ANALYSIS AND DESIGN OF OPTIMAL DISCONTINUOUS FINITE ELEMENT SCHEMES

A DISSERTATION
SUBMITTED TO THE DEPARTMENT OF AERONAUTICS AND ASTRONAUTICS
AND THE COMMITTEE ON GRADUATE STUDIES
OF STANFORD UNIVERSITY
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

Kartikey Asthana
June 2016

Acknowledgements

The only appropriate beginning to this dissertation is through the expression of my gratitude to Prof. Antony Jameson. Without his supervision and encouragement, this thesis would not have been possible. In fact, many of the ideas pursued herein owe themselves to the emphasis that Prof. Jameson places on independent thought. Even beyond academics, the experiences of his life, ranging from sports and sports cars to economics and airplane design, have been truly inspiring. It has been my pleasure and privilege to be his 50th doctoral student.

I am very grateful to my thesis readers, Prof. Sanjiva Lele and Prof. Peter Pinsky, for their time and patience. The scope of this work has benefited immensely from Prof. Lele's observations and continual guidance. The content, itself, has its foundations in the courses by Prof. Pinsky who provided me with the enriching opportunity to learn the art of teaching as a course assistant. Indeed, their approval of this thesis is a great honor for me.

I would like to sincerely thank Prof. Gerritsen and Prof. Papanicolaou for serving on my oral defense committee. There is no doubt that my appreciation of Applied Mathematics, in all its beauty and breadth, stems from the classrooms of such excellent teachers. It is fair to say that CME302 and MATH238 have been two of my favorite classes at Stanford.

My stay at the university was supported by the Stanford Graduate Fellowship. In particular, I would like to thank the family of Thomas V Jones for their generosity. The importance of this fellowship program for international students cannot be emphasized enough.

My time at the Aerospace Computing Laboratory was made very pleasant by the company of jovial colleagues. I begin by thanking Jerry Watkins for co-authoring not only research papers but also rock songs, and for teaching me how to drive. Thanks to Josh Romero for explaining traditions in the American culture; to Abhishek Sheshadri for his perspective on pure mathematics and the meaning of life; to Manuel for his positive outlook and practical advice; to Andy for introducing me to the Gewürztraminer and to Jacob, David, Jonathan, Zach and Freddy for begin awesome. Special thanks to Prof. David Williams, Jonathan Bull, Prof. Philipp Birken and Prof. Peter Vincent for their detailed comments and valuable feedback to my research.

Beyond the lab, much of my graduate life has been spent in the splendid company of Abhimanyu Banerjee, Akash Sarma, Akshay Subramaniam and Raunak Borker. I am thankful to them for all the food, drinks, movies, video games and late-night walks that we spent contemplating on matters of zero consequence.

Coming from a culturally entrenched Indian society, I recognize that it is still not common for one to pursue a doctoral degree in science from across the seas. I am eternally grateful to my parents, Kamal and Ranjana, for providing me with their unconditional love and support to break out from the norm. Their faith in my abilities is very empowering. I am blessed to have my elder brother, Siddharth, who has used the wisdom of his hard-earned experiences to serve as the pathfinder for my life. Finally, I am thankful to my girlfriend, Apoorva, for always lending a patient ear to my worries and my troubles in the several years that have gone by.

This acknowledgement cannot be complete without a very special mention of my undergraduate adviser, Prof. Tapan Sengupta who introduced me to computational fluid dynamics. Indeed, my fascination with von Neumann analysis is deeply rooted in my training at the High Performance Computing Laboratory in IIT Kanpur. Prof. Sengupta's lessons on integrity, diligence and persistence continue to motivate me through life. This thesis is dedicated to him.

Preface

The latter half of the twentieth century has witnessed the advent of computational methods for simulating fluid flow. Fueled by the rapid development of computer hardware, finite-volume approximations of conservation laws have relieved the applied physicist from the need for closed-form solutions. Modern researchers now find themselves in the easily accessible field of computational fluid dynamics (CFD), facilitated by the availability of myriad open-source as well as commercial software packages. Owing to classical results on gas dynamics for first-order [53] and second-order [136, 75] schemes, the finite-volume approach is now a mature technology. For the aerospace industry, modern computer codes provide a significantly cheaper, reproducible and reasonably robust alternative to physical experiments conducted in the wind-tunnel. Engineers of the last millennium have employed CFD to aid in the design [74] of most commercial airliners in operation today.

