \sum_{i=0}^{n-1} \mathbf{P}^{-1}(\mathbf{I} - \mathbf{A}\mathbf{P}^{-1})^{n+i} \mathbf{f} \rangle \\
&= \langle \mathbf{g}, -\sum_{i=0}^{n-1} (\mathbf{I} - \mathbf{P}^{-1}\mathbf{A})^i \mathbf{P}^{-1}\mathbf{f} \rangle - \langle \mathbf{g}, \sum_{i=0}^{n-1} (\mathbf{I} - \mathbf{P}^{-1}\mathbf{A})^{n+i} \mathbf{P}^{-1}\mathbf{f} \rangle \\
&= \langle \mathbf{g}, -[\mathbf{I} + (\mathbf{I} - \mathbf{P}^{-1}\mathbf{A})^n] \left(\sum_{i=0}^{n-1} (\mathbf{I} - \mathbf{P}^{-1}\mathbf{A})^i \mathbf{P}^{-1}\mathbf{f} \right) \rangle \\
&= \langle \mathbf{g}, \sum_{i=0}^{2n-1} (\mathbf{I} - \mathbf{P}^{-1}\mathbf{A})^i \mathbf{P}^{-1}\mathbf{f} \rangle \\
&= J^{\mathrm{pr}}(\mathbf{U}^{2n}).
\end{aligned} \tag{5.19}$$

That is, in this instance the corrected output estimate exactly equals the original sequence of output estimates at twice the iteration count. Thus, the correction term $\langle \boldsymbol{\Psi}^n, \mathbf{R}_n^{\mathrm{pr}} \rangle$ can either be used as an output error estimate or increase the convergence rate by exactly a factor of 2.

The above is a result relating the convergence of duality corrected output to the original values in a special situation. More generally, for output estimates obtained from arbitrary iterative solvers of *linearized* primal and dual systems, one has the following error expansion [57]:

$$J^{\mathrm{pr}}(\mathbf{U}) = J^{\mathrm{pr}}(\mathbf{U}^n) - \langle \boldsymbol{\Psi}^n, \mathbf{R}_n^{\mathrm{pr}} \rangle + \langle \boldsymbol{\Psi} - \boldsymbol{\Psi}^n, \mathbf{A}(\mathbf{U} - \mathbf{U}^n) \rangle. \tag{5.20}$$

Hence, if $\sigma_{\min}$ is the minimum singular value of $\mathbf{A}$, the following bound on the output

error is obtained:

$$|J^{\mathrm{pr}}(\mathbf{U}) - (J^{\mathrm{pr}}(\mathbf{U}^n) - \langle \mathbf{\Psi}^n, \mathbf{R}_n^{\mathrm{pr}} \rangle)| \leq \frac{1}{\sigma_{\min}} \|\mathbf{R}_n^{\mathrm{pr}}\| \|\mathbf{R}_n^{\mathrm{du}}\|. \qquad (5.21)$$

Therefore, if the primal and dual residual converge at the same rate, the corrected output estimate is expected to converge at twice the original rate. Alternatively, the correction term can be used as an effective error estimate due to incomplete iteration. For the concurrent flow-adjoint solver developed in this chapter the underlying problem (5.1) is nonlinear. In this situation, higher-order terms need to be considered in addition to those in the error expression (5.20). Nevertheless, the inner product $\langle \mathbf{\Psi}^n, \mathbf{R}_n^{\mathrm{pr}} \rangle$ constitutes a first-order error correction.

5.5   Convergence and timing results

All the computational results in this section are performed on a 2688 element mesh around a NACA 0012 airfoil. All test cases are subsonic with $M_\infty = 0.5$ and angle of attack $\alpha = 2.0°$. For the laminar test case, $Re = 5000$ is used. The line solver results are obtained with the preconditioner consisting of a set block-tridiagonal pieces of the Jacobian obtained from line elements as described in [25]. The block-tridiagonal subsystems are LU-decomposed to obtain the state updates. For all v-cycle multigrid solver results, line element Jacobi smoothing as described in [25, 67] is used. On each level, 8 pre- and post-smoothing iterations are carried out and 150 sweeps are performed on the coarsest, $p = 0$ level.
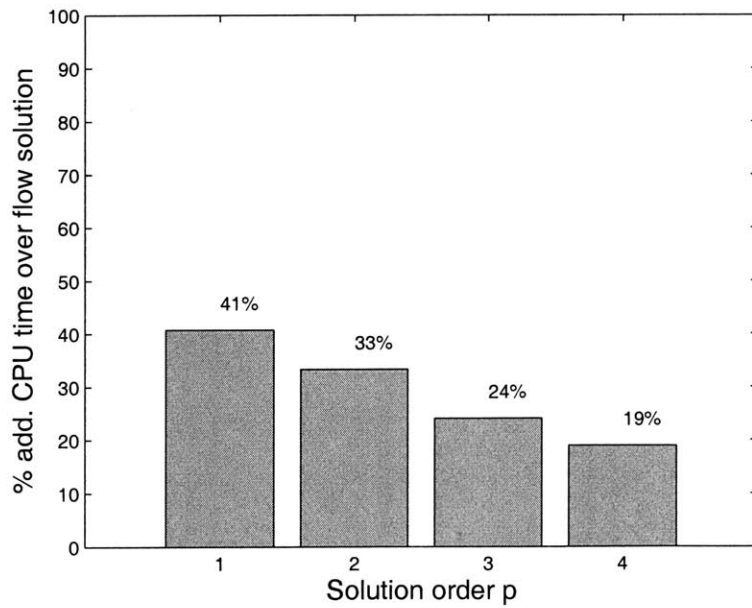
5.5.1   Inviscid Euler equations

Comparison of the compute time profiles for the baseline flow and flow-adjoint codes show that for all orders, about 2/3 of the additional CPU time required by the latter arises from computing the adjoint residual. Since for DG schemes the residual coupling exist only through nearest neighbor elements, as discussed in Section 5.3 the adjoint residual computation can be done in the same loops over the elements

105

and edges as carried out for the primal problem. By making use of the linearization already performed by the baseline flow solver, for all orders the adjoint residual is observed to require only about 1/2 the time taken to compute the flow residual and Jacobian. As the solution order increases, the compute time required to invert the preconditioner $\mathbf{P}$ becomes dominant. Using $\mathbf{P}^{-T}$ to obtain adjoint updates avoids another preconditioner inversion for the adjoint problem hence results in significant efficiency. In fact, the ratio of the time needed to invert the primal, block-tridiagonal element line Jacobian against the total additional time required for adjoint computation increases from about 0.8 to over 3 as the solution order increases from $p = 1$ to $p = 4$. This largely accounts for the decrease in the compute time required per adjoint solution in relation to the baseline flow solution from 31% to 13% over the same range of solution orders shown in Figure 5-1. The implications for multigrid solution are similar and demonstrates favorable scaling of the relative efficiency of adjoint solver with the finest solution order p.

Figure 5-2 shows the residual and output convergence for concurrent flow-adjoint line solver on an initial $p = 3$ solution using 10 inner iterations per step. The flow and adjoint residual rates reach the same asymptote and the drag and lift output corrections do indeed result in doubled convergence rate. Figure 5-3 shows the convergence of full multigrid (FMG) starting from a $p = 0$ solution. As can be seen from the top plot, the convergence of the primal and dual residuals again share the same asymptotic rates. For comparison, the convergence for the sequential adjoint solution approach of linearizing about a converged flow solution is shown in Figure 5-4, demonstrating that the adjoint convergence using the concurrent and sequential approaches are similar with the former somewhat affected by the convergence rate of the flow solution. However, as shown in the bottom plot of Figure 5-3, the former enables duality-based correction for outputs that results in full convergence within 15 multigrid iterations.
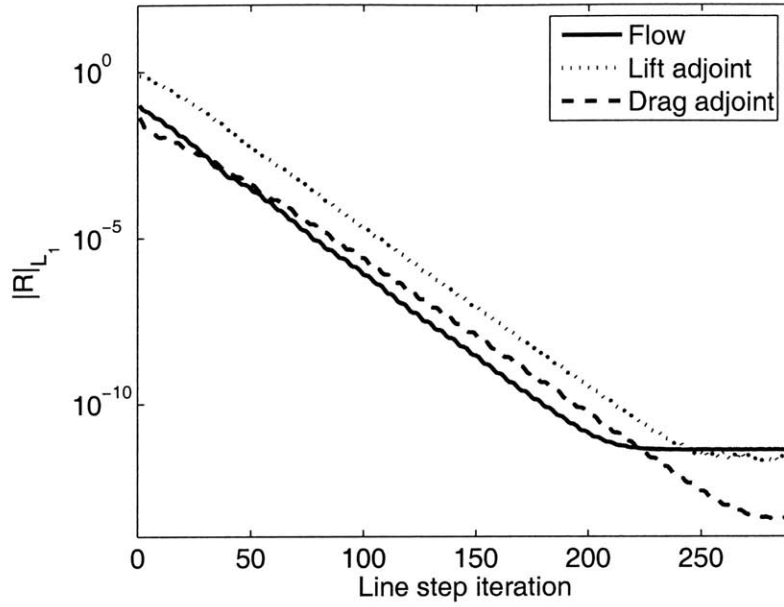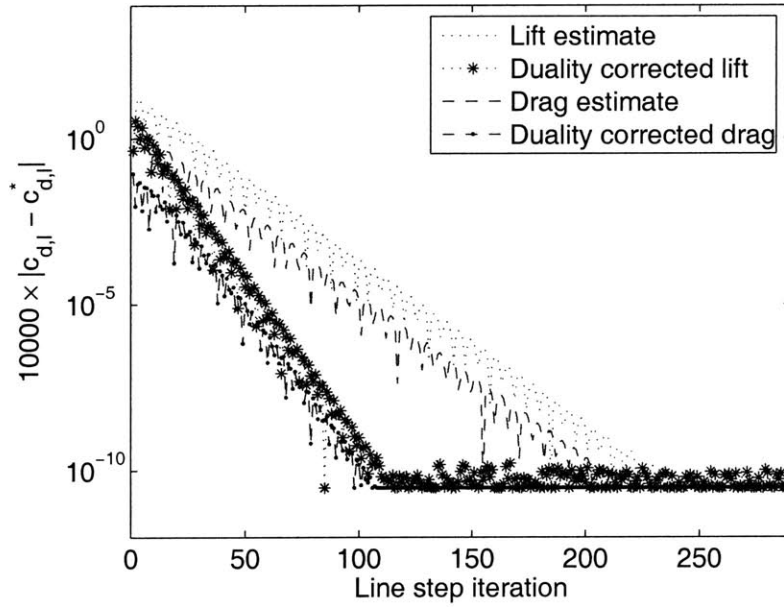
(a) Line solver

(b) MG solver

Figure 5-1: Relative cost of adjoint solution. Inviscid Euler equations.
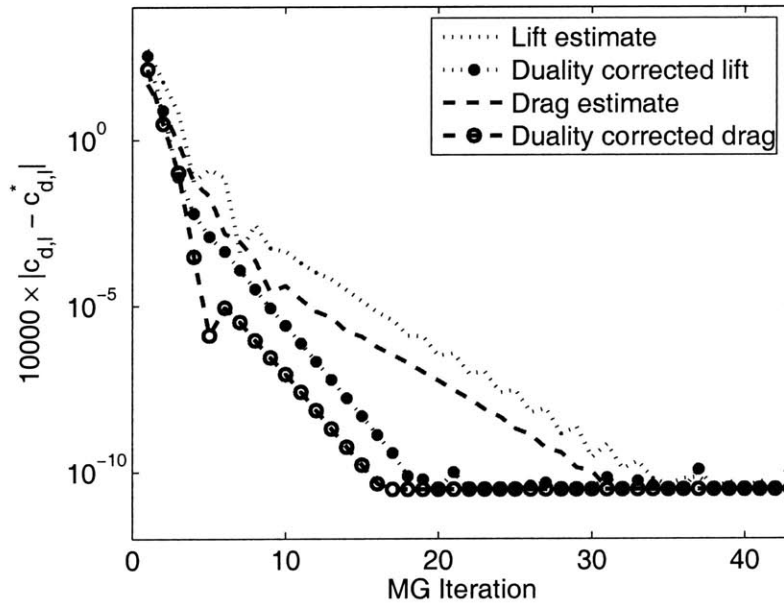
(a) Residual convergence

(b) Output convergence

Figure 5-2: Convergence of concurrent flow-adjoint line solver for $p = 3$ interpolation. Inviscid flow over NACA 0012 airfoil, $M_\infty = 0.5$, $\alpha = 2.0°$.

(a) Residual convergence

(b) Output convergence

Figure 5-3: Convergence of concurrent flow-adjoint multigrid for $p = 3$ interpolation. Inviscid flow over NACA 0012 airfoil, $M_\infty = 0.5$, $\alpha = 2.0°$.
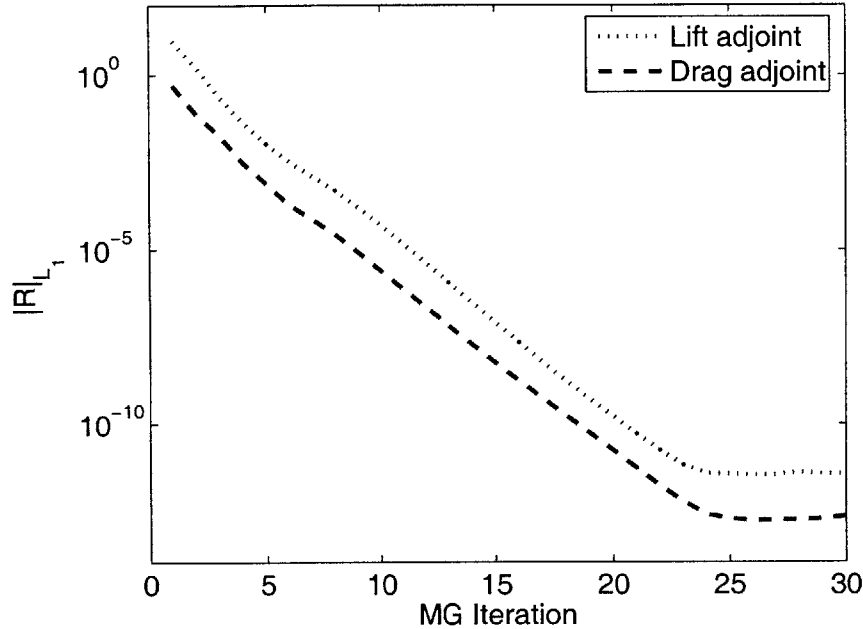
Figure 5-4: Convergence of sequential adjoint multigrid for $p = 3$ interpolation. Inviscid flow over NACA 0012 airfoil, $M_\infty = 0.5$, $\alpha = 2.0°$.

5.5.2 Compressible Navier-Stokes equations

In contrast to the inviscid results of Section 5.5.1, for the compressible Navier-Stokes case the computational effort in the current implementation of the baseline primal algorithm is dominated by the flow residual and Jacobian calculation rather than cost of inverting the preconditioner. In fact, for all of $p = 1$ to $p = 4$ solutions these computations turn out to take around 90% of the total solver time. Comparatively, the adjoint residual calculation is a much smaller proportion than the inviscid case: here, the adjoint residual requires less than 5% that of the primal residual/Jacobian. A factor that contributes to the relative efficiency of the former is that the discrete adjoint residual is defined algebraically, in particular it is independent of the complexity required to compute the various contributing terms and the number of quadrature points required to approximate the corresponding integrals both of which result in the increased cost of the primal residual. Given that the linearization is available from the primal routine, the discrete adjoint residual for the chosen DG discretiza-

110

tion can always be computed via a matrix vector multiplication that is independent of solution order and physics. This fact accounts for the minimal cost of computing the adjoint solution in relation to the flow which is clearly shown by Figure 5-5. The residual and output convergence for the concurrent, $p = 3$ line solver are given in Figure 5-6. Again, the flow and adjoint residual convergence reach the same asymptotic rate as guaranteed by the exact dual construction. The corrected outputs fully converge within half the number of solver iterations that needed for the flow residual and original outputs. Similar results are obtained for the $p = 0$ to 3 FMG solver, shown in Figure 5-7. Figure 5-8 shows the adjoint convergence for the sequential approach, where the adjoint iterations are only commenced after full convergence of the flow equations. Comparison of the concurrent and sequential approaches verify that the two converge with the same rate, differing only in the initial transient where the solution order in the former is increased from $p = 0$ to $p = 3$.

