

Solution-Acceleration Strategies for High-Order Unstructured Methods

Carlos A. Pereira

A Thesis
In the Department of
Mechanical, Industrial and Aerospace Engineering

Presented in Partial Fulfillment of the Requirements
For the Degree of
Doctor of Philosophy (Mechanical Engineering)

at Concordia University
Montréal, Québec, Canada

July, 2023
© Carlos A. Pereira, 2023

CONCORDIA UNIVERSITY
School of Graduate Studies

This is to certify that the thesis prepared

By: Carlos A. Pereira
Entitled: Solution-Acceleration Strategies for High-Order Unstructured Methods

and submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Mechanical Engineering)

complies with the regulations of the University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

_____	Chair
Dr. Xiupu Zhang	
_____	External Examiner
Dr. Krzysztof Fidkowski	
_____	Examiner
Dr. Charles Kiyanda	
_____	Examiner
Dr. Marius Paraschivoiu	
_____	Examiner
Dr. Tristan Glatard	
_____	Supervisor
Dr. Brian Vermeire	

Approved by

Dr. Muthukumaran Packirisamy, Graduate Program Director
Department of Mechanical, Industrial, and Aerospace Engineering

August 29, 2023

Dr. Mourad Debbabi, Dean
Faculty of Engineering and Computer Science

ABSTRACT

Solution-Acceleration Strategies for High-Order Unstructured Methods

Carlos A. Pereira, Ph.D.
Concordia University, 2023

The design of next-generation aircraft relies on computational fluid dynamics (CFD) to minimize testing requirements at reduced cost and risk. However, current industry reliance on Reynolds-averaged Navier-Stokes (RANS)-based CFD is limited in predicting transitional and turbulent flows. Large-eddy simulation (LES) offers accuracy where RANS methods fail, but can have prohibitive computational cost. To address this, we propose a high-order CFD framework to advance flux reconstruction (FR) methods toward industrial-scale simulations. FR is a family of high-order, unstructured schemes that provide accuracy at reduced cost per degree-of-freedom (DOF) compared to low-order methods, with proven potential for LES. We develop practical strategies to reduce the computational cost of FR methods for explicit and implicit formulations. Due to the low cost per time step, explicit time stepping is typically used in FR methods. However, stability constraints prohibitively limit time-step sizes in numerically stiff problems. Hence, implicit time stepping is preferred in these cases, but it requires solving large, nonlinear systems and can be computationally expensive.

This thesis introduces optimal Runge-Kutta methods to alleviate stability limits and reduce wall-clock times by approximately half in moderately low stiffness problems. For increased stiffness, we hybridize implicit FR methods using a trace variable, which allows a reduction of the implicit system via static condensation, decreasing implicit time stepping costs, especially at higher orders. Hybridization with both discontinuous (HFR) and continuous function spaces (EFR) is suitable for advection and advection-diffusion type problems within the FR method and enables significant speedup gains over standard FR. We incorporate polynomial adaptation to the hybridized framework, varying the solution polynomial's degree locally within each element, which results in an overall reduction in DOF and significant speedup gains in a two-dimensional problem against standard polynomial-adaptive formulations. Finally, we combine implicit-explicit (IMEX) time stepping with hybridization to tackle geometry-induced numerical stiffness. The resulting method reduces computational cost at least fifteen times over explicit methods in a multi-element airfoil problem at Reynolds 1.7 million. Our proposed framework enables substantial reductions in both moderate and high stiffness problems, thus advancing high-order methods toward large industrial-scale problems.

Acknowledgements

I would like to acknowledge the support I received from the Gina Cody School of Engineering through the Gina Cody Scholarship and the international tuition waiver award, and the Fonds de recherche du Québec- Nature et Technologie via a B2X Scholarship. This financial support relieved me of financial burdens, allowing me to focus on my studies. I am grateful for the opportunities provided by these scholarships.

I would like to express my deepest gratitude to my family for their support and encouragement throughout my Ph.D. journey. Their love and belief in me throughout my life have been vital in my personal and professional paths.

I would also like to acknowledge my current and past labmates. The collaborative environment we shared in the lab, filled with knowledge-sharing and enjoyable moments, made these years wonderful. Their support and friendship have been invaluable to me, and I am grateful for the memories we have created together. In particular, Hamid, Mohsen, Ramin, Siavash, and Abhishek embarked on their graduate journeys alongside me and have become lifelong friends. I also want to acknowledge more recent members, Amar, Marie-Pier, and Mohammad, who have provided me with insightful conversations and motivation during the last months of my Ph.D.

I want to acknowledge my friends Romulo and Divya for their continuous motivation and support. I am also deeply thankful to Nicolas, Ana, and Cynthia for their love, encouragement, and support during the challenges I faced throughout my studies. They have been a constant source of strength. I am truly fortunate to have their presence in my life, and I am grateful for their constant belief in me and for having so much love for Téo.

Finally, I would like to extend my heartfelt gratitude to my supervisor, Brian Vermeire, for his consistent guidance, trust, and invaluable mentorship throughout my graduate journey. His expertise, support, and constructive feedback have played a pivotal role in shaping my research and professional growth. I am genuinely grateful for all the invaluable conversations we had and the advice I received.

Table of Contents

List of Figures	viii
List of Tables	xiii
List of Abbreviations	xvi
List of Symbols	xviii
1 Introduction	1
1.1 Motivation	1
1.2 Numerical Approaches for Turbulent Flows	2
1.3 High-Order Methods	5
1.4 Solution Acceleration in High-Order Computational Fluid Dynamics	7
1.4.1 Time-Marching Methods	7
1.4.2 Hybridization	9
1.4.3 Polynomial Adaptation	10
1.5 Thesis Objectives and Contributions	10
1.6 Thesis Outline	12
2 Governing Equations	13
2.1 Linear Advection-Diffusion	13
2.2 Burgers Equation	14
2.3 Euler Equations	14
2.4 Navier-Stokes	16
2.5 Entropically-Damped Artificial Compressibility	16
3 The Flux Reconstruction Framework	18
3.1 Preliminaries	18
3.2 Implementation	19
3.2.1 Standard Formulations	23

3.2.2	Hybridized Formulations	24
4	Optimal Explicit Runge-Kutta Formulations	32
4.1	Overview	32
4.2	Stability Polynomials	32
4.2.1	Semidiscrete Analysis	34
4.2.2	Optimal Stability Polynomials	37
4.2.3	Numerical Results	48
4.3	Fully-Discrete Analysis	55
4.3.1	Fully-Discrete Properties of FR schemes using Optimal RK schemes .	58
4.3.2	Numerical Experiments	69
5	Implicit Hybridized Formulations	79
5.1	Overview	79
5.2	Advection-Type Problems	79
5.2.1	Spectral Analysis	79
5.2.2	Numerical Results	83
5.3	Advection-Diffusion-Type Problems	97
5.3.1	Stability Analysis	97
5.3.2	Local Post-Processing	107
5.3.3	Numerical Examples	108
6	Implicit Polynomial-Adaptive Hybridized Formulations	123
6.1	Overview	123
6.2	Implementation	123
6.3	Numerical Examples	126
6.3.1	Isentropic Vortex	126
6.3.2	Cylinder at $Re = 150$	129
6.3.3	NACA 0012 Airfoil at $Re = 10,000$	135
7	Implicit-Explicit HFR-FR Formulations	142
7.1	Overview	142
7.2	Formulation	144
7.3	Numerical Examples	148
7.3.1	Verification	148
7.3.2	Laminar Flow over a Circular Cylinder	150
7.3.3	Turbulent Flow over a Circular Cylinder	154
7.3.4	Turbulent Flow over a Multi-Element Airfoil	158

8	Conclusions	169
8.1	Practical Summary	171
8.2	Future Work	172
	References	173
	Appendix A HFR Grid Convergence Tables	188
	Appendix B Gauss-Lobatto Collocation in EFR	193
	Appendix C HFR Parallelization	197
	Appendix D IMEX Butcher Tableaus	199

List of Figures

1.1	Turbulent kinetic energy (E_k) cascade	2
1.2	Sketch of boundary layer and streamlines around an airfoil	4
4.1	Vector decomposition reference	35
4.2	Subdivision of element types for $d = 2$ and $d = 3$	35
4.3	Collection of eigenvalues for a fourth-order spatial discretization for two and three-dimensional elements	38
4.4	Optimal time-step size for OERK schemes for $p = 1 - 8$ spatial discretizations	39
4.5	Normalized optimal time-step size for OERK schemes for $p = 1 - 8$ spatial discretizations	40
4.6	Stability polynomial of classical RK schemes of degree 1, 2, 4 and 8 alongside eigenvalues of a $p = 4$ FR spatial discretization	41
4.7	Stability polynomial of OERK schemes of degree 1, 2, 4 and 8 shown in a dashed line alongside eigenvalues of a $p = 4$ semidiscretization	42
4.8	Plots of the scaled optimal time-step size as a function of the number of stages for different ($p = 4$) element types	44
4.9	Convergence plots for linear advection using the method of manufactured solutions using optimal RK schemes	50
4.10	Isosurfaces of Q-criterion coloured by velocity magnitude for the Taylor-Green vortex	52
4.11	Energy decay rate versus dimensionless time for the Taylor-Green vortex case	53
4.12	Enstrophy versus dimensionless time for the Taylor-Green vortex case	54
4.13	Spectral properties of $P_{16,1}$ schemes at different CFL numbers τ	59
4.14	Spectral properties of $P_{16,2}$ schemes at different CFL numbers τ	60
4.15	Spectral properties of $P_{3,3}$ schemes at different CFL numbers τ	62
4.16	Spectral properties of $P_{16,3}$ schemes at different CFL numbers τ	63
4.17	Spectral properties of $P_{4,4}$ schemes at different CFL numbers τ	64
4.18	Spectral properties of $P_{16,4}$ schemes at different CFL numbers τ	65
4.19	Spectral properties of $P_{8,8}$ schemes at different CFL numbers τ	67

4.20	Spectral properties of $P_{16,8}$ schemes at different CFL numbers τ	68
4.21	Solution for advection of a Gaussian bump using first and second-degree OERK schemes at $t = 1$ at different time-step sizes.	70
4.22	Solution at $t = 1$ for the advection of a Gaussian bump using OERK schemes	71
4.23	Solution of the advection of a Gaussian bump problem using first and second-degree OERK schemes at $t = 4000$	72
4.24	Advection of a Gaussian bump using OERK schemes at $t = 4000$	73
4.25	Energy spectrum for viscous Burger turbulence using classical and optimized fourth-order RK schemes	77
4.26	Energy spectrum for viscous Burger turbulence using classical and optimized third-order RK schemes	78
5.1	Element coupling involved in HFR and collocation-type EFR semidiscrete operators	81
5.2	Dissipation curves for different c -parameters at $\theta = \pi/6$ for different polynomial degrees	82
5.3	Dispersion curves for different c -parameters at $\theta = \pi/6$ for different polynomial degrees	83
5.4	Maximum resolved wavenumber $\bar{\kappa}_{1\%}$ against the wave incidence angle θ for different parameters c	84
5.5	Linear advection solution for initial Gaussian profile at time $t = 20$ and $p = 3$.	86
5.6	LHS matrix views for $p = 3$ FR, HFR and EFR methods for a 25×25 -element grid for the isentropic vortex case	88
5.7	Effects of the c -parameter on the performance of FR schemes. Results are shown for $p = 3$ for the isentropic vortex case. Values of c begin at $c_-/2$ and increase in the direction of the arrows	89
5.8	Wall-clock time against number of elements for FR, HFR, and EFR for the isentropic vortex case for different correction parameters	90
5.9	Unstructured mesh of 6830 quadrilateral elements for the airfoil-vortex interaction case	91
5.10	Pressure readings at $y = 0$ before the interaction in the airfoil-vortex interaction case with c_{DG}	92
5.11	Evolution of the lift coefficient for the airfoil-vortex interaction case with c_{DG}	94
5.12	Evolution of the moment coefficient for the airfoil-vortex interaction case with c_{DG}	95
5.13	Density residual against wall-clock time for the airfoil-vortex interaction case with c_{DG}	96

5.14	Wall-clock times for the airfoil-vortex interaction case for polynomial degrees 1 to 4 using all considered parameters c	97
5.15	Contours of the solution and post-processed solution for the steady-state linear advection-diffusion problem	110
5.16	Effects of the correction function on the grid convergence rate of hybridized FR methods with discontinuous traces for the steady-state linear advection-diffusion problem for multiple c -parameters	111
5.17	Performance metrics for the steady-state linear advection-diffusion problem.	112
5.18	Performance metrics for the unsteady linear advection-diffusion problem . . .	114
5.19	Normalized time to solve the linear system as a function of the time step size.	114
5.20	Time-evolution of the solution energy for hybridized FR methods for the Gaussian Pulse case with $\beta = 0$ for different values of $\bar{\lambda}$	116
5.21	Zoom of the time-evolution of the solution energy for hybridized FR methods for the Gaussian Pulse case with $\beta = 0$ for different values of $\bar{\lambda}$	117
5.22	Time-evolution of the solution energy for hybridized FR methods for the Gaussian Pulse case with $\beta = 0.01$ for different values of $\bar{\lambda}$	118
5.23	Zoom of the time-evolution of the solution energy for hybridized FR methods for the Gaussian Pulse case with $\beta = 0.01$ for different values of $\bar{\lambda}$	119
5.24	Performance metrics for the planar Couette problem	120
5.25	Temperature and density residual against wall-clock time in seconds for a $p = 4$ scheme with different correction functions on the 16×8 grid	121
5.26	Average time to solve the linear system per iteration ($p = 4$ finest grids) for the Couette problem	122
6.1	Portion of the isentropic vortex domain showing contours of (a) vorticity, and polynomial distribution for (b) $\epsilon = 0.2$ and (c) $\epsilon = 0.01$ after one time step .	128
6.2	Performance metrics for the isentropic vortex problem in terms of internal degrees of freedom DOF_v and runtime in seconds	129
6.3	Portion of the cylinder mesh composed of 3564 quadrilateral elements refined toward the cylinder walls to resolve the boundary layer	129
6.4	Instantaneous snapshots of the cylinder problem showing contours of vorticity for FR, HFR, and EFR at moments of minimum lift	131
6.5	Instant snapshots of the cylinder problem showing the adaptation algorithm for FR, HFR and EFR at moments of minimum lift. Contour lines of vorticity are superimposed	132
6.6	Runtime in seconds for the cylinder problem using p -adaptation	135
6.7	NACA 0012 airfoil mesh consisting of 8658 quadrilateral elements refined at the wall boundaries	135

6.8	Instant snapshots of the airfoil problem showing contours of vorticity for FR, HFR, and EFR at moments of minimum lift	137
6.9	Instant snapshots of the airfoil problem with superimposed contours of vorticity for FR, HFR, and EFR at moments of minimum lift	138
6.10	Averaged C_p coefficient for the adaptive schemes on the upper and lower walls of the airfoil. Reference is a $p = 5$ FR result	138
6.11	Overview of the runtime spent on the global and local solves for the airfoil problem in 100 iterations	139
7.1	Reference domain partitioning for IMEX schemes	145
7.2	Distribution of trace and solution points in this configuration for HFR and EFR IMEX methods for a $p = 2$ discretization	146
7.3	Distribution of explicit and implicit elements for the IMEX verification of linear advection	149
7.4	Mesh for the laminar cylinder case at $Re = 150$ consisting of 3090 quadrilateral elements	151
7.5	Distribution of element sizes for the laminar cylinder case	151
7.6	Maximum stable time-step size for multiple implicit fractions for the AIMEX _{10,2} scheme	152
7.7	Speedup factors for multiple implicit fractions for the AIMEX _{10,2} scheme	153
7.8	Evolution of drag and lift coefficients for the cylinder at $Re = 150$	155
7.9	Distribution of implicit (red) and explicit (blue) elements in the computational domain for the turbulent cylinder case	156
7.10	Spectra of the unsteady lift coefficient fluctuations for the turbulent cylinder problem	157
7.11	Side view of Q-criterion contours for the turbulent cylinder problem	159
7.12	Computational grid for the multi-element airfoil	161
7.13	Distribution of element sizes for the multi-element airfoil grid	162
7.14	Distribution of implicit and explicit elements after IMEX partitioning. Blue elements are solved explicitly, and red elements are solved implicitly	162
7.15	Vorticity contours for the multi-element airfoil problem	163
7.16	Q-criterion contours for the multi-element airfoil coloured by velocity magnitude	164
7.17	Q-criterion contours zoomed in the slat for the multi-element airfoil coloured by streamwise vorticity	165
7.18	Pressure coefficient C_p for the multi-element airfoil problem	166
7.19	Line plot locations in the vicinity of the slat	167
7.20	Plots of sampled normalized velocity magnitude for the multi-element airfoil across slat shear layer	168

B.1	Dissipation curves for different c -parameters at $\theta = \pi/6$ for different polynomial degrees with exact integration of the transmission conditions	195
B.2	Dispersion curves for different c -parameters at $\theta = \pi/6$ for different polynomial degrees with exact integration of the transmission conditions	196

List of Tables

3.1	Coefficients α_{DOF} and α_{NNZ} for quadrilateral elements	31
3.2	Coefficients α_{DOF} and α_{NNZ} for hexahedral elements	31
4.1	Ratio of the maximum time step for $p = 4$ two-dimensional elements using their optimal stability polynomial and s stages, relative to the maximum time step obtained using a stability polynomial optimized for a one-dimensional element	43
4.2	Ratio of the maximum time step for $p = 4$ three-dimensional elements using their optimal stability polynomial and s stages, relative to the maximum time step obtained using a stability polynomial optimized for a one-dimensional element	45
4.3	Time step ratio of $q = 3$ optimal RK schemes for different element types with $p = 4$ relative to RK _{3,3} , $\frac{\Delta t_{\text{opt}}}{\Delta t_{\text{RK}_{3,3}}}$	46
4.4	Speedup factor of $q = 3$ optimal RK schemes for different element types with $p = 4$ relative to RK _{3,3} , $(\frac{\Delta t_{\text{opt}}}{s} / \frac{\Delta t_{\text{RK}_{3,3}}}{3})$	47
4.5	Time step ratio of $q = 4$ optimal RK schemes for different element types with $p = 4$ relative to RK _{4,4} , $\frac{\Delta t_{\text{opt}}}{\Delta t_{\text{RK}_{4,4}}}$	47
4.6	Speedup factor of $q = 4$ optimal RK schemes for different element types with $p = 4$ relative to RK _{4,4} , $(\frac{\Delta t_{\text{opt}}}{s} / \frac{\Delta t_{\text{RK}_{4,4}}}{4})$	48
4.7	Convergence of the L_2 norm for the sine advection case	74
4.8	Maximum time-step sizes for viscous Burgers turbulence using $P_{s,4}$ schemes .	75
5.1	Energy $\ u^h\ ^2$ at $t = 20$ for the Gaussian profile case with $p = 3$	85
5.2	Average number of GMRES iterations for $p = 3$ schemes to solve the global system for the linear advection case. Results are shown for a 20×20 grid . .	86
5.3	Vortex pressure core readings at $t_c = 4$ for the airfoil-vortex interaction case	93
5.4	Average time (in seconds) per linear solve for the airfoil-vortex interaction case	93
5.5	Coefficients α_{DOF} and α_{NNZ} resulting from the airfoil-vortex interaction problems at different polynomial degrees	97

5.6	Speedup factors for the unsteady linear advection-diffusion problem with 1600 elements. Values are shown for two time step sizes and compared against the corresponding FR formulation with the same c -parameter	115
5.7	Speedup factors for the Couette problem with 32×16 elements for the HFR and EFR methods. These factors take into account the time to solve the global and local problems but exclude the assembly time of the Jacobian matrices	120
6.1	Resulting average number of degrees of freedom for the p -adaptive runs for the isentropic vortex problem	127
6.2	Averaged results for the cylinder problem for FR, HFR, and EFR schemes including both p -uniform and p -adaptive runs	133
6.3	Speedup factors for the cylinder problem. Results compare runtime ratios between the FR method and the corresponding scheme for all considered p -adaptive and p -uniform runs	134
6.4	Results for the airfoil problem for FR, HFR, and EFR methods with p -uniform and p -adaptive discretizations	140
6.5	Speedup factors for the airfoil problem. Results compare runtime ratios between the FR method and the corresponding scheme for all considered p -adaptive and p -uniform runs	141
7.1	General form of Butcher tableaus for IMEX schemes	143
7.2	Convergence table showing the L_2 -norm of the solution error and the order of accuracy for linear advection using multiple IMEX schemes	150
7.3	Summary of results for the cylinder at $Re = 150$ for dual and single scheme IMEX methods	154
7.4	Summary of results for the turbulent cylinder case	157
7.5	Summary of performance metrics for the turbulent cylinder case for 100 time steps	158
7.6	Summary of performance metrics for the multi-element airfoil problem for 100 time steps	167
A.1	L_2 norm of the error for the linear advection problem. Results are presented for HFR/FR and EFR methods	188
A.2	L_2 norm of the error for the steady-state linear advection-diffusion problem	189
A.3	L_2 -norm of the post-processed solution error for the steady-state linear diffusion problem	189
A.4	L_2 -norm of the solution error for the steady-state linear diffusion problem with consistent LDG switch	189

A.5	L_2 norm of the error for the unsteady advection-diffusion problem	190
A.6	L_2 -norm of the post-processed solution error for the unsteady linear diffusion problem	190
A.7	L_2 norm of the temperature error for the planar Couette flow case	191
A.8	L_2 norm of the isentropic vortex for FR, HFR, and EFR with c_{DG} , c_{SD} , c_{HU} for degrees $p = 1$ to $p = 5$	192
D.1	IMEX _{3,2} Butcher tableau	199
D.2	IMEX _{5,3} Butcher tableau	200

List of Abbreviations

AIMEX	Accelerated implicit-explicit
ACM	Artificial compressibility method
BR1	Bassi-Rebay I
BR2	Bassi-Rebay II
CFD	Computational fluid dynamics
CFL	Courant–Friedrichs–Lewy
DG	Discontinuous Galerkin
DGSEM	Discontinuous Galerkin spectral element method
DIRK	Diagonally-implicit Runge-Kutta
DNS	Direct numerical simulation
DOF	Degrees of freedom
EDAC	Entropically damped artificial compressibility
EDG	Embedded discontinuous Galerkin
EFR	Embedded flux reconstruction
ERK	Embedded Runge-Kutta
ESFR	Energy-stable flux reconstruction
FD	Finite difference
FR	Flux reconstruction
FV	Finite volume
GLL	Gauss-Lobatto-Legendre
GMRES	Generalized minimal residual method
HDG	Hybridized discontinuous Galerkin
HFR	Hybridized flux reconstruction
HLL	Harten-Lax-van Leer
HORUS	High-order unstructured solver

IEDG	Interior-embedded discontinuous Galerkin
ILES	Implicit large-eddy simulation
IMEX	Implicit-explicit
IF	Implicit factor
IP	Interior penalty
LDG	Local discontinuous Galerkin
LES	Large-eddy simulation
LHS	Left-hand side
LLF	Local Lax-Friedrichs
NNZ	Number of nonzeros
ODE	Ordinary differential equation
OERK	Optimal explicit Runge-Kutta
PIV	Particle image velocimetry
PSD	Power spectral density
RANS	Reynolds-averaged Navier-Stokes
RAS	Restricted additive Schwarz
RK	Runge-Kutta
SD	Spectral difference
SDIRK	Singly-diagonally implicit Runge-Kutta
SGS	Subgrid-scale
SV	Spectral volume
VCJH	Vincent-Castonguay-Jameson-Huynh

List of Symbols

In this thesis, the following notation is employed. Variables are expressed in this list in scalar form. When applicable, their vector/matrix forms are represented using bold formatting.

Roman Letters

\bar{c}_d	Time-averaged drag coefficient
c_d, C_D	Two and three-dimensional drag coefficients
c_l, C_L	Two and three-dimensional lift coefficients
$\bar{c}_{m,c/4}$	Time-averaged moment coefficient about the quarter-chord point
C_p	Pressure coefficient
c_P	Specific heat at constant pressure
c_V	Specific heat at constant volume
E_k	Kinetic energy
E_s	Stiffness indicator
\bar{f}	Face in global mesh
\mathfrak{F}	Interface flux
F	Flux variable
H	Normal flux jump
$H(\text{div}; \Omega)$	Sobolev space of vector functions with square-integrable divergence
J	Transformation Jacobian
ℓ	Length scale
\mathbb{M}	Trace polynomial space
\mathcal{M}	Transformation to physical space
M	Mach number
N	Number of elements

N_{ec}	Number of elements in a periodic cube
N_f	Number of faces in an element
N_g	Number of mapping points per element
N_{GMRES}	Number of GMRES iterations
N_r	Number of flux points in an element
$N_{r,f}$	Number of flux points in a face
N_s	Number of solution points in an element
N_v	Number of mesh vertices
\mathbb{P}	Solution polynomial space
p	Polynomial degree
p_{\max}	Cutoff polynomial degree in p -adaptation
Pr	Prandtl number
$P_{s,q}$	Stability polynomial of degree s and order q
Re	Reynolds number
St	Strouhal number
t	Time
t_c	Characteristic time
t_G	Time spent solving the global implicit system
\mathcal{T}^h	Set of elements in computational domain
$\partial\mathcal{T}^h$	Set of element boundaries in computational domain
t_J	Time spent assembling the Jacobian matrix
t_R	Time spent on the right-hand side residual evaluations
t_w	Wall-clock time
\mathfrak{U}	Interface solution
u	Solution variable
v_g	Group velocity
Z	Vorticity

Greek Letters

α	Linear advection velocity
α_{DOF}	Coefficient of degrees of freedom
α_{NNZ}	Coefficient of number of nonzeros

β	Diffusion coefficient
Δt	Time-step size
Δt_{opt}	Maximum time-step size in optimal RK methods
ϵ	Tolerance value
ε	Enstrophy
ε^h	Set of all unique faces
$\bar{\gamma}$	Ratio of specific heats
κ	Wavenumber
$\bar{\lambda}$	Linear advection scalar stabilization
$\tilde{\lambda}$	Linear advection scalar stabilization equivalent to $\bar{\lambda} = 0$
λ	Eigenvalue
μ	Coefficient of dynamic viscosity
Ω	Solution domain
Ω_k	Element in Ω
$\partial\Omega_k$	Element boundary
ω	Frequency
ϕ	Trace nodal basis function
φ	Solution nodal basis function
Φ	Vorticity indicator
Ψ	Legendre polynomial
$\bar{\tau}_{ij}$	Component of viscous stress tensor
τ_v	Viscous stabilization scalar in hybridized linear diffusion
Θ	Incompressibility factor
Υ	IMEX interface
ζ	Directional parameter in LDG/LDG-H formulations

Number Sets

\mathbb{C}	Set of complex numbers
\mathbb{R}	Set of real numbers

Operators

argmax	Index of maximum value
$\operatorname{Im}(\cdot)$	Imaginary part

\max	Maximum operator
$\bar{\mathbb{P}}$	Projection operator
$\text{Re}(\cdot)$	Real part
$\{\!\!\{\cdot\}\!\!\}$	Averaging operator, $\{\!\!\{\cdot\}\!\!\} = \frac{(\cdot)_+ + (\cdot)_-}{2}$
$\llbracket \cdot \rrbracket$	Jump operator, $\llbracket \cdot \rrbracket = (\cdot)_- \mathbf{n}_- + (\cdot)_+ \mathbf{n}_+$
$ \cdot $	Magnitude/complex modulus operator
$\ \cdot\ _p$	L_p norm

Subscripts, Superscripts, and Diacritics

\square_0	Interior quantity
$\hat{\square}$	Trace quantity
\square_∂	Domain boundary quantity
\square_\pm	Data from left/right side of interface
$\tilde{\square}$	Transformed quantity
$\square^{(c)}$	Convective component
\square^{ex}	Quantity from explicit formulation
\square^{FR}	Quantity associated with the FR method
\square^{HFR}	Quantity associated with the HFR method
\square^h	Discrete approximation
\square^{hC}	Continuous
\square^{hD}	Discontinuous
\square^{im}	Quantity from implicit formulation
\square_k	Variable associated with Ω_k
\square^r	Quantity associated with flux points
\square^s	Quantity associated with solution points
$\square^{(v)}$	Viscous component

Chapter 1

Introduction

1.1 Motivation

Over four billion passengers use air transportation annually, producing over 800 million tonnes of CO₂ and other greenhouse gas emissions [1]. These emissions are known to be two to four times more potent than those at sea level and are a significant driver of climate change [2]. By 2030, there will be a considerable increase in the number of new commercial aircraft, which requires a 20-25% reduction in fuel consumption to comply with current and forthcoming emission regulations [3]. Since 2012, programs such as Canada's Aviation Climate Action Plan [4] have initiated a fight against greenhouse emissions from aviation and expect net-zero emissions within the next few decades. In order to remain globally competitive, aerospace firms acknowledge the significance of investing in various sustainable technology domains, such as alternative propulsion systems and innovative aircraft design. Consequently, further actions must be taken to improve aerodynamic efficiency, reduce noise pollution, and develop new technologies for sustainable aviation fuels. Hence, the design of next-generation aircraft relies on our ability to understand and analyze the intricate flow dynamics throughout the full flight envelope.

The vast majority of flows in engineering applications within the aerospace industry do not show the organized and predictable motion of laminar flows, but instead are turbulent and have a time-dependent chaotic behaviour. The flow regime, whether laminar or turbulent, is characterized by the ratio of inertial or convective forces to viscous forces in the flow, known as the Reynolds number, Re . Since most flows in nature are turbulent, engineers and scientists need to understand their effects on the bodies they surround. While important, turbulent flows are challenging to predict, given the nonlinearity and complexities of the flow behaviour. This complexity is attributed to the presence of a vast range of turbulent structures within the flow, known as eddies. The size of these eddies may range from a macroscale, defined by a characteristic length ℓ such as the chord of an airfoil, to the Kolmogorov microscale [5],

which can be a small fraction of it. This microscale size η is well approximated [6] by

$$\eta \approx \text{Re}^{-3/4} \ell. \quad (1.1)$$

As these structures interact, the large eddies arising from the characteristic length transfer energy to smaller eddies, which in turn transfer energy to even smaller ones. At the level of the smallest scales, the viscosity of the fluid dissipates the remaining energy. The distribution of these scales can be seen in the turbulent energy cascade in Figure 1.1, showing the majority of energy contained in the largest structures. The higher the Reynolds number, the larger the separation between the largest and smallest eddies. This ratio can be estimated to be proportional to $\text{Re}^{3/4}$ [6]. To give some perspective, the Reynolds number in a typical cruising aircraft is in the order of 10 to 20 million for small aircraft such as turbofan airliners, and 45 to 80 million for large aircraft such as a Boeing 747 [7].