However, the impact of finite-volume schemes has been largely restricted to steady flows such as that during an aircraft in cruise. Other important flight configurations such as pitching, plunging, gusts, buffet and flutter have been tackled using conventional empirical techniques, flight experiments and proprietary wisdom. The presence of excessive numerical dissipation and dispersion in such low-order schemes has also impeded fundamental research in flow physics. For instance, in computational aero-acoustics [129], pressure disturbances originating at solid surfaces need to be propagated across long distances at the correct wave and group velocities and without spurious dissipation. Similarly, in direct numerical solution (DNS) of turbulent flows, numerical dissipation can have an unacceptable adulteration of the energy spectrum [97]. A similar problem manifests in the case of vortex-dominated flows such as that over helicopter blades or flapping wings.

Consequently, the last two decades in computational mechanics have been dedicated to the development of high-order methods. The principal argument in their support stems from the observation that increasing the order of convergence offers an exponential reduction in error. This is in contrast to the polynomial reduction offered by an increase in resolution of the computational domain. Even in regard to design, high-order methods have the potential to provide the same level of accuracy with considerably reduced computational effort [62] as compared to traditional second-order methods.

Several approaches have been adopted to achieve high-order accuracy. The conventional practice has relied upon compact finite difference schemes [58, 88] that are particularly well suited for Cartesian meshes. These schemes have been responsible for remarkable success in fundamental explorations on rectilinear geometries involving DNS of compressible shear flows [87], hypersonic boundary layers [159], receptivity studies regarding transition [115] etc. Although extensions to curvilinear and deforming meshes have been proposed [145], the applicability of such schemes to industrial problems remains limited due to the difficulty involved in generating structured meshes for complex geometries. Consequently, the development of discontinuous finite-element methods has been a focal point for recent efforts. These methods, which can operate on unstructured meshes, combine the high-order accuracy of continuous finite-element schemes [68] with the localized stencils of finite-volume schemes. Specifically, the numerical solution in each element can be arbitrarily refined on account of its representation in an appropriate, often hierarchal, basis. In contrast to continuous finite-element schemes, the solution is allowed to be discontinuous at the element interfaces which yields an element-local mass matrix that can be easily inverted. Finally, and perhaps most importantly, these methods are amenable to massively parallel algorithms that can be distributed to the level of individual elements, paving the way for accelerators such as multi-core processors and Graphical Processing Units [29, 92, 156].

The most popular example within this class of methods is the Discontinuous Galerkin (DG) [39, 89, 106] formulation. This computationally tractable, high-order method has been used to tackle conservation laws arising in a very wide variety of applications ranging across fluid mechanics [12, 39], solid mechanics [108, 158], electromagnetics [63], semiconductors [33], polymer processing [14], stochastic control [34] etc. Similar to continuous finite-element schemes, the DG formulation approximates the weak form of the conservation law. However, conservation of the flux is enforced at the element interfaces not by continuity of state variables but by the introduction of Riemann solutions as in finite-volume schemes. The DG formulation comes in two main flavors. The traditional one incorporates full-order integrations of fluxes within each element, often leading to significant computational costs, especially for nonlinear problems. The more recent flavor comprises the collocation-based nodal DG method which projects the flux function onto a set of nodes in the interior of the element. This helps to reduce computational cost while providing a very intuitive, nodal representation of the solution [62].

Another popular instance of discontinuous finite element methods is the Spectral Difference (SD) formulation, proposed initially as a staggered-grid Chebyshev multidomain method [81], and generalized later to simplexes [91]. Similar to DG, the numerical solution and flux in SD are both approximated locally through Lagrangian basis functions, but on different nodes warranting the name 'staggered-grid'. Another importance difference is that the SD formulation directly approximates the strong form of the conservation law leading to a simpler prescription and computational efficiency [90].

In 2007, Huynh [69] proposed a Flux Reconstruction (FR) approach to first-order hyperbolic systems that provides a generalized differential framework for discontinuous finite-element schemes on tensor-product elements. This formulation is unifying in that it recovers several existing high-order methods including the collocation-based nodal DG method and a version of the SD formulation. It also provides a broad scope for the derivation of new linearly stable schemes such as the g_2 scheme by Huynh himself, the Energy Stable FR (ESFR) family by Vincent et al. [142], and the CMFR family by Lopez et al. [93] which offers the promise of extending FR to elastodynamics.

In 2009, Huynh [71] extended the FR formulation to second-order hyperbolic systems on tensor-product elements, enabling the computation of problems with diffusion. Soon after, the ESFR family was extended to advection-diffusion problems [28] as well. The correction procedure was further generalized to triangular [70, 32, 152] and tetrahedral elements [154, 155], and to general non-linear, second-order conservation laws such as the Euler [31] and Navier-Stokes equations [151]. The surge in popularity of FR also motivated other frameworks such as Correction Procedure via Reconstruction (CPR) [49] and Lifting Collocation Penalty (LCP) [147].