(a) Line solver

(b) MG solver

Figure 5-5: Relative cost of adjoint solution. Compressible Navier-Stokes equations.

(a) Residual convergence

(b) Output convergence

Figure 5-6: Convergence of concurrent flow-adjoint line solver for $p = 3$ interpolation. Laminar flow over NACA 0012 airfoil, $M_\infty = 0.5$, $Re = 5000$, $\alpha = 2.0°$.

113

(a) Residual convergence

(b) Output convergence

Figure 5-7: Convergence of concurrent flow-adjoint multigrid for $p = 3$ interpolation. Laminar flow over NACA 0012 airfoil, $M_\infty = 0.5$, $Re = 5000$, $\alpha = 2.0°$.

Figure 5-8: Convergence of sequential adjoint multigrid for $p = 3$ interpolation. Laminar flow over NACA 0012 airfoil, $M_\infty = 0.5$, $Re = 5000$, $\alpha = 2.0°$.

Chapter 6

Adjoint approach to shape sensitivity

6.1 Introduction

There are two main approaches to adjoint-based shape sensitivity calculation, namely discrete and continuous formulations classified according to whether duality techniques are applied to the discrete or continuous system. The fully discrete approach allows for the computation of exact gradients of the discretized functional. However, this requires one to fully differentiate the motion of the mesh with respect to design as well as calculate the full primal residual for each design variable of interest. In many applications this is undesirable as the mesh differentiation may be unavailable or difficult to implement and the residual calculation can require a non-trivial compute time which scales linearly with the design dimension. Alternatively, the continuous approach can be used which has no dependence on interior mesh movement and can be computed via the primal and dual states at the control boundary [5, 50, 42]. However, this approach has the drawback of potentially large discrepancy between the computed gradient approximation and its exact value [5] which can slow down or halt the underlying nonlinear programming algorithm if the quality of approximation is not appropriately controlled by the underlying mesh refinement procedure.

Given the advantages and disadvantages of the above described approaches, the

best choice appears to be some compromise between the two. The proposed approach is developed in the context of adaptive precision framework as discussed in Chapter 2 where exact gradients for the finite dimensional models are not required. However, the computed gradients should be *pseudo-consistent* [70] in the sense that the error is controlled to be vanishingly small as the discretization parameter is increased. The proposed incomplete gradient is computed via the discrete approach with the assumption that only the boundary elements are perturbed. Thus, it is an exact gradient for the particular case where the interior mesh is held fixed when design is updated. For general mesh movement strategy, a discrepancy with the exact gradient exists. However, as will be demonstrated the error in the proposed incomplete gradient converges at the same rate as the discrete functional. In particular, by using higher order solutions in a dual-consistent formulation, incomplete gradients are demonstrated to be in close agreement with the full discrete results for smooth solutions of compressible Navier-Stokes equations.

In Section 6.2, discrete shape sensitivity formulation is reviewed. The explicit mesh movement strategy used in this work is described in Section 6.3. Finally, computational results are shown in Section 6.4 to demonstrate the accuracy of the proposed incomplete shape sensitivities in simple aerodynamic settings as well as verifying its convergence with respect to the full discrete values.

6.2   Shape sensitivity calculation

The discussions given in this section follow closely the standard treatments given in [21, 66]. Let $\mathbf{d}$ denote the vector design variables, $\mathbf{U}$ and $\mathbf{\Psi}$ the primal and dual degrees of freedom. With explicit dependence on the mesh $\mathbf{X}$, the discrete Lagrangian $\mathcal{L}(\mathbf{U}, \mathbf{\Psi}, \mathbf{X}, \mathbf{d})$ for the primal equations $\mathbf{A}(\mathbf{U}, \mathbf{X}, \mathbf{d}) = 0$ may be obtained from the discrete cost function $J(\mathbf{U}, \mathbf{X}, \mathbf{d})$ as in Chapter 3,

$$\mathcal{L}(\mathbf{U}, \mathbf{\Psi}, \mathbf{X}, \mathbf{d}) \equiv J(\mathbf{U}, \mathbf{X}, \mathbf{d}) - \langle \mathbf{A}(\mathbf{U}, \mathbf{X}, \mathbf{d}), \mathbf{\Psi} \rangle, \tag{6.1}$$

118

where again $\langle \cdot, \cdot \rangle$ denotes the Euclidean inner product. Differentiating with respect to design $\mathbf{d}$ and using the fact that $\boldsymbol{\Psi}$ satisfies the discrete adjoint equations (3.6), the following shape sensitivity expression is obtained:

$$
\begin{aligned}
\frac{dJ}{d\mathbf{d}} &= \frac{\partial \mathcal{L}}{\partial \mathbf{d}} \\
&= \frac{\partial J}{\partial \mathbf{d}} + \left\langle \left[\frac{\partial \mathbf{X}}{\partial \mathbf{d}} \right], \frac{\partial J}{\partial \mathbf{X}} \right\rangle - \left\langle \left[\frac{\partial \mathbf{A}}{\partial \mathbf{d}} \right] + \left[\frac{\partial \mathbf{A}}{\partial \mathbf{X}} \right] \left[\frac{\partial \mathbf{X}}{\partial \mathbf{d}} \right], \boldsymbol{\Psi} \right\rangle .
\end{aligned}
\tag{6.2}
$$

The above adjoint-based approach to shape sensitivity allows for calculating derivatives in a much more efficient manner than performing $dim(\mathbf{D})$ linear solves required for the forward differentiation approach. In particular, provided $\frac{\partial \mathbf{A}}{\partial \mathbf{X}}$ is not expensive to evaluate shape sensitivities can be computed at a cost essentially independent of the design dimension.

6.3 Mesh movement procedure

The domain deformation used in the current work is an explicit strategy described in [61]. Similar strategies have been used in a number of settings [23, 82]. More sophisticated implicit strategies are needed for robustness in more complex geometries [66], which remains an area for future work. Let $\Gamma_{\text{cont}} \in \partial \Omega$ denote the control surface being perturbed and M the set of all nodes in the mesh. For all element nodes m lying on the control surface, the mesh movement $\delta \mathbf{x}_m$ is completely specified by the design variable. Consider a boundary update $\Gamma_{\text{cont}}(\mathbf{x}) \to \Gamma_{\text{cont}}(\mathbf{x}) + \gamma(\mathbf{x})$ where $\gamma(\mathbf{x})$ is a prescribed shape change. Then necessarily,

$$
\delta \mathbf{x}_m = \gamma(\mathbf{x}_m), \quad \forall m \in M_{\Gamma_{\text{cont}}},
\tag{6.3}
$$

where the control node set is defined as $M_{\Gamma_{\text{cont}}} \equiv \{ m \in M : \mathbf{x}_m \cap \Gamma_{\text{cont}} \neq \emptyset \}$. For the remaining nodes, the deformation is chosen to contain a proportional factor

accounting for their distance to the control boundary being moved. That is,

$$\delta \mathbf{x}_m = \left(\frac{1}{\sum_{k \in M_{\Gamma_{\text{cont}}}} w_k \alpha_{k,m}} \right) \times \left(\sum_{k \in M_{\Gamma_{\text{cont}}}} w_k \alpha_{k,m} \gamma(\mathbf{x}_k) \right) , \quad \forall m \in M \backslash M_{\Gamma_{\text{cont}}} \quad (6.4)$$

where the first inverse sum is a normalization factor, w_k a weight factor taken to be the size of elemental boundary segment containing node k and $\alpha_{k,m} = 1/|\mathbf{x}_k - \mathbf{x}_m|^\beta$ with $\beta \geq 1$ a decay function. It has been found that $\beta = 4$ is a good choice [61] and hence this is used throughout the current work. The mesh movement formula (6.4) can be interpreted as the convolution of kernel $1/|\mathbf{x}|^\beta$ with the shape perturbation function $\gamma(\mathbf{x})$ on the boundary.

6.4 Computational verifications

The present computational study serves two main purposes. Firstly, it is a demonstration of the effect of high-order interpolation on the accuracy of shape gradients. Secondly, it serves to verify gradient convergence for problems of relevance to adaptive precision work in Chapter 7.

To study the convergence shape sensitivities, smooth Hicks-Henne sine bump functions [43] are used as perturbations. These functions are given analytically on the unit interval as:

$$b_{x_m}(x) = \sin^4(\pi x^{\log \frac{1}{2}/ \log x_m}), \quad x \in [0, 1], \quad (6.5)$$

where $x_m \in [0, 1]$ is a parameter specifying the peak of the bump. Figure 6-1 gives a plot of the bump functions (6.5) as x_m takes on a range of values from 0.025 to 0.95.

6.4.1 Inviscid design

The test case considered in this section is the same smooth, channel flow over a Gaussian bump whose convergence is verified in Section 3.5.1. The output of interest is similarly a weighted lift over the bump surface given in (3.59). Figure 6-2 is a plot

Figure 6-1: Samples of Hicks-Henne sine functions for various x_m parameters

of lift sensitivity computed on the fine, 9169 element mesh using $p = 5$ interpolation. Each data point in the plot represents the sensitivity with respect to a perturbation in the Hicks-Henne shape function whose peak is at the given x-coordinate. Figure 6-3 compares the full and incomplete discrete adjoint sensitivities, using uniformly refined meshes containing 587 and 9169 elements with solution orders $p = 1$ and $p = 3$. For both meshes and solution orders, the discrepancy between the two are small and nearly indistinguishable on the plots. Figure 6-4 plots the discrepancies directly for various discretization levels, showing the error decay with mesh refinement. Table 6.1 gives the values of Euclidean norm of the gradient discrepancies, normalized with respect to the values of the full discrete sensitivity. Convergence plot of the same quantities is given in Figure 6-5, showing that the error decreases at $\mathcal{O}(h^{2p+1})$ rate for the $p = 1$ and $p = 2$ solutions. For the $p = 3$ solution, the convergence does not reach the expected rate. However, this degradation in the sensitivity convergence rate is not in contradiction to the lift convergence result as shown in Figure 3-4.

121

Figure 6-2: Full discrete adjoint shape sensitivities (6.2) with respect to Hicks-Henne sine bump perturbations using $p = 5$ solution on fine, 9196 element mesh. Gaussian bump, inviscid subsonic flow ($M_\infty = 0.5$).
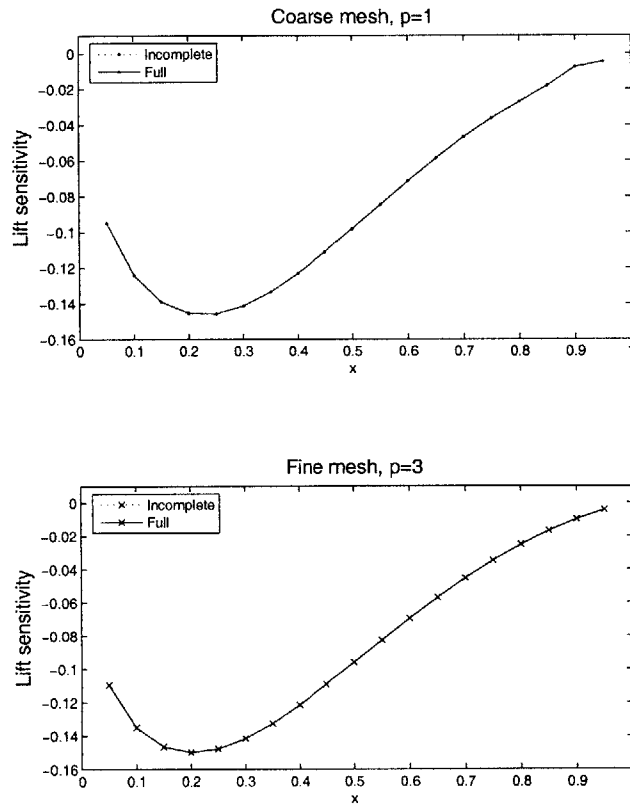

Figure 6-3: Comparison of incomplete and full discrete sensitivities. Gaussian bump, inviscid subsonic flow ($M_\infty = 0.5$).
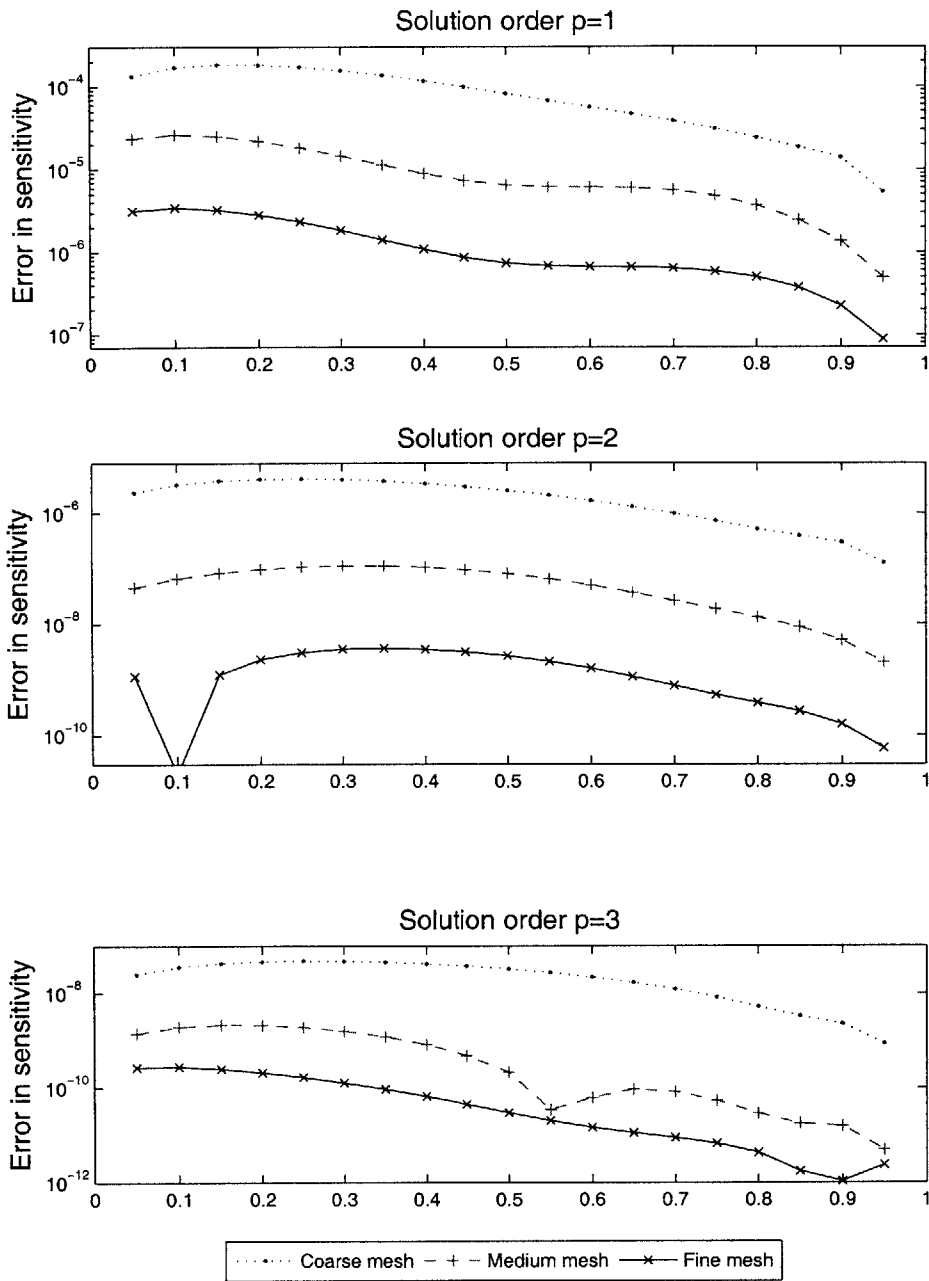
Figure 6-4: Convergence of incomplete to full discrete sensitivities. Gaussian bump, inviscid subsonic flow ($M_\infty = 0.5$).