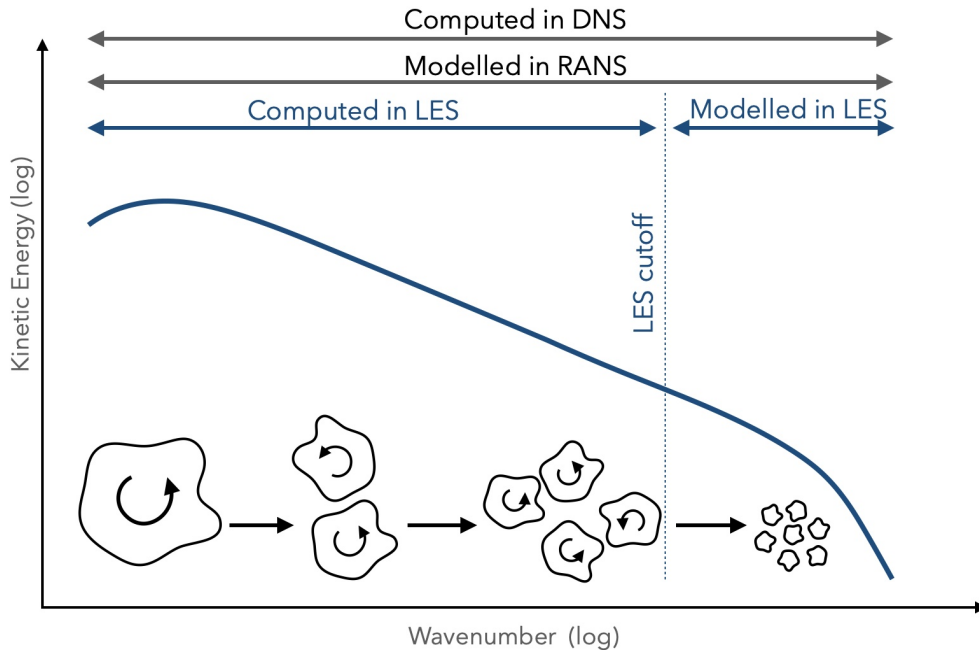


Figure 1.1. Turbulent kinetic energy (E_k) cascade

1.2 Numerical Approaches for Turbulent Flows

In general, the mechanical behaviour of fluids in engineering applications can be described by the Navier-Stokes equations. These equations are nonlinear and can describe phenomena from laminar flow to turbulence. Because no general analytical solution is available for the full range of phenomena involved in the Navier-Stokes equations, computational fluid

dynamics (CFD) focuses on devising numerical methods to efficiently approximate what would otherwise be analytically intensive or impossible to solve. CFD enables reductions in in-flight and ground-based testing requirements and provides valuable physical insight, allowing for better designs at reduced cost and risk. The potential of CFD in the aerospace industry is currently limited by its ability to predict transitional and turbulent flows [8].

To numerically solve the governing equations, we must discretize the equations using suitable numerical schemes. Common approaches include the finite volume (FV) methods, finite difference (FD) methods, and finite element (FE) methods. Due to its geometrical flexibility, the FV method has been the workhorse of industrial applications of flow problems over past decades. Unlike the FD method, the FV method can be easily applied to complex geometries and offers more flexibility and stability than continuous FE methods for convection-dominated problems [9]. However, standard second-order FV schemes generally introduce significant numerical error in the form of dissipation and dispersion. This presents challenges to many practical engineering applications containing high-frequency information that needs to be captured. For instance, highly separated flows, which are dominated by unsteady vortices over an extensive range of length and time scales; and aeroacoustics problems, which contain pressure fluctuations usually orders of magnitude smaller than mean flow quantities [10]. Furthermore, coupled with numerical discretizations, additional considerations need to be made to capture the effects of turbulent flows.

In addition to low-order FV methods, CFD implementations typically rely on the Reynolds-Averaged Navier-Stokes (RANS) equations to model turbulent flows. While commonly used in industrial-scale computations, the limitations of RANS methods in flows where the unsteadiness of the solution plays a major role in the physics confines them to the cruise flight portion of the aircraft operating design space [8]. RANS methods provide time-averaged solutions of the governing equations from an initial time t_0 to a final time T , i.e.

$$\bar{u} = \lim_{T \rightarrow \infty} \frac{1}{T} \int_{t_0}^{t_0+T} u dt, \quad (1.2)$$

and heavily rely on empirical parameters to model the effects of turbulence. Consequently, the spatiotemporal complexity of turbulent flows poses significant challenges to RANS equations. These challenges include flows at moderate Reynolds numbers and high angles of attack and many other complex applications where the inherent unsteadiness of the solution has an important effect on the results, such as transition to turbulence and vortex transport.

Scale-resolving approaches, such as direct numerical simulation (DNS) [11] and large-eddy simulation (LES), are known to be more accurate in regimes where RANS methods fail. These approaches directly model all or some length and temporal scales in the largely disparate range sizes that characterize turbulent flows. As shown in Figure 1.1, DNS aims to resolve

the full turbulence spectrum and requires domains large enough to cover the flow physics as well as fine grids to capture all scales. The cost of this high resolution is dictated by the Kolmogorov microscale. Most of the computational effort in DNS simulations goes to resolving these smallest sizes, which generally do not significantly contribute to the overall flow dynamics of engineering problems. The number of grid points required to perform DNS simulations can be approximated to scale with [6]

$$N \propto \text{Re}^{9/4}. \quad (1.3)$$

For typical applications in the aerospace industry, where Re can be in the range of tens of millions, the computational resources required by DNS outweigh the accuracy gains compared to LES. DNS is therefore confined primarily to academic studies due to this high cost, and is not expected to be feasible for many years [12].




Figure 1.2. Sketch of boundary layer and streamlines around an airfoil

On the other hand, LES methods resolve only the largest energy-containing eddies, which dominate the transfer of momentum and heat. As shown in Figure 1.1, this is done up to a cutoff scale, typically dictated by the grid size. The rest of the medium to smallest scales can be modelled explicitly via a subgrid-scale (SGS) model such as the dynamic Smagorinsky model [13]. In some spatial discretizations, the numerical error behaves similarly to an SGS model and is often used as an implicit LES (ILES) approach [14], which has shown significant potential [15, 16]. The ILES approach does not require additional model implementations beyond the governing equations and has been shown in some cases to provide more accurate results than explicit SGS models, which can introduce more dissipation than necessary [17]. Therefore, we make use of ILES throughout this thesis.

Despite the significant reduction in the number of resolved scales and a considerable cut in the computational cost over DNS, LES simulations are still computationally expensive. This cost is mainly attributed to the strict requirements on the grid cell sizing when resolving a boundary layer. Inside the boundary layer, the influence of viscosity is confined to a thin layer in contact with the walls. A velocity profile $u(\eta)$ is shown for an airfoil problem in Figure 1.2,

showing that the velocity goes to zero at the wall and rapidly increases to the freestream value U_∞ at the interface. This layer significantly decreases in thickness with increasing Reynolds number. Thus, in industrial-scale problems, it is typically very thin. Despite their small thickness, boundary layers are the source of vorticity and complex flow dynamics, and thus they are responsible for the primary flow dynamics. Hence, at large values of Re , the majority of the cost in LES is attributed to the small grid sizes to capture the large gradients within the boundary layer. This is due to two main reasons. First, the number of overall cells within this region grows toward the requirements of a DNS simulation [18], thereby expanding the size of the problem. Second, smaller time steps are necessary to maintain stability in typical explicit implementations, increasing computational time. These two aspects reduce the feasibility of using LES in industrial-scale design processes.

In order to design next-generation, quieter, and more sustainable aviation, further efforts still need to be made to enable scale-resolving LES simulations of industrial-scale problems at a reasonable computational cost. Hence, there is a requirement to develop novel efficient numerical methods that minimize numerical errors and are optimized for computational resource utilization.

1.3 High-Order Methods

The order of a numerical scheme influences the required grid resolution to attain a given error level. For smooth enough solutions, the numerical error of a spatial discretization of order p decreases proportionally to [19]

$$e \propto h^p, \tag{1.4}$$

where h is the cell size. High-order methods are able to provide improved accuracy with reduced computational cost per degree of freedom compared to conventional low-order methods [19]. Finite difference methods have existed for decades and provide high-order solutions. However, they are formulated on structured or curvilinear grids, confining their applications to simple geometries. Applications of engineering interest involve complex geometries such as landing gears and high-lift devices, which make the mesh generation process more cumbersome. These types of problems benefit from the flexibility of unstructured methods. High-resolution FV methods [20, 21] can provide high-order and are suitable for unstructured grids. However, they achieve this by increasing the number of neighbouring elements, which adds significant complexity and deteriorates the performance of modern parallel hardware implementations.

High-order compact unstructured methods rely on basis functions and locally represent

solutions inside each element using polynomials. Similar to the FV method, the solution is allowed to be discontinuous at the interfaces, and they incorporate Riemann solvers to control the stabilization mechanism properties. Compared to continuous FE methods, this added flexibility makes them ideal candidates for unstructured high-fidelity simulations. Among such high-order schemes are the discontinuous Galerkin (DG) [22, 23], spectral volume (SV) [24], and spectral difference (SD) [25] methods.

In 2007, Huynh introduced flux reconstruction (FR) as a unifying high-order method for advection problems on tensor-product elements [26]. Later, it was extended to the diffusion regime [27] and unstructured grids [28]. This framework is referred to as unifying because it can recover existing formulations of the abovementioned high-order methods via correction functions. Similarly, additional high-order methods have been identified, such as the range of energy-stable flux reconstruction (ESFR) schemes of Vincent et al. [29, 30]. In fact, the choice of the correction parameter introduces additional flexibility and influences the behaviour of the discretization. For instance, correction parameters have been identified to increase the maximum allowable time-step size when used with explicit time-stepping [31]. For implicit time-stepping, this parameter can influence the performance of linear solvers [32]. FR methods offer high-order accuracy on compact stencils, making their explicit formulations exceptionally well-suited for massively parallel computer architectures [33].

These characteristics make the FR approach appealing to model complex phenomena, including convection-dominated flows, where wave propagation dominates the dynamics of flow physics. In this sense, multiple spectral analyses of these methods [31, 34, 35, 36] have been carried out, demonstrating the potential of this approach in simulating turbulence. In fact, research has shown that the FR method provides appropriate numerical error to act as a simple subgrid-scale model in ILES computations [37, 38, 39].

Despite recent progress in computational power and algorithm development, scale-resolving simulations of high Reynolds number flow around complex aircraft configurations are still prohibitive and require very long computational times. Due to the repeated runs needed in the design and optimization processes, high-order methods must yield simulations with small turnaround times. Hence, the feasibility of applying high-order FR methods to complex problems of engineering interest relies on developing new efficient algorithms that maintain accuracy in space and time.

1.4 Solution Acceleration in High-Order Computational Fluid Dynamics

The following section provides a comprehensive literature review of current state of high-order CFD methods. Specifically, we discuss challenges in computing problems of industrial scale and how current solution-acceleration approaches still require further advancement to enable simulations at reduced turnover times. This work primarily concentrates on developing approaches that offer broad applicability. The following subsections encompass three main areas in high-order computational fluid dynamics, which will be the basis of the proposed framework developed in this thesis.

1.4.1 Time-Marching Methods

The behaviour of flow at industrial scales is inherently unsteady and complex, so we need appropriate temporal methods to advance the solution in time. These methods are generally classified as explicit, implicit, and implicit-explicit (IMEX).

Explicit methods, such as explicit Runge-Kutta (ERK) schemes, calculate the state of the system at a later time or stage from a known value at the current time or stage. They are generally easy to implement, do not require significant memory, and have a short computation time per time step. However, the maximum allowable time-step size is determined by the relationship between the spatial discretization and a constrained stability region that dictates the maximum stable Courants-Fredrichs-Lewys (CFL) number. Hence, they are generally suitable for nonstiff problems. On the other hand, implicit methods calculate the future state of the system via the solution of a coupled system of equations. Unlike explicit methods, implicit schemes can be unconditionally linearly stable, and the value of the time-step size is typically chosen based on accuracy rather than stability. In general industrial applications, the solution in an implicit method must be obtained via nonlinear solvers.

With advances in high-order methods, computations of more complex problems are becoming more feasible. Typical industrial applications of flow over aircraft wings are in the range of seven/eight-figure Reynolds numbers. The higher the value of this parameter, the denser the computational grid must be to capture the resulting thin boundary layer and smallest turbulent length scales. This poses an important challenge for the use of explicit methods due to their stability constraints, as the largest allowable time-step size becomes prohibitively small for these problems. In implicit methods, however, we are required to perform linearization, storage, preconditioning, and solution of large nonlinear systems, which scale very rapidly with the order of accuracy at $O(p^d)$, where p is the polynomial degree representing the solution, and d is the dimension of the problem. For typical problems of

engineering interest, wall-resolved LES simulations using purely implicit methods necessitate unattainable amounts of computational resources, especially memory.

IMEX methods combine the use of explicit and implicit methods to balance efficiency, stability, and accuracy, using the best features of each approach [40, 41, 42]. Generally, they are designed such that an implicit method is used to solve the stiff components of the underlying equations and an explicit method is used on the nonstiff portion. Then, the two are paired to enable time integration that maintains conservation and a desired order of accuracy. With IMEX methods, the maximum CFL condition can be relaxed in the explicit part, and the simulation can be run more efficiently. IMEX methods were initially devised to time-split the convection-diffusion-reaction operator into stiff and nonstiff components [40]. More recently, they have been applied to tackle so-called geometry-induced stiffness [43, 44]. This occurs when there exist largely disparate cell sizes in a single computational domain, as in, for instance, the cells within the boundary layer of an airfoil at high-Reynolds numbers compared to those in the far-field [43]. IMEX schemes have been shown to provide speedups over explicit methods by one to two orders of magnitude, thereby reducing the computational cost of LES simulations [44].

More recently, novel explicit and implicit-explicit methods have been introduced. These are obtained by reshaping stability polynomials of typical explicit Runge-Kutta (RK) methods to increase their maximum allowable time-step size for specific spatial discretization methods and conservation laws [45, 46, 47]. Optimal explicit RK schemes for the FR methods have been devised and analyzed in this thesis, and optimal IMEX RK schemes have been formulated in [48]. These methods have also been developed for specific applications such as in pseudo-time stepping for incompressible flows [49] and combined with p -multigrid methods [50, 51] for convergence acceleration.

The coupling of the temporal component with a semidiscretization yields a scheme with distinct spectral properties. A fully-discrete analysis of these properties was investigated by Yang et al. [52] for DG schemes coupled with classical RK schemes of second and third-order and contrasted with Lax-Wendroff schemes of the same orders. Their study was confined to the behaviour of well-resolved wavenumbers. Vermeire et al. [53] extended this analysis to the region of under-resolved wavenumbers to explore the dispersive and dissipative properties of the fully-discrete energy-stable FR schemes [30]. They considered conventional implicit and explicit RK schemes and found that RK schemes may change the semidiscrete spectral properties and provided recommendations for ILES, such as utilizing a time-step close to but not on the stability limit. Combining optimal explicit RK methods with implicit schemes leads to the accelerated IMEX (AIMEX) [48] methods, which in turn have been shown to yield speedup factors in excess of two over standard IMEX methods. While both IMEX and AIMEX methods have shown significant speedups over conventional explicit time-stepping,

it is known that these schemes spend the majority of their computation time solving the nonlinear systems that result from the governing equations.

1.4.2 Hybridization

Hybridization can reduce the cost of the implicit solver, which typically takes most of the computation time in implicit and IMEX formulations. It was introduced in the context of DG methods by Cockburn et al. [54], associated with the static condensation of finite-element methods of de Veubeke [55]. Hybridized methods define a new unknown on the faces of the elements, known as the trace variable, which effectively decouples interelement information generally observed in the Jacobian matrix. Then, with transmission conditions, the problem is globally defined in terms of the trace variable and later reduced via static condensation [56]. This leads to a Jacobian matrix, the size of which grows proportionally to $O(p^{d-1})$ instead of the typical $O(p^d)$ of standard implicit schemes. However, note that this is only the scaling, and at low values of p , this can result in a higher number of degrees of freedom. Hence, these methods are most suitable for higher orders.

Hybridized DG (HDG) methods are locally conservative and can exhibit superconvergence after postprocessing in diffusion-dominated problems [57, 58, 59]. This postprocessing recovers a globally continuous solution which converges at $p + 2$, as opposed to the conventional $p + 1$. However, this behaviour is only exhibited in a limited number of problems and is not a generalization of the method for typical industrial LES computations. In the standard HDG methods, the trace unknown can be chosen to be discontinuous. Subsets of these methods have been introduced by modifying the functional space of the trace variable. For example, by enforcing continuity on the skeleton of the domain, the number of globally coupled degrees of freedom decreases. This is known as the embedded discontinuous Galerkin (EDG) method, initially introduced in the context of linear shell problems [60]. An analysis of this method was performed in [61], where it was shown that although EDG results in smaller linear systems, local conservation and superconvergent properties are lost. Instead, Kamenetskiy [62] showed they are conservative on dual volumes. A combination of EDG and HDG methods gives rise to the interior-embedded discontinuous Galerkin (IEDG) methods, which Nguyen et al. [63] showed to be superior to EDG methods. These methods use continuous trace polynomials on interior faces and discontinuous on the domain boundary. This has been shown to provide flexibility and better accuracy when implementing boundary conditions. Hence, the IEDG implicit systems are smaller than for EDG.

HDG methods have been successfully implemented in the context of computational fluid dynamics for a range of problems, including convection [64], convection-diffusion [65, 59, 66], incompressible [67, 68, 69] and compressible [70, 71, 72] flows, as well as turbulent flows via ILES [38, 70, 73]. However, the use of hybridization has been confined to discontinuous

Galerkin methods.

1.4.3 Polynomial Adaptation

In practical engineering applications of CFD, the complexity of the flow physics may not be known *a priori*. This poses a challenge in generating grids to capture important unsteady phenomena that may evolve in space and time. Furthermore, even when *a priori* information is known, high-order methods may result in a large number of unnecessary degrees of freedom. In typical aerospace LES simulations of complex geometries, the domain size is chosen sufficiently large to capture the flow and reduce the effects of boundary conditions. In these cases, the most complex physics associated with the boundary layer and vortical structures are concentrated in relatively small fractions of the full domain. In practice, only these areas require high resolution.

Problems within the FR framework are locally high-order. This allows a natural application of local p -adaptation algorithms, which enable the use of variable degrees of the solution polynomial from one element to another. With this approach, it is possible to reduce the computational cost of high-order simulations while maintaining accuracy, which is suitable for problems with confined regions of large gradients. Applications of p -adaptive methods in the FR family of schemes have been done by implementing interface elements [74], using mortar elements [75] and have been shown to outperform h -adaptation techniques for relatively smooth problems [76]. Adaptive strategies require the use of error estimation indicators that identify regions where higher numerical resolution is desired. These can be truncation-error based [77], adjoint-based [78, 79, 80, 81] and feature-based [82, 83]. A review of these estimators can be found in [84, 85, 86]. Application of p -adaptation can also be performed within hybridized schemes.

Previous work on p -adaptive hybridizable methods [87, 88, 89] has focused on the HDG method where the trace variable is allowed to be globally discontinuous. As we will observe in Chapter 5, the embedded method is, in fact, competitive in terms of efficiency [63, 90] and can be a promising approach where p -adaptation can be incorporated.

1.5 Thesis Objectives and Contributions

The main objective of this thesis is to develop solution-accelerating strategies within the family of flux reconstruction methods toward industrial-scale scale-resolving simulations. We seek to answer: *How can we perform LES simulations of industrial-scale at reduced computational cost?* To achieve this objective, efficient explicit and implicit forms of the FR method will be developed.

For problems of moderate stiffness, such as moderately low Reynolds numbers, optimal explicit Runge-Kutta methods will be developed and analyzed in one, two, and three dimensions. These schemes will then be validated via benchmark cases involving turbulent flow.

Then, hybridization will be introduced for the family of flux reconstruction methods. This includes development, implementation and spectral analysis against standard formulations. A series of numerical examples will verify this implementation and demonstrate the performance and accuracy benefits of this approach in implicit formulations. This implementation will be available within the in-house High-ORder Unstructured Solver (HORUS) software.

In combination with hybridization, polynomial adaptation will be introduced to FR methods to further reduce overall number of degrees of freedoms in implicit formulations. Performance and accuracy measurements will be showcased via a series of numerical examples.

Finally, hybridized and standard FR methods will be combined to obtain partially hybridized IMEX schemes suitable for moderately high Reynolds numbers. A comparison with standard IMEX schemes and explicit methods will be made in terms of performance. The new algorithm will be validated with laminar and turbulent numerical examples, including a multi-element airfoil at $Re = 1.7 \times 10^6$.

The primary contributions of this work are

- The development and analysis of optimal RK methods for FR schemes in one and multiple dimensions with speedups of $2\times$ over standard explicit formulations.
- The development, verification, and validation of a multi-core hybridization framework for the FR approach into the HORUS solver exhibiting speedup factors of an order of magnitude against standard implicit FR.
- The development of a polynomial-adaptive hybridized FR framework capable of reducing the overall computational cost exceeding $6\times$ over conventional polynomial-adaptive standard FR.
- The development, verification, and application of partially hybridized IMEX schemes to solve industrial-scale Reynolds number multi-element airfoil problem. Demonstrated speedup factors of $4\times$ over standard IMEX formulations and at least $15\times$ over explicit methods.

1.6 Thesis Outline

This thesis is organized as follows

- Chapter 2 presents the sets of governing equations considered throughout this work, from the compressible Navier-Stokes equations to simplified systems used in the analysis of the developed methods.
- Chapter 3 presents the flux reconstruction framework, its formulation, hybridization, and implementation details. For the hybridized portion, we discuss the block formulation and static condensation approach and its form for implicit time-marching methods.
- Chapter 4 contains the first solution-accelerating strategy considered in this work, i.e., optimal explicit Runge-Kutta methods. We discuss the computation of optimal stability polynomials in one and multiple dimensions. We also present a fully-discrete analysis characterizing the numerical error for a range of CFL numbers. Results are validated with a series of numerical examples.
- Chapter 5 analyzes hybridized flux reconstruction methods in the advection and advection-diffusion regimes. Here, we discuss the performance, stability, and accuracy of these methods for a range of correction functions.
- Chapter 6 presents polynomial adaptation of hybridized FR methods using a feature-based vorticity indicator. Via a projection and a new algorithm, we perform dynamically adaptive simulations of vortex-dominated flows and present the performance gains from this approach.
- In Chapter 7, implicit-explicit formulations of combined hybridized and standard flux reconstruction methods are developed to tackle geometry-induced stiffness. Specifically, a conservative formulation is presented. Performance and validation results are showcased via LES simulations at moderately high Reynolds numbers.
- Finally, Chapter 8 presents conclusions and immediate future directions of the present work.

Chapter 2

Governing Equations

The numerical experiments considered in this thesis are governed by systems of conservation laws applied on a domain Ω . We can describe a general form of these equations in the form of a convection-diffusion equation

$$\frac{\partial \mathbf{u}}{\partial t} + \sum_{i=1}^d \frac{\partial}{\partial x_i} \left[\mathbf{F}_i^{(c)}(\mathbf{u}) - \mathbf{F}_i^{(v)}(\mathbf{u}, \nabla \mathbf{u}) \right] = \mathbf{s} \quad \text{in } \Omega, \quad (2.1)$$

where $\mathbf{F}^{(c)}$ is the convective flux, associated with the transport effects, $\mathbf{F}^{(v)}$ is the viscous flux, associated with the diffusive phenomena and d is the problem dimension. In some cases, we will make use of a source term vector \mathbf{s} to introduce additional forcing contribution into the physics. We will describe each of the systems that we have used throughout this thesis, starting with the scalar equations such as linear advection-diffusion and the Burgers equation, followed by the nonlinear compressible Euler and Navier-Stokes equations. Finally, we will present an incompressible approach for Navier-Stokes, which allows for explicit formulations.

2.1 Linear Advection-Diffusion

The linear advection-diffusion equation is used several times throughout this work. It describes the transport and diffusion of a scalar u in a medium. The quantity moves at a velocity denoted $\boldsymbol{\alpha}$, and the diffusion process occurs at a rate β . Hence, the convective and viscous fluxes are defined as

$$\mathbf{F}_i^{(c)} = \alpha_i u, \quad \mathbf{F}_i^{(v)} = \beta \frac{\partial u}{\partial x_i}, \quad i = 1, \dots, d. \quad (2.2)$$

When $\beta = 0$, we obtain linear advection associated with wave propagation phenomena, which is commonly used for verification of codes since its exact solution is readily available and

given by

$$u(\mathbf{x}, t) = u^0(\mathbf{x} - \boldsymbol{\alpha}t), \quad (2.3)$$

where u^0 is the initial condition, $u^0 = u(\mathbf{x}, t = 0)$. In addition, we make use of this equation to analyze the spectral properties of semi and fully-discrete systems. In some cases, we also use the method of manufactured solution for verification of both convective and viscous implementations.

2.2 Burgers Equation

Burgers equation is the nonlinear analogue of the previous equation. We define its fluxes

$$\mathbf{F}_i^{(c)} = \frac{u^2}{2}, \quad \mathbf{F}_i^{(v)} = \beta \frac{\partial u}{\partial x_i}, \quad i = 1, \dots, d. \quad (2.4)$$

This equation describes the transport of u with local advective speed equal to itself. Hence, portions with larger values of u will move faster than portions with lower values. With this equation, we generally observe the formation of shocks, especially when the diffusion coefficient β is zero. When $\beta \neq 0$, this equation can mimic turbulent behaviour in a one-dimensional setting by including an appropriate source term s .

2.3 Euler Equations

The Euler equations describe nonlinear behaviour of fluid flow without the effects of viscosity. They consist of equations for conservation of mass, momentum, and energy. These equations can be used to study various phenomena, including shock waves, rarefaction waves, and other compressibility effects in fluids. The conserved variables in three dimensions are

$$\mathbf{u} = \begin{bmatrix} \rho \\ \rho v_x \\ \rho v_y \\ \rho v_z \\ \rho E \end{bmatrix}, \quad (2.5)$$

where ρ is density, ρv_i is the momentum component in the i -th direction, and E is the specific total energy, which we will define shortly. Since Euler equations do not include any dissipative

effects, the viscous flux term is set to zero $\mathbf{F}_i^{(v)} = \mathbf{0}$, while the convective flux is defined

$$\mathbf{F}_i^{(c)} = \begin{bmatrix} \rho v_i \\ \rho v_i v_x + P \delta_{ix} \\ \rho v_i v_y + P \delta_{iy} \\ \rho v_i v_z + P \delta_{iz} \\ v_i (\rho E + P) \end{bmatrix}, \quad (2.6)$$

where P is the static pressure and δ_{ij} is the Kronecker delta. For an ideal gas, the pressure and temperature T are related via the equation of state

$$P = \rho R T, \quad (2.7)$$

where R is the specific gas constant, which can be shown to be related to the specific heat capacities at constant volume c_V and constant pressure c_P as well as the specific heat capacity ratio $\bar{\gamma} = c_P/c_V$ by

$$c_V = \frac{R}{\bar{\gamma} - 1}, \quad c_P = \frac{\bar{\gamma} R}{\bar{\gamma} - 1}. \quad (2.8)$$

In addition, the specific internal energy e can be defined

$$e = c_V T, \quad (2.9)$$

which allows us to compute the specific total energy via

$$E = e + \frac{|\mathbf{v}|^2}{2}, \quad (2.10)$$

where $|\mathbf{v}|$ is the magnitude of the velocity vector. With these equations, the system is closed with five equations and five unknowns. By solving the Euler equations, we can study various phenomena such as shock waves, rarefaction waves, and other nonlinear wave interactions in compressible fluids. Compressibility is governed by the nondimensional Mach number M , defined as

$$M = \frac{|\mathbf{v}|}{c}, \quad (2.11)$$

where c is the speed of sound given by

$$c = \sqrt{\frac{\bar{\gamma} P}{\rho}}. \quad (2.12)$$

For $M < 0.3$, compressibility effects can be considered negligible. The Euler equations do not capture the effects of viscosity, and other dissipative processes, which are significant in many real-world situations. In these cases, we need to resort to the Navier-Stokes equations.

2.4 Navier-Stokes

Industrial applications of computational fluid dynamics involve flows where the effects of viscosity cannot be neglected. These include problems where the influence of boundary layers, shear forces, and thermal gradients are important such as high Reynolds numbers flows. These equations inherit the vector of conservative quantities, convective fluxes, thermodynamic properties, and the Mach number from the Euler equations. Hence, we define $\mathbf{F}^{(c)}$ as in the Euler equations and define viscous flux vectors of the form

$$\mathbf{F}_i^{(v)} = \begin{bmatrix} 0 \\ \bar{\tau}_{xi} \\ \bar{\tau}_{yi} \\ \bar{\tau}_{zi} \\ \bar{\boldsymbol{\tau}}_i \cdot \mathbf{v} + k \frac{\partial T}{\partial x_i} \end{bmatrix}, \quad (2.13)$$

where k is the thermal conductivity coefficient and $\bar{\tau}_{ij}$ is a component of the viscous stress tensor $\bar{\boldsymbol{\tau}}$, which can be written

$$\bar{\tau}_{ij} = \mu \left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} \right) - \lambda_\mu \nabla \cdot \mathbf{v} \delta_{ij}, \quad (2.14)$$

where μ is the dynamic viscosity coefficient and λ_μ is the bulk viscosity coefficient, typically set to $-\frac{2\mu}{3}$.

2.5 Entropically-Damped Artificial Compressibility

As the Mach number decreases, the disparity between the entropy and the acoustic wave speeds in the flow becomes significant. Hence, for these problems, solutions via compressible Navier-Stokes become quite challenging. Artificial compressibility methods (ACM) [91] provide mechanisms toward the divergence-free condition of the velocity field via pseudo-time [92], which reduces the stiffness of the pressure and velocity field toward a divergence-free result. These methods enable explicit time stepping, contrary to the typical solutions of the incompressible Navier-Stokes equations [93], which requires solving a Poisson problem to

obtain the pressure field.