Despite the aforementioned advances in FR, which have been concerned primarily with implementation, progress along the lines of fundamental analysis has been rather limited. For instance, there are no published results guaranteeing that the FR formulation is indeed consistent [86]! Along the same line, convergence of the numerical solution, or even stability as per Lax's criterion, is not assured. While these might get dismissed by a casual practitioner as merely theoretical details, there are some serious voids from a practical point of view as well. There are no analytical estimates, published as yet, for the rate of convergence (order of accuracy) of FR for either time-dependent or steady-state problems. Researchers in the past [30, 150] have relied on canonical numerical experiments as a proxy for these rates. Most often, in these experiments, the numerical solution is marched for a fixed amount of time using the DG scheme with fully-upwinded fluxes. Consequently, the corresponding observations fail to capture the long-time super-convergence exhibited by FR schemes. Similarly, the effect of interface flux, location of solution points and the choice of correction function on the accuracy of the scheme is not well understood.

Another major challenge faced by the FR formulation is in regard to stability for non-linear fluxes. The nodal nature of the formulation naturally introduces aliasing errors [76] when a general flux is projected onto the polynomial space of the numerical solution. These errors can lead to instabilities especially if the weak form of the conservation law admits discontinuous solutions such as in transonic or supersonic flow. Jameson et al. have suggested the use of Gauss-Legendre solution points, or more generally quadrature points [150], to minimize the magnitude of aliasing errors. However, this is often insufficient by itself when the polynomial order of the scheme is high. Consequently, some of the most important problems in aerospace design are still out of reach of high-order discontinuous finite-element schemes.

In this dissertation, we undertake these three major challenges in a rigorous study supported by adequate numerical evidence. Correspondingly, the discussion is distributed in three parts as follows.

We begin with an eigensolution analysis for linear time-dependent problems to prove that the FR formulation is indeed consistent. Moreover, any FR scheme which is stable in the von-Neumann sense is provably convergent. Interestingly, the rate of convergence for such problems is a function of time, starting from a short-time rate associated with polynomial interpolation and eventually asymptoting to a long-time rate associated with the spatial derivative operator. Both these rates can be inferred from a simple eigenvalue computation for any given FR scheme.

The eigensolution analysis also allows us to examine the effect of polynomial order, correction function, interface flux and solution points on the dispersion and dissipation properties of the formulation. Using this framework, a new set of linearly stable high-order FR schemes is proposed that minimizes wave propagation errors for the range of resolvable wavenumbers. The resulting optimal schemes, designated Optimal ESFR (OESFR) and Optimal FR (OFR), provide a higher resolving efficiency in comparison to the Discontinuous Galerkin (DG) scheme as well as standard high-order compact finite difference schemes.

The second part of the dissertation derives analytical estimates for the rate of convergence of FR for steady-state problems. This has the lead to the derivation of a special class of schemes, designated Super-convergent FR (SFR), which exhibits an enhanced rate of convergence. The DG scheme is recovered as one member of this family while the SD scheme is not. We also show that the rate of convergence for steady-state problems is identical to the short-time rate of convergence for time-dependent problems.

In the third and final part of the dissertation, we prove that stability of the FR solution can be ensured for non-linear fluxes through the addition of adequate artificial dissipation. Unfortunately, a direct application of dissipation would severely limit the rate of convergence of the solution. Towards this end, we pose viscosity as a spectral filtering operation implemented in the physical space via a strictly local convolution integral. Coupled with a discontinuity sensor, this approach provides a computationally efficient method that captures shock discontinuities while preserving accuracy in smooth regions of the solution, even for very, very high polynomial orders such as $P = 119$. The filtered solution provides reduced total variation, reduced maximum overshoot/undershoot, and allows sub-element shocks to be localized in the interior of an element.

Note that the results derived herein for the general FR formulation are indeed directly applicable to DG. This is valuable from the point of view of verification since the DG scheme has been thoroughly investigated in the last two decades. We also note the scope of this work has been restricted to tensor-product elements for analytical simplicity. However, this is hardly a limitation since any triangular element can either be completely replaced by quadrilaterals [68] or derived from degenerated quadrilaterals with collapsed edges [110]. All computations recorded in this work have been performed using computer codes developed by members of the Aerospace Computing Laboratory.