Figure 6-5: Convergence of Euclidean norm of error in incomplete discrete sensitivities. Gaussian bump, inviscid subsonic flow ($M_\infty = 0.5$).

Order	$\left\| \frac{dJ}{dd}\|_{\text{full}} - \frac{dJ}{dd}\|_{\text{incomp.}} \right\| / \left\| \frac{dJ}{dd}\|_{\text{full}} \right\|$		
	Coarse mesh	Medium mesh	Fine mesh
p=1	1.11E-3	1.34E-4	1.72E-5
p=2	2.67E-5	7.09E-7	2.09E-8
p=3	3.15E-7	1.12E-8	1.20E-9

Table 6.1: Euclidean norm of error in incomplete discrete sensitivities. Gaussian bump, inviscid subsonic flow ($M_\infty = 0.5$).

124

6.4.2 Viscous design

Here, the lift and drag sensitivities are examined for a subsonic, laminar flow ($M_\infty = 0.5$, $Re = 5000$) over the NACA 0012 airfoil at $2°$ angle of attack for which the drag convergence is studied in Section 3.5.2. Figure 6-6 shows the computed, fully discrete sensitivities with respect to Hicks-Henne shape perturbations on the fine, 10752 element mesh with $p = 3$ solution. Figures 6-7, 6-8 and 6-9 each compare the discrete sensitivities obtained using full mesh and surface movement on a set of uniformly refined meshes containing 672, 2688 and 10752 elements. For the $p = 1$ result given in Figure 6-7, the discrepancies are significant on coarser meshes which only become small on the finest mesh. Moreoever, on the two coarser meshes there are large regions where the incomplete and full sensitivities are of the opposite sign. For the $p = 2$ result given in Figure 6-8, the discepancies are significantly smaller and very good agreement is obtained on the medium mesh. For the $p = 3$ result given in Figure 6-9, the discrepancy is small even on the coarsest mesh and is within plotting accuracy on the two finer meshes. The mesh convergence for each order is given in Figure 6-10. Tables 6.2 and 6.3 show the Euclidean norm of the discrepancy for lift and drag sensitivities respectively. Convergence plots of the respective values are given in Figures 6-11 and 6-12. For $p = 1$ and $p = 2$ solutions, the convergence for both lift and drag sensitivities are at the rates $\sim \mathcal{O}(h^2)$ and $\mathcal{O}(h^4)$ respectively. For the $p = 3$ solution, the convergence rate for lift is $\mathcal{O}(h^6)$. In the case of drag, the $p = 3$ rate is lower than expected. Again, this may be a consequence of the fact that $q = 3$ boundary representation is used and higher-order boundary interpolation may be needed to acheieve the full convergence order.

Figure 6-6: Full discrete adjoint shape sensitivities (6.2) on 10752 element fine mesh using $p = 3$ solution. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).
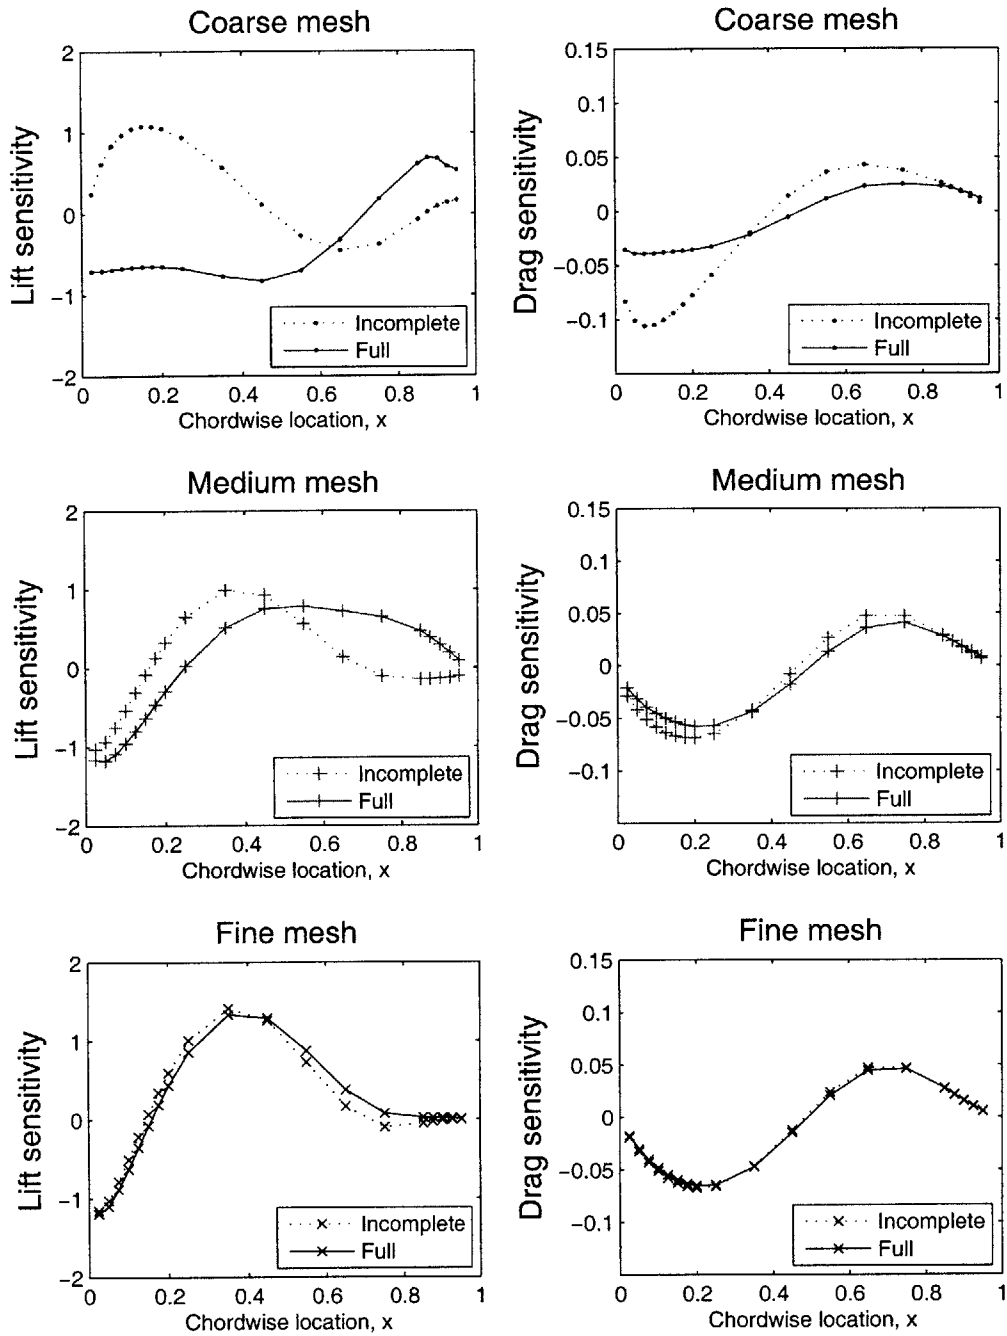
Figure 6-7: Comparison of incomplete and full discrete sensitivities for $p = 1$ solution. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).

Figure 6-8: Comparison of incomplete and full discrete sensitivities for $p = 2$ solution. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).
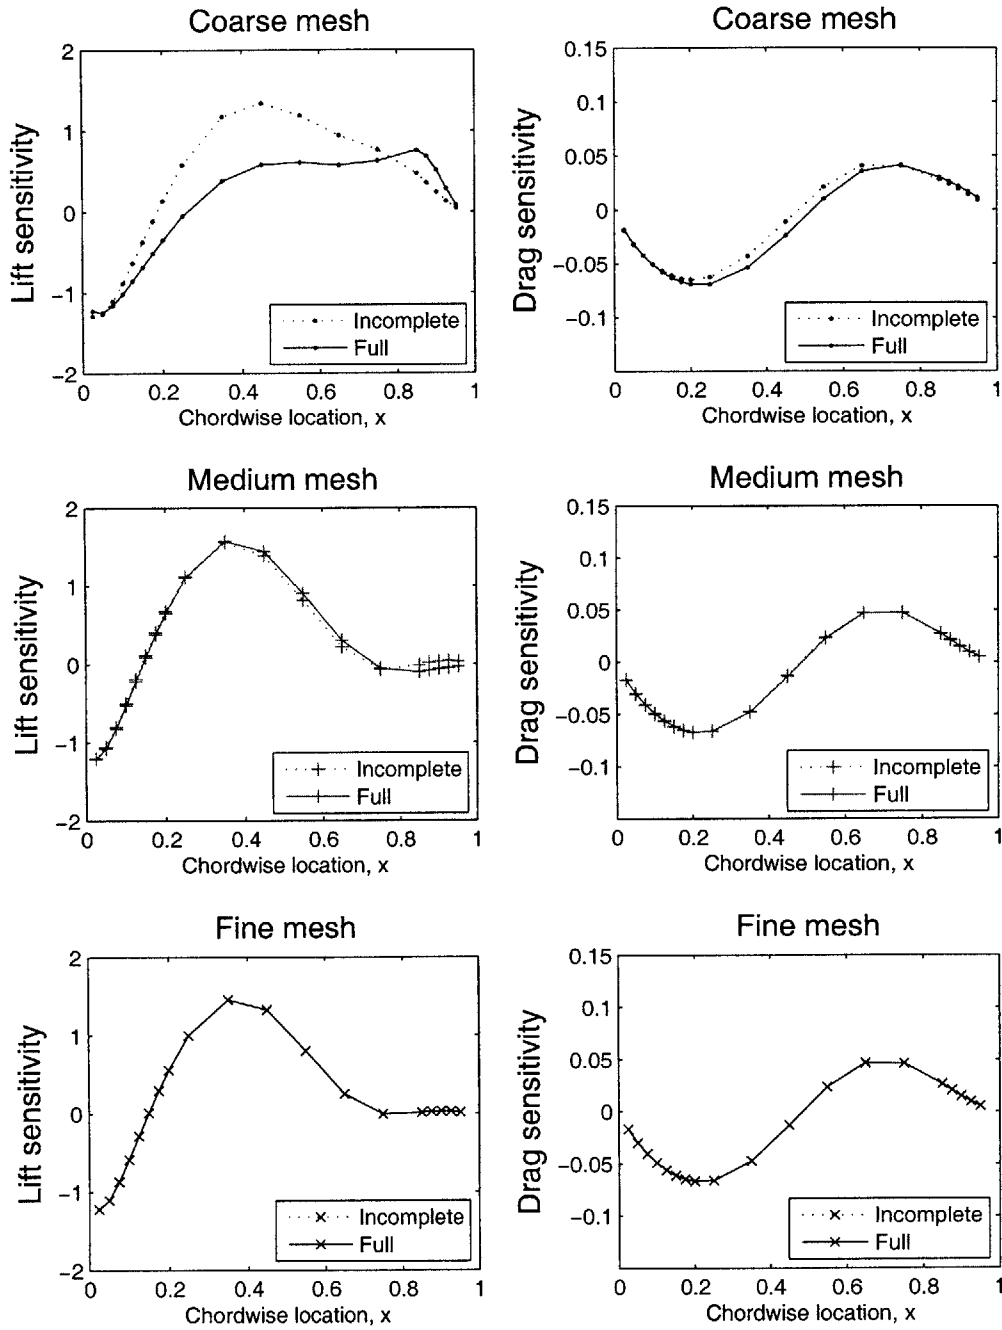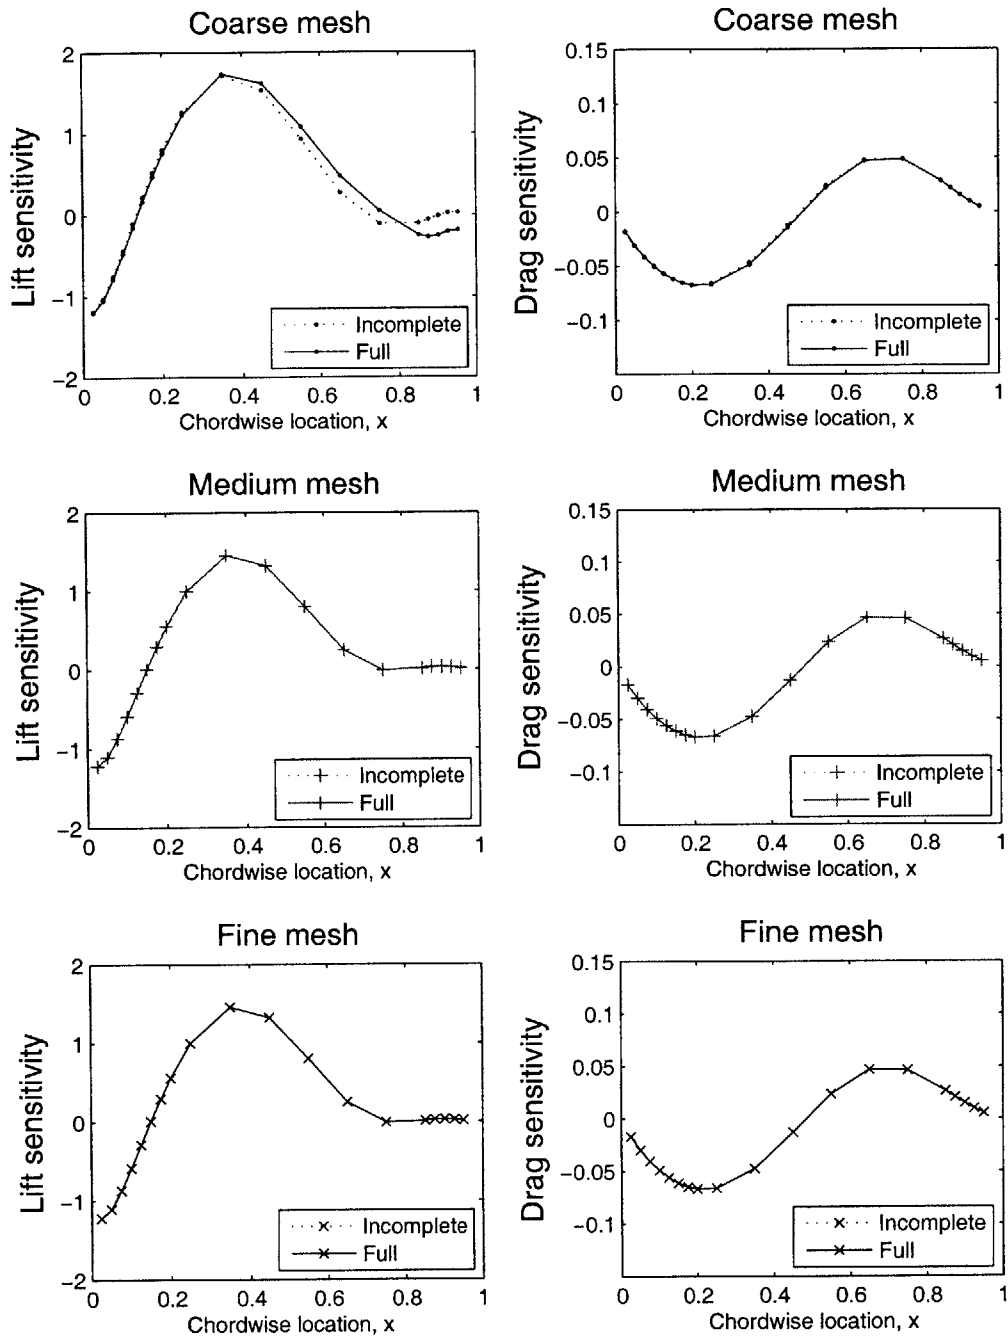
Figure 6-9: Comparison of incomplete and full discrete sensitivities for $p = 3$ solution. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).
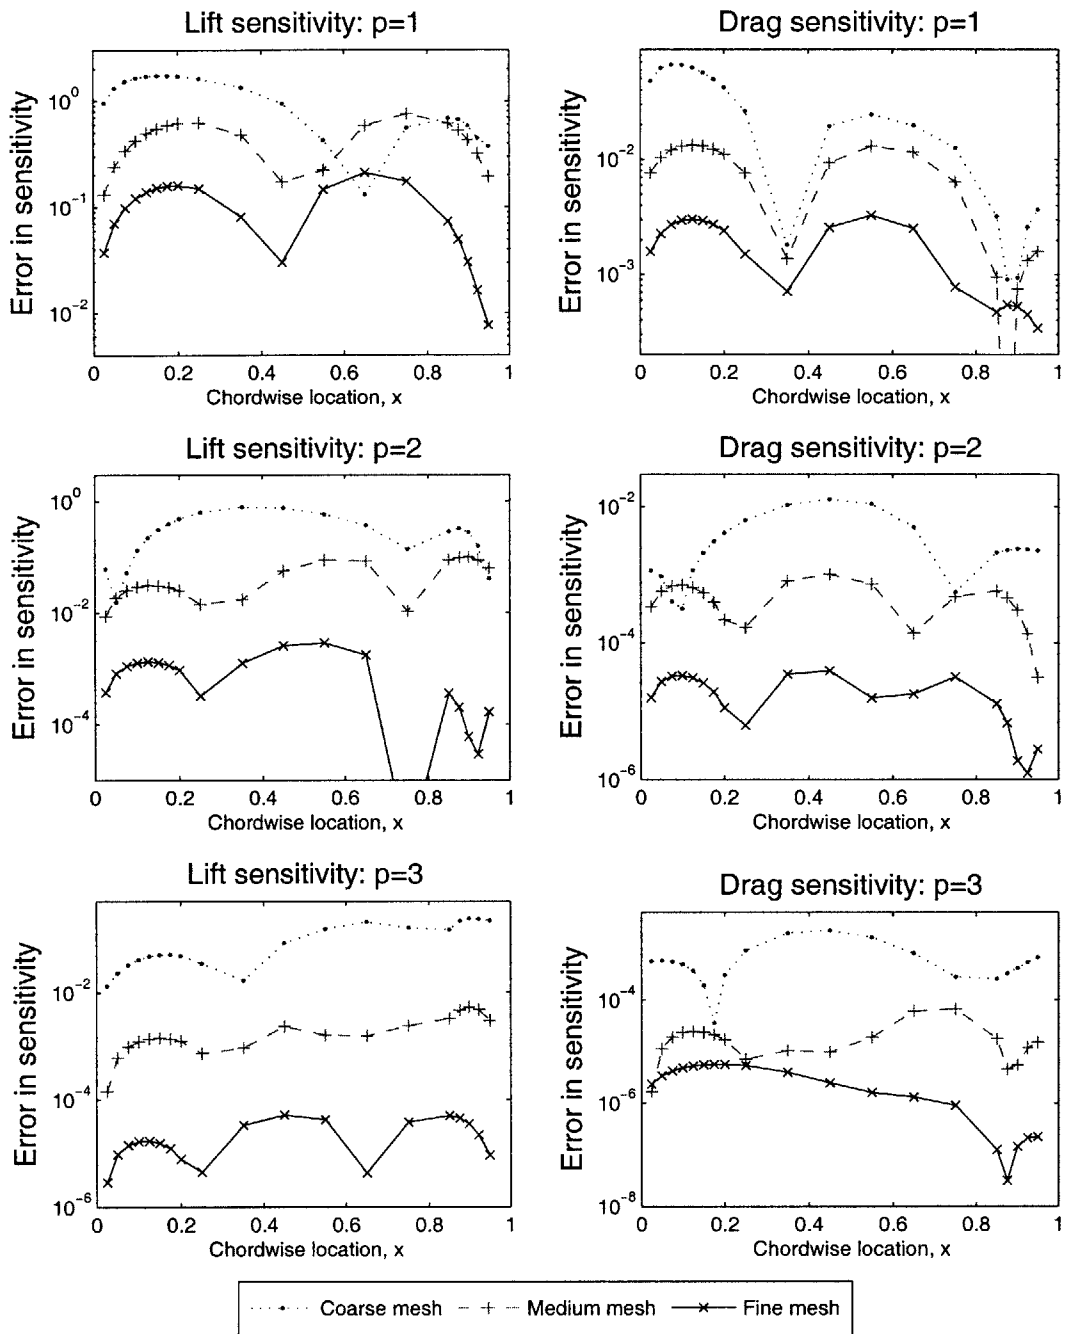
Figure 6-10: Convergence of incomplete to full discrete sensitivities. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).