The entropically-damped artificial compressibility (EDAC) method of Clausen [94] achieves closure by minimizing density variations rather than setting a constant entropy constraint, such as in the ACM method. The resulting equation has a pressure diffusion operator. The pressure evolution is hence dictated by

$$\frac{\partial P}{\partial t} + \mathbf{v} \cdot \nabla P + \frac{1}{M^2} \nabla \cdot \mathbf{v} - \frac{1}{\text{Re}} \nabla^2 P = 0. \quad (2.15)$$

This leads to a vector \mathbf{u} of the form

$$\mathbf{u} = \begin{bmatrix} P \\ v_x \\ v_y \\ v_z \end{bmatrix}, \quad (2.16)$$

as well as convective and viscous flux vectors defined by

$$\mathbf{F}_i^{(c)} = \begin{bmatrix} v_i(P + \Theta) \\ v_i v_x + P \delta_{ix} \\ v_i v_y + P \delta_{iy} \\ v_i v_z + P \delta_{iz} \end{bmatrix}, \quad \mathbf{F}_i^{(v)} = \nu \frac{\partial}{\partial x_i} \begin{bmatrix} P \\ v_x \\ v_z \\ v_y \end{bmatrix}, \quad (2.17)$$

where $\Theta = 1/M^2$ with M an artificial Mach number, and $\nu = \frac{1}{\text{Re}}$. These equations have been successfully applied in the context of FR methods [95], showing that they are an effective alternative to the ACM method, providing faster results and simpler implementation. Higher values of Θ introduce stiffness in the problem, but provide a more accurate approximation of the divergence-free condition. The maximum stable time step size is also influenced by this parameter in explicit numerical schemes. We make use of the EDAC equations in Chapter 7, where we take advantage of their explicit form and reduced size to solve a multi-element airfoil problem at a relatively high Reynolds number.

Chapter 3

The Flux Reconstruction Framework

3.1 Preliminaries

Consider the following conservation law

$$\frac{\partial u}{\partial t} + \nabla \cdot \mathbf{F}(u, \nabla u) = 0 \quad \text{in } \Omega, \quad (3.1)$$

where Ω is a bounded subset of \mathbb{R}^d with boundary $\partial\Omega \in \mathbb{R}^{d-1}$ and d dimensions, u is the conserved quantity, $\mathbf{F} = \mathbf{F}(u, \nabla u)$ is the flux and t is time. To discretize this problem, we rewrite it as a system of first-order ordinary differential equations

$$\frac{\partial u}{\partial t} + \nabla \cdot \mathbf{F}(u, \mathbf{q}) = 0, \quad (3.2a)$$

$$\mathbf{q} - \nabla u = 0, \quad (3.2b)$$

where \mathbf{q} is an auxiliary variable referring to the gradient of the conserved variable.

We define \mathcal{T}^h to be the partition of Ω into N nonoverlapping, conforming elements Ω_k , each with boundary $\partial\Omega_k = \{f\}$, where f is an element face. Define also $\partial\mathcal{T}^h$ to be the collection of all element borders such that $\partial\mathcal{T}^h = \{\partial\Omega_k : \Omega_k \in \mathcal{T}^h\}$. Here, every face is counted from the point of view of every element. Furthermore, consider $\varepsilon^h = \varepsilon_\partial^h \cup \varepsilon_0^h = \bar{f}$ to be the union of all boundary and interior faces in the computational domain. Two interior faces from $\partial\mathcal{T}^h$ have a single corresponding face or vice-versa $\bar{f} \in \varepsilon^h$, $\bar{f} = (f \in \partial\Omega_k) \cap \varepsilon^h$. Furthermore, we define the jump operator of a vector \mathfrak{F} and a scalar u to be defined at face f

$$\llbracket \mathfrak{F} \rrbracket = \mathfrak{F}_+ \cdot \mathbf{n}_+ + \mathfrak{F}_- \cdot \mathbf{n}_-, \quad (3.3)$$

$$\llbracket u \rrbracket = u_+ \mathbf{n}_+ + u_- \mathbf{n}_-, \quad (3.4)$$

respectively, and similarly, the averaging operators by

$$\{\{u\}\} = \frac{u_- + u_+}{2}, \quad (3.5)$$

$$\{\{\mathfrak{F}\}\} = \frac{\mathfrak{F}_- + \mathfrak{F}_+}{2}. \quad (3.6)$$

The first step in the implementation is to map each Ω_k to a reference element $\tilde{\Omega}$ and each \bar{f} to a reference face. Then, we can make use of invertible one-to-one mapping functions $\mathcal{M}_k(\tilde{\mathbf{x}})$ to convert quantities between physical and reference space within each element. These are obtained via

$$\mathbf{x} = \mathcal{M}_k(\tilde{\mathbf{x}}) = \sum_{i=1}^{N_g} M_i(\tilde{\mathbf{x}}) \mathbf{x}_i^g, \quad (3.7)$$

where \mathbf{x} is the physical coordinate of a given point in Ω_k , and M_i is a shape function associated with one of the N_g mapping points $\{\mathbf{x}_i^g\}$. Define the Jacobian matrix of these transformations by $\mathbf{J}_k(\tilde{\mathbf{x}})$ and its determinant by $J_k(\tilde{\mathbf{x}})$. These geometric parameters allow us to rewrite the conservation law considering the approximated physical solution $u^h = \cup_{k=1}^N u_k^h$ and flux $\mathbf{F}^h = \cup_{k=1}^N \mathbf{F}_k^h$ in reference space such that

$$\tilde{u}_k^h = \tilde{u}_k^h(\tilde{\mathbf{x}}, t) = J_k u_k^h(\mathcal{M}_k(\tilde{\mathbf{x}}), t), \quad (3.8)$$

$$\tilde{\mathbf{F}}_k^h = \tilde{\mathbf{F}}_k^h(\tilde{\mathbf{x}}, t) = J_k \mathbf{J}_k^{-1} \mathbf{F}_k^h(\mathcal{M}_k(\tilde{\mathbf{x}}), t), \quad (3.9)$$

$$\tilde{\mathbf{q}}_k^h = \tilde{\mathbf{q}}_k^h(\tilde{\mathbf{x}}, t) = \mathbf{J}_k^T \mathbf{q}_k^h(\mathcal{M}_k(\tilde{\mathbf{x}}), t), \quad (3.10)$$

so that the evolution of the physical solution within each element satisfies

$$\frac{\partial u_k^h}{\partial t} + \frac{1}{J_k} \tilde{\nabla} \cdot \tilde{\mathbf{F}}_k^h = 0, \quad (3.11)$$

$$\tilde{\mathbf{q}}_k^h - \tilde{\nabla} u_k^h = 0, \quad (3.12)$$

where $\tilde{\nabla}$ is the divergence operator in reference space.

3.2 Implementation

To obtain a discretization of order $p + 1$, we place N_s solution points $\{\tilde{\mathbf{x}}_i^s\}_{i=1}^{N_s}$ inside each element and $N_{r,f}$ flux points $\{\tilde{\mathbf{x}}_{f,i}^r\}_{i=1}^{N_{r,f}}$ on each of the N_f element faces. A diagram is provided in Figure 3.1 for two neighbouring quadrilateral elements showcasing these points. At each face, we define functions such that outward unit normal vectors can be obtained via $\tilde{\mathbf{n}}_f^m = \tilde{\mathbf{n}}_f(\tilde{\mathbf{x}}_{f,m}^r)$ and $\mathbf{n}_{k,f}^m = \mathbf{n}_{k,f}(\tilde{\mathbf{x}}_{f,m}^r)$ in reference and physical space, respectively.

The conserved variable can be represented within a given element $\tilde{\Omega}_k$ via interpolation of




Figure 3.1. Solution and flux point locations for two neighbouring quadrilateral elements for a $p = 2$ FR method

node values $\{U_{k,i}\}_{i=1}^{N_s}$ with nodal basis functions $\{\varphi_i(\tilde{\mathbf{x}})\}_{i=1}^{N_s}$

$$u_k^h = \sum_{i=1}^{N_s} U_{k,i}(t) \varphi_i(\tilde{\mathbf{x}}), \quad (3.13)$$

resulting in degree- p discontinuous polynomials. Note that a discontinuous flux $\tilde{\mathbf{F}}_k^{hD}$ in the reference element can also be represented using the same basis functions

$$\tilde{\mathbf{F}}_k^{hD} = \sum_{i=1}^{N_s} \tilde{\mathbf{F}}_{k,i}(t) \varphi_i(\tilde{\mathbf{x}}). \quad (3.14)$$

For second-order partial differential equations, we first compute an approximation to the gradient, which we denote \mathbf{q}^h . This is obtained by taking the gradient of a globally continuous scalar variable. In the FR method, this is carried out by performing a reconstruction procedure, which penalizes the interface jumps via correction functions. In other words, we add the following correction term to the discontinuous solution from Equation (3.13)

$$u_k^{hC} = \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \tilde{\mathbf{n}}_f^m \cdot \mathbf{g}_f^m(\tilde{\mathbf{x}}) [\mathfrak{U}_{k,f} - u_{k,f}^h]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_{f,m}^r}, \quad (3.15)$$

with correction functions $\mathbf{g}_f^m(\tilde{\mathbf{x}})$ satisfying

$$\tilde{\mathbf{n}}_f^m \cdot \mathbf{g}_l^n(\tilde{\mathbf{x}}_f^m) = \delta_{fl} \delta_{mn}. \quad (3.16)$$

Here, subscripts indicate the face number, superscripts the node number within the face, and δ is the Kronecker delta. The correction functions we consider in this work define energy-stable

FR methods. They can be defined as follows in one dimension

$$g_L = \frac{(-1)^p}{2} \left[\Psi_p - \frac{\eta_p \Psi_{p-1} + \Psi_{p+1}}{1 + \eta_p} \right], \quad g_R = \frac{1}{2} \left[\Psi_p + \frac{\eta_p \Psi_{p-1} + \Psi_{p+1}}{1 + \eta_p} \right], \quad (3.17)$$

where Ψ_p is a Legendre polynomial of degree p and

$$\eta_p = \frac{c(2p+1)(a_p p!)^2}{2}, \quad (3.18)$$

with a_p the leading coefficient in Ψ_p

$$a_p = \frac{(2p)!}{2^p (p!)^2}, \quad (3.19)$$

and c a free parameter which can recover existing methods. For instance, $c = c_{DG} = 0$ recovers the discontinuous Galerkin method, and c_{SD} recovers the spectral difference method. For completeness, we also consider the c_{HU} method, as defined in [26]. These functions are shown in Figure 3.2. They can be directly extended to tensor product elements [96] and have been identified for simplices in [30, 97]. The specific values of c_{SD} and c_{HU} can be written as a function of the polynomial degree [29]

$$c_{SD} = \frac{2p}{(2p+1)(p+1)(a_p p!)^2}, \quad (3.20)$$

$$c_{HU} = \frac{2(p+1)}{(2p+1)p(a_p p!)^2}. \quad (3.21)$$




Figure 3.2. Left and right correction functions for c_{DG} , c_{SD} and c_{HU} FR $p = 2$ schemes

Equation (3.15) requires a common value of the solution $\mathfrak{U}_{k,f}$, which we will discuss in the next subsections. The auxiliary variable \mathbf{q} is a vector polynomial of degree p that results from taking the gradient of the solution and its correction. For each element, we can write it as

$$\tilde{\mathbf{q}}_k^h = \sum_{i=1}^{N_s} U_{k,i}(t) \tilde{\nabla} \varphi_i(\tilde{\mathbf{x}}) + \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \tilde{\mathbf{n}}_f^m \cdot \tilde{\nabla} \cdot \mathbf{g}_f^m(\tilde{\mathbf{x}}) [\mathfrak{U}_{k,f} - u_{k,f}^h]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_{f,m}^r}, \quad (3.22)$$

and then map it to physical space via Equation (3.10). Following a similar procedure to computing the corrected gradient, we formulate a continuous flux

$$\tilde{\mathbf{F}}_k^h = \tilde{\mathbf{F}}_k^{hD} + \tilde{\mathbf{F}}_k^{hC}, \quad (3.23)$$

where the term $\tilde{\mathbf{F}}_k^{hC}$ uses the same correction function and can be written

$$\tilde{\mathbf{F}}_k^{hC} = \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \mathbf{g}_f^m(\tilde{\mathbf{x}}) \left[\tilde{H}(\tilde{\mathbf{x}})_{k,f} \right]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_{f,m}^r}, \quad (3.24)$$

and the normal jump of the flux at the face is defined as follows

$$\tilde{H}_{k,f}(\tilde{\mathbf{x}}) = \tilde{\mathfrak{F}}_{k,f} \cdot \tilde{\mathbf{n}}_f - \tilde{\mathbf{F}}_{k,f}^{hD} \cdot \tilde{\mathbf{n}}_f. \quad (3.25)$$

$\tilde{\mathbf{F}}_{k,f}^{hD}$ is the transformed discontinuous flux polynomial interpolated to face f , and the relationship between physical and reference space for the common flux is [96]

$$\tilde{\mathfrak{F}}_{k,f} \cdot \tilde{\mathbf{n}}_f = J_{k,f} \hat{\mathfrak{F}}_{k,f} \cdot \mathbf{n}_{k,f}, \quad (3.26)$$

where $J_{k,f}$ is the Jacobian determinant of the face. Note that there is no constraint on applying different correction functions for the gradient and the flux in Equations (3.15) and (3.24), but we choose to make use of the same in this work. There remains to define both the corrected solution and flux interface variables, which lead to different types of discretization. Depending on the choice of the common solution and fluxes, the numerical method will have different stability properties, and the stencil will also vary. These choices are well-documented and lead to the well-known interior penalty (IP) [98], Bassi and Rebay I (BR1) [99] and II (BR2) [100] and the so-called local DG (LDG) [101]. We will refer to these as standard formulations. Other approaches rely on adding a new unknown into the system, which we will refer to as the hybridized formulations. Up to this point, both methods share the same methodology. We now branch out into the formulation of each of these approaches.

3.2.1 Standard Formulations

In standard formulations, the numerical flux at the interface typically takes the following form for the inviscid or convective component

$$\hat{\mathfrak{F}}^{(c)}(u_-, u_+) \cdot \mathbf{n}_- = \frac{1}{2} [\mathbf{F}^{(c)}(u_-) + \mathbf{F}^{(c)}(u_+)] \cdot \mathbf{n}_- + \frac{s^{\text{FR}}}{2}(u_- - u_+), \quad (3.27)$$

where s^{FR} is a stabilization parameter in the standard formulation. This parameter is chosen depending on the physics of the problem. Typical values (or matrices in systems of equations) lead to the Rusanov or local Lax-Friedrichs (LLF) [102, 103], Roe [104], and HLL methods [105]. Throughout this work, we employ the Rusanov/LLF fluxes with $s^{\text{FR}} = \bar{\lambda}^{\text{FR}}$, with $\bar{\lambda}^{\text{FR}}$ the maximum wave speed in the system. The viscous flux is typically computed by evaluating $\mathbf{F}^{(v)}$ using a common gradient $\hat{\mathbf{q}}$ and a common solution \mathfrak{U} , i.e.

$$\hat{\mathfrak{F}}^{(v)}(\mathfrak{U}, \hat{\mathbf{q}}) \cdot \mathbf{n}_- = \mathbf{F}^{(v)}(\mathfrak{U}, \hat{\mathbf{q}}) \cdot \mathbf{n}_-. \quad (3.28)$$

Note that the computation of the total interface flux is the sum of the viscous and inviscid fluxes. At the interface, conservation is strongly enforced in the standard formulation by setting

$$\hat{\mathfrak{F}} \cdot \mathbf{n}_- = -\hat{\mathfrak{F}} \cdot \mathbf{n}_+. \quad (3.29)$$

In this work, we consider the LDG and BR2 formulations, which provide explicit definitions for the values of the common solution and the common gradient at the interface. For the LDG method, these can be computed via

$$\mathfrak{U} = \{\{u\}\} - \zeta \llbracket u \rrbracket, \quad (3.30)$$

$$\hat{\mathbf{q}} = \{\{\mathbf{q}\}\} + \zeta \llbracket \mathbf{q} \rrbracket - \theta \llbracket u \rrbracket, \quad (3.31)$$

where ζ is a scalar value in one dimension, and θ is a penalty term. In the case of the BR2 method

$$\mathfrak{U} = \{\{u\}\}, \quad (3.32)$$

$$\hat{\mathbf{q}} = \{\{\nabla u\}\} + s^{\text{BR2}} \{\{r^f(\llbracket u \rrbracket)\}\}, \quad (3.33)$$

with s^{BR2} a stabilization parameter generally set to unity for flow problems, and r^f is a lifting operator satisfying

$$\int_{\Omega} \mathbf{r}^f(\llbracket u \rrbracket) \cdot \boldsymbol{\psi} d\Omega = \int_f \llbracket u \rrbracket \cdot \{\{\boldsymbol{\psi}\}\} ds. \quad (3.34)$$

Hence, a modified gradient that considers only the effect of the discontinuities at the interface is considered in the BR2 method. In BR2, the common gradient involves a lifting term, which results in a more compact stencil than the LDG method. It is the form of these fluxes that may lead to a certain degree of coupling between elements, decreasing or increasing the sparsity of the implicit systems and resulting in more or less expensive computations [106].

After choosing a suitable interface formulation, we apply the divergence on the corrected flux in Equation (3.23) and sum over all elements, which yields the following system for the standard flux reconstruction method

$$\sum_{\Omega_k \in \mathcal{T}^h} \tilde{\mathbf{q}}_k^h - \sum_{i=1}^{N_s} U_{k,i} \tilde{\nabla} \varphi_i(\tilde{\mathbf{x}}) - \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \tilde{\mathbf{n}}_f^m \cdot \tilde{\nabla} \cdot \mathbf{g}_f^m(\tilde{\mathbf{x}}) [\mathbf{u}_{k,f} - u_{k,f}^h]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_{f,m}^r} = 0, \quad (3.35a)$$

$$\sum_{\Omega_k \in \mathcal{T}^h} \frac{\partial \tilde{u}_k^h}{\partial t} + \sum_{i=1}^{N_s} \tilde{\mathbf{F}}_{k,i} \cdot \tilde{\nabla} \varphi_i(\tilde{\mathbf{x}}) + \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \tilde{\nabla} \cdot \mathbf{g}_f^m(\tilde{\mathbf{x}}) [\tilde{H}(\tilde{\mathbf{x}})_{k,f}]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_{f,m}^r} = 0. \quad (3.35b)$$

Since the definitions of the interface variables are explicitly defined, this system can directly be solved both implicitly and explicitly. Similar to this methodology, we now demonstrate the steps for the hybridized method. In our implicit implementation, the nonlinear system is formed for Equation (3.35b). Equation (3.35a) is added via the chain rule.

3.2.2 Hybridized Formulations

In addition to the conserved variable, we introduce an approximation to u^h on the skeleton of the computational grid such that at any face $\bar{f} \in \varepsilon^h$, a degree- p polynomial can be obtained via

$$\hat{u}_{\bar{f}}^h(\tilde{\mathbf{x}}, t) = \sum_{i=1}^{N_{r,f}} \hat{U}_{\bar{f},i}(t) \phi_i(\tilde{\mathbf{x}}), \quad (3.36)$$

which is the so-called trace variable. Here, $N_{r,f}$ is the number of flux points in face f , which we consider equal to the number of trace points at a given face, and ϕ is a $d - 1$ -dimensional trace basis function. Furthermore, the common flux can be computed by adding contributions from the convective and diffusive components. Then, we can write

$$\hat{\mathfrak{F}}_{k,f} = \hat{\mathfrak{F}}_{k,f}^{(c)} + \hat{\mathfrak{F}}_{k,f}^{(v)}. \quad (3.37)$$

Hybridization is then achieved by considering the following form of the common fluxes

$$\hat{\mathfrak{F}}_{k,f}^{(c)} = \mathbf{F}^{(c)}(\hat{u}_{\bar{f}}^h) + s^{(c)}(u_{k,f}^h - \hat{u}_{\bar{f}}^h) \mathbf{n}_{k,f}, \quad (3.38)$$

$$\hat{\mathfrak{F}}_{k,f}^{(v)} = \mathbf{F}^{(v)}(\hat{u}_{\bar{f}}^h, \mathbf{q}_{k,f}^h) + s^{(v)}(u_{k,f}^h - \hat{u}_{\bar{f}}^h) \mathbf{n}_{k,f}, \quad (3.39)$$

where $s^{(c)}$ and $s^{(v)}$ are convective and viscous stabilization parameters or matrices. Note that with this form of fluxes, the definition of the common flux is given for each element. These fluxes depend on information from a single element and the corresponding trace. Contrary to the standard FR implementation, we do not strongly enforce conservation at the interface since $\hat{\mathfrak{F}}_-$ need not be equal to $\hat{\mathfrak{F}}_+$ in the general case. Hence, we seek to enforce an additional statement satisfying global conservation, i.e.,

$$[[\hat{\mathfrak{F}}]]_{\varepsilon^h} = 0, \quad (3.40)$$

which can be discretely written

$$\sum_{\bar{f} \in \varepsilon_0^h} \int_{\bar{f}} [[\hat{\mathfrak{F}}]]_{\bar{f}} \phi ds + \sum_{\bar{f} \in \varepsilon_\partial^h} \int_{\bar{f}} \mathfrak{F}_{\bar{f}}^{\text{BC}} \phi ds = 0, \quad (3.41)$$

and provides closure to the system. In these equations, we have separated the interior and normal boundary fluxes, the latter of which we denote $\mathfrak{F}_{\bar{f}}^{\text{BC}}$, and the jump operator is defined as in (3.4). After summing over all elements, we can state the hybridized form of the flux reconstruction approach for convection-diffusion type problems as follows

$$\sum_{\Omega_k \in \mathcal{T}^h} \tilde{\mathbf{q}}_k^h - \sum_{i=1}^{N_s} U_{k,i} \tilde{\nabla} \varphi_i(\tilde{\mathbf{x}}) - \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \tilde{\mathbf{n}}_f^m \cdot \tilde{\nabla} \cdot \mathbf{g}_f^m(\tilde{\mathbf{x}}) [\mathfrak{U}_{k,f} - u_{k,f}^h]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_{f,m}^r} = 0, \quad (3.42a)$$

$$\sum_{\Omega_k \in \mathcal{T}^h} \frac{\partial \tilde{u}_k^h}{\partial t} + \sum_{i=1}^{N_s} \tilde{\mathbf{F}}_{k,i} \cdot \tilde{\nabla} \varphi_i(\tilde{\mathbf{x}}) + \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \tilde{\nabla} \cdot \mathbf{g}_f^m(\tilde{\mathbf{x}}) [\tilde{H}(\tilde{\mathbf{x}})_{k,f}]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_{f,m}^r} = 0, \quad (3.42b)$$

$$\sum_{\bar{f} \in \varepsilon_0^h} \int_{\bar{f}} [[\hat{\mathfrak{F}}]]_{\bar{f}} \phi ds + \sum_{\bar{f} \in \varepsilon_\partial^h} \int_{\bar{f}} \mathfrak{F}_{\bar{f}}^{\text{BC}} \phi ds = 0, \quad (3.42c)$$

where we have readily taken the divergence of the flux and its correction to arrive at (3.42b). Typically, hybridized methods make use of discontinuous or globally continuous function spaces for the trace variable, which can be respectively defined by

$$\mathbb{M}_p^h = \{\phi \in L_2(\varepsilon^h) : \phi|_{\bar{f}} \in \mathbb{P}^p(\bar{f}), \forall \bar{f} \in \varepsilon^h\}, \quad (3.43a)$$

$$\bar{\mathbb{M}}_p^h = \mathbb{M}_p^h \cap C^0(\varepsilon^h). \quad (3.43b)$$

The choice of nodal basis functions for the trace variable leads to different types of hybridizations. In this work, we consider two types of hybridization. First, we consider discontinuous trace nodal basis functions. This leads to the hybridized flux reconstruction (HFR) method. We also define globally continuous nodal basis functions on interior faces ε_0^h and discontinuous on boundary faces ε_∂^h . This leads to the interior-embedded flux reconstruction (IEFR) scheme,

which we will denote EFR for simplicity. A diagram of the resulting solution and trace points is shown for two neighbouring elements Ω_- and Ω_+ in Figure 3.3 for a $p = 2$ scheme, i.e., a scheme representing the solution using a polynomial of degree 2. An important feature of HFR methods over EFR is that HFR methods are pointwise conservative, as opposed to EFR methods, which are only conservative on the dual volumes [62]. This is a direct consequence of the function space choice and will be further discussed later in this thesis. Note that the EFR method is an HFR scheme with reduced function space for the trace. Hence, on some occasions, we use the term HFR as a universal term to refer to all hybridized FR methods. Furthermore, throughout this thesis, we typically choose the solution points inside the element to be those of the Gauss-Legendre quadrature. In the case of EFR, Gauss-Lobatto-Legendre (GLL) points are used at the faces. When different sets of points are used in the EFR method for the volume and faces, the correction functions are still generated using a tensor-product formulation, and the lifting operation includes an additional interpolation operator. See [32].




Figure 3.3. Trace variable location for two neighbouring quadrilateral elements in an HFR (left) and EFR (right) discretization considering a $p = 2$ scheme

3.2.2.1 The Global System

Hybridized unsteady problems can be written

$$\frac{\partial \mathbf{u}}{\partial t} + f(\mathbf{u}, \hat{\mathbf{u}}) = 0, \quad (3.44a)$$

$$g(\mathbf{u}, \hat{\mathbf{u}}) = 0, \quad (3.44b)$$

where $\mathbf{u} \in \mathbb{R}^{N \times N_s}$, f is the hybridized spatial discretization function associated with Equation (3.42b), and g is the residual associated with the flux conservation statement in Equation (3.42c) [107]. Note that Equation (3.42a) is added via the chain rule in the flux terms, as will be discussed in the implementation in the following section. Here, we have dropped the

superscript h for simplicity in the notation. In conjunction with high-order spatial methods, we can use s -stage singly-diagonal implicit Runge Kutta (SDIRK) schemes [108], for which the solution at every stage can be written as a linear combination of the off-diagonal and diagonal coefficients of the associated Butcher tableau, which has for general form for all RK methods

$$\begin{array}{c|ccc} & c_1 & a_{1,1} & \cdots & a_{1,s} \\ \mathbf{c} & \mathbf{A} & \vdots & \ddots & \vdots \\ \hline & \mathbf{b} & c_s & a_{s,1} & \cdots & a_{s,s} \\ \hline & & b_1 & \cdots & b_s \end{array} \quad (3.45)$$

SDIRK time integration methods can be efficiently implemented as they only need to be preconditioned once per iteration since they have a constant diagonal in their Butcher tableaux and are lower diagonal, i.e., $a_{i,j} = 0$, $j > i$. In the majority of this thesis, unless otherwise indicated, we employ the two-stage SDIRK method for implicit time stepping, which can be written

$$\begin{array}{c|cc} \mathbf{c} & \mathbf{A} & \gamma \\ \hline & \mathbf{b} & 1 \quad \begin{array}{cc} \gamma & \\ 1 - \gamma & \gamma \end{array} \\ \hline & & 1 - \gamma \quad \gamma \end{array} \quad (3.46)$$

with $\gamma = 1 - \frac{\sqrt{2}}{2}$. We seek to advance the current solution u^t by a time step Δt to obtain u^{t+1} . This requires computing a number of intermediate stages. For SDIRK methods, the solution of the hybridized problem at the i -th stage is given by the system

$$\mathbf{u}^i - \mathbf{u}^t - \sum_{j=1}^{i-1} a_{ij} \Delta t f(\mathbf{u}^j, \hat{\mathbf{u}}^j) - a_{ii} \Delta t f(\mathbf{u}^i, \hat{\mathbf{u}}^i) = 0, \quad (3.47a)$$

$$g(\mathbf{u}^i, \hat{\mathbf{u}}^i) = 0. \quad (3.47b)$$

Here, information associated with the diagonal coefficients is unknown, and with the off-diagonal coefficients is known. We can rewrite the system in a simplified manner by introducing

$$\mathbf{u}^{t*,i} = \mathbf{u}^t + \sum_{j=1}^{i-1} a_{ij} \Delta t f(\mathbf{u}^j, \hat{\mathbf{u}}^j), \quad (3.48)$$

which yields

$$\mathbf{u}^i - \mathbf{u}^{t*,i} - a_{ii} \Delta t f(\mathbf{u}^i, \hat{\mathbf{u}}^i) = 0, \quad (3.49a)$$

$$g(\mathbf{u}^i, \hat{\mathbf{u}}^i) = 0. \quad (3.49b)$$

Finally, the solution at the next time step is found from the computed values of the solution and trace at the intermediate stages

$$\mathbf{u}^{t+1} = \mathbf{u}^t + \Delta t \sum_{j=1}^s b_j f(\mathbf{u}^j, \hat{\mathbf{u}}^j). \quad (3.50)$$

Now, the system resulting from Equation (3.49) is larger than a typical FR method, which can be formulated similarly by omitting Equation (3.47b) from the Newton method. However, after hybridization, we can solve instead a reduced problem via static condensation, which we discuss next.