Figure 6-11: Convergence of Euclidean norm of error in incomplete discrete sensitivities for lift output. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).

	$\left\| \frac{dJ}{d\mathbf{d}}\right\|_{full} - \left.\frac{dJ}{d\mathbf{d}}\right\|_{incomp.} \left\| \right. / \left\| \left.\frac{dJ}{d\mathbf{d}}\right\|_{full}\right\|$		
Order	Coarse mesh	Medium mesh	Fine mesh
p=1	1.84E-0	6.82E-1	1.67E-1
p=2	5.43E-1	7.58E-2	1.71E-3
p=3	1.61E-1	3.47E-3	3.86E-5

Table 6.2: Euclidean norm of error in incomplete discrete sensitivities for lift output. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).

131

Figure 6-12: Convergence of Euclidean norm of error in incomplete discrete sensitivities for drag output. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).

	$\left\| \frac{dJ}{d\mathbf{d}}\big|_{\text{full}} - \frac{dJ}{d\mathbf{d}}\big|_{\text{incomp.}} \right\| / \left\| \frac{dJ}{d\mathbf{d}}\big|_{\text{full}} \right\|$		
Order	Coarse mesh	Medium mesh	Fine mesh
p=1	1.35E-0	2.42E-1	5.03E-2
p=2	1.21E-1	1.26E-2	5.38E-4
p=3	2.10E-2	6.07E-4	8.33E-5

Table 6.3: Euclidean norm of error in incomplete discrete sensitivities for drag output. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).

132

Chapter 7

Applications

In this chapter, the Adaptive Precision Algorithm is applied to a number of aero-dynamic design cases. In Section 7.1, drag minimization using Hicks-Henne sine functions on the NACA 0012 airfoil is carried out. To avoid thin airfoils, a quadratic penalty term is added to maintain constant airfoil volume, as is commonly used in aerodynamic design [5, 66, 52]. In Section 7.2, an additional penalty term is added to enforce a constraint on the lift and the angle of attack is further introduced as a design variable. In Section 7.3, an inverse design problem is performed to minimize the interference effects of a nearby airfoil.

7.1 Drag minimization at constant volume

The flow case utilized here is a subsonic, laminar flow over NACA 0012 airfoil with $M_\infty = 0.5$, $Re = 5000$, at an angle of attack $\alpha = 2.0°$. As a function of design $\mathbf{d}$, the objective is a drag squared functional combined with a quadratic penalization term to maintain constant airfoil volume:

$$\mathcal{J}(\mathbf{u},\mathbf{d}) \equiv \left(\frac{c_d(\mathbf{u},\mathbf{d})}{c_d^0}\right)^2 + 10 \times \left(\frac{\mathrm{Vol}(\mathbf{d}) - \mathrm{Vol}^0}{\mathrm{Vol}^0}\right)^2, \qquad (7.1)$$

where $\mathrm{Vol}(\mathbf{d})$ denotes the airfoil volume, $c_d(\mathbf{u},\mathbf{d})$ the drag coefficient consisting of pressure and viscous stress contributions and the superscript 0 denotes the values

Figure 7-1: Hicks-Henne sine functions used for airfoil design.

corresponding to the initial, NACA 0012 airfoil discretized with $p = 1$ interpolation on a 2688 element mesh, for which drag convergence with adaptive solution order is studied in Section 4.6.

The design space used in this test case consists of 30 Hicks-Henne functions on each of the upper and lower surfaces of the aifoil, plotted in Figure 7-1. The underlying optimizer is the Broyden-Fletcher-Goldfarb-Shanno (BFGS) quasi-Newton algorithm with Armijo linear-search [34]. The concurrent flow-adjoint line-solver is used to iteratively obtain the state estimates. The Adaptive Precision Algorithm of Section 2.3 with parameters $\gamma = 2$, $\zeta = 0.2$, $\tau = 0.5$, $\omega = 1$, $E = 0.8$, $j_{max} = 3$ is applied to the initial, NACA 0012 airfoil. Figure 7-2, 7-3, 7-4 show the pressure, friction profiles as well as airfoil shapes at various stages of the optimization procedure, plotted against those of the initial NACA 0012 airfoil. Owing to the coarseness of the mesh, error exists in the pressure and friction coefficient on this initial NACA 0012 solution as is evident by the irregular profiles. By the 4th iteration, p-refinement has been carried out to ensure the validity of design updates and as a result the pressure and friction coefficients are smoother. By the 12th iteration, the errors are small and the design has improved. Over the course of the optimization and adaptation steps, the drag decreases by 4.5% of its original value while the airfoil volume increases by only 0.03%.

134

(a) Optimization step 4

(b) Optimization step 6

(c) Optimization step 12

(d) Optimization step 20

Figure 7-2: Results of drag minimization at constant volume. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

(a) Optimization step 4

(b) Optimization step 6

(c) Optimization step 12

(d) Optimization step 20

Figure 7-3: Results of drag minimization at constant volume. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

(a) Optimization step 4



(b) Optimization step 6

(c) Optimization step 12

(d) Optimization step 20

Figure 7-4: Results of drag minimization at constant volume. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.
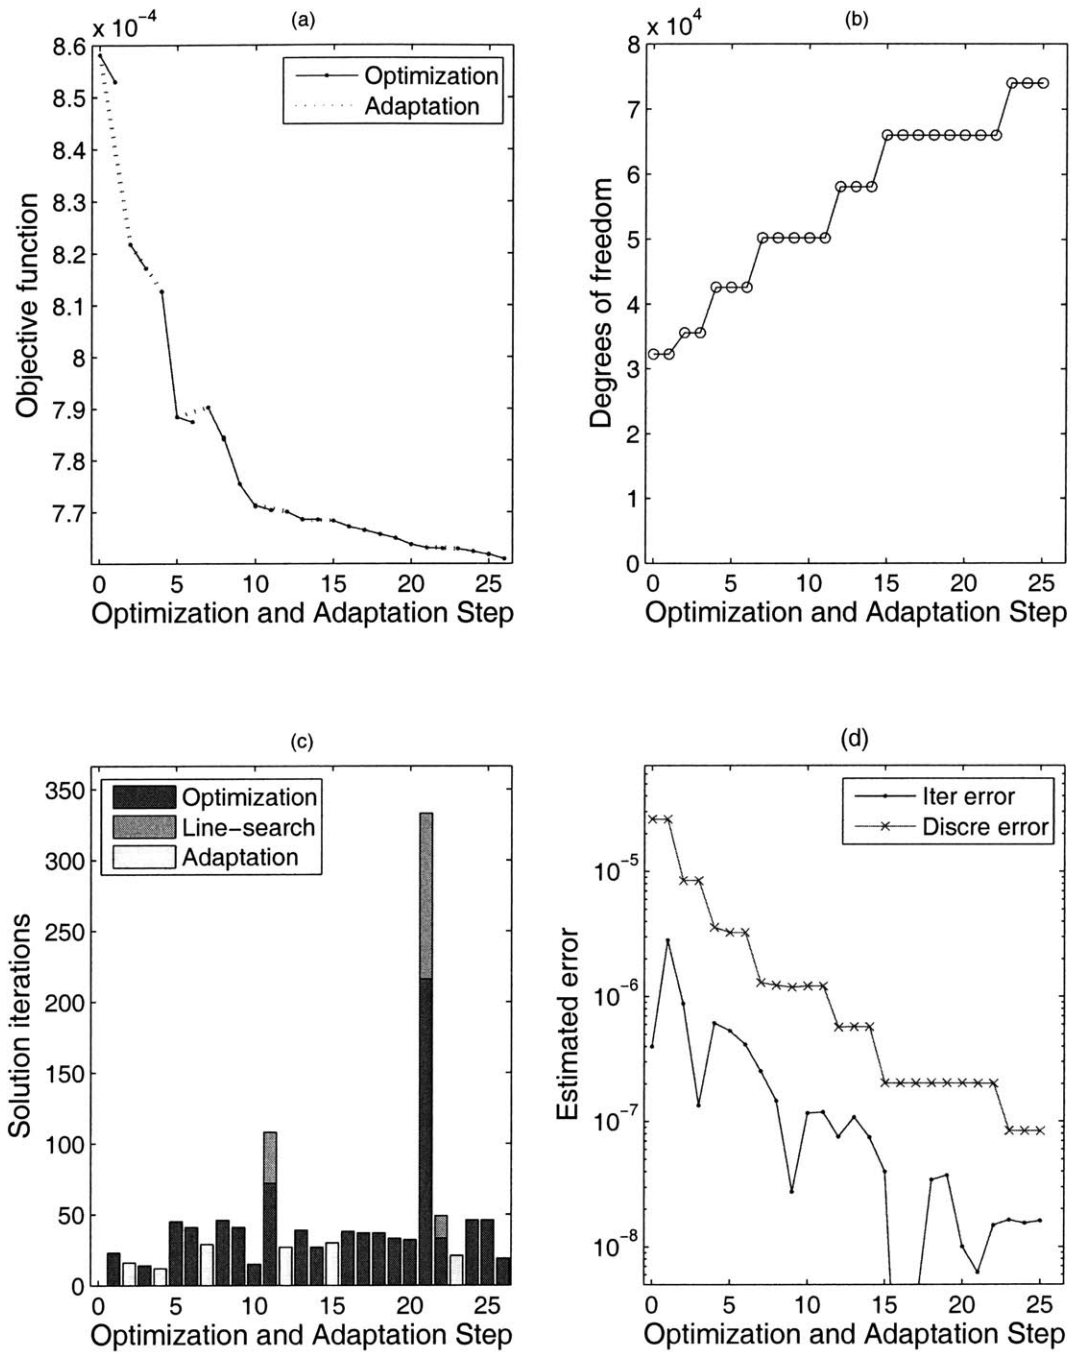
Figure 7-5: Convergence of adaptive precision optimization for drag minimization at constant volume. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

Figure 7-5 shows the convergence result obtained by the Adaptive Precision Algorithm, where successful design updates are plotted together with the adaptive steps. Plot 7-5 (a) shows the values of the computed objective function as optimization or adaptation steps are carried out. The solid line from step 0 to 1 shows that an optimization step is attempted. However, the resulting change in the value of the objective function is smaller than the estimated error. Hence, an adaptation step is carried out on the initial design and the change in objective function from carrying out p-refinement is shown by the dotted line from step 0 to step 2. The result shows that adaptation gives rise to significant change in the computed objective and confirms that the attempted initial design step does not constitute a valid design improvement. At step 2, another optimization iteration is attempted as shown by the solid line from step 2 to 3. However, it is again invalid and another adaptation step is carried out, shown by the dotted line from step 2 to 4. This adaptation again results in a large change in the computed objective function. Next, the design change is successful and the objective function improves significantly, as shown by the solid line from step 4 to 5. The rest of the convergence plot similarly follows. Plot 7-5 (b) shows the solution degrees of freedom at every step, which increases at relatively regular intervals in constant increments. The p-adaptation steps decrease the estimated error in the objective function, which is shown in plot 7-5 (d). It is seen that during optimization steps, the error stays virtually constant. Every time adaptation is carried out, the estimated error drops by a constant fraction, confirming that p-refinement is effective for this problem. By the choice of parameter $\zeta = 0.2$, as the plot confirms the iterative error estimate is always controlled to be less than 0.2 that of the discretization error estimate. Plot 7-5 (c) shows the number of solution iterations carried out for each optimization or adaptation step. In the first 4 steps, the number of solution iterations carried out are small owing to the low tolerance required. From step 5 onwards, more solution iterations are needed to meet the more stringent error requirement. The solution iterations labelled as line-search are those additional solution iterations carried out when the size of the design steps are decreased so as to satisfy Armijo's criterion. The data for the last few design steps

139

show that as the size of the design changes become small, it is no longer necessary to carry out a large number of solution iterations to meet the error criterion.

In demonstrating the adaptive precision approach, an important question is its potential reliability benefits over a fixed precision strategy where a single model is used. This is illustrated here via an example where the fixed precision approach is carried out on a coarse model. In particular, the chosen model consists of a uniform $p = 1$ solution on the same 2688 element mesh. The solutions are tightly converged, with flow residual dropping to lower than 10^{-9} its initial value. Figure 7-6 compares the airfoil shapes, pressure and friction coefficient profiles for adaptive and fixed precision approaches. Although the lower surfaces of the airfoil computed using either approach are very similar, discrepancies in the upper airfoil surfaces exist at certain locations. This leads to a difference in the friction coefficient profiles close to the leading edge of the top airfoil surface.

To compare the quality of optimized results given by the fixed and adaptive precision strategies, Table 7.1 shows certain quantities of interested computed via the approximation models and the corresponding "exact" results. The latter is computed from a $p = 4$ solution on the same 2688 element mesh. From columns 2 and 3 of Table 7.1, it can be seen that the drag is over predicted by the initial, $p = 1$ solution. This is also reflected in the fixed precision result, where the computed value is a pessimistic prediction of the optimal drag. In fact, in the current test case the true drag obtained from the fixed precision approach is essentially identical to that attained by the adaptive precision approach. However, the latter approach provides the designer with a confidence in design optimality that is lacking in the former approach. Owing to the p-adaptation carried out, the computed adaptive precision results (which has solution orders consisting mainly of $p = 2$ and $p = 3$) closely match the exact values. Furthermore, the remaining error is effectively captured by the estimator: computing the ratio of the discretization error estimate over the "true" error in the objective function shows that $I_{eff} = 0.98$.

Figure 7-6: Comparison of optimized airfoils computed using fixed $p = 1$ interpolation and Adaptive Precision Algorithm for drag minimization at constant volume. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

	Initial		Adapt. Prec.		Fixed Prec.	
	Computed	Exact	Computed	Exact	Computed	Exact
c_d	5.7769E-2	5.6829E-2	5.5171E-2	5.5168E-2	5.6224E-2	5.5169E-2
c_l	5.1590E-2	3.7035E-2	2.7280E-2	2.7210E-2	3.5906E-2	3.8896E-2
$\mathcal{J}$	3.3384E-3	3.2296E-3	3.0439E-3	3.0436E-3	3.1622E-3	3.04460E-3

Table 7.1: Computed and exact results for drag minimization at constant volume. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).

7.2 Drag minimization at constant volume and lift

The subsonic, laminar flow test case used here is the same as that used in Section 7.1. Here, the objective function takes the form:

$$\mathcal{J}(\mathbf{u},\mathbf{d}) \equiv \left(\frac{c_d(\mathbf{u},\mathbf{d})}{c_d^0}\right)^2 + 10 \times \left(\frac{c_l(\mathbf{u},\mathbf{d}) - c_l^0}{c_l^0}\right)^2 + 10 \times \left(\frac{\mathrm{Vol}(\mathbf{d}) - \mathrm{Vol}^0}{\mathrm{Vol}^0}\right)^2, \quad (7.2)$$

where in addition to terms given in (7.1) is a penalty functional involving the lift coefficient $c_l(\mathbf{u},\mathbf{d})$ and its value computed on the initial NACA 0012 airfoil, c_l^0. The results shown here is based on Adaptive Precision Algorithm using BFGS optimizer, with the same choice of parameters as that used in Section 7.1. Also, the same Hicks-Henne sine functions with the addition of the angle of attack are chosen as the design variables, with the latter implemented by rotating the chordline of the airfoil rather than modifying the farfield BC.

The presence of an additional penalty functional to maintain a specified lift provides a further test of the adaptive precision algorithm. An effect that arises in this case is that discretization error in the lift penalty functional can be significantly larger at far-off design points where the lift does not match its desired value, than at close-to-optimal design points where it does. In comparison to the test case examined Section 7.1, in this case there is expected to be a stronger influence between optimization and adaptation. If too much optimization is carried out before adaptation, then at a given design the computed lift on the coarse mesh is in fact far away from

142

its true value and hence the angle of attack has to be significantly updated when adaptation is finally carried out. In the reverse case where significant adaptation is performed before carrying out optimization, the former may result in a highly refined mesh that is necessary to accurately estimate the lift penalty functional when in fact if some optimization steps were carried out to decrease the penalty term, the demand on the accuracy of the predicted lift can be decreased.

Figures 7-7, 7-8, 7-9 show pressure, friction profiles as well as airfoil shapes at various stages of the optimization procedure, plotted against those of the initial NACA 0012 airfoil. Here, apparently due to the lift constraint, the design changes are smaller than the those shown in Section 7.1. However, design objectives are successfully achieved. At the end of optimization and adaptation, the lift and volume discrepancies are only 0.3% and 0.01% of its respective desired values while the drag has decreased by 4.1% over its initial value. Figure 7-10 shows the convergence result of the Adaptive Precision Algorithm. The computed objective function as optimization or adaptation is carried out is shown in Plot 7-10 (a). The solid line from step 0 to 1 shows that an initial design update is attempted. However it is too small for the design update to be valid, hence an adaptation step is carried out on the original design as denoted by the dotted line from step 0 to 2. Because the lift is poorly predicted on the initial mesh, the objective function increases drastically after adaptation. Given that the lift is poorly predicted, the algorithm is consistent with the heuristic that adaptation should be chosen over optimization. After adaptation, two successful design updates are carried out, denoted by solid lines from steps 2 to 3 and 3 to 4. Owing to the closer matching of computed to the desired lift, the objective function decreases significantly. A design update is attempted as denoted by the flat solid line from step 4 to 5. However, the discretization error is large in relation to the design change and hence adaptation is carried out, denoted by the dotted line from step 4 to 6. The large size of discretization error is confirmed by the relatively large change in objective function brought about by adaptation. Similar "zig-zag" effect in the objective function can be observed till step 9. Plot 7-10 (b) shows the increase in the solution degrees of freedom brought about by adaptation. It can be seen that the solution size increases

in a relatively uniform manner. Plot 7-10 (d) shows the discretization and iteration error estimates during the optimization procedure. Of particular interest is that, in contrast to Plot 7-5 (d), the discretization error is significantly non-monotone for the initial steps. In particular, the discretization error estimate can increase after adaptation steps which is a result of the computed lift on the more refined mesh being further away from the desired value than previously. Also, the discretization error can change significantly with optimization steps depending on whether the updates lift is closer to the desired value or not.

Figure 7-11 demonstrates the difference in the optimal results produced using a fixed $p = 1$ interpolation and adaptive precision approach. Some difference in trends can be observed in the optimized airfoil profiles. In particular, for the top surface airfoil shape shown in the bottom plot, near the leading edge the adaptive precision result is fairly close to the initial NACA 0012 whereas the $p = 1$ result is thinner in this region. This thinning is compensated by a slightly thicker profile in the middle of the airfoil. The discrepancy in the airfoils at the leading edge is also reflected in the computed pressure profiles shown in the top plot.

As a comparison of the outputs computed using the fixed and adaptive precision strategies, Table 7.2 shows the corresponding values for drag and lift outputs as well as the objective function. The "exact" values given in the tables are again obtained from the uniform, $p = 4$ solution computed on the same underlying, 2688 element mesh. As columns 2 and 3 of Table 7.2 show, the lift predicted by the initial, $p = 1$ solution is too large by almost 1/3 of the true value, hence resulting in a large discrepancy between the computed and true objective function values. For the optimized drag, the computed $p = 1$ result is $c_d = 5.6401E\text{-}2$ whereas the adaptive precision result is $c_d = 5.5402E\text{-}2$. Since the exact result for the fixed precision optimized airfoil is $c_d = 5.5435E\text{-}2$, the discrepancy arises as a result of inaccuracy in the drag prediction. The c_l for the fixed precision result computed using $p = 1$ interpolation is $5.3893E\text{-}2$, which is significantly different from the desired value of $c_l^0 = 5.1590E\text{-}2$. Thus, the lack of precision in the $p = 1$ solution results in an optimized design that deviates from an underlying design goal of maintaining a specified lift. Upon evaluation on a fine

144

(a) Optimization step 3

(b) Optimization step 12

(c) Optimization step 16

(d) Optimization step 20

Figure 7-7: Results of drag minimization at constant volume and lift. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0^\circ$.

(a) Optimization step 3

(b) Optimization step 12

(c) Optimization step 16

(d) Optimization step 20

Figure 7-8: Results of drag minimization at constant volume and lift. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

146

(a) Optimization step 3

(b) Optimization step 12

(c) Optimization step 16

(d) Optimization step 20

Figure 7-9: Results of drag minimization at constant volume and lift. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.
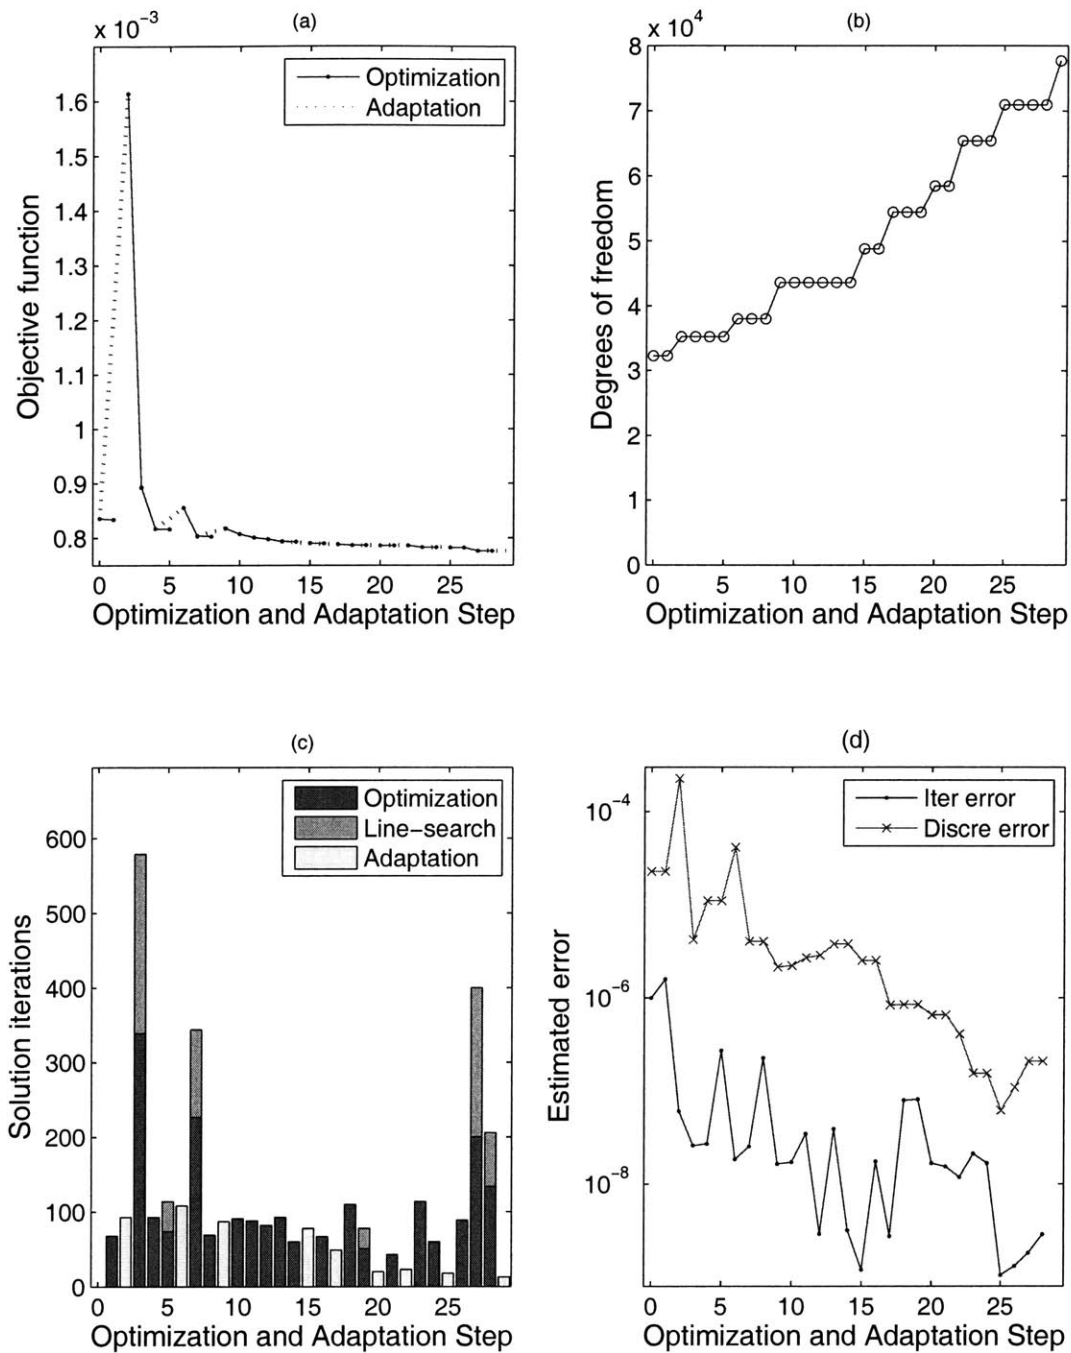
Figure 7-10: Convergence of adaptive precision optimization for drag minimization at constant volume and lift. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

mesh, redesign has to be carried out so that the lift requirement is met. Finally, for the adaptive precision results the small discrepancy between the exact and computed objectives (where again the adapted solution order is mostly of $p = 2$ or $p = 3$) confirms the effectiveness of error control strategy within the optimization procedure. The effectivity is computed to be $I_{eff} = 1.55$, thereby demonstrating the correctness of the estimated error at convergence.

Figure 7-11: Comparison of optimized airfoils computed using fixed $p = 1$ interpolation and Adaptive Precision Algorithm for drag minimization at constant volume and lift. NACA 0012 airfoil, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

150

	Initial		Adapt. Prec.		Fixed Prec.	
	Computed	Exact	Computed	Exact	Computed	Exact
c_d	5.7769E-2	5.6829E-2	5.5402E-2	5.5398E-2	5.6401E-2	5.5435E-2
c_l	5.1590E-2	3.7035E-2	5.1447E-2	5.1458E-2	5.1341E-2	5.3893E-2
$\mathcal{J}$	3.3438E-3	1.3370E-2	3.1048E-3	3.1043E-3	3.1827E-3	3.1272E-3

Table 7.2: Computed and exact results for drag minimization at constant volume and lift. NACA 0012, laminar subsonic flow ($M_\infty = 0.5$, $Re = 5000$).