3.2.2.2 Static Condensation

Static condensation is a dimensionality reduction approach typically used in finite-element problems [109, 110]. After hybridization, FR methods become amenable to this approach, as we can express the interior solution as a function of the trace variable. The system (3.49) can be written for a general nonlinear problem

$$h(\mathbf{u}, \hat{\mathbf{u}}) = 0, \quad (3.51a)$$

$$g(\mathbf{u}, \hat{\mathbf{u}}) = 0, \quad (3.51b)$$

which we linearize about the interior and trace solutions $\mathbf{u}^t, \hat{\mathbf{u}}^t$ to obtain a linear system for the n -th Newton iteration

$$\begin{bmatrix} \mathbf{A}^n & \mathbf{B}^n \\ \mathbf{C}^n & \mathbf{D}^n \end{bmatrix} \begin{bmatrix} \delta \mathbf{u}^n \\ \delta \hat{\mathbf{u}}^n \end{bmatrix} = \begin{bmatrix} \mathbf{r}^n \\ \mathbf{s}^n \end{bmatrix}, \quad (3.52)$$

where $\delta \mathbf{u}^n, \delta \hat{\mathbf{u}}^n$ refer to the update vectors for the solution and trace variable at the n -th Newton iteration. Due to the discontinuous nature of the interior solution and the direct interelement decoupling resulting from the hybridized numerical fluxes, \mathbf{A} is block-diagonal, with each block associated with an element in \mathcal{T}^h . Consequently, it is possible to compute its inverse locally at each element and rewrite the system via static condensation. To do this, we eliminate the interior solution \mathbf{u} from Equation (3.52) by substituting its definition in the first row of the system into the second row. Finally, we obtain a reduced system that depends only on the trace variable, i.e.

$$\mathbf{L}^n \delta \hat{\mathbf{u}}^n = \mathbf{t}^n, \quad (3.53)$$

where $\mathbf{L}^n = \mathbf{D}^n - \mathbf{C}^n(\mathbf{A}^n)^{-1}\mathbf{B}^n$ and $\mathbf{t}^n = \mathbf{s}^n - \mathbf{C}^n(\mathbf{A}^n)^{-1}\mathbf{r}^n$. Then, the solution can be obtained from

$$\delta\mathbf{u}^n = (\mathbf{A}^n)^{-1}(\mathbf{r}^n - \mathbf{B}^n\delta\hat{\mathbf{u}}^n). \quad (3.54)$$

These matrices can be built efficiently on a per-element basis and then globally assembled, allowing the solution to be recovered instead for each element Ω_k via

$$\delta\mathbf{u}_k^n = (\mathbf{A}_k^n)^{-1}(\mathbf{r}_k^n - \mathbf{B}_k^n\delta\hat{\mathbf{u}}_k^n), \quad (3.55)$$

which contains operators of smaller sizes and are local to each element. These local solves scale linearly in parallel implementations since they are local operations to every element. Dropping the subscript n for simplicity, the global operators can be formulated by

$$\mathbf{L}_{ij} = \mathbf{L}_{ij} + \mathbf{L}_{k,\bar{i}\bar{j}}, \quad (3.56)$$

$$\mathbf{t}_i = \mathbf{t}_i + \mathbf{t}_{k,\bar{i}}, \quad (3.57)$$

where the elemental matrices \mathbf{L}_k and \mathbf{t}_k are defined by

$$\mathbf{L}_k := \mathbf{D}_k - \mathbf{C}_k\mathbf{A}_k^{-1}\mathbf{B}_k, \quad (3.58)$$

$$\mathbf{t}_k := \mathbf{s}_k - \mathbf{C}_k\mathbf{A}_k^{-1}\mathbf{r}_k, \quad (3.59)$$

and the indices \bar{i}, \bar{j} are associated with a many-to-one mapping of the element's flux points to the global trace points. The elemental blocks $\mathbf{A}^k, \mathbf{B}^k$ can be specifically defined as follows

$$\mathbf{A}_{k,ij} = \delta_{ij} - \Delta ta_{ii} \frac{1}{J_{k,i}} \left[\sum_{g=1}^{N_s} \tilde{\nabla} \varphi_g(\tilde{\mathbf{x}}) \cdot \frac{\partial \tilde{\mathbf{F}}_{k,g}}{\partial u_{k,j}} + \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \tilde{\nabla} \cdot \mathbf{g}_f^m(\tilde{\mathbf{x}}) \frac{\partial \tilde{H}_{k,f}}{\partial u_{k,j}}(\tilde{\mathbf{x}}_{f,m}^r) \right]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_i^s}, \quad (3.60a)$$

$$\mathbf{B}_{k,il} = \frac{\Delta ta_{ii}}{J_{k,i}} \left[\sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}} \tilde{\nabla} \cdot \mathbf{g}_f^m(\tilde{\mathbf{x}}) \frac{\partial \tilde{H}_{k,f}}{\partial \hat{u}_{k,l}}(\tilde{\mathbf{x}}_{f,m}^r) \right]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_i^s}, \quad (3.60b)$$

and the remaining blocks can be further structured face-by-face within each element and then assembled into a single elemental block. This way, we can easily accommodate faces with different numbers of trace points, which will arise in the p -adaptive formulation of Chapter 6. Note that p -adaptation is only considered in two dimensions in this thesis. Further challenges appear for EFR methods in three-dimensions but are not addressed in this work. That is, \mathbf{C}_k

and \mathbf{D}_k can be trivially assembled from

$$\bar{\mathbf{C}}_{k,qj}^f = \frac{\partial \tilde{\mathfrak{F}}_{k,f}}{\partial u_{k,j}}(\tilde{\mathbf{x}}_{f,q}^r) \cdot \tilde{\mathbf{n}}_f^q, \quad (3.61a)$$

$$\bar{\mathbf{D}}_{k,qt}^f = \frac{\partial \tilde{\mathfrak{F}}_{k,f}}{\partial \hat{u}_{k,f,t}}(\tilde{\mathbf{x}}_{f,q}^r) \cdot \tilde{\mathbf{n}}_f^q, \quad (3.61b)$$

such that elementwise matrices

$$\mathbf{C}_k = (\mathbf{M}_k^0 \bar{\mathbf{C}}_k^0, \dots, \mathbf{M}_k^{N_f} \bar{\mathbf{C}}_k^{N_f})^T, \quad (3.62a)$$

$$\mathbf{D}_k = \text{diag}(\mathbf{M}_k^0 \bar{\mathbf{D}}_k^0, \dots, \mathbf{M}_k^{N_f} \bar{\mathbf{D}}_k^{N_f}), \quad (3.62b)$$

are obtained. The indices in Equations (3.60) and (3.61) are defined by $i, j = 1, \dots, N_s$, $l = 1, \dots, N_r$ and $q, t = 1, \dots, N_{r,f}$, $N_r = \sum_f^{N_f} N_{r,f}$. Note that $\hat{u}_{k,l}$ refers to the trace living on the l -th point of element Ω_k , and $\hat{u}_{k,f,t}$ is the trace at the t -th point of the f -th face in element Ω_k . In addition, \mathbf{M}_k^f is the face mass matrix $\mathbf{M}_{k,qt}^f = \int_{\bar{f}} \phi_q \phi_t d\bar{f}$. The contributions from the viscous gradients are merged into the flux Jacobians via the chain rule

$$\frac{\partial \mathbf{F}}{\partial \mathbf{u}} = \frac{\partial \mathbf{F}}{\partial \mathbf{q}} \frac{\partial \mathbf{q}}{\partial \mathbf{u}}, \quad (3.63)$$

$$\frac{\partial \mathbf{F}}{\partial \hat{\mathbf{u}}} = \frac{\partial \mathbf{F}}{\partial \mathbf{q}} \frac{\partial \mathbf{q}}{\partial \hat{\mathbf{u}}}. \quad (3.64)$$

Similarly, the vectors \mathbf{r}_k , \mathbf{s}_k evaluate the right-hand side functions in the Newton algorithm

$$\mathbf{r}_k = -h(\mathbf{u}_k, \hat{\mathbf{u}}_k), \quad (3.65a)$$

$$\mathbf{s}_k = -g(\mathbf{u}_k, \hat{\mathbf{u}}_k). \quad (3.65b)$$

In order to solve Equation (3.53), we employ the generalized minimum residual (GMRES) iterative solver with restricted additive Schwarz (RAS) preconditioning [111] in PETSc [112] following [113]. Assembling the Jacobian matrix can be expensive for large problems, especially at high orders. We consider instead a quasi-Newton method where the Jacobian matrix is frozen for a given number of iterations to balance the time spent on the linear solver and the assembly of the Jacobian blocks. Appendix C discusses some implementation details for parallel computations.

To see the benefits of using hybridized methods, assume an infinitely large structured mesh of quadrilateral elements in two dimensions and hexahedra in three dimensions. Nguyen

et al. [63] estimated the number of degrees of freedom associated with the global system by

$$N_{\text{DOF}} = N_v l \alpha_{\text{DOF}}, \quad (3.66)$$

where N_v is the number of mesh nodes, l is the number of conserved variables or unknowns, and α_{DOF} can be found in Tables 3.1-3.2. In addition, due to the sparse nature of the Jacobian matrices, the number of nonzeros NNZ is associated with the memory storage requirements and general performance of the linear solver and can be estimated by

$$\text{NNZ} = N_v l^2 \alpha_{\text{NNZ}}, \quad (3.67)$$

where α_{NNZ} can be obtained from Tables 3.1-3.2 for the aforementioned elements [114]. These values have been used to describe the performance improvement potential of these methods and can be found for simplex elements in [63]. In the presence of boundary conditions, the values in Tables 3.1-3.2 should be preceded by a $<$ operator for EFR since boundary trace points do not take part in the size of the \mathbf{L} matrix. From these tables, we can see the advantages of using hybridized methods, especially at high orders. For a low-order method such as $p = 1$, HFR and FR have the same N_{DOF} in the linear solve, which can make HFR more expensive due to the overhead computations of the local Equations (3.54). In three dimensions, the EFR method shows benefits for all polynomial degrees, but HFR only for $p \geq 3$.

Table 3.1. Coefficients α_{DOF} and α_{NNZ} for quadrilateral elements

p	α_{DOF}						α_{NNZ}					
	1	2	3	4	5	6	1	2	3	4	5	6
FR	4	9	16	25	36	49	48	189	512	1125	2160	3773
HFR	4	6	8	10	12	14	56	126	224	350	504	686
EFR	1	3	5	7	9	11	9	47	113	207	329	479

Table 3.2. Coefficients α_{DOF} and α_{NNZ} for hexahedral elements

p	α_{DOF}						α_{NNZ}					
	1	2	3	4	5	6	1	2	3	4	5	6
FR	8	27	64	125	216	343	256	2187	10240	34375	93312	218491
HFR	12	27	48	75	108	147	528	2673	8448	20625	42768	79233
EFR	1	7	19	37	61	91	27	459	2415	7803	19323	40467

Chapter 4

Optimal Explicit Runge-Kutta Formulations

4.1 Overview

Conventional explicit Runge-Kutta methods have been the go-to schemes in high-order implementations of the FR method. Typical RK methods are only suitable for nonstiff problems [115, 116]. For solutions of ordinary differential equations, the time step must be chosen in accordance with their maximum stable CFL condition, which is dictated by a stability function associated with the time-stepping scheme. This is so that the scaled eigenspectrum of the spatial discretization lies inside this stability region [45].

This chapter demonstrates the computational benefits and numerical analysis of optimal Runge-Kutta methods for explicit formulations of the standard FR discretization. The work described here is partially the result of a collaboration with Siavash Hedayati. The semi and fully-discrete analysis, as well as optimal RK polynomials for one dimension, were undertaken by Pereira. The computation and validation of optimal RK polynomials for multidimensions were undertaken by Hedayati. A summarized version is presented in this chapter. Further details can be found in [117, 118].

4.2 Stability Polynomials

Explicit RK schemes take the value of the solution at the current time t and predict its value after a given time-step size Δt using intermediate stages. This can be expressed for equations of the form

$$u' = \lambda u, \tag{4.1}$$

as

$$u(t + \Delta t) = P_{s,q}(z)u(t), \quad (4.2)$$

where $P_{s,q}(z)$ is the characteristic degree s polynomial of the RK scheme, which is of order q . The stability of RK temporal schemes depends on the eigenvalues λ^h of the system of equations to solve scaled by the time-step size Δt , i.e., $z = \Delta t \lambda^h$. More specifically, these eigenvalues must lie within the region of absolute stability S

$$\Delta t \lambda^h \subseteq S, \quad (4.3)$$

where S is defined as

$$S = \{z \in \mathbb{C} : |P_{s,q}| \leq 1\}, \quad (4.4)$$

and $|\cdot|$ is the complex modulus operator. Generally, a stability polynomial is given by the monomial expansion

$$P_{s,q}(z) = \sum_{j=0}^s \gamma_j z^j, \quad (4.5)$$

where $\{\gamma_0, \dots, \gamma_s\}$ are the aforementioned polynomial coefficients. These coefficients are required to be [119, 45]

$$\gamma_j = \frac{1}{j!}, \quad j = 0, \dots, q, \quad (4.6)$$

to guarantee a temporal scheme of, at least, order q . The remaining $s - q$ coefficients are free parameters, which can be optimized to more accurately fit the eigenvalues of the system of ODEs. Hence, we can significantly increase the time-step size by adding new coefficients to $P_{s,q}(z)$. The result is an optimal explicit Runge-Kutta (OERK) method.

Generating an optimal ERK stability polynomial is generally cast as an optimization problem. We seek the set of coefficients $\{\gamma_{q+1}, \dots, \gamma_s\}$ that yields the maximum stable time-step size for a given semidiscretization. In other words

$$\begin{aligned} & \underset{\gamma_{q+1}, \dots, \gamma_s}{\text{maximize}} && \Delta t \\ & \text{subject to} && |P_{s,q}(\Delta t \lambda^h)| - 1 \leq 0, \quad \forall \lambda^h. \end{aligned} \quad (4.7)$$

Finding these optimal free parameters can be challenging due to the nonconvex characteristic of the problem for $s > 2$, which may lead to finding sub-optimal local minima [45, 119]. Ketcheson et al. [45] proposed recasting the problem in Equation (4.7) into a generalization of the classical least absolute problem, in which we look for how small the maximum modulus

of $P_{s,q}(z)$ can be for the set of $\boldsymbol{\lambda}^h$ at a given time-step size Δt

$$\underset{\gamma_{q+1}, \dots, \gamma_s}{\text{minimize}} \quad \max_{\boldsymbol{\lambda}^h \in \boldsymbol{\Lambda}^h} (|P_{s,q}(\Delta t \boldsymbol{\lambda}^h)| - 1). \quad (4.8)$$

Denote the solution of this problem as $r_{s,q}(\Delta t \boldsymbol{\lambda}^h)$. We can evaluate the minimization problem using the bisection method to find the maximum stable time-step size, which we will refer to as Δt_{opt} , such that

$$\begin{aligned} & \underset{\gamma_{q+1}, \gamma_{q+2}, \dots, \gamma_s}{\text{maximize}} \quad \Delta t \\ & \text{subject to} \quad r_{s,q}(\Delta t, \boldsymbol{\lambda}^h) \leq 0. \end{aligned} \quad (4.9)$$

Using Algorithm 1 [46], we are able to solve Equation (4.9) to find Δt_{opt} . We solve the

Algorithm 1: Finding Δt_{opt} [119]

```

Select  $\Delta t_{max}$ 
 $\Delta t_{min} = 0$ 
while  $\Delta t_{max} - \Delta t_{min} > \epsilon$  do
     $\Delta t = \frac{\Delta t_{min} + \Delta t_{max}}{2}$ 
    Solve Equation (4.8)
    if  $r_{s,q}(\Delta t, \boldsymbol{\lambda}^h) \leq 0$  then
        |  $\Delta t_{min} = \Delta t$ 
    else
        |  $\Delta t_{max} = \Delta t$ 
return  $\Delta t_{opt} = \Delta t_{min}$ 

```

optimization problem in Equation (4.8) using the package for modelling convex optimization problems CVXPY 1.0.24 by Diamond and Boyd [120, 121]. We have converged our results using a tolerance ϵ of 10^{-12} . Our results agree well with those provided by Vermeire [119] for the second-order RK polynomials. Butcher tableaus of a consistent number of stages can be generated via the stability polynomial of the RK scheme.

4.2.1 Semidiscrete Analysis

The spectral properties of high-order semidiscretizations have been thoroughly studied in the context of linear advection via von Neumann analysis in one [122, 31] and higher dimensions [123]. Fully-discrete analyses have shown the effects of the temporal schemes on spatial discretizations. Yang et al. [52] analyzed different first and second-order RK methods on DG discretizations. Vermeire et al. [53] investigated the properties of conventional explicit and implicit RK methods in the context of ILES for FR_{DG} schemes and ESFR schemes [124].

In order to generate the aforementioned optimal stability polynomials of d -dimensional spatial FR schemes, we perform von Neumann analysis to obtain the eigenvalues $\boldsymbol{\lambda}^h$ of the

semidiscretization. The analysis consists of evaluating a scheme's wave propagation properties by considering the d -dimensional linear advection equation

$$\frac{\partial u}{\partial t} + \nabla \cdot (\boldsymbol{\alpha}u) = 0, \quad (4.10)$$

where $u = u(\mathbf{x}, t)$ is the scalar solution variable, $\boldsymbol{\alpha}u$ is the linear advection flux with $\boldsymbol{\alpha}$ the advection velocity, t is time and \mathbf{x} is the spatial coordinate. The advection direction is defined by the angles θ_1, θ_2 , with $\theta_2 = 0$ in the case of $d = 2$ (see Figure 4.1). We divide a




Figure 4.1. Vector decomposition reference

computational domain Ω into N elements. We consider triangular and quadrilateral element types for $d = 2$ and hexahedral, prismatic, and tetrahedral elements for $d = 3$. This division must be performed to ensure that Ω is periodic in all directions. In this sense, N_{ec} nonperiodic elements are agglomerated into a referential periodical element with edge length h , as shown in Figure 4.2. We consider h to be unity for all edges.




Figure 4.2. Subdivision of element types for $d = 2$ and $d = 3$

Next, we seek plane wave solutions of the form

$$u(\mathbf{x}, t) = e^{I(\boldsymbol{\kappa} \cdot \mathbf{x} - \omega t)}, \quad (4.11)$$

where $\boldsymbol{\kappa} = |\boldsymbol{\kappa}|(\beta_1, \beta_2)$ defines the wavenumber, ω is the frequency of the wave and $I = \sqrt{-1}$ is the imaginary number. We consider $|\boldsymbol{\alpha}| = 1$ and the exact dispersion relation is simply $\omega = |\boldsymbol{\kappa}|$. After a simple projection, we may write for any given element

$$\mathbf{u}_i^h(\mathbf{x}, t) = e^{I(\boldsymbol{\kappa} \cdot \mathbf{x}_i - \omega^h t)} \bar{\mathbf{u}}, \quad (4.12)$$

where $\bar{\mathbf{u}}$ is an unknown vector that contains the amplitudes of the wave in numerical space, and \mathbf{x}_i is the center coordinate of the element. Note that the solution within any given element is the same, and the choice of the subscript $i \in [1, N]$ is arbitrary. After discretizing the linear advection equation using FR with an upwind Riemann flux and the c_{DG} correction parameter, we can rewrite the conservation law at the discrete level

$$\frac{d\mathbf{u}^h}{dt} = \mathbf{L}\mathbf{u}^h, \quad (4.13)$$

where \mathbf{L} is the sparse semidiscrete matrix of dimension $N_{ec}N_s \times N_{ec}N_s$, which depends on the wavevector $\boldsymbol{\kappa}$, advection direction, c parameters, Riemann flux choice, and element type. The matrix \mathbf{L} contains the contributions from the neighbouring elements and can be written in general for any d number of dimensions

$$\mathbf{L} = \left[\mathbf{C}^0 + \sum_{i=1}^d \left(\mathbf{C}_u^i e^{-I\boldsymbol{\kappa} \cdot \bar{\mathbf{x}}^i} + \mathbf{C}_d^i e^{+I\boldsymbol{\kappa} \cdot \bar{\mathbf{x}}^i} \right) \right], \quad (4.14)$$

where $\bar{\mathbf{x}}^i$, $i = 1, \dots, d$ is a unit vector, \mathbf{C}^0 is a matrix associated with the element in study, and \mathbf{C}_u^i , \mathbf{C}_d^i are respectively the upstream element and the downstream element Jacobian matrices along the i -th direction. Note that the second term inside the summation in Equation (4.14) vanishes when an upwind Riemann flux is implemented. After substituting Equation (4.12) into (4.13), we obtain

$$-I\omega^h \bar{\mathbf{u}} = \mathbf{L}\bar{\mathbf{u}}, \quad (4.15)$$

which is clearly a classical eigenvalue problem with $N_{ec} \times N_s$ eigenvalues $\boldsymbol{\lambda}^h \in \mathbb{C}$. Then, these eigenvalues are related to numerical frequencies by

$$\omega^h = I\boldsymbol{\lambda}^h. \quad (4.16)$$

The spectrum of the eigenvalues considering all wavenumbers and wavevector orientations define the stability properties of the spatial discretization and are scaled by the time-step size to fit within the stability regions of the explicit temporal schemes. In addition, the imaginary part of the numerical frequencies must be nonpositive to ensure boundedness [123, 125, 31]. Element types with larger eigenspectra typically require smaller time-step sizes, similar to the effects of increasing the polynomial degree [119]. However, this condition is also influenced by the number of solution points $N_{ec}N_s$ within the referential element. Figure 4.3 displays the collection of eigenvalues for all considered element types using a solution polynomial degree $p = 4$ spatial discretization for resolvable wavenumbers and orientations. We note that using even double precision in the computation of eigenvalues may yield numerical frequencies with positive spurious real components on the order of machine precision. As a consequence, these

values were removed from the results in the computation of the optimal stability polynomials, discussed in the following sections.

4.2.2 Optimal Stability Polynomials

4.2.2.1 One-Dimensional Elements

By solving the optimization problem described in Section 4.2, it is possible to generate optimal RK polynomials for schemes of arbitrary order. In this work, we present results for spatial schemes of order $p = 1$ to $p = 8$. Two sets of parameters are obtained using the optimization algorithm, namely the optimal time-step Δt_{opt} and the coefficients $\gamma_{q+1}, \dots, \gamma_s$ for each $P_{s,q}$ scheme. Figure 4.4 shows the maximum stable step sizes that can be used for schemes of degrees 1, 2, 4, and 8. Results are shown for polynomials with $s = q$ to $s = 16$. It is clear that Δt_{opt} increases considerably with the first added coefficient, after which the growth becomes relatively monotonic. It can be argued that this linear increase is balanced with a larger number of residual evaluations. In Figure 4.5, we show the time step size divided by the number of residual evaluations, which we assume equal to s . It is clear that the low-order temporal schemes $P_{s,1}, P_{s,2}$ only show improvement until $s = 4$. After this point, the benefit of optimizing the RK polynomial may be lost due to the large number of required residual evaluations. However, higher-order schemes $P_{s,4}, P_{s,8}$ generally show improvement in computational cost, at least until the maximum degree considered, $s = 16$. The use of OERK schemes with a large number of optimized coefficients can bring additional advantages. For instance, Vermeire proposed the use of Paired-Explicit RK schemes (P-ERK) to reduce computational cost by pairing different ERK schemes based on the distribution of numerical stiffness across the computational domain [119].

The second set of parameters regards the coefficients of the optimal RK polynomials. Classical ERK stability regions can be seen in Figure 4.6 for temporal schemes of order 1, 2, 4, and 8. The eigenvalues of a $p = 4$ FR spatial discretization are scaled by Δt_{opt} and shown as black dots on top of the stability contours. We note that the circular shape of the $P_{1,1}$ scheme cannot fit the eigenvalues of any FR discretization of $p \geq 1$ [126]. Hence, this low-order scheme is unstable for higher-order advection problems. We observe a closer representation of the semidiscretization for the classical high-order temporal schemes. As a consequence, they allow larger Δt_{opt} . In contrast, Figure 4.7 shows the OERK schemes with $s = 16$. These optimal RK polynomials contain $s - q$ additional coefficients relative to their classical analogues. Hence, the eigenspectra of the semidiscrete system are more closely represented by $P(z)$. It is important to note that due to the stiff, flat region of eigenvalues near the imaginary axis, the optimization for both the first and second-order temporal schemes moves in the same direction. As a result, their coefficients are very close. Due to the less stiff shape







Figure 4.3. Collection of eigenvalues for a fourth-order spatial discretization for two and three-dimensional elements




(a) $q = 1$



(b) $q = 2$



(c) $q = 4$



(d) $q = 8$

Figure 4.4. Optimal time-step size Δt_{opt} for OERK schemes from Algorithm 1 with $p = 1 - 8$ spatial discretizations showing a larger Δt_{opt} with increasing s

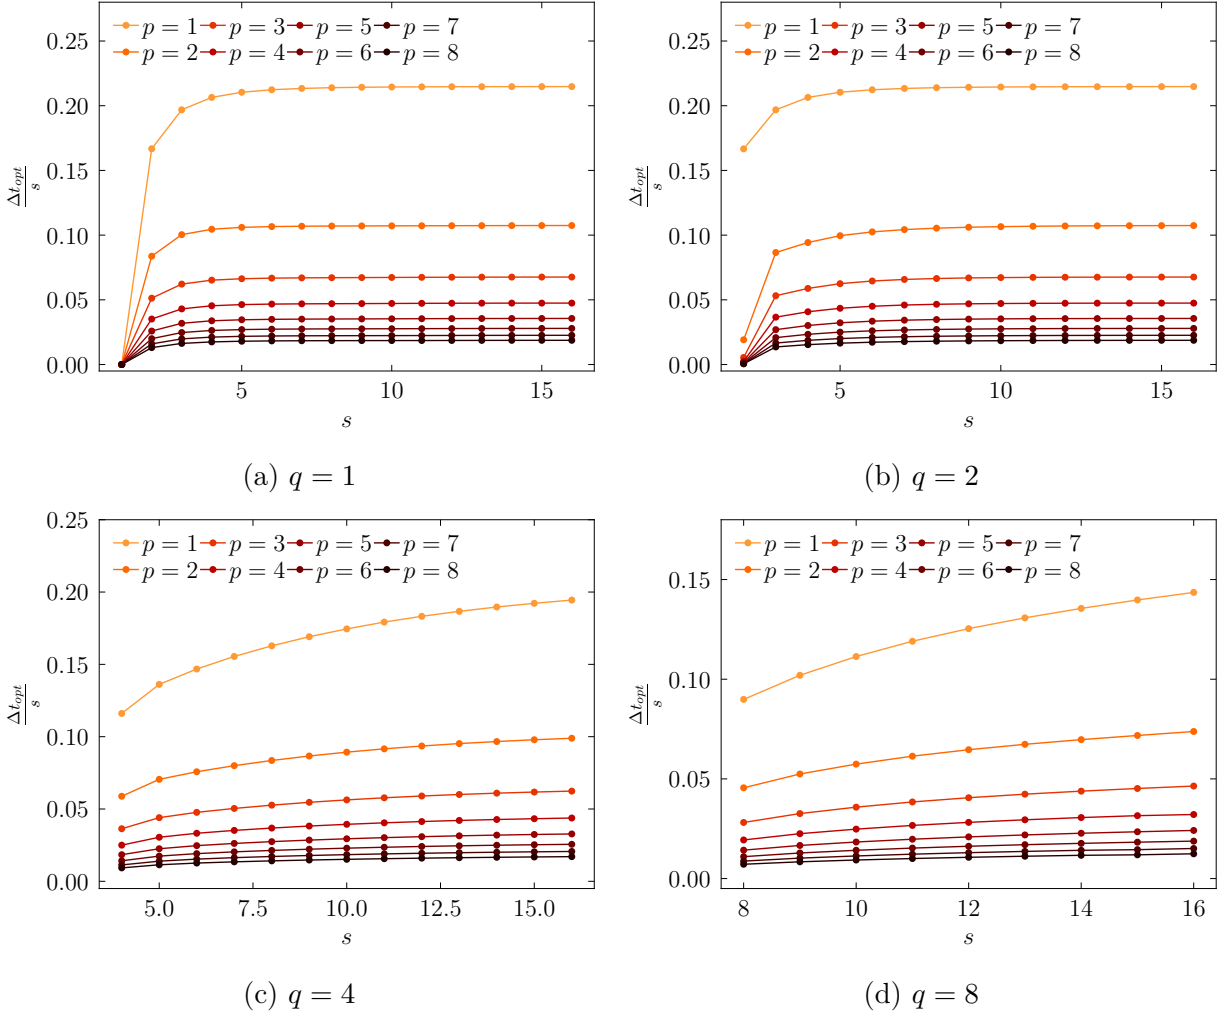


Figure 4.5. Optimal time-step size for OERK schemes normalized by the degree of the corresponding optimized polynomial for $p = 1 - 8$ spatial discretizations. Increasing $\Delta t_{opt}/s$ can be observed particularly for the first increases of s with lower-order temporal schemes and in general for all considered s with higher-order temporal schemes




Figure 4.6. Stability polynomial of classical RK schemes of degree 1, 2, 4 and 8 shown in a dashed line alongside eigenvalues of a degree $p = 4$ FR spatial discretization. Lower-order temporal schemes require smaller time-step sizes to fit the eigenvalues. The $q = 1$ temporal scheme is unstable for spatial schemes $p \geq 1$




Figure 4.7. Stability polynomial of OERK schemes of degree 1, 2, 4 and 8 shown in a dashed line alongside eigenvalues (black circles) of the semidiscretization of degree $p = 4$. Oscillatory behaviour can be observed for higher-order temporal schemes, such as $p = 4$ and $p = 8$, around the eigenvalues of the FR scheme

of the low-order schemes, they do not require as large s as high-order ones to represent the region of eigenvalues. At $s = 16$, both $P_{16,4}(z)$ and $P_{16,8}(z)$ present some oscillatory behaviour around the eigenspectra of the FR scheme.


4.2.2.2 Multidimensional Elements

We repeat the optimization procedure for multidimensional elements. Figure 4.8 shows the obtained optimal Δt_{opt} . We compute a number of metrics to compare against optimized polynomials for line elements as well as to standard RK methods. Tables 4.1 and 4.2 show the ratio of Δt_{opt} using the optimal stability polynomial of each multidimensional element type, relative to using the optimal one-dimensional stability polynomial with the multidimensional element. This is shown for all considered orders of accuracy in time and for all stage counts.


Table 4.1. Ratio of the maximum time step for $p = 4$ two-dimensional elements using their optimal stability polynomial and s stages, relative to the maximum time step obtained using a stability polynomial optimized for a one-dimensional element

		Quadrilaterals				Triangles			
$s \backslash q$	q	1	2	3	4	1	2	3	4
2		1.00	-	-	-	1.00	-	-	-
3		1.00	1.00	-	-	1.01	1.00	-	-
4		1.00	1.00	1.00	1.00	1.10	1.06	1.01	-
5		1.00	1.00	1.00	1.00	1.10	1.10	1.10	1.02
6		1.00	1.00	1.00	1.00	1.11	1.10	1.05	1.11
7		1.00	1.00	1.00	1.00	1.11	1.12	1.11	1.08
8		1.00	1.00	1.00	1.00	1.12	1.10	1.08	1.12
9		1.00	1.00	1.00	1.00	1.12	1.12	1.12	1.08
10		1.00	1.00	1.00	1.00	1.12	1.12	1.09	1.12
11		1.00	1.00	1.00	1.00	1.12	1.12	1.12	1.08
12		1.00	1.00	1.00	1.00	1.12	1.12	1.10	1.12
13		1.00	1.00	1.00	1.00	1.12	1.12	1.12	1.09
14		1.00	1.00	1.00	1.00	1.12	1.12	1.11	1.12
15		1.00	1.00	1.00	1.00	1.12	1.12	1.12	1.09
16		1.00	1.00	1.00	1.00	1.12	1.12	1.11	1.12


From these results, it is clear that optimizing the stability polynomial for the quadrilateral and hexahedral elements provides no additional benefits relative to using the optimal one-dimensional stability polynomial. However, minor performance improvements can be obtained for triangular, tetrahedral, and prismatic elements with additional speedup factors of up to 1.12 observed.