7.3 Interference inverse design

The inverse design problem examined here is motivated by practical applications where it is desired to minimize the aerodynamic interference effect on a component whose performance in isolation has been separately optimized. This includes situations such as the problem of redesigning the wing so as to minimize the aerodynamic impact of the engine nacelle. For the combined configuration, by inverse designing the wing so that the pressure profile matches its original values as far as possible, the effect of the engine is hopefully minimized.

The particular case examined here is that of inviscid, subsonic flow ($M_\infty = 0.5$, $\alpha = 0.0°$) over two NACA 0012 airfoils in close proximity. The two airfoils are placed in parallel, with the airfoil undergoing redesign placed one chord in front and above the airfoil whose interference effect is to be minimized. The airfoils and the underlying 3382 element mesh are shown in Figure 7-12. The goal of the inverse design is to match the target pressure profile of a NACA 0012 airfoil placed in isolation with the same flow condition. For the computations shown here, the desired pressure profile $p_{target}(\mathbf{x})$ is taken to be the value computed using a $p = 4$ interpolation on a uniform mesh containing 7344 elements. That is, the objective function is given by:

$$\mathcal{J}(\mathbf{u}, \mathbf{d}) \equiv \int_{\Gamma_{airfoil}} \left(p(\mathbf{u}, \mathbf{d}) - p_{target}(\mathbf{x}) \right)^2 ds_x, \qquad (7.3)$$

where s_x is the chordwise distance along the airfoil. The design variables are the 30

151

Figure 7-12: Mesh for interference inverse design. NACA 0012 airfoils in close proximity, $M_\infty = 0.5, \alpha = 0.0°$.

Hicks-Henne sine functions (shown in Figure 7-1) used to design the mean camber line of the airfoil, as well as the angle of rotation pivoted about the mid-span of the airfoil. Thus, the volume of the airfoil is not affected by the design variables. The optimization problem is initiated using a uniform, $p = 1$ solution.

Figure 7-13 and 7-14 show the pressure profiles and airfoil shapes respectively at various steps of adaptive precision optimization. At the first step, significant tilt downward is introduced at the leading edge while the trailing edge is deflected upwards. This results in the pressure at the leading edge closely matching that of the desired distribution but an oscillatory pressure profile is produced near the trailing edge. By step 3, the trailing edge pressure oscillation is removed and is significantly closer to the desired profile. By step 6, the upper and lower pressure profiles are essentially overlapping except near the leading edge. For the final design, in the front section of the airfoil close matching to the desired pressure profile is attained. In the rear portion of the airfoil, interference effect of the trailing airfoil remains.

Figure 7-15 shows the convergence of the Adaptive Precision Algorithm. From Plot 7-15 (a) it can be seen that most of the design improvements take place within the first 5 optimization steps, after which improvements are significantly decreased and adaptation steps are needed to ensure that discretization error is small in comparison. Plot 7-15 (d) shows significant decrease in the discretization error error estimates when adaptation is carried out at steps 11, 18 and 20. Finally, Plot 7-15 (c) shows that from step 20 onwards, as the design converges the number of solution iterations necessary to meet the accuracy tolerance decreases.

(a) Optimization step 1

(b) Optimization step 3

(c) Optimization step 6

(d) Optimization step 25

Figure 7-13: Results of interference inverse design. NACA 0012 airfoils in close proximity, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

(a) Optimization step 1

(b) Optimization step 3

(c) Optimization step 6

(d) Optimization step 25

Figure 7-14: Results of interference inverse design. NACA 0012 airfoils in close proximity, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.
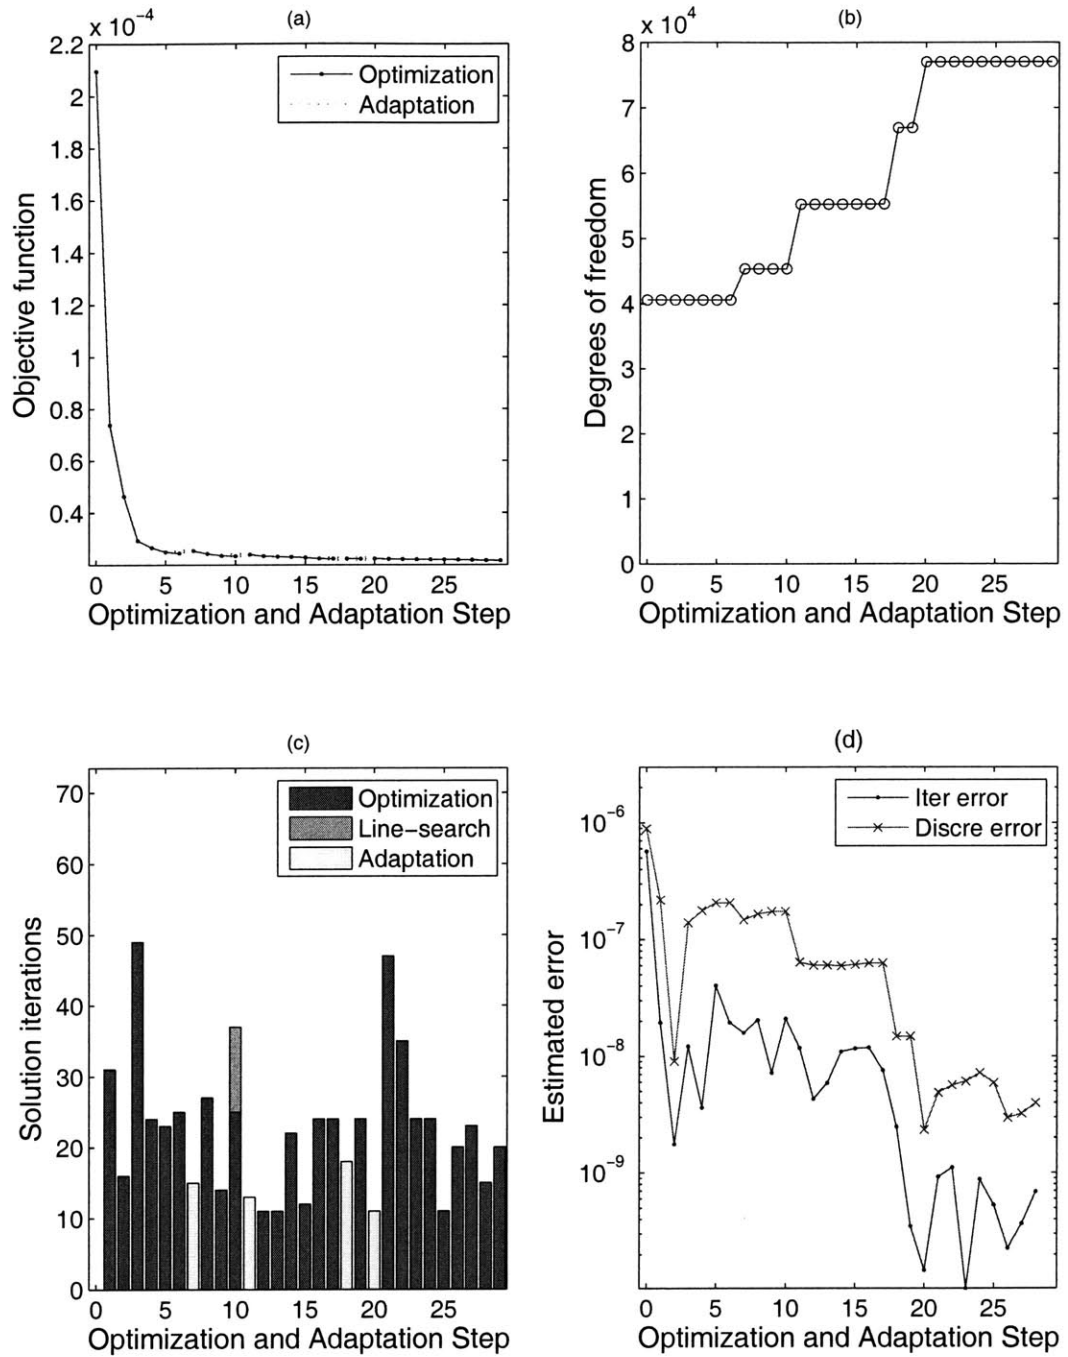
Figure 7-15: Convergence of adaptive precision optimization for interference inverse design. NACA 0012 airfoils in close proximity, $M_\infty = 0.5, Re = 5000, \alpha = 2.0°$.

Chapter 8

Conclusions

A general methodology has been demonstrated to incorporate *a posteriori* error estimates to give the designer confidence in the validity of design updates computed on approximation models. By introducing a precision control framework based on rigorous error estimates, approximation models can be used in computational optimization without compromising reliability while ensuring that the limit points obtained by the algorithm are stationary points for the underlying continuous system of interest. The error estimation ingredients for the proposed methodology can be realized in an inexpensive manner by utilizing a single adjoint solution that is already computed for shape sensitivities and hence does not require any additional global solves. In terms of the computational cost benefit, by allowing the model precision to increase adaptively when necessary and effecting error control on the objective function of interest via local refinements, the time needed for carrying out computational optimization can be significantly decreased in comparison to the fixed precision approach where computations may have to be carried out on a fine mesh constructed *a priori*.

In the aerodynamic optimization context, higher-order DGFEM is demonstrated to be an effective way to realize the adaptive precision methodology. From its consistency and dual-consistency properties, discretization error for functional outputs can be accurately estimated and controlled. Due to its nearest-neighbor stencil, DGFEM allows for an efficient implementation of concurrent flow-adjoint solver necessary for iterative error estimation. These attributes of DGFEM also make it attractive for

general control applications, independent of the benefits within the adaptive precision framework. Also of interest to general shape optimization applications is the demonstration that exact shape sensitivities can be accurately estimated from higher-order solutions via an incomplete formulation, where the contribution from the potentially expensive or unavailable interior mesh motion is ignored.

Based on DGFEM implementation developed, the work presented here represents a demonstration of adaptive precision methodology in simple but relevant aerodynamic optimization settings using inviscid Euler and compressible Navier-Stokes equations. From the numerical results, a number of conclusions may be drawn. Firstly, provided discretization error can be reliably assessed, the model precision can be effectively controlled in the optimization context. The optimization procedure can be initiated on a coarse solution and adaptively refined along the optimization trajectory so that a fine mesh that represents the true optimal solution to a sufficient accuracy is obtained at convergence. Secondly, within the optimization context the iterative error estimates for the objective function can be used as an effective stopping criterion. By requiring an increasingly stricter tolerance on the iterative error, the number of solution iterations carried out is appropriately increased.

Future work

The methodology as formulated is restricted to the context of unconstrained minimization. Hence, in the computational results demonstrated the desired equality constraints have been weakly enforced via the use of penalty functions. Future research can be carried out to extend the methodology to classes of optimization algorithms that handle the constraints in an explicit manner. In particular, for cases where inequality constraints are enforced through functionals of the state, numerical errors for both the objective and constraints need to be considered.

Significant further work is necessary to apply the proposed duality-based tools to practical, large-scale aerodynamic problems. In the area of error estimation and mesh adaptation for functional outputs, more work is necessary to study the quality

of discretization error estimates in the presence of flow discontinuities. Furthermore, to effectively control the discretization error, hp-adaptation strategies need to be developed. In the area of shape sensitivities, it remains to show the accuracy of incomplete gradient computed from hp-adapted meshes when sharp discontinuities exist in the solutions computed around complex geometries.

Some theoretical issues also remain to be addressed. Firstly, although finite element discretizations of linear elliptic systems have been shown to form consistent approximations, the same is not known for convection dominated, nonlinear problems relevant to many applications. In certain situations, theoretical analysis may be needed to choose the objective function appropriately in order for the discretized problems to consistently approximate the continuous counterpart. Secondly, given consistent approximations obtained by some chosen discretization, a question that arises from the use of *a posteriori* error estimates is how does the discretization error depend parametrically on the design in a given neighborhood and if the error can be effectively estimated. This is an issue that may need to be better understood in situations involving more complex geometries or the use of shape design functions of less regularity.

Appendix A

Review of Discretizations for Adjoint of Inviscid Euler Equations

In this appendix, a review of previous work on the discretization of the adjoint PDE is given, first for non-conservative finite-volume formulations and then for a recently proposed conservative, DG formulation. Although the latter formulation is consistent for smooth flows, it is not so for general shocked flows as is demonstrated in the context of quasi-one dimensional Euler equations.

A.1 Non-conservative formulations

The adjoint discretization adopted by Jameson *et al.* [75, 76, 51, 63] is a second-order central difference scheme with blended first and third-order artificial dissipation of the same form as that used for the flow equations to remove high frequency oscillatory errors.

In [5] Anderson and Venkatakrishnan implements a vertex-based finite volume discretization of continuous adjoint for compressible Naveris-Stokes equations. The derivation of the scheme is guided by the discrete adjoint formulation, so that for the first order accurate version, direct correspondence with discrete adjoint is achieved. This is in part motivated by the desire to maintain good agreement of adjoint-based sensitivities with the finite difference results. Let N_i be the set of vertex neighbors of

vertex i in a given triangulation T_h of the flow domain. Taking the median duals of T_h to obtain control volumes Ω_i, the inviscid discretization for the adjoint is:

$$\frac{\partial}{\partial t} \int_{\Omega_i} \psi d\mathbf{x} - \sum_{j \in N_i} \mathbf{G}_{ij} l_{ij} = 0, \tag{A.1}$$

where l_{ij} is the length of the dual edge for nodes i, j, $\mathbf{G}_{ij}$ is given by the following expression,

$$\mathbf{G}_{ij} = \frac{1}{2}\left[\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}_i]^T (\psi_i + \psi_j) \right.$$
$$\left. + \frac{\partial}{\partial \mathbf{u}_l}\left(|\hat{\mathbf{n}} \cdot \tilde{\mathbf{A}}(\mathbf{u}_r, \mathbf{u}_l)|(\mathbf{u}_r - \mathbf{u}_l)\right)^T (\psi_r - \psi_l)\right]$$

and $\tilde{\mathbf{A}}(\mathbf{u}_r, \mathbf{u}_l)$ is the Jacobian matrix evaluated at the Roe-averaged state. The left and right flow, adjoint states are obtained from extrapolation:

$$\mathbf{u}_l = \mathbf{u}_i + \frac{\phi}{2}\nabla \mathbf{u} \cdot (\mathbf{x}_j - \mathbf{x}_i), \qquad \mathbf{u}_r = \mathbf{u}_j + \frac{\phi}{2}\nabla \mathbf{u} \cdot (\mathbf{x}_i - \mathbf{x}_j)$$
$$\psi_l = \psi_i + \frac{\phi}{2}\nabla \psi \cdot (\mathbf{x}_j - \mathbf{x}_i), \qquad \psi_r = \psi_j + \frac{\phi}{2}\nabla \psi \cdot (\mathbf{x}_i - \mathbf{x}_j),$$

where $\phi = 0, 1$ for first- and second-order discretizations respectively and $\mathbf{x}_i, \mathbf{x}_j$ are nodal locations. The scheme has the property that adjoint flux is not conservative: $\mathbf{G}_{ij} \neq -\mathbf{G}_{ji}$; however this is not in contradiction with the fact that the adjoint of a conservative system is not in conservative form.

In [10], Baysal and Ghayour propose a cell-based scheme for the discretization of continuous adjoint for the inviscid Euler's equations. The underlying flow scheme is similarly a second-order discretization using Roe's flux-difference splitting, taking as left and right states the Taylor expanded cell-centered flow states to the faces of each triangular cell. The equations for the adjoint are again given by (A.1), but now with l_{ij} denoting the length of the side lying between cell i and j. The expression for $\mathbf{G}_{ij}$

is given by

$$\mathbf{G}_{ij} = \frac{1}{2}\left[\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}_i]^T(\psi_l + \psi_r) + \frac{\partial}{\partial \mathbf{u}_l}\left(|\hat{\mathbf{n}} \cdot \tilde{\mathbf{A}}(\mathbf{u}_r, \mathbf{u}_l)|(\mathbf{u}_r - \mathbf{u}_l)\right)^T(\psi_r - \psi_l)\right]$$

where the extrapolated adjoint average $\psi_l + \psi_r$ is used.