(a) Line




(b) Quadrilaterals




(c) Triangles



(d) Hexahedra



(e) Prisms



(f) Tetrahedra

Figure 4.8. Plots of the scaled optimal time-step size as a function of the number of stages for different ($p = 4$) element types

Table 4.2. Ratio of the maximum time step for $p = 4$ three-dimensional elements using their optimal stability polynomial and s stages, relative to the maximum time step obtained using a stability polynomial optimized for a one-dimensional element

$s \backslash q$	Hexahedra				Prisms				Tetrahedra			
	1	2	3	4	1	2	3	4	1	2	3	4
2	1.00	-	-	-	1.00	-	-	-	1.09	-	-	-
3	1.00	1.00	-	-	1.00	1.00	-	-	1.06	1.00	-	-
4	1.00	1.00	1.00	-	1.08	1.04	1.00	-	1.04	1.02	1.02	-
5	1.00	1.00	1.00	1.00	1.07	1.08	1.08	1.04	1.04	1.03	1.02	1.00
6	1.00	1.00	1.00	1.00	1.09	1.06	1.04	1.08	1.05	1.03	1.03	1.02
7	1.00	1.00	1.00	1.00	1.09	1.09	1.09	1.05	1.05	1.05	1.05	1.03
8	1.00	1.00	1.00	1.00	1.09	1.08	1.05	1.09	1.05	1.06	1.06	1.05
9	1.00	1.00	1.00	1.00	1.09	1.09	1.09	1.05	1.06	1.06	1.06	1.06
10	1.00	1.00	1.00	1.00	1.09	1.08	1.06	1.09	1.06	1.06	1.07	1.06
11	1.00	1.00	1.00	1.00	1.09	1.09	1.09	1.05	1.06	1.06	1.07	1.07
12	1.00	1.00	1.00	1.00	1.09	1.09	1.07	1.09	1.06	1.06	1.07	1.07
13	1.00	1.00	1.00	1.00	1.09	1.09	1.09	1.06	1.07	1.07	1.07	1.07
14	1.00	1.00	1.00	1.00	1.09	1.09	1.08	1.09	1.06	1.07	1.07	1.07
15	1.00	1.00	1.00	1.00	1.09	1.09	1.09	1.07	1.06	1.06	1.06	1.07
16	1.00	1.00	1.00	1.00	1.09	1.09	1.08	1.09	1.06	1.06	1.06	1.07

4.2.2.3 Comparison with Classical Runge-Kutta Schemes

To investigate the utility of the stability polynomials optimized for each multidimensional element type, we can compare them to the classical RK_{3,3} and RK_{4,4} methods. Table 4.3 shows the ratio of Δt_{opt} relative to the maximum stable time step size using RK_{3,3}. This shows that a larger time step relative to RK_{3,3} can be taken for all elements as the number of stages increases, which is to be expected. More importantly, Table 4.4 shows that the speedup factor relative to RK_{3,3} also increases with the number of stages, which is measured by the ratio of optimal schemes $\frac{\Delta t_{opt}}{s}$ relative to that of RK_{3,3}. By 16 stages, speedup factors of 1.46 to 1.74 are observed, depending on the element type. Similar to Table 4.3, Table 4.5

Table 4.3. Time step ratio of $q = 3$ optimal RK schemes for different element types with $p = 4$ relative to RK_{3,3}, $\frac{\Delta t_{opt}}{\Delta t_{RK_{3,3}}}$

s	Hexahedra	Prisms	Tetrahedra	Quadrilaterals	Triangles	Lines
3	1.00	1.00	1.00	1.00	1.00	1.00
4	1.53	1.69	1.48	1.53	1.73	1.53
5	2.11	2.27	2.02	2.11	2.32	2.11
6	2.70	2.97	2.57	2.70	3.05	2.70
7	3.29	3.57	3.11	3.29	3.65	3.29
8	3.87	4.24	3.65	3.87	4.36	3.87
9	4.44	4.84	4.17	4.44	4.96	4.44
10	5.01	5.46	4.71	5.01	5.63	5.01
11	5.57	6.08	5.22	5.57	6.23	5.57
12	6.13	6.67	5.73	6.13	6.86	6.13
13	6.68	7.29	6.25	6.68	7.48	6.68
14	7.23	7.87	6.76	7.23	8.08	7.23
15	7.77	8.47	7.26	7.77	8.70	7.77
16	8.31	9.07	7.77	8.31	9.29	8.32

shows the ratio of Δt_{opt} relative to the maximum stable time step size using RK_{4,4}. Again, this shows that a larger time step relative to RK_{4,4} can be taken for all elements as the number of stages increases. In addition, similar to Table 4.4, Table 4.6 demonstrates that the speedup factor relative to classical RK_{4,4} also increases with the number of stages, which is again measured by the ratio of the optimal schemes $\frac{\Delta t_{max}}{s}$ relative to that of RK_{4,4}. By 16 stages, speedup factors of 1.63 to 1.97 are observed, depending on the element type.

Table 4.4. Speedup factor of $q = 3$ optimal RK schemes for different element types with $p = 4$ relative to RK_{3,3}, $(\frac{\Delta t_{opt}}{s} / \frac{\Delta t_{RK_{3,3}}}{3})$

s	Hexahedra	Prisms	Tetrahedra	Quadrilaterals	Triangles	Lines
3	1.00	1.00	1.00	1.00	1.00	1.00
4	1.15	1.26	1.11	1.15	1.30	1.15
5	1.27	1.36	1.21	1.27	1.39	1.27
6	1.35	1.48	1.28	1.35	1.52	1.35
7	1.41	1.53	1.33	1.41	1.57	1.41
8	1.45	1.59	1.37	1.45	1.63	1.45
9	1.48	1.61	1.39	1.48	1.65	1.48
10	1.50	1.64	1.41	1.50	1.69	1.50
11	1.52	1.66	1.42	1.52	1.70	1.52
12	1.53	1.67	1.43	1.53	1.71	1.53
13	1.54	1.68	1.44	1.54	1.73	1.54
14	1.55	1.69	1.45	1.55	1.73	1.55
15	1.55	1.69	1.45	1.55	1.74	1.55
16	1.56	1.70	1.46	1.56	1.74	1.56

Table 4.5. Time step ratio of $q = 4$ optimal RK schemes for different element types with $p = 4$ relative to RK_{4,4}, $\frac{\Delta t_{opt}}{\Delta t_{RK_{4,4}}}$

s	Hexahedra	Prisms	Tetrahedra	Quadrilaterals	Triangles	Lines
4	1.00	1.00	1.00	1.00	1.00	1.00
5	1.52	1.68	1.47	1.52	1.72	1.52
6	1.99	2.17	1.88	1.99	2.21	1.99
7	2.46	2.72	2.32	2.46	2.79	2.46
8	2.94	3.22	2.77	2.94	3.30	2.94
9	3.43	3.78	3.23	3.44	3.89	3.44
10	3.93	4.32	3.69	3.94	4.42	3.94
11	4.44	4.87	4.17	4.45	5.02	4.45
12	4.95	5.44	4.63	4.96	5.58	4.95
13	5.47	5.99	5.11	5.47	6.15	5.47
14	5.97	6.56	5.58	5.98	6.73	5.98
15	6.48	7.11	6.05	6.49	7.29	6.49
16	6.99	7.67	6.52	7.00	7.88	7.00

Table 4.6. Speedup factor of $q = 4$ optimal RK schemes for different element types with $p = 4$ relative to RK_{4,4}, $(\frac{\Delta t_{opt}}{s} / \frac{\Delta t_{RK_{4,4}}}{4})$

s	Hexahedra	Prisms	Tetrahedra	Quadrilaterals	Triangles	Lines
4	1.00	1.00	1.00	1.00	1.00	1.00
5	1.22	1.34	1.18	1.22	1.37	1.22
6	1.33	1.45	1.25	1.33	1.48	1.33
7	1.41	1.56	1.32	1.41	1.60	1.41
8	1.47	1.61	1.38	1.47	1.65	1.47
9	1.53	1.68	1.44	1.53	1.73	1.53
10	1.57	1.73	1.48	1.58	1.77	1.58
11	1.61	1.77	1.52	1.62	1.82	1.62
12	1.65	1.81	1.54	1.65	1.86	1.65
13	1.68	1.84	1.57	1.68	1.89	1.68
14	1.71	1.87	1.60	1.71	1.92	1.71
15	1.73	1.90	1.61	1.73	1.95	1.73
16	1.75	1.92	1.63	1.75	1.97	1.75

4.2.3 Numerical Results

4.2.3.1 Verification

For verification, a Butcher tableau was generated for each stability polynomial using the least truncation method [45] (available in [117]), and was used to solve a linear advection test case. To amplify the temporal error beyond machine precision, a prescribed source term was added using the method of manufactured solutions. The linear advection equation with this added source term is

$$\frac{\partial u}{\partial t} + \nabla \cdot (\boldsymbol{\alpha}u) = S(t), \quad (4.17)$$

where $u(\mathbf{x}, t)$ is a scalar, $\boldsymbol{\alpha}$ is the advection velocity, and $S(t)$ is the source term. Assuming periodic boundaries and an initial condition $u(\mathbf{x}, 0) = e^{-\kappa|\mathbf{x}|^2}$, where $\kappa = 0.4$, the exact solution is

$$u = e^{\kappa|\mathbf{x}-\boldsymbol{\alpha}t|^2} + \int S(t)dt. \quad (4.18)$$

Taking the exact solution to be

$$u(\mathbf{x}, t) = e^{\kappa|\mathbf{x}-\boldsymbol{\alpha}t|^2} + 100\sin(10\pi t), \quad (4.19)$$

the resulting source term is

$$S(t) = 1000\pi \times \cos(10\pi t). \quad (4.20)$$

We used a 20×20 two-dimensional and $20 \times 20 \times 20$ three-dimensional domain Ω , with periodic boundary conditions in all directions. Two sets of two-dimensional simulations, one with 80×80 quadrilateral elements, and another with 160×160 triangular elements, and three sets of three-dimensional simulations, one with $80 \times 80 \times 80$ hexahedral elements, and two sets with $160 \times 160 \times 160$ tetrahedral and prismatic elements were run to evaluate the order of accuracy of optimal schemes. The solution was represented using $p = 6$ degree polynomial on each element to minimize spatial error. A set of simulations was run using different optimal stability polynomials with $s = 16$ and $q = 1, 2, 3, 4$. Each simulation was run with unit advection velocity in all directions to a final simulation time of $t = 20$ to allow the flow to complete a full cycle through the periodic domain. The accuracy of each simulation is evaluated using the L_2 norm of the error at the end of each simulation, defined as

$$E_{L_2}(\Omega^h) = \sqrt{\int_{\Omega^h} (u^h(\mathbf{x}) - u(\mathbf{x}))^2 d\Omega^h}. \quad (4.21)$$

A summary of the error for each temporal scheme and time-step size, including observed orders of accuracy, are plotted in Figure 4.9 for all element types. Importantly, we observe that all schemes achieve their designed order of accuracy.

4.2.3.2 Navier-Stokes Equations

To explore the utility of the optimized stability polynomials for multidimensional nonlinear problems, we consider DNS of the Taylor-Green vortex. The initial flow field for this test case is specified as [19]

$$v_x = +U_0 \sin\left(\frac{x}{\ell}\right) \cos\left(\frac{y}{\ell}\right) \cos\left(\frac{z}{\ell}\right), \quad (4.22)$$


$$v_y = -U_0 \cos\left(\frac{x}{\ell}\right) \sin\left(\frac{y}{\ell}\right) \cos\left(\frac{z}{\ell}\right), \quad (4.23)$$

$$v_z = 0, \quad (4.24)$$


$$P = P_0 + \frac{\rho_o U_0^2}{16} \left(\cos\left(\frac{2x}{\ell}\right) + \cos\left(\frac{2y}{\ell}\right) \right) \left(\cos\left(\frac{2z}{\ell}\right) + 2 \right), \quad (4.25)$$

$$\rho = \frac{P}{RT_0}, \quad (4.26)$$


where v_x , v_y , and v_z are the velocity components, P is the pressure, P_0 is the background pressure, ρ is the density, and T_0 and U_0 are constants specified such that the flow Mach number based on U_0 is $M = 0.1$, effectively incompressible. The domain is a periodic cube with the dimensions $-\pi\ell \leq x, y, z \leq +\pi\ell$. For the current study, we consider a Reynolds number $\text{Re} = 1600$ based on the length scale ℓ and velocity scale U_0 . The test case is run to




(a) Quadrilateral element




(b) Triangular element



(c) Hexahedral element



(d) Tetrahedral element



(e) Prismatic element

Figure 4.9. Convergence plots for linear advection using the method of manufactured solutions using optimal RK schemes


a final nondimensional time of $t = 20t_c$ where $t_c = \ell/U_0$, and reference spectral DNS data is available from van Rees et al. [127]. Of primary interest is the temporal evolution of the total kinetic energy, which can be found via

$$E_k = \frac{1}{\rho_0 \Omega} \int_{\Omega} \rho \frac{\mathbf{v} \cdot \mathbf{v}}{2} d\Omega, \quad (4.27)$$


and, more specifically, its dissipation rate $-dE_k/dt$. Furthermore, the temporal evolution of enstrophy can be computed via

$$\varepsilon = \frac{1}{\rho_0 \Omega} \int_{\Omega} \rho \frac{\mathbf{Z} \cdot \mathbf{Z}}{2} d\Omega, \quad (4.28)$$


where \mathbf{Z} is the vorticity. For incompressible flows, the dissipation rate can be related to the enstrophy by $-dE_k/dt = 2\frac{\mu}{\rho_0}\varepsilon$ [127, 19]. Consequently, enstrophy is a direct measure of the expected physical dissipation, and the dissipation rate computed from the kinetic energy is typically higher due to the additional numerical dissipation of the scheme. A total of six simulations were run using $p = 5$ and a nominal 256^3 solution points using hexahedral, prismatic, and tetrahedral elements with the classical RK_{4,4} scheme and each element types corresponding optimal fourth-order temporal scheme with $q = 4$ and $s = 16$. The maximum stable time step sizes for each element type and temporal scheme were determined via bisection, and each simulation was run to completion using a time step 80% of this maximum size. This resulted in speedup factors of 1.68, 1.64, and 1.57 for hexahedral, prismatic, and tetrahedral elements, respectively, when using the optimal stability polynomials. Importantly, these speedup factors were obtained with negligible code modification, simply the substitution of the coefficients in the Butcher tableau. Qualitative results in terms of isosurfaces of Q-criterion coloured by velocity magnitude are shown in Figure 4.10 for all element types and both sets of temporal schemes at a nondimensional simulation time of $t_c = 15$, which is beyond the turbulent transition time. From these images, it is apparent that the flow is qualitatively indistinguishable between the two temporal schemes for all element types. Furthermore, quantitative results in terms of the total kinetic energy dissipation rate and enstrophy are shown in Figure 4.11 and Figure 4.12, respectively. These results demonstrate a negligible difference in accuracy for both quantitative measures throughout each simulation, where the plots for the classical and optimized Runge-Kutta methods are visually indistinguishable. We observe that the optimized stability polynomials are able to yield significant reductions in simulation time with negligible influence on both qualitative and quantitative results.




(a) Hexahedra, RK_{4,4}




(b) Hexahedra, OERK_{16,4}




(c) Prisms, RK_{4,4}



(d) Prisms, OERK_{16,4}




(e) Tetrahedra, RK_{4,4}




(f) Tetrahedra, OERK_{16,4}


Figure 4.10. Isosurfaces of Q-criterion coloured by velocity magnitude for the Taylor-Green vortex



(a) Hexahedra




(b) Tetrahedra




(c) Prisms


Figure 4.11. Energy decay rate versus dimensionless time for the Taylor-Green vortex case



(a) Hexahedra



(b) Tetrahedra



(c) Prisms

Figure 4.12. Enstrophy versus dimensionless time for the Taylor-Green vortex case

4.3 Fully-Discrete Analysis

We now analyze the properties of the fully-discrete discretization using optimal RK schemes. We focus only on one-dimensional formulations, as they typically provide general insights into the overall numerical errors of a scheme. To analyze the methods, we consider again the linear advection equation

$$\frac{\partial u}{\partial t} + \frac{\partial u}{\partial x} = 0, \quad (4.29)$$

which is nondissipative and admits solutions of the form

$$u(x, t) = e^{I(\kappa x - \omega t)}, \quad (4.30)$$

where κ is the wavenumber, ω is the frequency, and $I = \sqrt{-1}$. Substituting the plane wave into the advection equation yields the exact dispersion relation $\omega = \kappa$. Divide a computational domain into a series of N nonoverlapping elements of equal length h using periodic boundary conditions. We can discretize Equation (4.29) using the standard FR approach described in Chapter 3 and write it in the form

$$\frac{d\mathbf{u}^h}{dt} = \mathbf{S}\mathbf{u}^h, \quad (4.31)$$

where \mathbf{u}^h is a column vector containing the nodal expansion coefficients of each element, i.e.

$$\mathbf{u}^h = [\mathbf{u}_1^h, \mathbf{u}_2^h, \dots, \mathbf{u}_N^h]^T,$$

and \mathbf{S} is the semidiscrete operator of size $N(p+1) \times N(p+1)$. Equation (4.31) can be expanded for the upwind scheme considered here as

$$\begin{bmatrix} \frac{d\mathbf{u}_1^h}{dt} \\ \frac{d\mathbf{u}_2^h}{dt} \\ \vdots \\ \frac{d\mathbf{u}_N^h}{dt} \end{bmatrix} = \begin{bmatrix} \mathbf{C}_0 & 0 & \dots & \mathbf{C}_{-1} \\ \mathbf{C}_{-1} & \mathbf{C}_0 & & \\ & \ddots & \ddots & \\ & & \mathbf{C}_{-1} & \mathbf{C}_0 \end{bmatrix} \begin{bmatrix} \mathbf{u}_1^h \\ \mathbf{u}_2^h \\ \vdots \\ \mathbf{u}_N^h \end{bmatrix},$$

where \mathbf{C}_0 and \mathbf{C}_{-1} are the $(p+1) \times (p+1)$ square linear transformation matrices of the spatial discretization

$$\mathbf{C}_0 = -\frac{2}{h} [\mathbf{D} - \mathbf{g}_{L,\tilde{x}} \mathbf{l}^T], \quad \mathbf{C}_{-1} = -\frac{2}{h} \mathbf{g}_{L,\tilde{x}} \mathbf{r}^T, \quad (4.32)$$

where $\mathbf{g}_{L,\tilde{x}}$ is a column vector containing the derivative of the left DG correction function at each node

$$\mathbf{g}_{L,\tilde{x}_j} = \left. \frac{dg_L}{d\tilde{x}} \right|_{\tilde{x}_j}, \quad (4.33)$$

\mathbf{D} is the differentiation matrix and \mathbf{l} , \mathbf{r} are the left and right interpolation vectors. These operators are respectively given by

$$D_{jl} = \left. \frac{d\varphi_l}{d\tilde{x}} \right|_{\tilde{x}_j}, \quad (4.34)$$

$$l_j = \varphi_j \Big|_{\tilde{x}=-1}, \quad (4.35)$$

$$r_j = \varphi_j \Big|_{\tilde{x}=1}. \quad (4.36)$$

A fully-discrete operator \mathbf{Q} can be obtained after application of a temporal method such that

$$\mathbf{Q} = \sum_{j=0}^s \gamma_j (\Delta t \mathbf{S})^j, \quad (4.37)$$

where γ_j are the coefficients of the stability polynomial of the time-stepping scheme. This yields a matrix of size $N(p+1) \times N(p+1)$, which has the following form

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q}_{1,1} & \mathbf{Q}_{1,2} & \cdots & \mathbf{Q}_{1,N} \\ \mathbf{Q}_{2,1} & \mathbf{Q}_{2,2} & \cdots & \mathbf{Q}_{2,N} \\ \vdots & & \ddots & \vdots \\ \mathbf{Q}_{N,1} & \mathbf{Q}_{N,2} & \cdots & \mathbf{Q}_{N,N} \end{bmatrix}. \quad (4.38)$$

The solution at the next time step can be obtained for any given element by

$$\mathbf{u}_n^h(t + \Delta t) = \sum_{j=1}^N \mathbf{Q}_{n,j} \mathbf{u}_j^h. \quad (4.39)$$

In this analysis, we seek solutions of the form

$$\mathbf{u}_n^h(t) = e^{I(\bar{\kappa}x_n - \bar{\omega}^h t)} \boldsymbol{\alpha}, \quad (4.40)$$

where $\bar{\kappa} = \kappa h$ is a prescribed nondimensional wavenumber $\bar{\kappa} \in [0, 2\pi]$, $\bar{\omega}^h$ is the numerical frequency of the fully-discrete scheme, x_n is the center of the element and $\boldsymbol{\alpha}$ is a nonzero

vector of unknowns. Substituting Equation (4.40) into (4.39) yields

$$e^{-i\bar{\omega}^h \Delta t} \boldsymbol{\alpha} = \sum_{j=1}^N \mathbf{Q}_{n,j} e^{i\bar{\kappa}(j-n)} \boldsymbol{\alpha}, \quad (4.41)$$

which we rewrite as a classical eigenvalue problem

$$\lambda \boldsymbol{\alpha} = \mathbf{M}(\bar{\kappa}) \boldsymbol{\alpha}, \quad (4.42)$$

where $\lambda = e^{-I\bar{\omega}^h \Delta t}$ and

$$\mathbf{M} \equiv \mathbf{M}_n = \sum_{j=1}^N \mathbf{Q}_{n,j} e^{I\bar{\kappa}(j-n)}. \quad (4.43)$$

Note that due to the equally-spaced grid, the $(p+1) \times (p+1)$ matrix \mathbf{M}_n is equivalent in every element. Hence the choice of $d \in [1, N]$ is arbitrary. The $p+1$ eigenvalues of \mathbf{M} are complex and are related to numerical frequencies by

$$\bar{\omega}^h = -\frac{\ln(\lambda)}{I\Delta t}. \quad (4.44)$$

The imaginary part of $\bar{\omega}^h$ is associated with dissipation error, whereas the real part is associated with the numerical dispersion of the scheme. A single mode dominates the spectral properties of the numerical scheme at each given wavenumber. The other k eigenvalues contain large numerical dissipation and will be rapidly damped. Given the periodicity of Equation (4.42) with period 2π , each eigenvalue can be interpreted as corresponding to a true wavenumber $\bar{\kappa} + 2\pi l$, where l is an integer [123]. To sort ω^h , we implemented the algorithm proposed by Vincent et al. [31].

Based on the fact that \mathbf{M} is continuous in $\bar{\kappa}$, and via the theory of perturbations [128], they are subsequently permuted to form a continuous function $\lambda : [0, 2(p+1)\pi] \rightarrow \mathbb{C}$. We validate the accuracy of the fully-discrete scheme by computing the frequency error

$$E(\bar{\kappa}) = |\bar{\kappa} - \bar{\omega}_\kappa^h|, \quad (4.45)$$

where $\bar{\omega}_\kappa^h$ is associated with the dominant eigenvalue at the prescribed $\bar{\kappa}$. Then, the local order of the frequency error can be found by [53]

$$A_\kappa = \frac{\ln[E(\bar{\kappa})/E(\bar{\kappa}/2)]}{\ln(2)}, \quad (4.46)$$

and the global order is [26]

$$A_g = A_\kappa - 1. \quad (4.47)$$

For the DG scheme considered here, the expected global order of accuracy is $2p + 1$, which is superconvergent [129]. In addition, we explore the group velocity, at which features at a given wavenumber propagate in space, given by

$$v_g = \frac{d}{d\bar{\kappa}} [\text{Re}(\bar{\omega}^h)]. \quad (4.48)$$

This analysis will give insights into the spectral properties of the fully-discrete scheme. In this study, we focus on FR_{DG} schemes coupled with their corresponding OERK schemes. The introduction of optimized stability polynomials modifies the type of error that the resulting numerical scheme introduces to the solution. In the following section, we analyze the spectral properties of fully-discrete FR-OERK schemes. Due to the unconditional unstable properties of the $P_{1,1}$ scheme and the poor properties of $P_{2,2}$ [126], results for these two temporal methods will be omitted in further discussions.

4.3.1 Fully-Discrete Properties of FR schemes using Optimal RK schemes

The properties of fully-discrete schemes do not rely only on their spatial order of accuracy but also the time-stepping method. We now explore the spectral properties of fully-discrete FR-OERK schemes according to the analysis presented in Section 4.3. In particular, we look at the dissipation, dispersion, group velocity, error, and the local order of fifth-order spatial discretizations ($p = 4$) coupled with high and low-order temporal schemes. Results are shown for different ratios of Δt relative to Δt_{opt} , defined by τ until the CFL condition reaches its stability limit at $\tau = 1.0$. We use the spectral properties of the semidiscretization as a reference, shown by a dashed line in all plots.

Figures 4.13 and 4.14 show the spectral properties for the $P_{16,1}$ and $P_{16,2}$ schemes, respectively. As we have previously mentioned, the stability polynomials of these two schemes contain very similar coefficients due to the shape of the spatial eigenspectra. As a consequence, their spectral properties are similar. We first look at the dispersion and dispersion error plots. Dispersion error is generally observed to increase with the time-step size for all wavenumbers and can be significant even at low wavenumbers. Dissipation error increases for larger τ for wavenumbers $0 < \bar{\kappa} \lesssim \pi$ but decreases after this point. Good agreement of v_g can be seen until at least $\bar{\kappa} = 2\pi$. After this, large values of group velocity appear for unresolved features, even for $\tau = 0.5$. The error plots include the slope of $E(\bar{\kappa})$, which represents the local order. This is shown for both the fully-discrete scheme on the left and also for the semidiscretization on the right. As previously observed by Yang et al. [52] and Vermeire et al. [53], the superconvergence property of the FR_{DG} scheme with local order $A_\kappa = 2p + 2$ is generally downgraded to the order of the temporal scheme in the region of small wavenumbers,

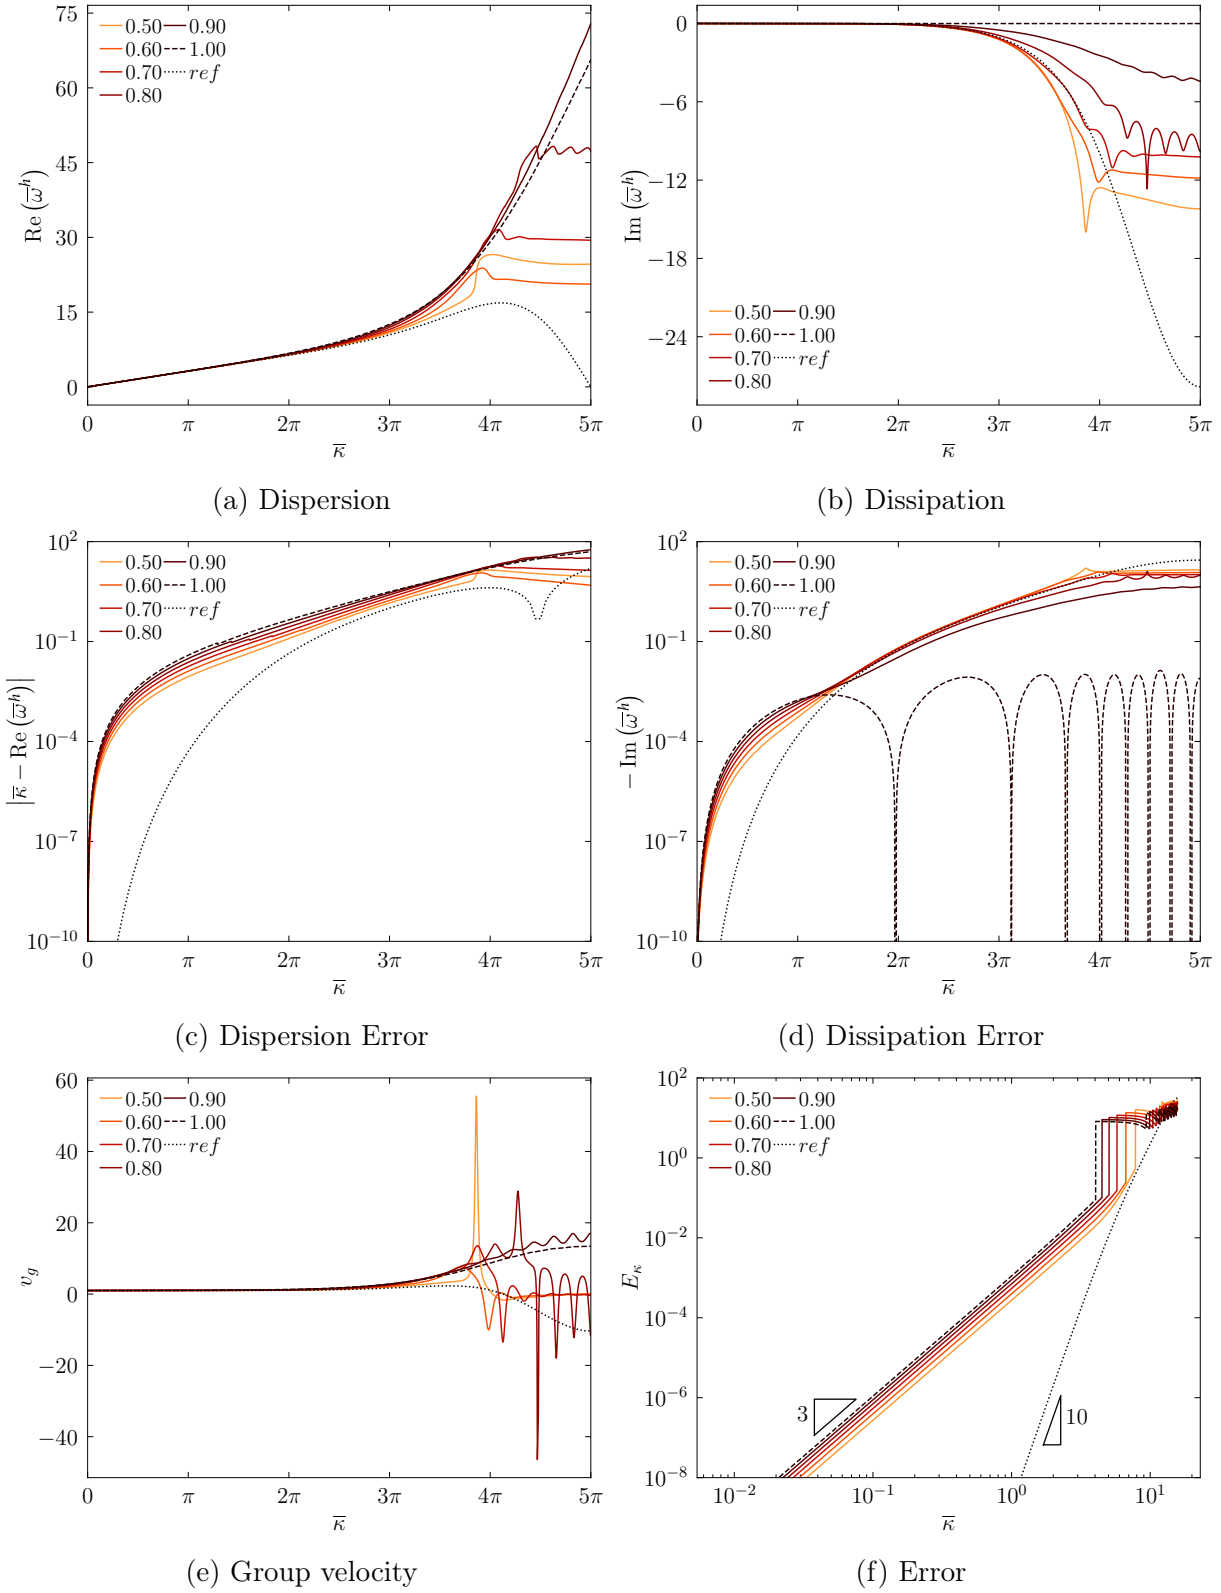


Figure 4.13. Spectral properties of $P_{16,1}$ schemes at different CFL numbers τ . Properties of the semidiscretization are shown in a dotted line