A.2 Conservative formulations

In [16], Chen and Collis propose a discontinuous Galerkin method for optimal control of flows governed by the viscous Burgers and compressible Navier-Stokes equations. The inviscid discretization is derived via rewriting the adjoint PDE into a divergence form plus a source term. That is,

$$\mathcal{F}'[\mathbf{u}]^T \cdot \nabla\psi = \nabla \cdot \left(\mathcal{F}'[\mathbf{u}]^T\psi\right) - \left(\nabla \cdot \mathcal{F}'[\mathbf{u}]^T\right)\psi. \tag{A.2}$$

Then the DG formulation for the inviscid adjoint term is chosen to be

$$\sum_{\kappa \in T_h} \int_\Omega \left\{\mathbf{v}_h^T(\mathbf{u}_h)_t + \nabla\mathbf{v}_h^T \cdot \left(\mathcal{F}'[\mathbf{u}]^T\psi_h\right) - \mathbf{v}_h^T\left(\nabla \cdot \mathcal{F}'[\mathbf{u}]^T\right)\psi_h\right\}dx$$

$$-\int_{\Gamma_i}(\mathbf{v}_h^+ - \mathbf{v}_h^-)^T\mathcal{H}_i^{du}(\psi_h^+, \psi_h^-, \hat{\mathbf{n}})ds$$

$$-\int_{\partial\Omega}(\mathbf{v}_h^+)^T\mathcal{H}_i^{b,du}(\psi_h^+, \psi_h^b, \hat{\mathbf{n}})ds = 0, \tag{A.3}$$

where the Lax-Friedrichs method based on the maximum eigenvalue λ_{max} of the Euler Jacobian is used for interelement adjoint flux,

$$\mathcal{H}_i^{du}(\psi_h^+, \psi_h^-, \hat{\mathbf{n}}) = \frac{1}{2}\left(\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}^+]^T\psi^+ + \hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}^-]^T\psi^- - \lambda_{max}(\psi^- - \psi^+)\right), \tag{A.4}$$

and similarly for $\mathcal{H}_i^{b,du}$ on the domain boundary.

If both the underlying flow and adjoint PDE solutions are smooth, the above adjoint formulation is consistent. However in situations where flow discontinuities are present, the formulation is in general inconsistent as may be shown by examining

the case of analytical adjoint solution for quasi-one-dimensional Euler equations. It has been shown that [32, 85] across the shock location, the analytical adjoint solution is continuous. Suppose the shock is located at a node of the mesh, $\mathbf{x}_s$. Since $\psi^+(\mathbf{x}_s) = \psi^-(\mathbf{x}_s) \equiv \psi^s$, $\mathcal{H}_i^{du} = \frac{1}{2}\left(\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}^+]^T + \hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}^-]^T\right)\psi^s$. Given arbitrary $\mathbf{v}_h$, denote its values at the shock as $\mathbf{v}_h(\mathbf{x}_s)^+, \mathbf{v}_h(\mathbf{x}_s)^-$. Then the consistency error consists of the jump of flux Jacobian across the shock:

$$
\begin{aligned}
&\psi^s \mathbf{v}_h(\mathbf{x}_s)^+ \left(\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}^+] - \left\{\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}]^T\right\}\right) \\
-\ &\psi^s \mathbf{v}_h(\mathbf{x}_s)^- \left(\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}^-] - \left\{\hat{\mathbf{n}} \cdot \mathcal{F}'[\mathbf{u}]^T\right\}\right),
\end{aligned}
\tag{A.5}
$$

where the notation $\{\cdot\}$ denotes the average across the shock. Hence, the consistency error does not vanish if $\psi^s \neq 0$.

164

Appendix B

Derivation of local error indicator expression

In this appendix, derivations of local error indicator expressions for BR2 discretization of second-order systems [9] are given. This is first done for primal residual contributions in Section B.1 followed by the corresponding dual residual contributions in Section B.2.

B.1 Primal local error indicator

Using an integration by parts formula on discontinuous functions as given in [6], obtain:

$$
\begin{aligned}
&\sum_{\kappa \in T_h} \int_\kappa \nabla \mathbf{v}_h^T \cdot (\mathcal{A}_v(\mathbf{u}_H)\nabla \mathbf{u}_H) dx \\
&= \sum_{\kappa \in T_h} - \int_\kappa \mathbf{v}_h^T \nabla \cdot (\mathcal{A}_v(\mathbf{u}_H)\nabla \mathbf{u}_H) dx \\
&\quad + \int_{\Gamma_i} [\![\mathbf{v}_h]\!]^T \cdot \{\mathcal{A}_v(\mathbf{u}_H)\nabla \mathbf{u}_H\} ds + \int_{\partial\Omega} \mathbf{v}_h^{+T} \hat{\mathbf{n}} \cdot \mathcal{A}_v(\mathbf{u}_H^+)\nabla \mathbf{u}_H^+ ds \\
&\quad + \int_{\Gamma_i} \{\mathbf{v}_h\}^T [\![\mathcal{A}_v(\mathbf{u}_H)\nabla \mathbf{u}_H]\!] ds.
\end{aligned}
\tag{B.1}
$$

Using the above, the BR2 semilinear form (3.50) may be written as:

$$
\begin{aligned}
R_h(\mathbf{u}_H, \mathbf{v}_h) &= \sum_{\kappa \in T_h} \int_\kappa \mathbf{v}_h^T \nabla \cdot (\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H) \nabla \mathbf{u}_H) d\mathbf{x} \\
&+ \int_{\Gamma_i} [\![\mathbf{u}_H]\!]^T \cdot \{\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H)^T \nabla \mathbf{v}_h\} ds \\
&- \int_{\Gamma_i} [\![\mathbf{v}_h]\!]^T \cdot \eta_f \{\boldsymbol{\delta}_f\} ds \\
&- \int_{\Gamma_i} \{\mathbf{v}_h\}^T [\![\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H) \nabla \mathbf{u}_H]\!] ds \\
&+ \int_{\partial\Omega} \mathbf{v}_h^{+T} [\mathbf{F}_h^b(\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_H^b) \nabla \mathbf{u}_H - \eta_f \hat{\mathbf{n}} \cdot \boldsymbol{\delta}_f^b) - \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_H^+) \nabla \mathbf{u}_H^+] ds \\
&- \int_{\partial\Omega} (\mathbf{u}_H^b - \mathbf{u}_H^+)^T [\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H^b)^T \nabla \mathbf{v}_h^+] ds.
\end{aligned}
\tag{B.2}
$$

From the definition of the jump operator $[\![\cdot]\!]$, the third line of above involving $\boldsymbol{\delta}_f$ may be rewritten as a sum over element boundaries:

$$
\int_{\Gamma_i} [\![\mathbf{v}_h]\!]^T \cdot \eta_f \{\boldsymbol{\delta}_f\} ds = \sum_{\kappa \in T_h} \int_{\partial\kappa \backslash \partial\Omega} \mathbf{v}_h^{+T} \eta_f \hat{\mathbf{n}} \cdot \{\boldsymbol{\delta}_f\} ds.
\tag{B.3}
$$

Finally, (4.24) is obtained by using (B.3) in (B.2) to write the resulting expression as a sum over all elements and making the substitution $\mathbf{v}_h \rightarrow \tilde{\boldsymbol{\psi}}_h - \boldsymbol{\psi}_H$.

B.2 Dual local error indicator

Since the dual local error indicator involves many terms, the derivation is broken into two steps. Firstly, in Section B.2.1 the assumption is made that the coefficient matrix has no dependence on the state, $\boldsymbol{\mathcal{A}}_v = \boldsymbol{\mathcal{A}}_v(\mathbf{x})$. Subsequently, in Section B.2.2 the contribution from state dependence of the coefficient matrix $\boldsymbol{\mathcal{A}}_v$ is given.

B.2.1 State-independent coefficient matrix

Using integration by parts, obtain:

$$
\sum_{\kappa \in T_h} \int_\kappa \nabla \boldsymbol{\psi}_H^T \cdot (\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H) \nabla \mathbf{v}_h) d\mathbf{x}
$$

$$
= \sum_{\kappa \in T_h} - \int_\kappa \mathbf{v}_h^T \nabla \cdot (\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H)^T \nabla \psi_H) d\mathbf{x}
$$
$$
+ \int_{\Gamma_i} [\![\mathbf{v}_h]\!]^T \cdot \{\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H)^T \nabla \psi_h\} ds + \int_{\partial \Omega} \mathbf{v}_h^{+T} \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_H^+)^T \nabla \psi_H^+ ds
$$
$$
+ \int_{\Gamma_i} \{\mathbf{v}_h\}^T [\![\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H)^T \nabla \psi_H]\!] ds. \tag{B.4}
$$

Using the above and the generic form of boundary functional given by (3.55), the expression for adjoint residual may be written as:

$$
R'_h[\mathbf{u}_H](\mathbf{v}_h, \psi_H) - \mathcal{J}'_h[\mathbf{u}_H](\mathbf{v}_h)
$$
$$
= \sum_{\kappa \in T_h} \int_\kappa \mathbf{v}_h^T \nabla \cdot (\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H)^T \nabla \psi_H) d\mathbf{x}
$$
$$
- \int_{\Gamma_i} \{\mathbf{v}_h\}^T [\![\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H)^T \nabla \psi_H]\!] ds + \int_{\Gamma_i} [\![\psi_H]\!]^T \cdot \{\boldsymbol{\mathcal{A}}_v(\mathbf{u}_H) \nabla \mathbf{v}_h\} ds
$$
$$
- \int_{\Gamma_i} [\![\mathbf{v}_h]\!] \cdot \eta_f \{\boldsymbol{\delta}_f^{du}\} ds
$$
$$
- \int_{\partial \Omega} \left(\left[\frac{\partial \mathbf{u}_H^b}{\partial \mathbf{u}_H^+} \right] \mathbf{v}_h^+ \right)^T \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v(\mathbf{u}_H^+)^T \nabla \psi_H^+ ds + \int_{\partial \Omega} \left(\left[\frac{\partial \mathbf{u}_H^b}{\partial \mathbf{u}_H^+} \right] \mathbf{v}_h^+ - \mathbf{v}_h^+ \right)^T \eta_f \hat{\mathbf{n}} \cdot \boldsymbol{\delta}_f^{b,du} ds
$$
$$
+ \int_{\partial \Omega} \psi_H^{+T} \left(\left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{v}_h^+ \right) ds
$$
$$
- \int_{\Gamma_{\text{output}}} J'_I[\mathbf{u}_H^b] \left(\left[\frac{\partial \mathbf{u}_H^b}{\partial \mathbf{u}_H^+} \right] \mathbf{v}_h^+ \right) ds
$$
$$
- \int_{\Gamma_{\text{output}}} J'_{II}[\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}_H] \left(\left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{v}_h^+ \right) ds, \tag{B.5}
$$

where, to rewrite the boundary integrals of the inner product of the adjoint state with the linearization of the primal auxiliary variables $\boldsymbol{\delta}_f$, the dual auxiliary variables $\boldsymbol{\delta}_f^{du}$ are introduced which are defined on interior faces by,

$$
\int_\Omega \boldsymbol{\tau}_h \cdot \boldsymbol{\delta}_f^{du} d\mathbf{x} = \int_{\sigma_f} [\![\psi_H]\!] \cdot \{\boldsymbol{\mathcal{A}}_v \boldsymbol{\tau}_h\} ds, \quad \forall \boldsymbol{\tau}_h \in [\mathcal{V}_h^p]^2, \tag{B.6}
$$

and for boundary faces $\boldsymbol{\delta}_f^{du,b}$ are defined by: $\forall \boldsymbol{\tau}_h \in [\mathcal{V}_h^p]^2$,

$$
\int_\Omega \boldsymbol{\tau}_h \cdot \boldsymbol{\delta}_f^{b,du} d\mathbf{x} = \int_{\sigma_f} (\psi_H^{+T} - J'_{II}[\hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \nabla \mathbf{u}_H]) \left(\left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \hat{\mathbf{n}} \cdot \boldsymbol{\mathcal{A}}_v \boldsymbol{\tau}_h^+ \right) ds. \tag{B.7}
$$

By breaking up the expression (B.5) into elemental volume and surface integral contributions and making the substitution $\mathbf{v}_h \rightarrow \tilde{\mathbf{u}}_h - \mathbf{u}_H$, the dual form of the local error indicator (4.25) is obtained.

B.2.2 State-dependent coefficient matrix

From the expression for the semilinear form (3.50), it is straightforward to compute:

$$
\begin{aligned}
R'_h[\mathbf{u}_H](\mathbf{v}_h, \psi_H) &- \mathcal{J}'_h[\mathbf{u}_H](\mathbf{v}_h) = (B.5) \\
&+ \sum_{\kappa \in T_h} \int_\kappa \nabla \psi_H^T \cdot (\mathcal{A}'_v[\mathbf{u}_H](\mathbf{v}_h) \nabla \mathbf{u}_H) \, d\mathbf{x} \\
&+ \int_{\Gamma_i} [\![\mathbf{u}_H]\!] \cdot \{\mathcal{A}'_v[\mathbf{u}_H](\mathbf{v}_h)^T \nabla \mathbf{v}_h\} ds + \int_{\Gamma_i} [\![\psi_H]\!] \cdot \{\mathcal{A}'_v[\mathbf{u}_H](\mathbf{v}_h) \nabla \mathbf{u}_H\} ds \\
&- \int_{\Gamma_i} \eta_f [\![\psi_H]\!] \cdot \left\{ \frac{\partial \delta_f}{\partial \mathbf{u}_H} \mathbf{v}_h \right\} ds \\
&- \int_{\partial\Omega} (\mathbf{u}_H^b - \mathbf{u}_H^+)^T (\hat{\mathbf{n}} \cdot \mathcal{A}'_v[\mathbf{u}_H^b] \left(\left[\frac{\partial \mathbf{u}_H^b}{\partial \mathbf{u}_H^+} \right] \mathbf{v}_h^+ \right)^T \nabla \psi_H^+) ds \\
&+ \int_{\partial\Omega} \psi_H^{+T} \left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \left(\hat{\mathbf{n}} \cdot \mathcal{A}'_v[\mathbf{u}_H^b] \left(\left[\frac{\partial \mathbf{u}_H^b}{\partial \mathbf{u}_H^+} \right] \mathbf{v}_h^+ \right) \nabla \mathbf{u}_H^+ - \eta_f \hat{\mathbf{n}} \cdot \frac{\partial \delta_f^b}{\partial \mathbf{u}_H^+} \left[\frac{\partial \mathbf{u}_H^b}{\partial \mathbf{u}_H^+} \right] \mathbf{v}_h^+ \right) ds \\
&- \int_{\Gamma_{\text{output}}} J'_{II}[\hat{\mathbf{n}} \cdot \mathcal{A}_v \nabla \mathbf{u}_H] \left[\frac{\partial \mathbf{F}_h^b}{\partial \mathbf{F}_h^+} \right] \left(\hat{\mathbf{n}} \cdot \mathcal{A}'_v[\mathbf{u}_H^b] \left(\left[\frac{\partial \mathbf{u}_H^b}{\partial \mathbf{u}_H^+} \right] \mathbf{v}_h^+ \right) \nabla \mathbf{u}_H^+ \right. \\
&\left. - \eta_f \hat{\mathbf{n}} \cdot \frac{\partial \delta_f^b}{\partial \mathbf{u}_H^+} \left[\frac{\partial \mathbf{u}_H^b}{\partial \mathbf{u}_H^+} \right] \mathbf{v}_h^+ \right) ds,
\end{aligned}
$$

where $\partial \delta_f / \partial \mathbf{u}_H$ is understood to denote the derivative of δ_f with respect to $\mathbf{u}_H$ through the state dependence of coefficient matrix, $\mathcal{A}_v = \mathcal{A}_v(\mathbf{u}_H)$. The contribution of state-dependent coefficient matrix to the local error indicator is again obtained by a decomposition into elemental contributions and replacing $\mathbf{v}_h \rightarrow \tilde{\mathbf{u}}_h - \mathbf{u}_H$.

Bibliography

[1] M. Ainsworth and T. J. Oden. *A posteriori error estimation in finite element analysis.* Pure and Applied Mathematics. John Wiley and Sons, New York, 2000.

[2] N. M. Alexandrov, R. M. Lewis, C. R. Gumbert, L. L. Green, and P. A. Newman. Approximation and model management in aerodynamic optimization with variable-fidelity models. *Journal of Aircraft*, 38(5):1093–1101, 2001.

[3] N. M. Alexandrov, E. J. Nielsen, R. M. Lewis, and W. K. Anderson. First-order model management with variable-fidelity physics applied to multi-element airfoil optimization. AIAA Paper 2000-4886, 2000.

[4] W. K. Anderson and D. L. Bonhaus. Aerodynamic design on unstructured grids for turbulent flows. Technical Report 112867, NASA, 1997.

[5] W. K. Anderson and V. Venkatakrishnan. Aerodynamic design optimization on unstructured grids with a continuous adjoint formulation. *Computers and Fluids*, 28:443–480, 1999.

[6] D. N. Arnold, F. Brezzi, B. Cockburn, and L. D. Marini. Unified analysis of discontinuous Galerkin methods for elliptic problems. *SIAM Journal on Numerical Analysis*, 39(5):1749–1779, 2002.

[7] W. Bangerth and R. Rannacher. *Adaptive finite element methods for differential equations.* Lectures in Mathematics ETH Zürich. Birkhäuser, Basel, 2003.

[8] T. J. Barth and M. G. Larson. A-posteriori error estimation for higher order Godunov finite volume methods on unstructured meshes. In R. Herban and

D. Kröner, editors, *Finite Volumes for Complex Applications III*, pages 41–36. Hermes Penton, London, 2002.

[9] F. Bassi and S. Rebay. GMRES discontinuous Galerkin solution of the compressible Navier-Stokes equations. In B. Cockburn, G. Karniadakis, and C.-W. Shu, editors, *Discontinuous Galerkin Methods: Theory, Computation and Applications*, pages 197–208. Springer, Berlin, 2000.

[10] O. Baysal and K. Ghayour. Continuous adjoint sensitivities for optimization with general cost functionals on unstructured meshes. *AIAA Journal*, 39(1):48–55, 2001.

[11] R. Becker. Mesh adaptation for stationary flow control. *Journal of Mathematics in Fluid Mechanics*, 3:317–341, 2001.

[12] R. Becker. Mesh adaptation for Dirichlet flow control via Nitsche's method. *Communications in Numerical Methods in Engineering*, 18:669–680, 2002.

[13] R. Becker and R. Rannacher. An optimal control approach to *a posteriori* error estimation in finite element methods. *Acta Numerica*, 10:1–102, 2001.

[14] G. Biros and O. Ghattas. Inexactness issues in the Lagrange-Newton-Krylov-Schur method in PDE-constrained optimization. In L. T. Biegler, O. Ghattas, M. Heinkenschloss, and B. van Bloemen Waanders, editors, *Large-scale PDE-constrained optimization*, pages 148–164. Springer, Berlin, 2003.

[15] A. N. Brooks and T. J. R. Hughes. Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations. *Computer Methods in Applied Mechanics and Engineering*, 32:199–259, 1982.

[16] G. Chen and S. S. Collis. Towards optimal control of aeroacoustic flows using the discontinuous Galerkin discretizations. AIAA Paper 2004-0364, 2004.

[17] S. S. Collis and M. Heinkenschloss. Analysis of the streamline upwind/Petrov Galerkin methods applied to the solution of optimal control problems. Technical Report 02/01, Rice University Department of Computational and Applied Mathematics, 2002.

[18] A. Dadone and B. Grossman. Progressive optimization of inverse fluid dynamic design problems. *Computers and Fluids*, 29(1):1–32, 2000.

[19] A. Dadone and B. Grossman. Fast convergence of viscous airfoil design problems. *AIAA Journal*, 40(10):1997–2005, 2002.

[20] A. Dadone and B. Grossman. Fast convergence of inviscid fluid dynamic design problems. *Computers and Fluids*, 32(4):607–627, 2003.

[21] J. Elliot and J. Peraire. Aerodynamic design using unstructured meshes. AIAA Paper 96-1941, 1996.

[22] J. Elliot and J. Peraire. Practical 3D aerodynamic design and optimization using unstructured grids. AIAA Paper 96-4170, 1996.

[23] J. Elliot and J. Peraire. Aerodynamic optimization on unstructured meshes with viscous effects. AIAA Paper 97-1849, 1997.

[24] A. Ern and J.-L. Guermond. *Theory and practice of finite elements*. Applied Mathematical Sciences. Springer-Verlag, New York, 2004.

[25] K. Fidkowski. A high-order discontinuous Galerkin multigrid solver for aerodynamic applications. Master's thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, June 2004.

[26] L. P. Franca, S. L. Frey, and T. J. R. Hughes. Stabilized finite element methods: I. Application to the advective-diffusive model. *Computer Methods in Applied Mechanics and Engineering*, 95(2):253–276, 1992.

[27] O. Ghattas and J.-H. Bark. Optimal control of two- and three-dimensional incompressible Navier-Stokes flows. *Journal of Computational Physics*, 136:231–244, 1997.

[28] M. Giles, M. G. Larson, J. M. Levenstam, and E. Süli. Adaptive error control for finite element approximations of the lift and drag coefficients in viscous flow. Technical Report NA 97/06, Oxford University Computing Lab Numerical Analysis Group, 1997.

[29] M. B. Giles. On the iterative solution of adjoint equations. In *Automatic Differentiation: From Simulation to Optimization*, pages 145–152. Springer-Verlag, 2001.

[30] M. B. Giles, M. C. Duta, J-D Muller, and N. A. Pierce. Algorithm developments for discrete adjoint methods. *AIAA Journal*, 41(2):1–32, 2003.

[31] M. B. Giles and N. A. Pierce. Adjoint equations in CFD: duality, boundary conditions and solution behaviour. AIAA Paper 97-1850, January 1997.

[32] 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.

[33] M. B. Giles and E. Süli. Adjoint methods for PDEs: a posteriori error analysis and postprocessing by duality. *Acta Numerica*, 11:145–236, 2002.

[34] P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization*. Academic Press, New York, 1986.

[35] A. Griewank. *Evaluating derivatives: principles and techniques of algorithmic differentiation*. Society for Industrial and Applied Mathematics, Philadelphia, 2000.

[36] A. Griewank and C. Faure. Reduced functions, gradients and Hessians from fixed point iterations for state equations. *Numerical Algorithms*, 30(2):113–139, 2002.

172

[37] A. Griewank and C. Faure. Piggyback differentiation and optimization. In L. T. Biegler, O. Ghattas, M. Heinkenschloss, and B. van Bloemen Waanders, editors, *Large-scale PDE-constrained optimization*, pages 148–164. Springer, Berlin, 2003.

[38] 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 NA 04/18, Oxford University Computing Lab Numerical Analysis Group, 2004.

[39] K. Harriman, P. Houston, B. Senior, and E. Süli. hp-version discontinuous Galerkin methods with interior penalty for partial differential equations with nonnegative characteristic form. In S. Y. Cheng, C.-W. Shu, and T. Tang, editors, *Recent advances in scientific computing and partial differential equations*, pages 89–109. American Mathematical Society, 2003.

[40] R. Hartmann and P. Houston. Adaptive discontinuous Galerkin finite element methods for the compressible Euler equations. *Journal of Computational Physics*, 183(2):508–532, 2002.

[41] R. Hartmann and P. Houston. Adaptive discontinuous Galerkin finite element methods with interior penalty for the compressible Navier-Stokes equations. In M. Feistauer, V. Dolejsi, P. Knobloch, and K. Najzar, editors, *Numerical Mathematics and Advanced Applications: Proceedings of ENUMATH 2003 the 5th European Conference on Numerical Mathematics and Advanced Applications*, pages 410–419. Springer, Berlin, 2004.

[42] J. Haslinger and R. A. E. Mäkinen. *Introduction to shape optimization: theory, approximation and computation*. Society for Industrial and Applied Mathematics, Philadelphia, 2003.

[43] R. M. Hicks and P. A. Henne. Wing design by numerical optimization. *Journal of Aircraft*, 15:407–412, 1978.

[44] P. Houston, R. Rannacher, and E. Süli. A posteriori error analysis for stabilized finite element approximations of transport equations. *Computer Methods in Applied Mechanics and Engineering*, 190:1483–1508, 2000.

[45] J. C. Newman III and A. C. Taylor III. Three-dimensional aerodynamic shape sensitivity analysis and design optimization using the Euler equations on unstructured grids. AIAA Paper 96-2464, 1996.

[46] J. C. Newman III, A. C. Taylor III, R. W. Barnwell, P. A. Newman, and G. J.-W. Hou. Overview of sensitivity analysis and shape optimization for complex aerodynamic configurations. *Journal of Aircraft*, 36(1):87–96, 1999.

[47] J. C. Newman III, A. C. Taylor III, and G. W. Burgreen. An unstructured grid approach to sensitivity analysis and shape optimization using the Euler equations. AIAA Paper 95-1646, 1995.

[48] A. Iollo, G. Kuruvila, and S. Ta'asan. Pseudo-time method for optimal shape design using the Euler equations. ICASE Report 1995-59, 1995.

[49] A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3:233–260, 1988.

[50] A. Jameson and S. Kim. Reduction of the adjoint gradient formula in the continuous limit. AIAA Paper 2003-0040, 2003.

[51] A. Jameson, L. Martinelli, and N. A. Pierce. Optimum aerodynamics design using the Navier-Stokes equations. *Theoretical Computational Fluid Dynamics*, 10:213–237, 1998.

[52] S. Kim, J. J. Alonson, and A. Jameson. Efficient aerodynamic design method using a tightly coupled algorithm. *Journal of Aircraft*, 41(5):1082–1097, 2004.

[53] C. Kirjner and E. Polak. On the use of consistent approximations for the optimal design of beams. *SIAM Journal on Control and Optimization*, 34(6):1891–1913, 1996.

[54] G. Kuruvila, S. Ta'san, and M. D. Salas. Airfoil design and optimization by the one-shot method. AIAA Paper 1995-0478, 1995.

[55] R. Li, W. Liu, H. Ma, and T. Tang. Adaptive finite element approximation for distributed elliptic optimal control problems. *SIAM Journal on Control and Optimization*, 41(5):1321–1349, 2002.

[56] W. Liu and N. Yan. A posteriori error estimates for convex boundary control problems. *SIAM Journal on Numerical Analysis*, 39(1):73–99, 2001.

[57] J. Lu and D. L. Darmofal. A quasi-minimal residual method for simultaneous primal-dual solutions and superconvergent functional estimates. *SIAM Journal on Scientific Computing*, 24(5):1693–1709, 2003.

[58] J. Lu and D. L. Darmofal. Adaptive precision methodology for flow optimization via discretization and iteration error control. AIAA Paper 2004-1096, 2004.

[59] H. Melbø and T. Kvamsdal. Goal oriented error estimators for Stokes equations based on variationally consistent postprocessing. *Computer Methods in Applied Mechanics and Engineering*, 192:613–633, 2003.

[60] B. Mohammadi. Optimal shape design, reverse mode of automatic differentiation and turbulence. AIAA Paper 97-0099, 1997.

[61] B. Mohammadi and O. Pironneau. *Applied shape optimization for fluids*. Oxford University Press, Oxford, 2001.

[62] A. L. Moigne and N. Qin. Variable-fidelity aerodynamic optimization for turbulent flows using a discrete adjoint formulation. *AIAA Journal*, 42(7):1281–1292, 2004.

[63] S. Nadarajah and A. Jameson. Studies of the continuous and discrete adjoint approaches to viscous automatic aerodynamic shape optimization. AIAA Paper 2001-2530, 2001.

[64] M. Nemec, D. W. Zingg, and T. H. Pulliam. Multipoint and multi-objective aerodynamic shape optimization. *AIAA Journal*, 42(6):1057–1065, 2004.

[65] E. Nielsen, J. Lu, M. Park, and D. L. Darmofal. An implicit, exact dual adjoint solution method for turbulent flows on unstructured grids. *Computers and Fluids*, 33:1131–1155, 2004.

[66] E. J. Nielsen and W. K. Anderson. Aerodynamic design optimization on unstructured meshes using the Navier-Stokes equations. *AIAA Journal*, 37(11):1411–1419, 1999.

[67] T. A. Oliver. Multigrid solution for high-order discontinuous Galerkin discretizations of the compressible Navier-Stokes equations. Master's thesis, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, August 2004.

[68] A. T. Patera and E. M. Rønquist. A general output bound result : application to discretization and iteration error estimation and control. *Mathematical Models and Methods in Applied Sciences*, 1:658–712, 2001.

[69] O. Pironneau. *Optimal shape design for elliptic systems*. Springer-Verlag, New York, 1984.

[70] O. Pironneau and E. Polak. Consistent approximations and approximate functions and gradients in optimal control. *SIAM Journal on Control and Optimization*, 41(2):487–510, 2002.

[71] E. Polak. *Optimization : algorithms and consistent approximations*. Springer-Verlag, New York, 1997.

[72] A. Quarteroni and A. Valli. *Numerical approximation of partial differential equations*. Springer-Verlag, New York, 1994.

[73] R. Radespiel and R.C. Swanson. An investigation of cell centered and cell vertex multigrid schemes for the Navier-Stokes equations. AIAA Paper Number 89-0453, 1989.

176

[74] R. Rannacher. Adaptive Galerkin finite element methods for partial differential equations. *Journal of Computational and Applied Mathematics*, 128:205–233, 2001.

[75] J. J. Reuther, A. Jameson, J. J. Alonso, M. J. Rimlingers, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part I. *Journal of Aircraft*, 36(1):51–60, 1999.

[76] J. J. Reuther, A. Jameson, J. J. Alonso, M. J. Rimlingers, and D. Saunders. Constrained multipoint aerodynamic shape optimization using an adjoint formulation and parallel computers, part II. *Journal of Aircraft*, 36(1):61–74, 1999.

[77] A. M. Sauer-Budge. *Computing upper and lower bounds on linear functional outputs from linear coercive partial differential equations*. PhD dissertation, Massachusetts Institute of Technology, Department of Aeronautics and Astronautics, 2003.

[78] A. Schwartz and E. Polak. Consistent approximations for the optimal control problems based on Runge-Kutta integration. *SIAM Journal on Control and Optimization*, 34(4):1235–1269, 1996.

[79] A. Shenoy, M. Heinkenschloss, and E. M. Cliff. Airfoil design by an all-at-once method. *International Journal of Computational Fluid Dynamics*, 11:3–25, 1998.

[80] G. Strang and G. J. Fix. *An analysis of the finite element method*. Series in Automatic Computation. Prentice-Hall, New Jersey, 1973.

[81] C.-H. Sung and J. K. Kwon. Aerodynamic design optimization using the Navier-Stokes and adjoint equations. AIAA Paper 2001-0266, 2001.

[82] C.-H. Sung and J. K. Kwon. Efficient aerodynamic design method using a tightly coupled algorithm. *AIAA Journal*, 40(9):1839–1845, 2002.

[83] D. A. Venditti and D. L. Darmofal. Grid adaptation for functional outputs : application to two-dimensional inviscid flows. *Journal of Computational Physics*, 176(1):40–69, 2002.

[84] 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.

[85] L. Xie. *Gradient-based optimum aerodynamic design using adjoint methods*. PhD dissertation, Virginia Polytechnic Institute and State University, Department of Aerospace Engineering, 2002.