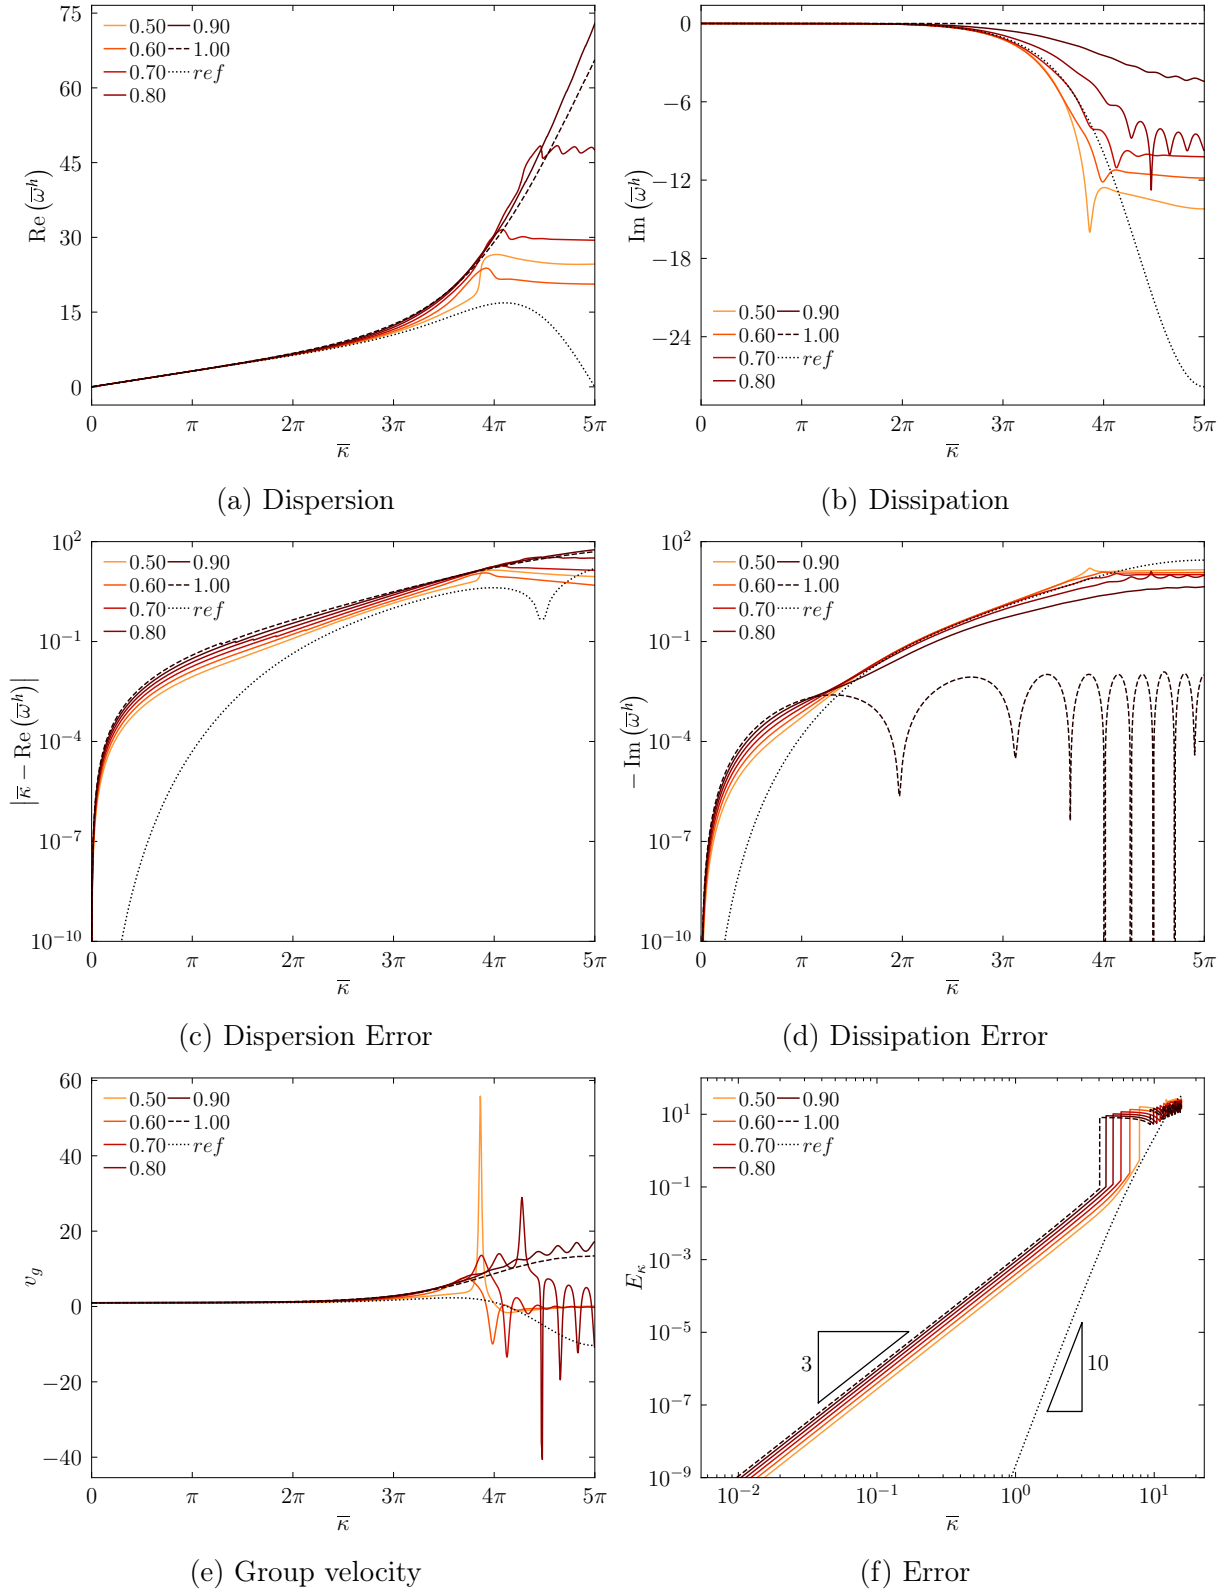


Figure 4.14. Spectral properties of $P_{16,2}$ schemes at different CFL numbers τ . Properties of the semidiscretization are shown in a dotted line

decreasing it to $A_\kappa = p + 1$. Hence, we observe $A_\kappa = 10$ for the semidiscretization, and a reduction to $A_\kappa = 3$ in the region of well-resolved wavenumbers. Note that the value of A_κ is the same for both $P_{16,1}$ and $P_{16,2}$ as a consequence of both the order conditions and the optimization.

We now compare the spectral properties of the classical and optimal RK schemes of third and fourth order. The behaviour of the $P_{4,4}$ and $P_{3,3}$ schemes was already presented in [53]. We show these results for completeness, which, in addition, serves as validation of the study. Figure 4.15 shows a $P_{3,3}$ scheme. Dispersion curves agree well with the reference until at least 3π , with additional dispersion error introduced for the smallest wavenumbers proportional to the times-step size. After this, we observe considerable increases in dispersion for higher τ , similar to the properties of the semidiscretization. Dissipation plots show higher damping at small wavenumbers when τ is increased, and the dissipation reduces for larger τ in the high-wavenumber region. Group velocity is underpredicted right before 3π and may be substantially larger for the highest wavenumbers. The local order is reduced from $A_\kappa = 10$ to $A_\kappa = 4$ in the well-resolved range.

In contrast, Figure 4.16 shows the optimized $P_{16,3}$ scheme. Increased dispersive behaviour can be seen starting at wavenumbers $\approx \pi$, in particular for large τ . Dissipation increases for small τ for the mid-range of wavenumbers and decreases after 4π relative to the semidiscrete properties. When $\tau = 1.0$, dissipation is drastically reduced for all wavenumbers with small amounts added at given wavenumbers after $\approx 3\pi$. Group velocity is overpredicted after 2π . Large magnitudes of v_g can be seen for certain high values of τ in the high-wavenumber region. Compared to the conventional scheme, the maximum (absolute) group velocity is generally smaller. The convergence of the error for all τ considered decreases to $A_\kappa = p + 1$ due to the large time-step size. Compared to the classical scheme, the optimal third-order RK scheme generally introduces more error, particularly at high wavenumbers. However, they are also capable of achieving a significantly larger time-step size.

Properties of the $P_{4,4}$ scheme can be seen in Figure 4.17. Compared to the third-order conventional RK scheme, we observe considerable agreement with the spectral properties of the semidiscretization for a higher range of wavenumbers until $\approx 3\pi$. The dispersion error behaves similarly to the semidiscretization at $\tau < 0.6$ with increased error after 4π . However, it is clear that even the semidiscrete spectral behaviour in this region includes large amounts of error. The dissipation properties show that increasing the time-step means additional dissipation for the mid-range of wavenumbers, but behaves in the opposite manner as it approaches $\bar{\kappa} = 5\pi$. High-wavenumber features travel at large, incorrect speeds for all considered τ at high wavenumbers. The error convergence behaves similarly to the previously discussed schemes, in which the temporal order dominates for the well-resolved range of wavenumbers, where the slope decreases by a factor of 2, specifically from $A_\kappa = 10$ to $A_\kappa = 5$.

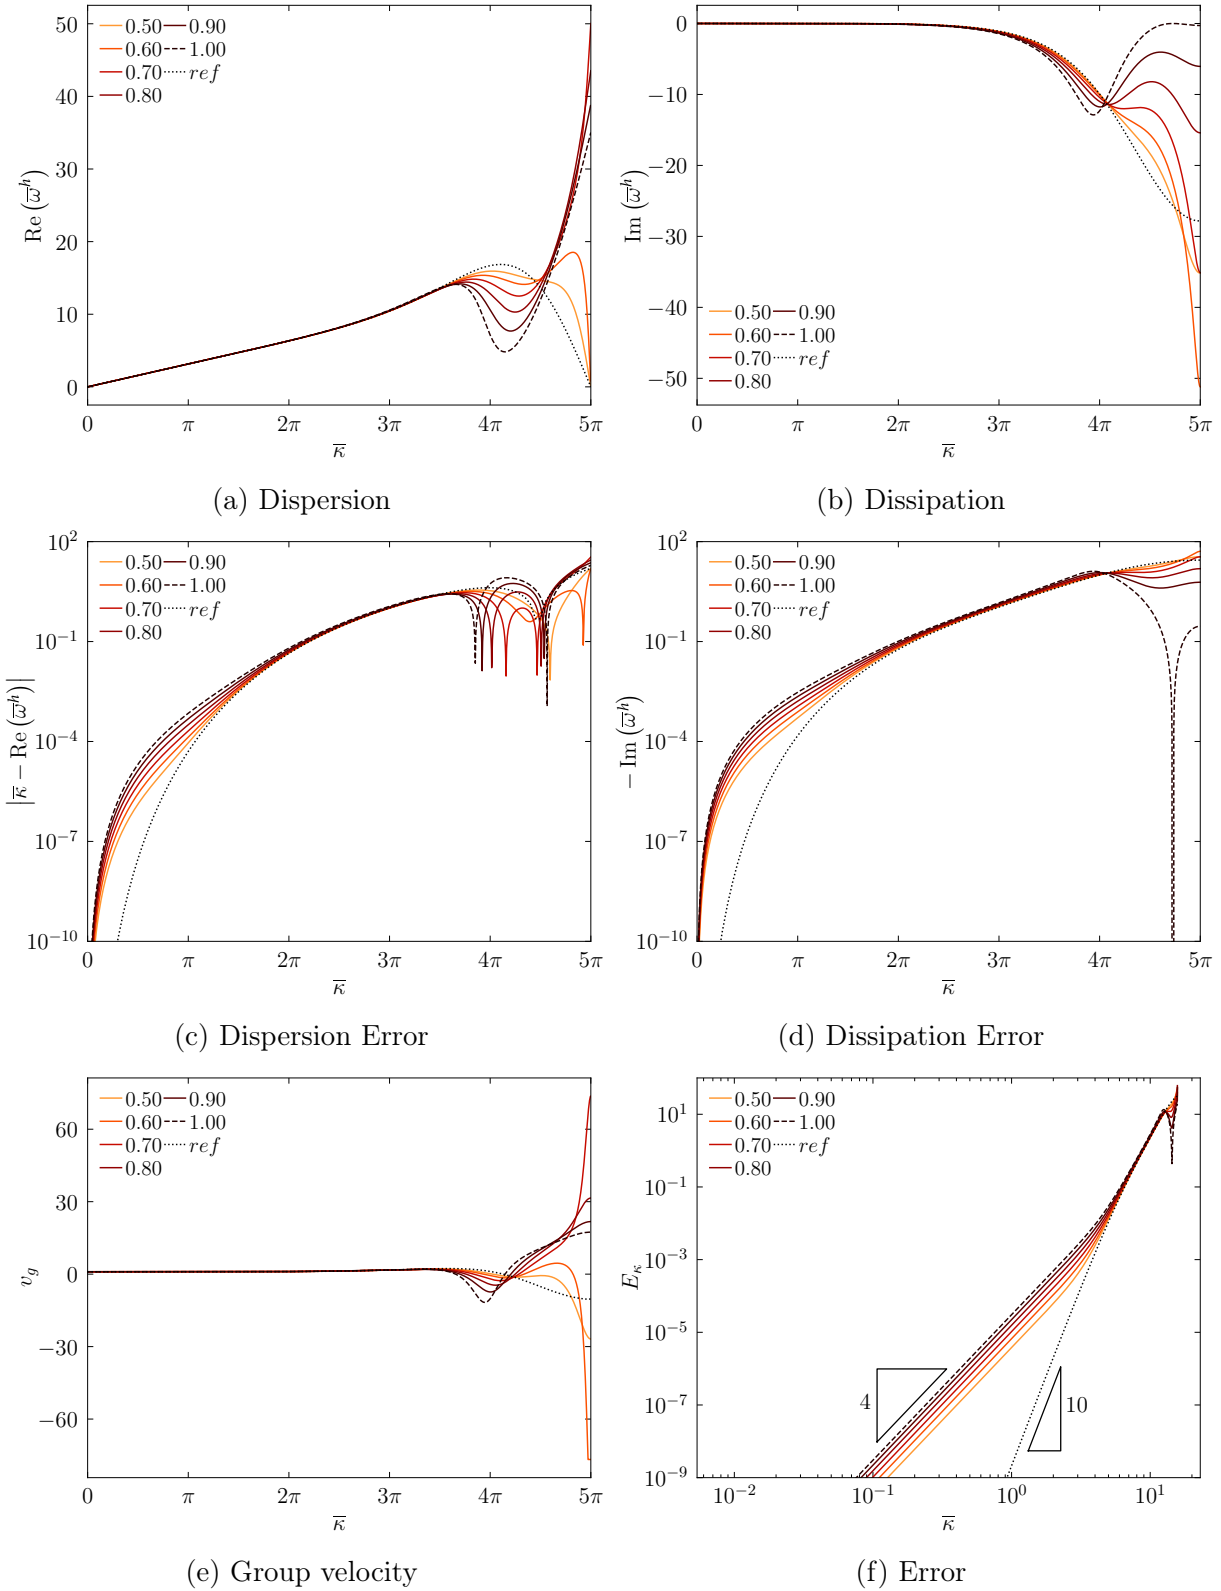


Figure 4.15. Spectral properties of $P_{3,3}$ schemes at different CFL numbers τ . Properties of the semidiscretization are shown in a dotted line

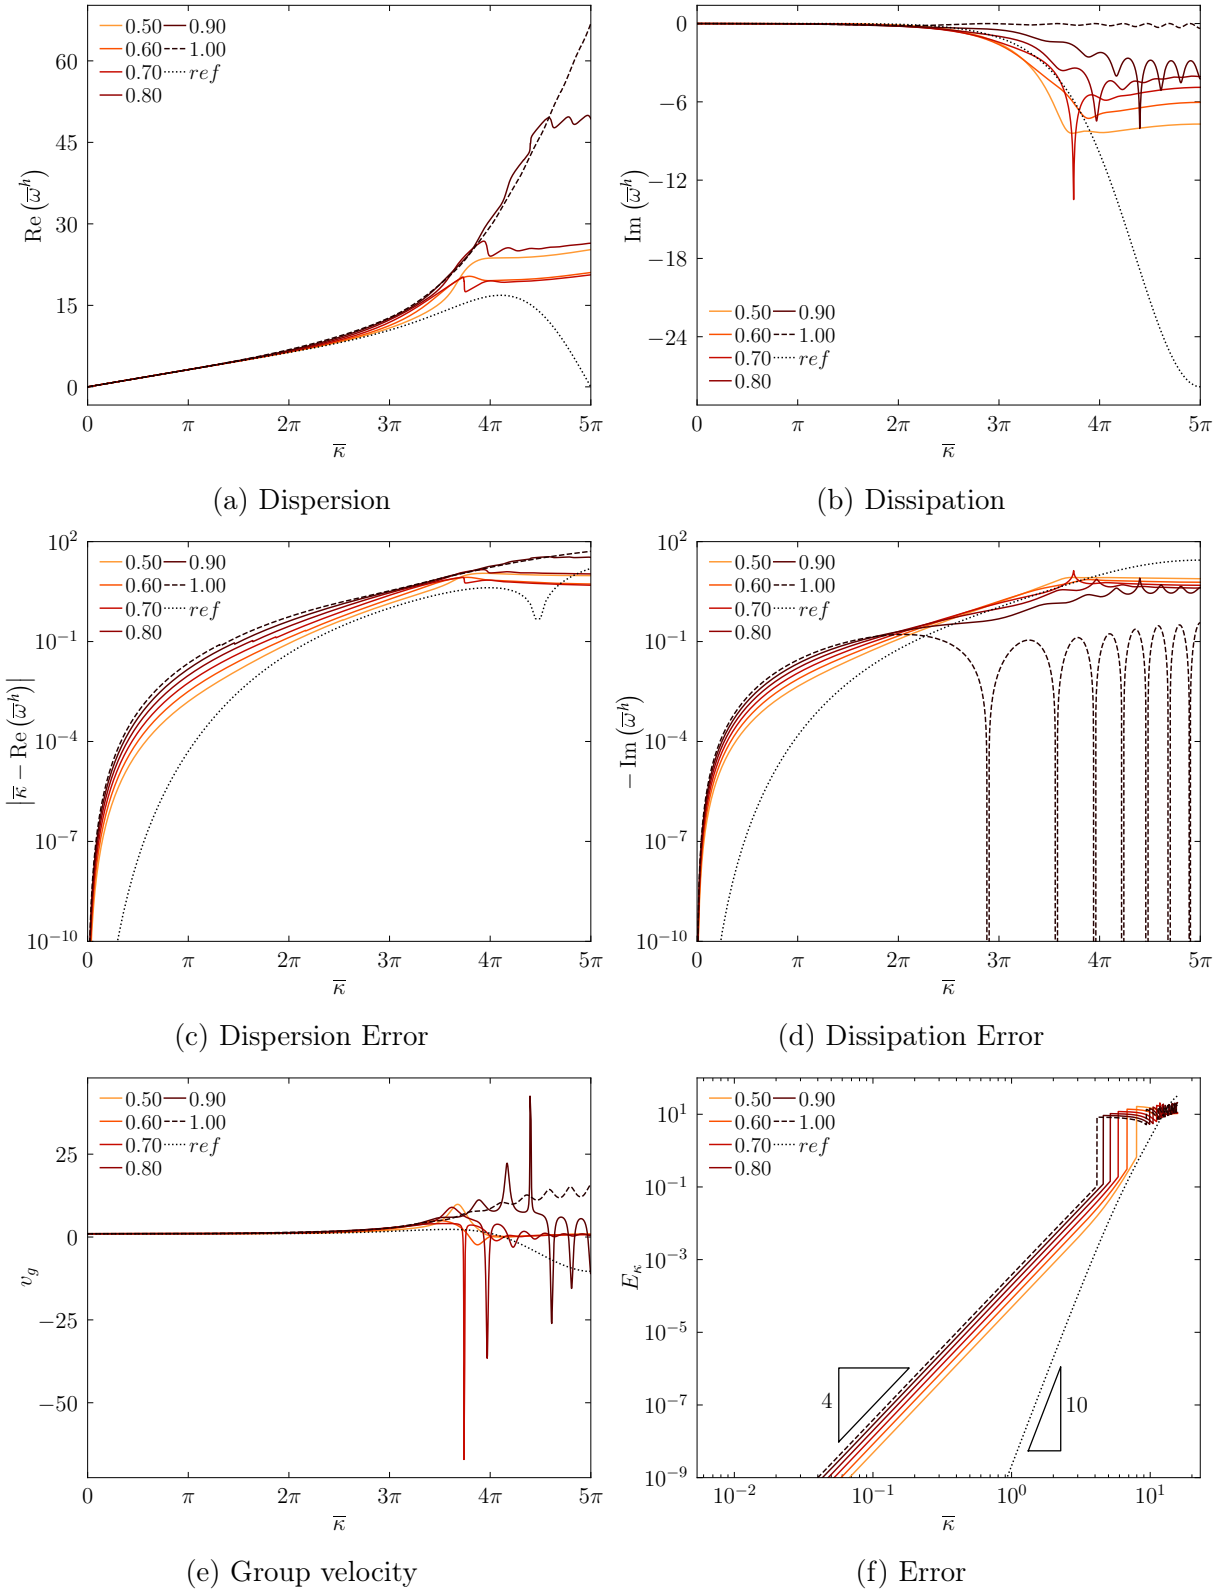


Figure 4.16. Spectral properties of $P_{16,3}$ schemes at different CFL numbers τ . Properties of the semidiscretization are shown in a dotted line

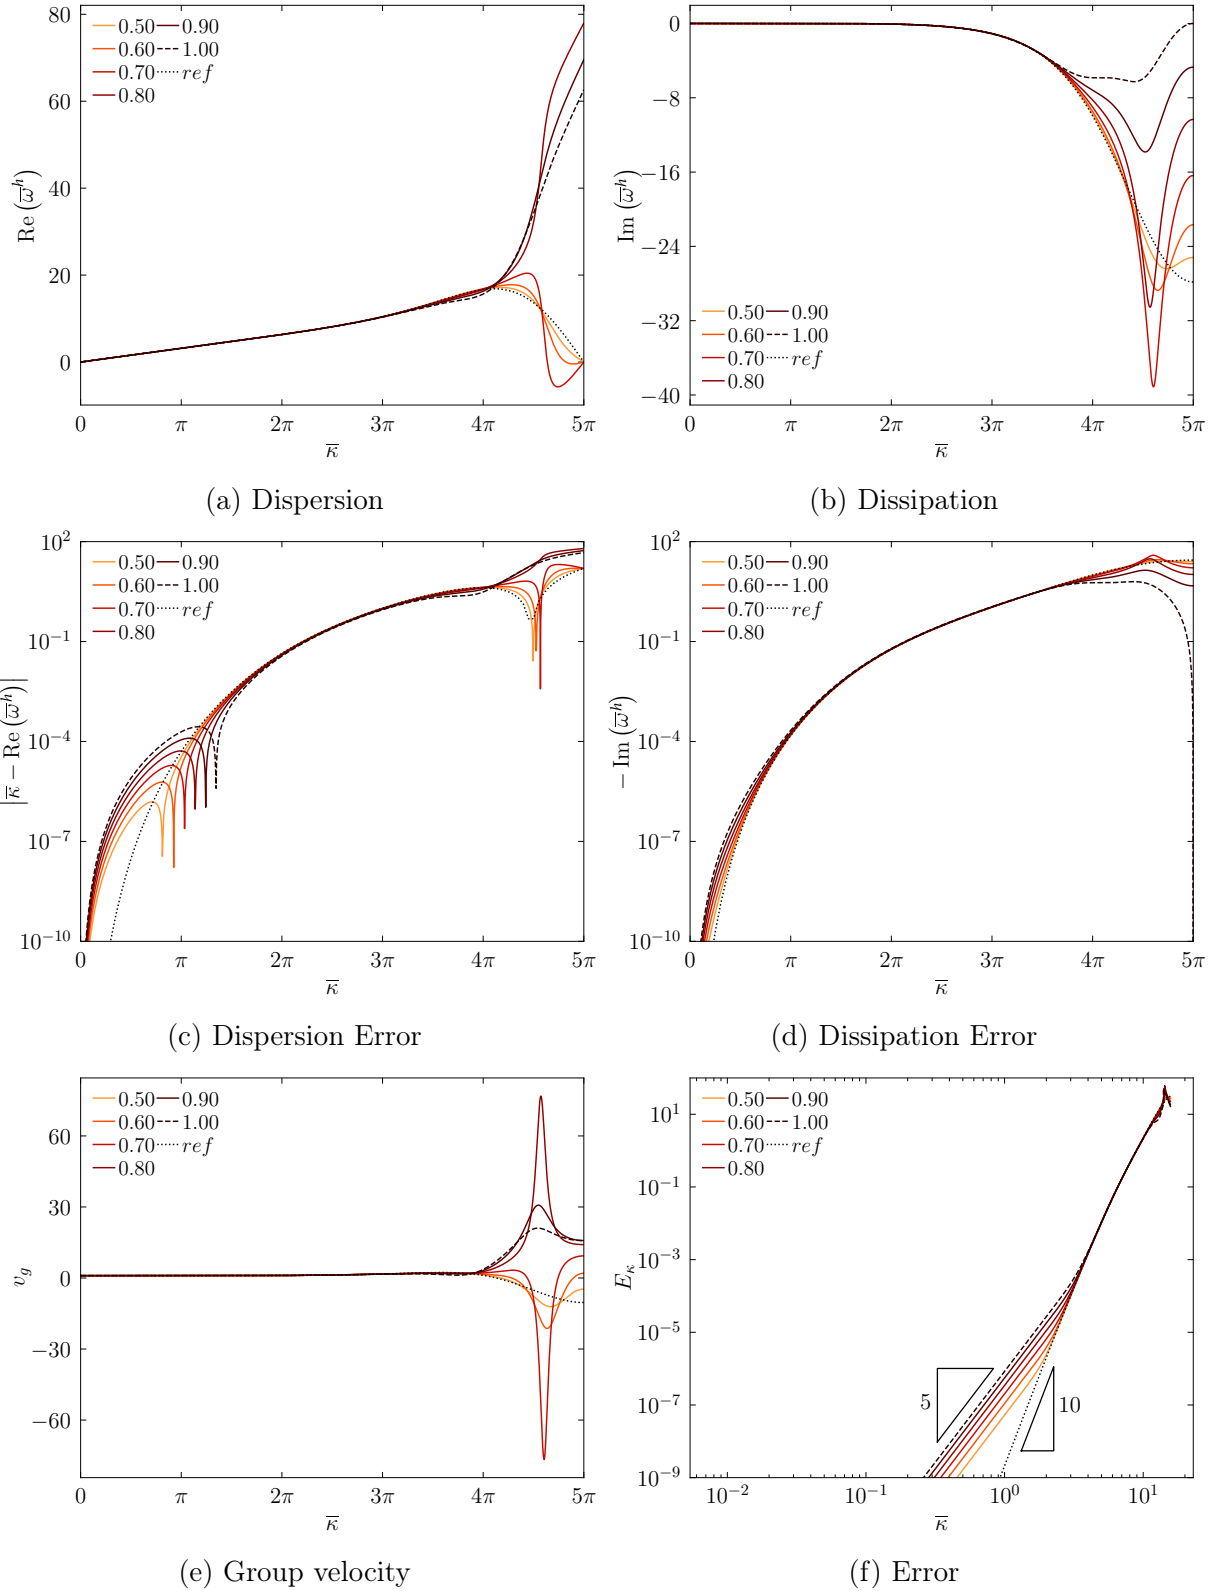


Figure 4.17. Spectral properties of $P_{4,4}$ schemes at different CFL numbers τ . Properties of the semidiscretization are shown in a dotted line

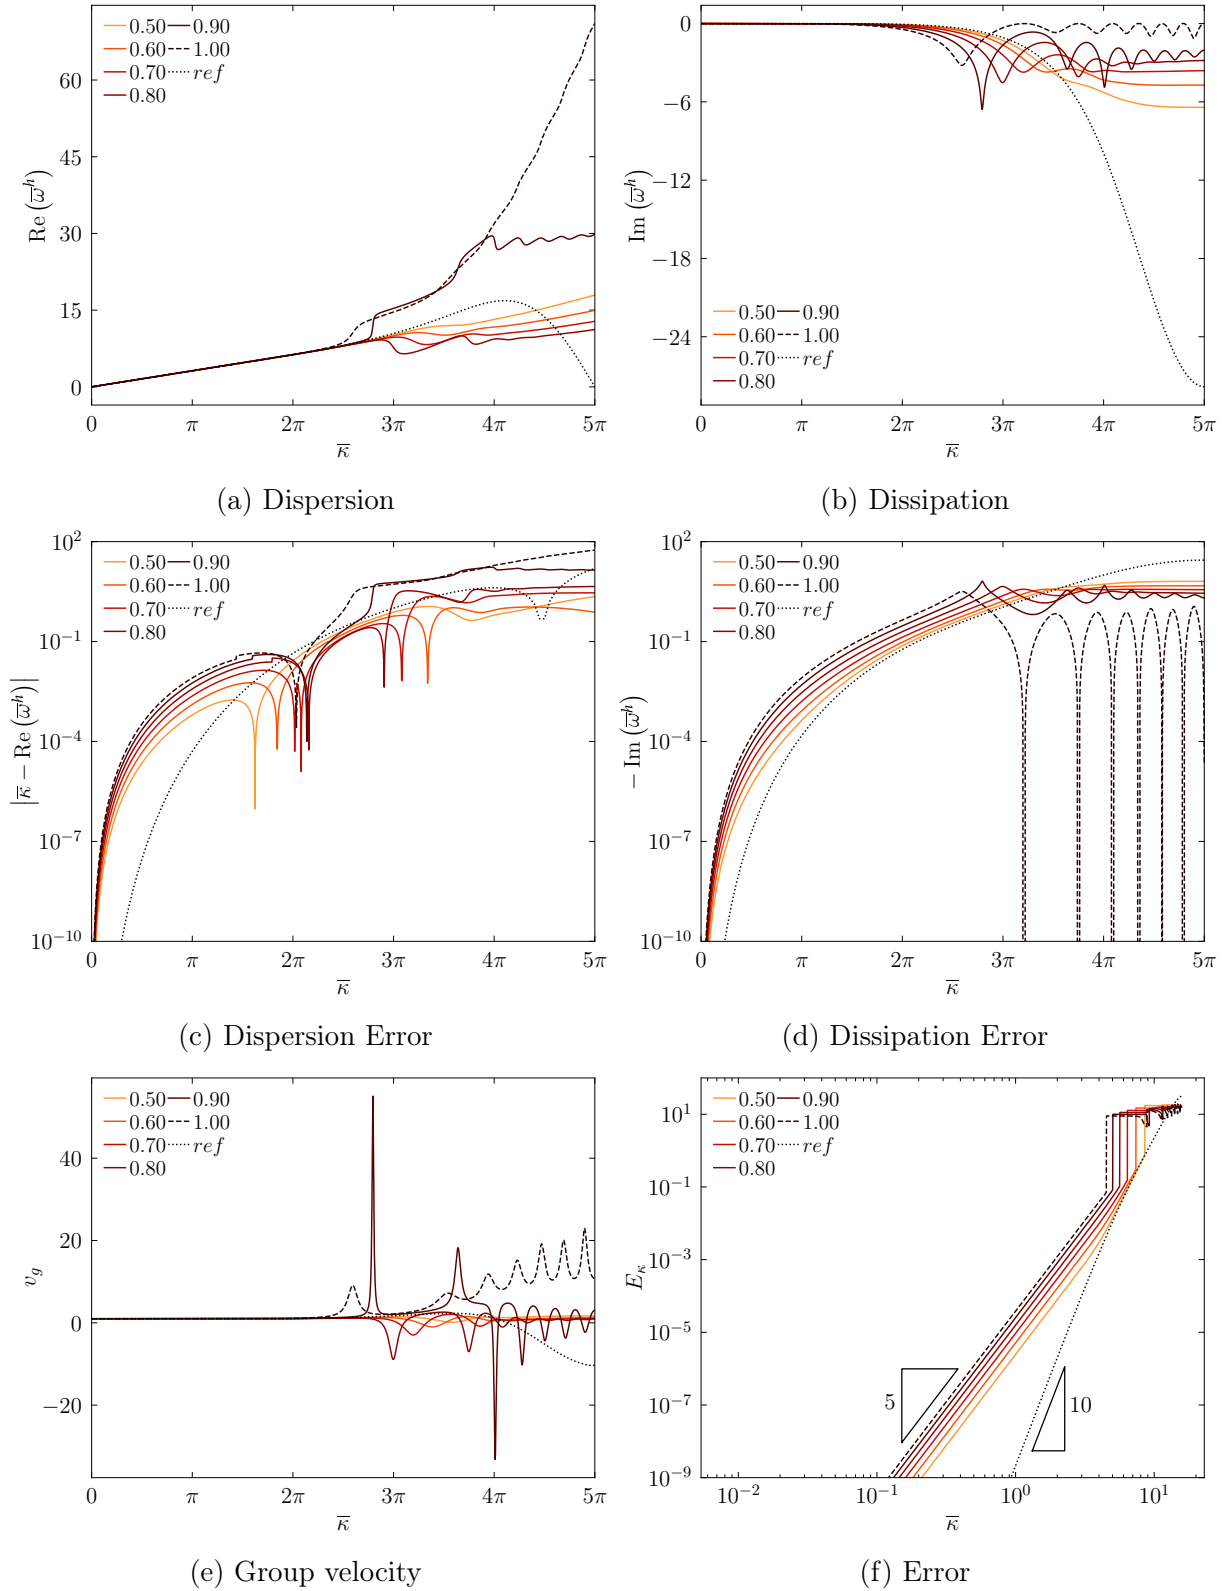


Figure 4.18. Spectral properties of $P_{16,4}$ schemes at different CFL numbers τ . Properties of the semidiscretization are shown in a dotted line

By optimizing the shape of the fourth-order polynomial to $P_{16,4}$, its spectral properties are presented in Figure 4.18. We observe that the dispersion properties are similar to those of the classical scheme but introduce significant additional error after $\approx 2\pi$. This point becomes closer to π as the time-step size increases. The shape of the $P_{16,4}$ method does not fully represent the shape of the eigenspectra, as previously shown in Figure 4.7. Hence, similar to the third-order optimized scheme, different amounts of dissipation are introduced for $\tau = 1.0$ in the high-wavenumber range. In general, the dissipation becomes significant after π for $\tau = 1.0$ and later for smaller Δt_{opt} , which have the opposite effect for wavenumbers $\bar{\kappa} > 4\pi$. We observe that small changes in the time-step size may cause large oscillatory behaviour in the group velocity of high-wavenumber features after $\bar{\kappa} = 2\pi$, but with smaller magnitude than the conventional scheme. The local order of the fully-discrete scheme is reduced to the temporal order in most of the wavenumber range from 10 to 5. Generally, the behaviour of the $P_{4,16}$ scheme can be compared with the classical fourth-order RK scheme in the range of low wavenumbers, especially as we move away from the maximum stable τ .

Figure 4.19 shows an eighth-order $P_{8,8}$ temporal scheme coupled with a fifth-order spatial discretization. The dispersion behaviour better follows the semidiscrete properties for a larger range of wavenumbers in comparison with the semidiscrete properties. This behaviour is maintained for all values of τ until at least 4π . Dissipation increases until about 4π and then gradually vanishes for $\tau = 1.0$ close to the Nyquist criterion. Group velocity features travel at the reference speed until close to $\bar{\kappa} = 4\pi$, after which large velocities can be expected at certain ranges of high wavenumbers. Contrary to the previously discussed temporal methods, the superconvergence property of the semidiscretization appears to be maintained. This may be due to temporal convergence beyond machine precision.

The spectral properties of $P_{16,8}$ schemes continue the trend, as shown in Figure 4.20. Dispersive behaviour occurs faster for higher values of τ , and the behaviour follows that of the semidiscretization up until about 2π for $\tau = 1.0$. Similar to the fourth-order OERK scheme, dissipative error is reduced in the high wavenumber regions for large τ , and the previously observed oscillatory behaviour at high wavenumbers is also present in this case. The group velocity follows the reference values for about 3π , after which large values may be predicted at given high wavenumbers.

We have observed that OERK schemes have a significant impact on the spectral properties of spatial discretizations. As $P(z)$ is optimized, the dissipative behaviour can be drastically reduced, with substantial introductions of dispersive behaviour. However, this is the case for time steps close to $\tau = 1.0$. As τ is reduced slightly below $\tau = 1.0$, the spectral properties approach those of the semidiscretization and contain similar characteristics to classical RK methods. Furthermore, we have shown that increasing the order of the temporal method improves the spectral properties of the fully-discrete scheme.

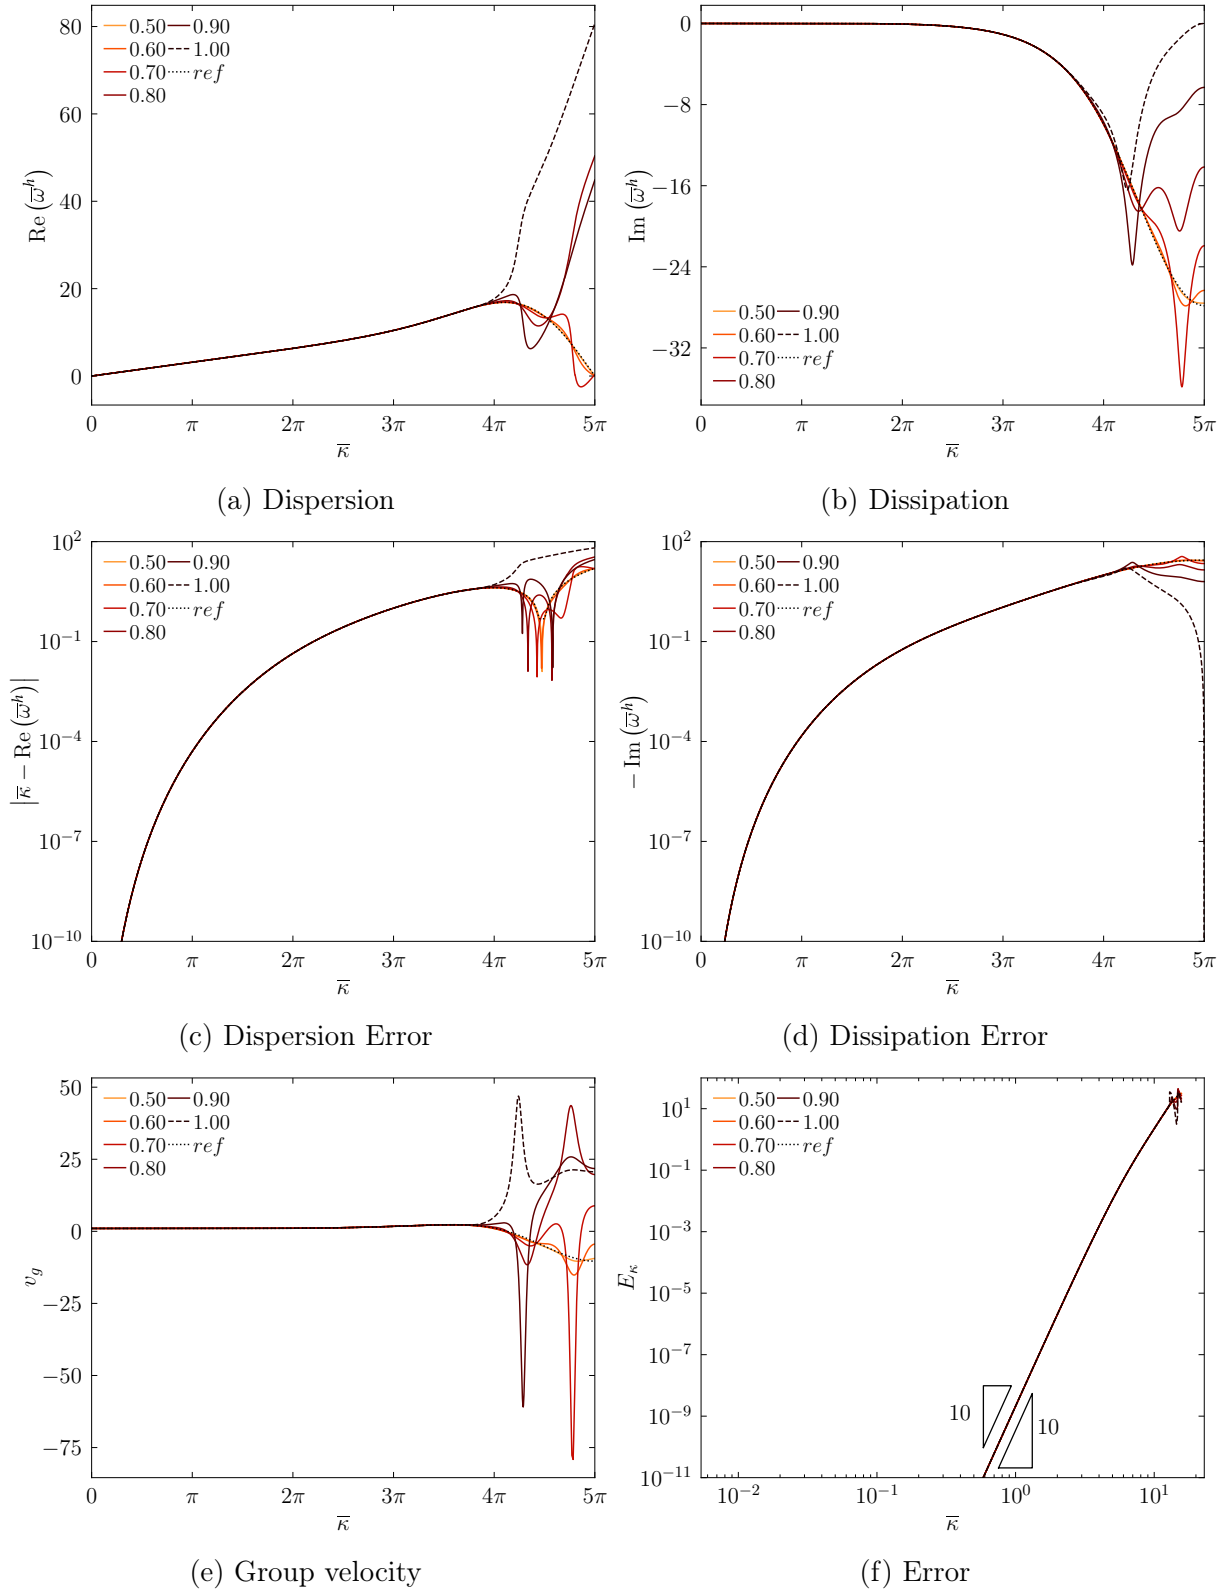


Figure 4.19. Spectral properties of $P_{8,8}$ schemes at different CFL numbers τ . Properties of the semidiscretization are shown in a dotted line

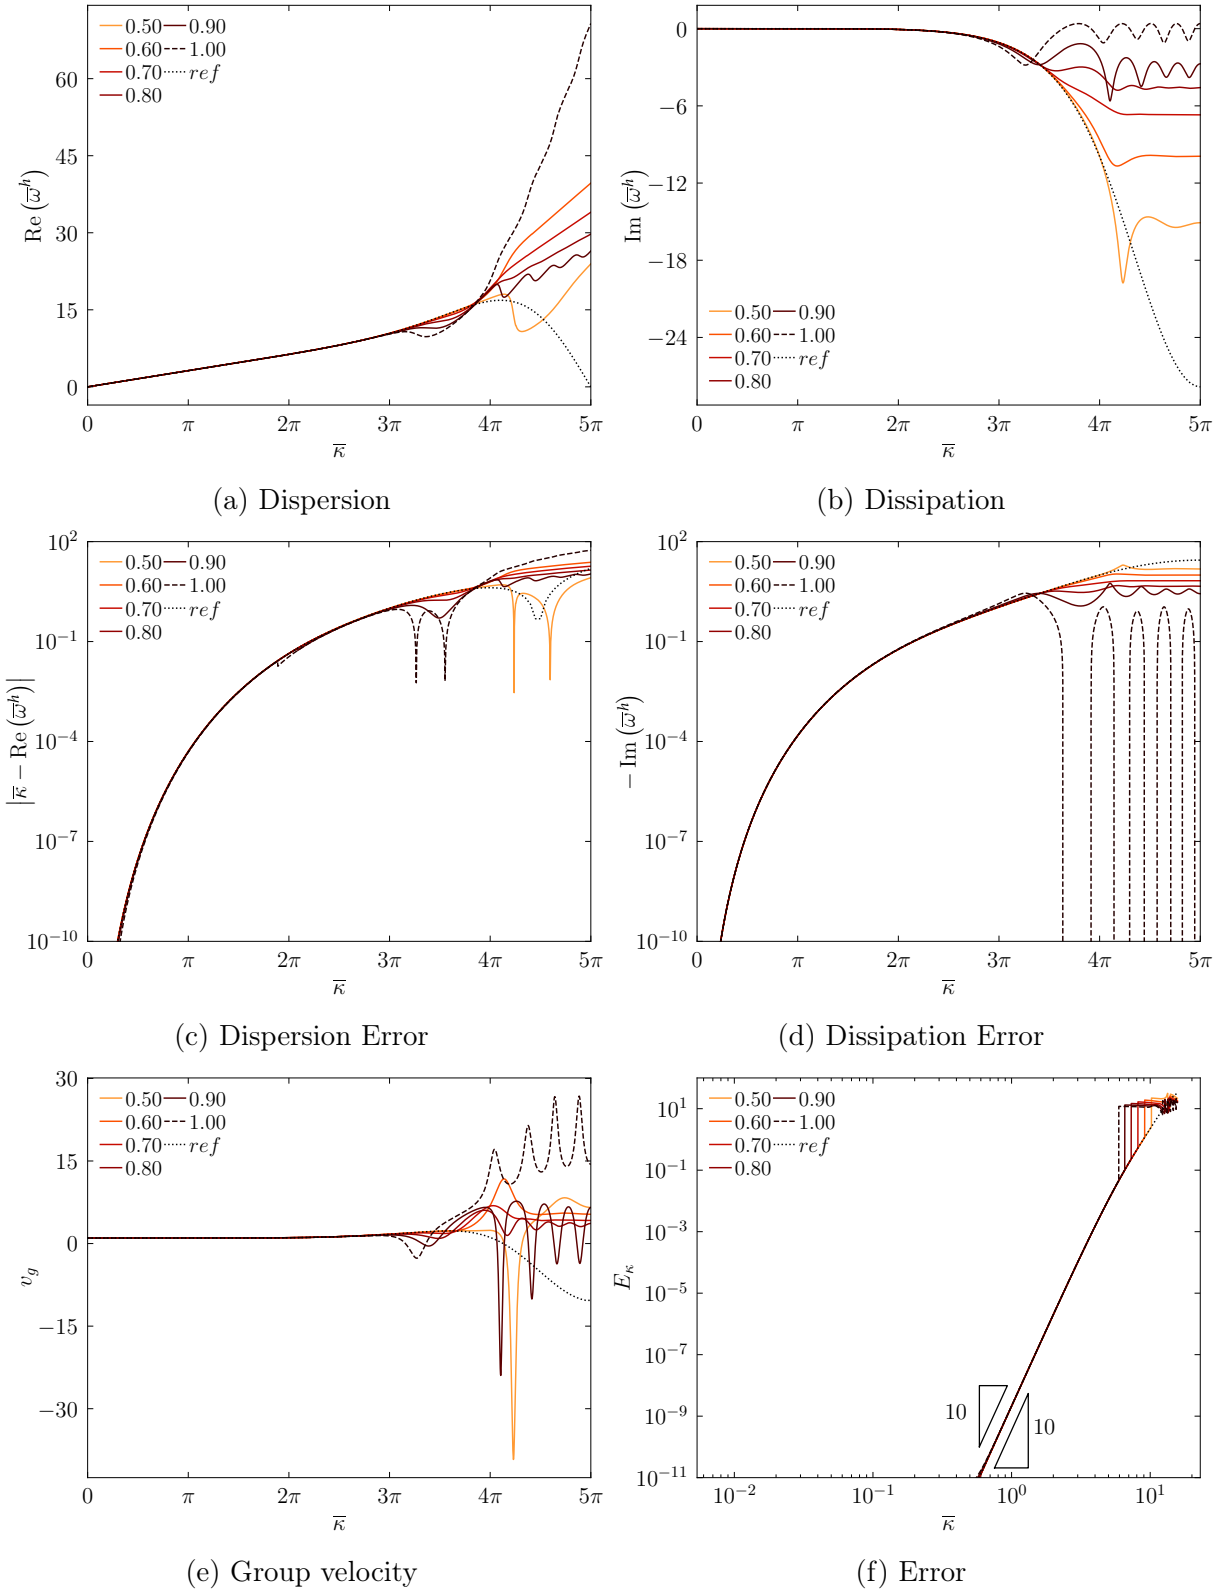


Figure 4.20. Spectral properties of $P_{16,8}$ schemes at different CFL numbers τ . Properties of the semidiscretization are shown in a dotted line

4.3.2 Numerical Experiments

Following the recommended numerical experiments in [53], we present advection of a Gaussian bump, advection of a sine wave, as well as one-dimensional viscous Burgers turbulence. We generated Butcher tableaux of orders 1, 2, 3, and 4, available in [118]. Generating Butcher tableaux for high-order temporal schemes can be particularly challenging due to the involved large number of order conditions and were omitted for the nonlinear Burgers test case.

4.3.2.1 Advection of a Gaussian Bump

We first evaluate the influence of the time-step size and number of optimized coefficients in advection problems. We simulate advection of a Gaussian bump on a computational domain of $x \in [-20, 20]$ divided into $N = 40$ elements. An initial condition of a superposition of wavenumbers is prescribed as $u(x, 0) = e^{-10x^2}$. Periodic boundary conditions, an upwind flux and Gauss points are used with a fifth-order spatial FR discretization ($p = 4$). Hence, we are able to see the behaviour of different OERK schemes for low and high wavenumber features. Since $h = 1$, we evaluate the spectral properties using ratios of the maximum time-step size τ_{CFL} , defined by $\tau = \Delta t / \tau_{CFL}$. We compare our results against the behaviour of the semidiscretization, using a very small time step size. Initially, results are shown at $t \approx 1$, and then for $t = 4000$. We take advantage of the linearity of the advection equation to advance the solution using the temporal matrix described in Equation (4.43). Hence, no Butcher tableau was required for this case.

It is known that spurious high-wavenumber features may quickly dissipate due to numerical dissipation. First, we look at the results for the advected solution after $t \approx 1$ for the low-order schemes $P_{16,1}$ and $P_{16,2}$ in Figure 4.21. Fast-travelling nonphysical features can be observed after the bump, which correspond to high-frequency information. We now increase the order of the temporal scheme to 3, 4, and 8. Figure 4.22 shows the solution for both classical and optimal RK schemes on the left and right sides, respectively. In general, the amount of dissipation added to the low-wavenumber features appears to be higher, more so for large values of τ , as shown in our fully-discrete results. Compared to the conventional methods, optimal RK schemes show fewer dissipated nonphysical features with slower-travelling group velocity. This agrees well with the findings in the spectral analysis, where the maximum value of v_g for large wavenumbers is generally smaller for the optimized schemes.

After $t = 4000$, we plot the results again. This time, we observe that the error from the semidiscretization alone, represented by the dashed black line, is significant. The results are shown in Figure 4.23 for the first and second-order schemes. High-frequency nonphysical modes can still be observed for $\tau = 1.0$. As shown in our spectral results, the amount of




Figure 4.21. Solution for advection of a Gaussian bump using first and second-degree OERK schemes at $t = 1$ at different time-step sizes. Reference solution is plotted using a dashed line. Results between these two schemes are very close due to the optimization problem

dissipation for these schemes at large τ is almost negligible. On the other hand, increased low-wavenumber dissipation has been added for the smaller time-step sizes. Higher-order schemes are shown at $t = 4000$ in Figure 4.24. First, we note that spurious high-frequency information has been dissipated. In addition, low-wavenumber features appear to be more damped for optimal and lower-order schemes in relation to the reference solution. This dissipation is dependent on the time-step size and is reduced with τ . For the optimal fourth-order scheme, nonphysical oscillations corresponding to the middle range of wavenumbers appear on the left side of the bump, different from the classical $P_{4,4}$ scheme, which shows these features both in front and after, similar to the reference solution. This is due to the underpredicted group velocity at lower wavenumbers with the optimal scheme, as shown in the spectral analysis.

4.3.2.2 Sine Wave

We perform a convergence study to compare the accuracy properties presented in the fully-discrete analysis. This test case has been used previously [130, 53, 31] to analyze the properties of the numerical scheme, independent of the initial error projection. To achieve this, we consider the simple advection of a sine wave on a computational domain $x \in [0, 20]$. We implement a time-dependent boundary condition at $x = 0$ of $u^h(0, t) = \sin(\pi t/2)$. No boundary condition is required at $x = 20$ since the scheme is fully upwinded. The initial conditions on the domain are set to zero. After four units of time, a single cycle has formed in the domain with wavelength $\lambda_s = 4$. Due to the cyclical nature of the boundary condition, it is expected that every four units in space contain a repeated complete sinusoid. The element

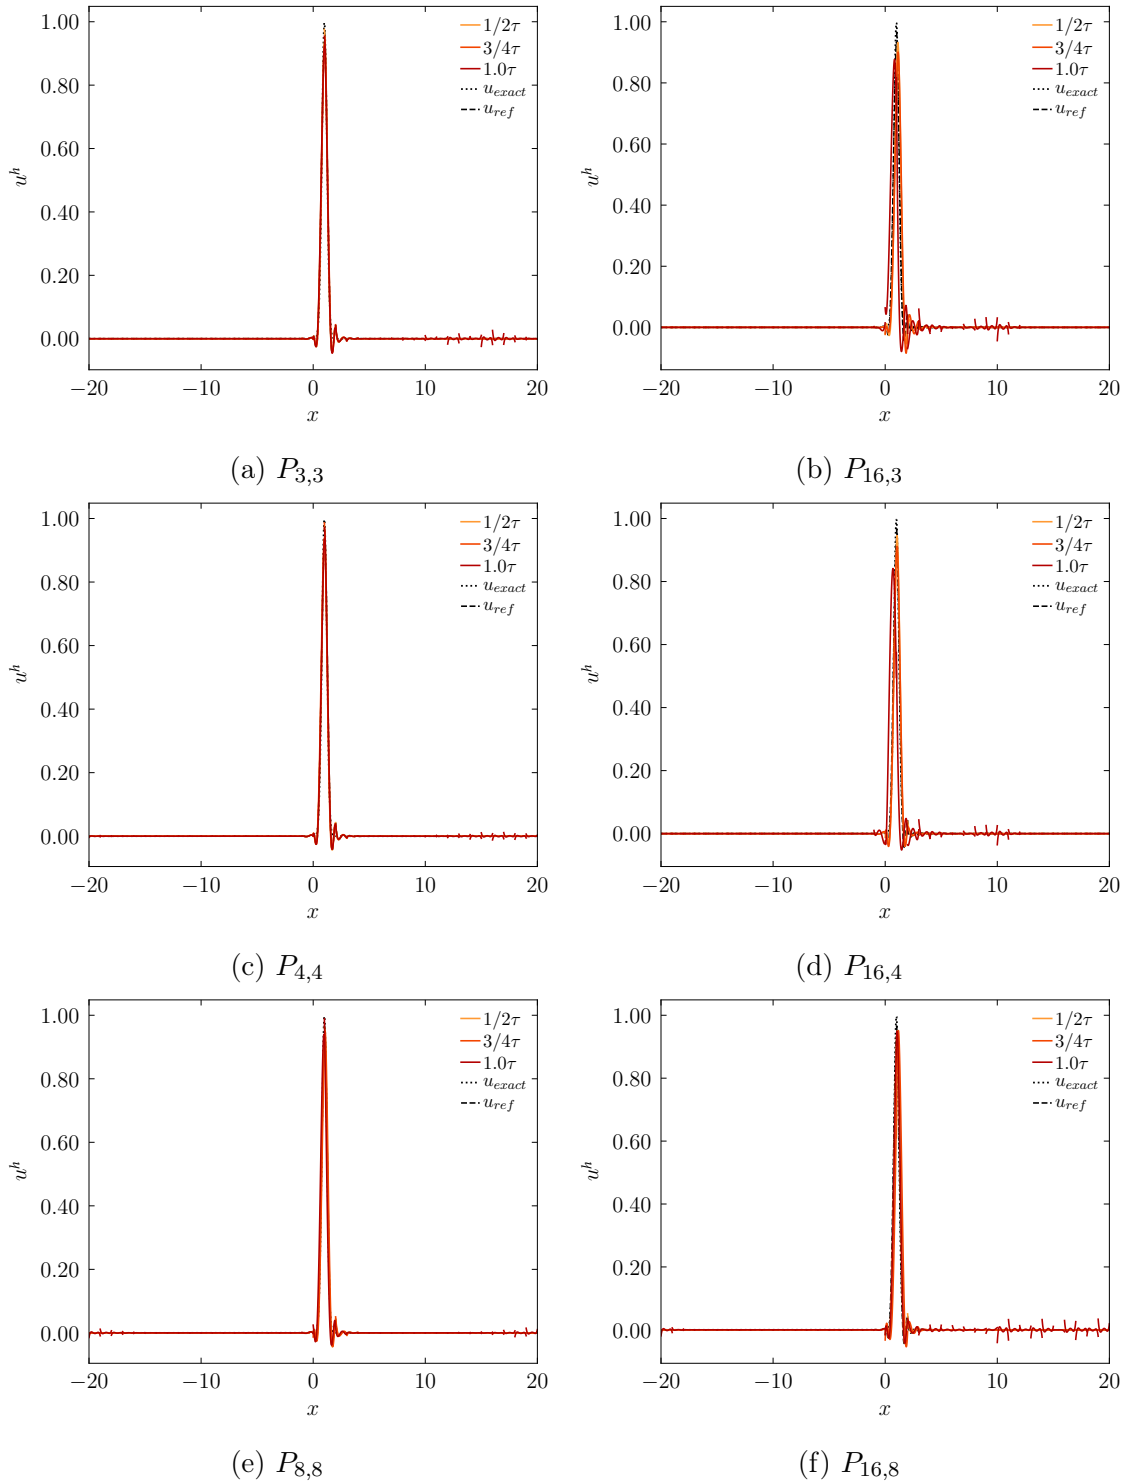


Figure 4.22. Advection of a Gaussian bump using OERK schemes at $t = 1$ using high-order temporal schemes. Results are shown at three different time-step sizes and contrasted with the semidiscrete solution in a dashed line. Differences in group velocities at high wavenumbers can be observed between classical and optimal schemes




Figure 4.23. Advection of a Gaussian bump using first and second-degree OERK schemes at $t = 4000$. Results are shown for different values of τ in continuous lines and the semidiscrete solution in a dashed line. Solution is significantly dissipated relative to the semidiscrete solution

length h is determined from the exact dispersion relation $\bar{\kappa} = \omega h$, where $\omega = \pi/2$. We present results for wavenumbers $\bar{\kappa} = 2\pi/n^2$, $n = 1, \dots, 5$. After $t = 100$, we compute the L_2 norm between the first and third sinusoids by

$$L_2 = \sqrt{\frac{h}{2} \sum_{i=1}^{\lambda_s/h} \int_{-1}^1 \left[u_i^h(\tilde{x}, t) - u_{i+\lambda_s/h}^h(\tilde{x}, t) \right]^2 d\tilde{x}}. \quad (4.49)$$

Table 4.7 shows the global convergence of the error for the FR $p = 4$ scheme coupled with temporal OERK schemes of order $q = 1, 2, 3$, and 4 . Similar to the spectral results, we note that the order of accuracy of the fully-discrete scheme is reduced from $A_g \approx 2p + 1$ to q for the range of well-resolved wavenumbers. Superconvergence is observed for all considered schemes with global orders of ≈ 7 or higher. Results for the $P_{16,1}$ and $P_{16,2}$ show similar behaviour for the reasons mentioned in the previous sections. Our results for the third and fourth-order classical schemes agree well with those provided in [53]. We show that optimal polynomials also converge to the expected orders of accuracy and can still achieve superconvergence when a sufficiently small time-step size is used.

4.3.2.3 Burgers Turbulence

Forced viscous Burgers turbulence is generally used as a simple one-dimensional test for numerical schemes for ILES [53, 131]. It is known that the FR scheme inherently adds the necessary dissipation for the high-frequency scales present in turbulent flows [53, 39,

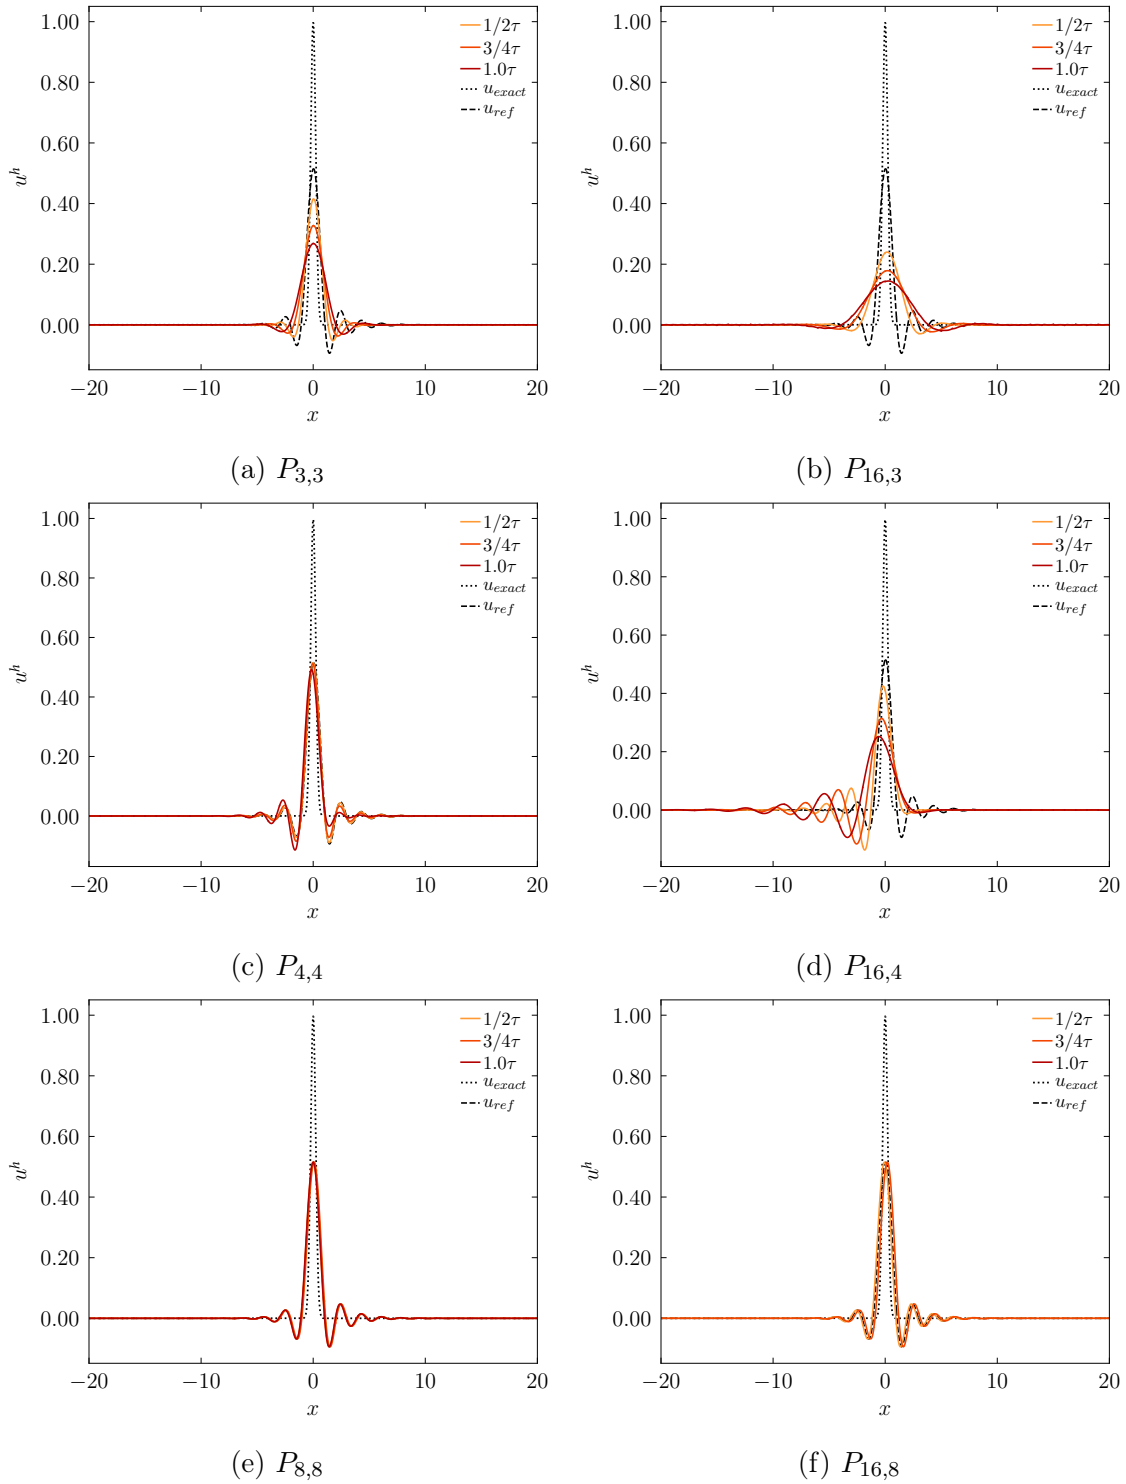


Figure 4.24. Advection of a Gaussian bump using OERK schemes at $t = 4000$. Results are shown for different values of τ in continuous lines and the semidiscrete solution in a dashed line. Optimizing ERK schemes can modify the spectral properties of the semidiscretization. Higher-order schemes show better agreement with semidiscrete solution

Table 4.7. Convergence of the L_2 norm for the sine advection case

$P_{16,1}$		$P_{16,2}$		
$\bar{\kappa}$	Error	A_g	Error	A_g
2π	2.60×10^{-1}	-	2.60×10^{-1}	-
π	2.10×10^{-3}	6.95	2.10×10^{-3}	6.95
$\pi/2$	3.11×10^{-4}	2.75	3.11×10^{-4}	2.76
$\pi/4$	7.76×10^{-5}	2.00	7.76×10^{-5}	2.00
$\pi/8$	1.94×10^{-5}	2.00	1.94×10^{-5}	2.00
$\pi/16$	4.85×10^{-6}	2.00	4.85×10^{-6}	2.00
$P_{3,3}$		$P_{16,3}$		
$\bar{\kappa}$	Error	A_g	Error	A_g
2π	2.58×10^{-1}	-	2.58×10^{-1}	-
π	1.34×10^{-3}	7.58	1.34×10^{-3}	7.58
$\pi/2$	3.72×10^{-6}	8.50	3.20×10^{-6}	8.71
$\pi/4$	7.38×10^{-8}	5.65	8.02×10^{-9}	8.64
$\pi/8$	8.42×10^{-9}	3.13	2.01×10^{-10}	5.32
$\pi/16$	1.05×10^{-9}	3.00	3.37×10^{-11}	2.60
$P_{4,4}$		$P_{16,4}$		
$\bar{\kappa}$	Error	A_g	Error	A_g
2π	2.57×10^{-1}	-	2.58×10^{-1}	-
π	1.32×10^{-3}	7.60	1.34×10^{-3}	7.59
$\pi/2$	9.67×10^{-6}	7.09	3.15×10^{-6}	8.72
$\pi/4$	6.27×10^{-7}	3.95	1.46×10^{-8}	7.80
$\pi/8$	4.37×10^{-8}	3.84	9.38×10^{-10}	3.93
$\pi/16$	3.20×10^{-9}	3.77	5.99×10^{-11}	3.97

132]. We evaluate the effects of the numerical scheme on the characteristic behaviour of the semidiscretization with this simple model. The viscous Burgers equation is given by

$$\frac{\partial u}{\partial t} + \frac{1}{2} \frac{\partial u^2}{\partial x} = \beta \frac{\partial^2 u}{\partial x^2} + s(x, t), \quad (4.50)$$

where β is a viscosity coefficient and is set to 10^{-4} . A forcing function $s(x, t)$ appears on the right-hand side and is specified as in [53, 17]. It is defined in wavenumber space

$$s(\kappa) = A \frac{e^{i\iota}}{\sqrt{\kappa} \sqrt{\Delta t}}, \quad (4.51)$$

where κ is the wavenumber, A is an amplification factor, ι is a randomly chosen number, $\iota \in [-\pi, \pi]$, computed at every time step for every wavenumber. In this work, we use $A = 0.01$.

The computational domain is set on $x \in [-\pi, \pi]$ using 20 elements and periodic boundary conditions. Initially, the solution is prescribed as $u(x, 0) = 0$, and the source term is allowed to introduce random behaviour into the results. We use the LDG approach for the viscous fluxes and a Roe scheme for the Burgers flux. Anti-aliasing is introduced to the Burgers flux to eliminate aliasing errors at high wavenumbers (see Hesthaven [9]). We compute the energy spectra according to

$$E(\kappa) = \frac{1}{2} |u(\kappa)|^2, \quad (4.52)$$

where $u(\kappa)$ can be obtained by computing the FFT of the solution. We generated third and fourth-order Butcher tableaux for $P_{s,q}$ schemes of degrees $s = p + i$, $i = 0, \dots, 4$ and s number of stages. Table 4.8 shows the time-step size τ_{max} that was used for each scheme. This number represents the time-step size that maintained at least 50% of stable simulations. We note that there is not a unique tableau for a given $P_{s,q}$ scheme, and then τ_{max} and results hereafter correspond to our generated Butcher tableaux, which can be found in the electronic supplementary material of [118]. For all schemes, we compare the results for simulations performed at $\tau = 0.9, 0.92, 0.94, 0.96, 0.98$ and $1.0\tau_{max}$ against a reference result computed with a small time-step size, which is shown in a dashed line.


Table 4.8. Maximum time-step sizes for viscous Burgers turbulence using $P_{s,4}$ schemes

s	4	5	6	7	8
τ_{max}	1.88	2.70	3.81	4.62	5.56


Figure 4.25 shows the energy spectrum for a fourth-order optimal RK scheme. We include an example solution, where the turbulent behaviour, including discontinuities, can be observed. Here, we explore the effects of the time-step size and the degree of the RK

polynomials. It is known that FR provides the necessary dissipation to resolve turbulent flows using ILES [37]. It is clear from the spectral plots in previous sections that the $P_{4,4}$ scheme is almost nondissipative for the highest wavenumber features. This is observed in the concentration of energy at these scales. By increasing the number of optimized coefficients, the spectral behaviour changes. In particular, it was observed that, on average, the dissipation decreases for the high-wavenumber region and increases for the low to mid-range. However, the amount of dissipation in the wavenumber space varies. This is consistent with the oscillatory behaviour observed in the fully-discrete analysis section. Furthermore, we observe that using a time-step size of at least 90% of the stable maximum introduces clear improvements in the dissipation properties. By moving away from the stability limit by a small percentage, it is expected that the optimal RK schemes will provide appropriate spectral characteristics to be implemented as FR-OERK schemes for ILES. For completeness, we include the results for $P_{s,3}$ schemes in Figure 4.26.


With these numerical examples, we have demonstrated the suitability of optimal RK methods for FR schemes in the context of turbulent simulations. With speedup factors close to two, optimal RK methods enable faster simulations of problems with moderately low stiffness while maintaining the characteristic spectral properties of FR methods. This is true so long as the time step size is chosen close to but not at the limit of stability. In the next chapters, we will focus on developing solution acceleration strategies for implicit discretizations.




(a) Example solution for $P_{8,4}$




(b) $P_{4,4}$




(c) $P_{5,4}$



(d) $P_{6,4}$



(e) $P_{7,4}$



(f) $P_{8,4}$

Figure 4.25. Energy spectrum for viscous Burger turbulence using classical and optimized fourth-order RK schemes. Results using a very small time-step size are included as reference data (dashed line). A dependence on s can be observed in terms of energy dissipation at high wavenumbers

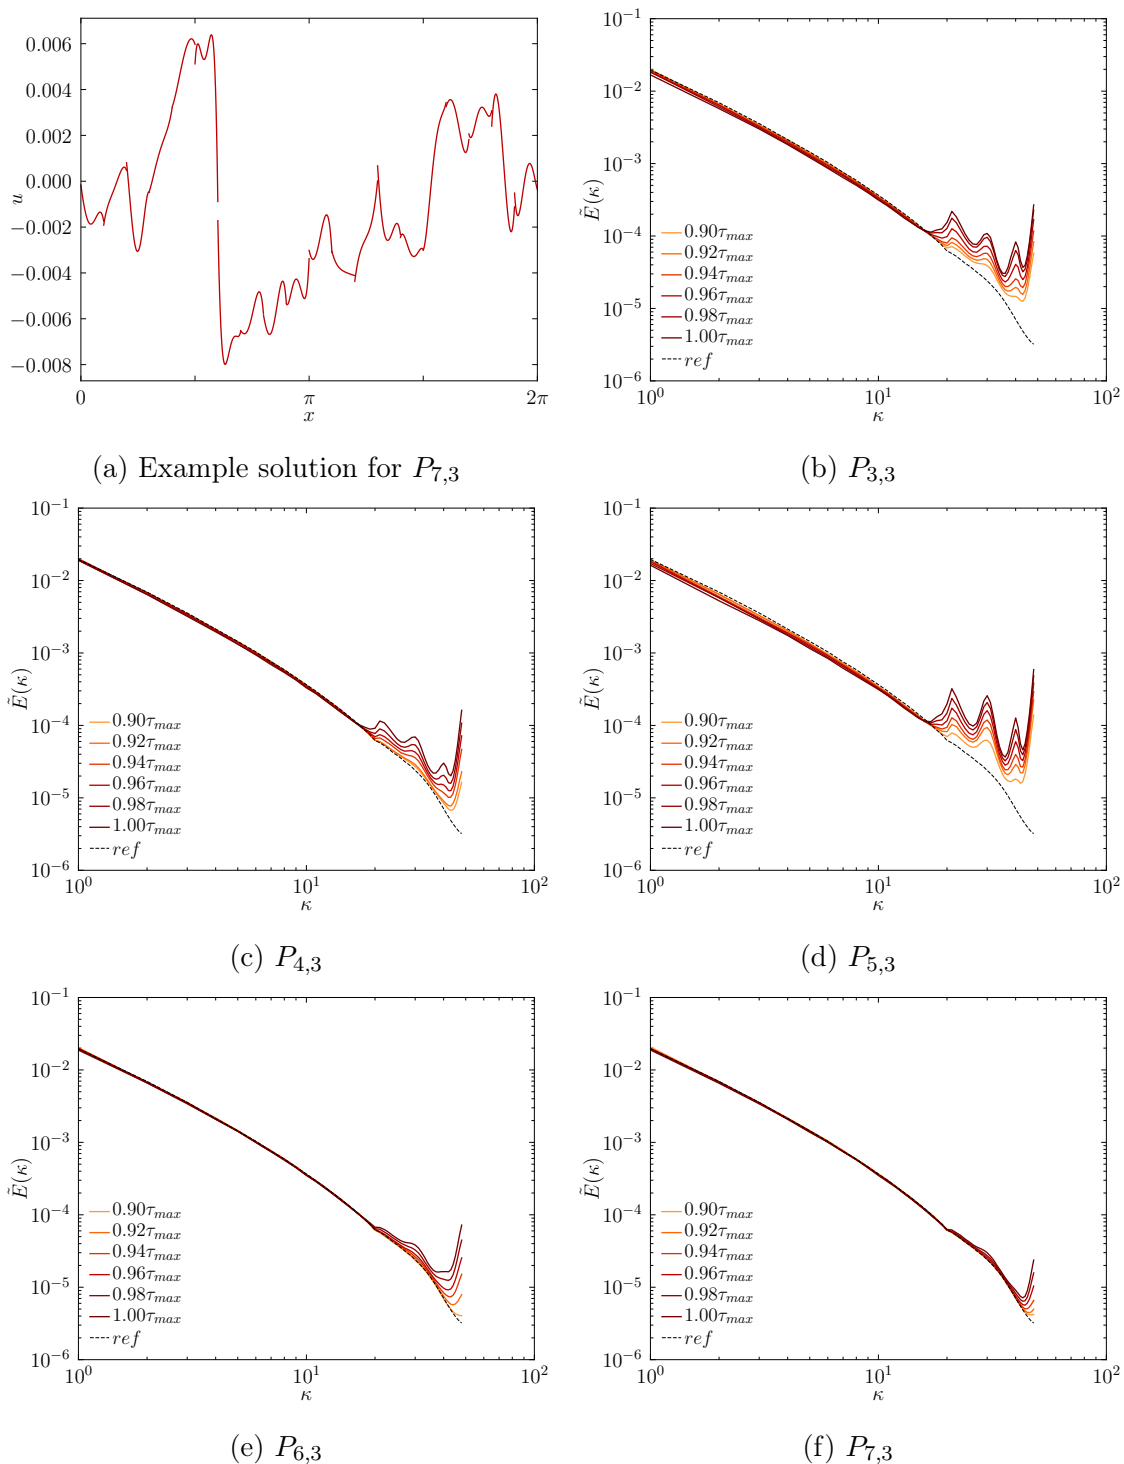


Figure 4.26. Energy spectrum for viscous Burger turbulence using classical and optimized third-order RK schemes. Results using a very small time-step size are included as reference data (dashed line). The behaviour of classical and OERK schemes is generally similar away from the stability limit

Chapter 5

Implicit Hybridized Formulations

5.1 Overview

In the previous chapter, we developed optimal explicit Runge-Kutta methods, where computations of nonstiff problems were accelerated to yield accurate and stable simulations with a speedup factor of two at most. While a factor of two is significant, we need to enable faster computations, especially for problems that may contain disparate scales that introduce stiffness, such as wall-resolved turbulence, for which explicit methods are not suitable. We continue the work in this thesis by investigating performance and accuracy of hybridization in the context of implicit formulations of FR methods. The description and implementation details of the schemes considered here can be found in Chapter 5.

5.2 Advection-Type Problems

5.2.1 Spectral Analysis

In this section, we consider the spectral behaviour of hybridized methods. These properties have been obtained for conventional FR methods [31, 36] and have provided important insights into the dispersion and dissipation properties of numerical schemes. This analysis will establish the behaviour of HFR and EFR methods for a range of c -parameters. We will show that HFR methods recover the conventional FR schemes for linear advection, and that EFR methods are identical to HFR methods in one-dimensional flow orientations. To properly characterize the behaviour of EFR methods, we perform a two-dimensional von Neumann analysis. The procedure is similar to that described in the previous chapter for the fully-discrete analysis, but additional steps are required to see the effects of hybridization. In the case of EFR, we consider a collocation of Gauss-Lobatto points for integration. A discussion on the influence

of this integration choice is included in Appendix B. We consider the linear advection equation

$$\frac{\partial u}{\partial t} + \boldsymbol{\alpha} \cdot \nabla u = 0, \quad (5.1)$$

which is nondispersive and nondissipative, where $\boldsymbol{\alpha}$ is the advection velocity with a prescribed modulus $|\boldsymbol{\alpha}| = 1$ and direction θ . The above equation admits solutions of the form

$$u(\mathbf{x}, t) = e^{I(\kappa \mathbf{1}_\theta \cdot \mathbf{x} - \omega t)}, \quad (5.2)$$

where κ is a wavenumber with orientation vector $\mathbf{1}_\theta = [\cos \theta, \sin \theta]$, ω is a temporal frequency and $I = \sqrt{-1}$. Here, we have chosen to align the advection and wave orientations for the sake of simplicity. For this analysis, a computational domain is subdivided into N quadrilateral elements of equal side length $h = 1$, and periodic conditions have been applied in all directions. Within each element, we seek numerical solutions of the form

$$u_k^h(\mathbf{x}, t) = e^{I(\bar{\kappa} \mathbf{1}_\theta \cdot \mathbf{x}_k^c - \bar{\omega}^h t)} \bar{\mathbf{u}}, \quad (5.3)$$

where $\bar{\kappa} \in [0, \pi]$ is a prescribed nondimensional wavenumber, \mathbf{x}_k^c is the center coordinate of Ω_k and $\bar{\mathbf{u}}$ is a vector of unknowns. Considering a Lax-Friedrichs [103, 102] type Riemann solver of the form

$$\hat{\mathfrak{F}}_{k,f}(u_{k,f}^h, \hat{u}_f^h) = \boldsymbol{\alpha} \hat{u}_f^h + |\boldsymbol{\alpha} \cdot \mathbf{n}_f^m| (u_{k,f}^h - \hat{u}_f^h) \mathbf{n}_f^m, \quad (5.4)$$

the problem can be written after discretization as the following linear system

$$\frac{d\mathbf{u}}{dt} + \bar{\mathbf{A}}\mathbf{u} + \bar{\mathbf{B}}\hat{\mathbf{u}} = \mathbf{0}, \quad (5.5)$$

$$\bar{\mathbf{C}}\mathbf{u} + \bar{\mathbf{D}}\hat{\mathbf{u}} = \mathbf{0}, \quad (5.6)$$

where $\bar{\mathbf{A}}, \bar{\mathbf{B}}$ and $\bar{\mathbf{C}}, \bar{\mathbf{D}}$ are the Jacobian matrices resulting from the hybridized discretization of Equation (5.1). We are interested in the effects of the semidiscretization on the interior solution. Finally, we can eliminate the trace variable $\hat{\mathbf{u}}$ and obtain a condensed system in terms of \mathbf{u} only

$$\frac{d\mathbf{u}}{dt} = \mathbf{M}\mathbf{u}, \quad (5.7)$$

where $\mathbf{M} = \bar{\mathbf{A}} - \bar{\mathbf{B}}\bar{\mathbf{D}}^{-1}\bar{\mathbf{C}}$. The block-circulant structure of \mathbf{M} allows us to consider a single element to analyze the scheme's properties. Inspection of this global semidiscrete operator demonstrates that the coupling between elements is stronger for the EFR method. From Figure 5.1, an element Ω_k is coupled with its neighbours via the faces in HFR methods, whereas an element in EFR is coupled with all elements sharing a vertex. In addition, HFR methods exactly recover the FR semidiscrete operator in (5.7), and hence it is equivalent in

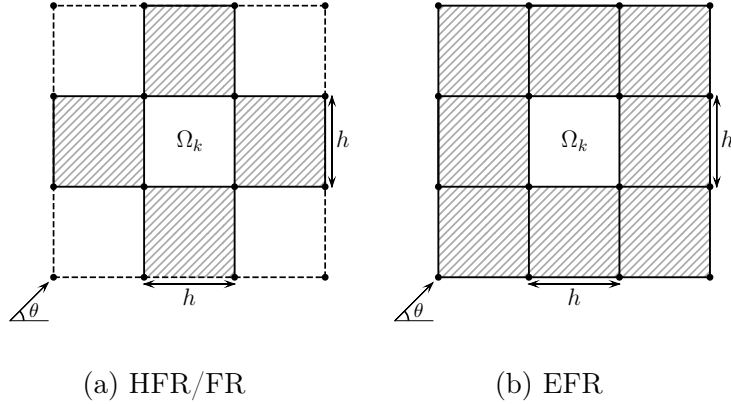


Figure 5.1. Element coupling involved in HFR and collocation-type EFR semidiscrete operators

the context of linear advection [133]. Inserting (5.3) in (5.7) yields an eigenvalue problem for any element Ω_k

$$-I\omega^h \bar{\mathbf{u}} = \bar{\mathbf{M}} \bar{\mathbf{u}}, \quad (5.8)$$

where

$$\bar{\mathbf{M}} = \sum_{i=1}^{N_b} \mathbf{C}_i e^{I(\bar{\kappa} \Delta \mathbf{x}_{i-k}^c \cdot \mathbf{1}_\theta)}. \quad (5.9)$$

In (5.9), \mathbf{C}_i is a Jacobian block associated with the i -th element involved in the coupling, and $\Delta \mathbf{x}_{i-k}^c$ is a vector holding the distance from \mathbf{x}_i^c to \mathbf{x}_k^c . $N_b = 5$ for HFR and $N_b = 9$ for EFR. The $(p+1)^2$ eigenvalues of $\bar{\mathbf{M}}$, $\lambda_j \in \mathbb{Z}$, are a function of κ , θ , $\boldsymbol{\alpha}$, c -parameter, and $\{\phi_i(\tilde{x})\}$. These eigenvalues are associated with numerical frequencies $-I\omega^h = \lambda$, which allow us to determine the dispersion ($\text{Re}(\omega^h) \neq \kappa$) and dissipation ($\text{Im}(\omega^h) < 0$) properties of the scheme. We follow [31, 36] for the sorting of these eigenvalues and to identify the dominant mode, which we use to make our analysis. We nondimensionalize figures in this section considering $\bar{\kappa} = \frac{\kappa}{(p+1)} \cos \theta$.

Of interest in this work is the comparison between the accuracy of HFR and EFR methods. Figures 5.2 and 5.3 show the dissipation and dispersion relations from $p = 1$ to $p = 4$ for different values of $c \in [\frac{c_-}{2}, c_{HU}]$ considering a wave incidence angle of $\theta = \frac{\pi}{6}$. Plots for other values of θ are relatively similar and are omitted for the sake of brevity. Increasing the value of c introduces more numerical dissipation for low to moderate wavenumbers for all methods. In the range of high wavenumbers, increasing values of c have less dissipation instead. Results for c_{DG} are shown in a continuous line, c_{SD} in a dashed line, and c_{HU} in a dash-dotted line. EFR methods introduce significantly more numerical error at $p = 1$ compared to HFR. As the polynomial degree is increased, it can be observed that these dissipation and dispersion mechanisms begin to behave similarly.




Figure 5.2. Dissipation curves for different c -parameters at $\theta = \pi/6$ for different polynomial degrees. Results are highlighted for c_{DG} (—), c_{SD} (---) and c_{HU} (-.-)

Whether one hybridization introduces more or less dissipation at given wavenumbers is also a function of the incidence angle. When $\theta = 0$, the numerical traces located at the top and bottom faces of the quadrilaterals do not play a role in the EFR formulation anymore, and the scheme becomes equivalent to HFR [134, 107]. To visualize the behaviour of these methods at multiple values of θ , let us consider the maximum resolved wavenumber $\bar{k}_{1\%}$, at which a wave decay of 1% is expected. This parameter is a reliable indicator of the resolution capabilities of a scheme [122]. Figure 5.4 shows the maximum resolved wavenumber against the wave orientation for multiple values of c increasing in the direction of the arrows displayed in the plots. For $p = 1$, the resolution power of EFR is greatly affected for $\theta > 0$. For example, EFR with c_{DG} shows significantly more dissipation than its HFR counterpart. As we increase p , results show that the behaviour of EFR, which is less computationally expensive, starts





Figure 5.3. Dispersion curves for different c -parameters at $\theta = \pi/6$ for different polynomial degrees. Results are highlighted for c_{DG} (—), c_{SD} (---) and c_{HU} (···)


to behave similarly to conventional FR or HFR schemes, especially at higher orders. Note that the maximum resolved wavenumber varies by less than 1% between the hybridization methods at $p = 4$. Despite having more or less dissipation at given values of θ than HFR, EFR is generally more dissipative.

5.2.2 Numerical Results


In this section, we explore the performance and accuracy of hybridized flux reconstruction methods and compare them against conventional implicit FR formulations. All numerical experiments were carried out serially on a 3.2GHz Intel Core i5-5600 processor with 16Gb of RAM. We present results for linear and nonlinear cases.




(a) $p = 1$



(b) $p = 2$



(c) $p = 3$



(d) $p = 4$

Figure 5.4. Maximum resolved wavenumber $\bar{\kappa}_{1\%}$ against the wave incidence angle θ for different parameters $c \in [c_-/2, c_{HU}]$. Arrows point toward increasing values of c . Specific values of c are highlighted for c_{DG} (—), c_{SD} (- - -) and c_{HU} (- · - ·)

5.2.2.1 Linear Advection

First, we analyze the behaviour of hybridized FR methods in the context of linear advection. As stated in the spectral analysis section, HFR schemes exactly recover the solution from the conventional FR methods for linear flux functions. For this problem, we consider a periodic ($\mathfrak{F}^{\text{BC}} = 0$) domain $\Omega \in [0, 2]^2$ with four levels of refinement of 5×5 , 10×10 , 20×20 , 40×40 quadrilateral elements. An initial Gaussian function is used to initialize the simulation, given by

$$u(\mathbf{x}, t) = e^{-20(x-1)^2 - 20(y-1)^2}. \quad (5.10)$$

Considering $\boldsymbol{\alpha} = [1, 1]$ and the Riemann solver in Equation (5.4), the L_2 norm of the solution error is computed for verification after the Gaussian bump has completed one advective cycle. Results are shown in Table A.1 for HFR and EFR with different c -parameters, c_{DG} , c_{SD} and c_{HU} . Both HFR and EFR methods display the expected $p + 1$ order of accuracy. As expected, all simulations using c_{DG} were able to obtain the smallest errors compared to higher values of c , which, when increasing, introduce additional numerical error. This behaviour occurs for both HFR and EFR simulations. When comparing both hybridization methods, HFR obtained the smaller L_2 norm except for a few configurations. As the polynomial degree increases, the errors become closer to each other, and both approaches behave in a more similar manner, with EFR generally having larger error.

We can measure the profile decay caused by numerical dissipation by measuring the L_2 norm of the solution. In order to do this, we carry out an additional set of simulations on a 9×9 grid considering the same initial condition. On quadrilateral elements, the numerical error propagates symmetrically [36]. Looking at a slice of the solution will be enough to analyze this Gaussian profile. Figure 5.5 shows a plot over $y = 1$ of the solution at $t = 20$ for the c_{DG} , c_{SD} and c_{HU} values. A slice of the exact solution is included in a dashed line for reference. When $c = c_{DG}$, the numerical solution is the closest to the reference data. For the other two methods c_{SD} and c_{HU} , higher numerical dissipation and dispersion are present, similar to the observations for the one-dimensional problem in [29]. Both the solutions for HFR and EFR present the same error characteristics, with the latter resulting in a slightly more damped Gaussian profile. However, EFR with c_{DG} is more accurate than HFR with c_{SD} and with c_{SD} . A measure of the solution energy $\|u^h\|^2$ at the final time is shown in Table 5.1 confirming the previous observations.

Table 5.1. Energy $\|u^h\|^2$ at $t = 20$ for the Gaussian profile case with $p = 3$

	c_{DG}	c_{SD}	c_{HU}
HFR	0.077939	0.0770379	0.0762666
EFR	0.077793	0.0768750	0.0761222




Figure 5.5. Linear advection solution for initial Gaussian profile at time $t = 20$ and $p = 3$. Results are shown for HFR, and EFR methods with c_{DG} (—), c_{SD} (---) and c_{HU} (-·-·)

Furthermore, we investigate the influence of c on the implicit solver performance. Table 5.2 shows the average number of GMRES iterations for convergence with and without preconditioning. These results belong to a mesh of 20×20 elements and a solution polynomial of degree 3. While the c_{DG} parameter is able to obtain a smaller numerical error, it can be seen that a larger number of GMRES iterations per linear solve are needed in comparison with c_{SD} and c_{HU} . The latter introduces the most numerical error from the three considered methods but improves the GMRES convergence. Interestingly, EFR methods require 1.5 to 2.5 times more GMRES iterations compared to FR when no preconditioner is used. However, when using RAS, FR requires ~ 1.4 times additional iterations per linear solve. Furthermore, all methods significantly improve the convergence rate when RAS preconditioning is used. Results for other polynomial degrees and refinement levels are generally consistent and are omitted for brevity.

Table 5.2. Average number of GMRES iterations for $p = 3$ schemes to solve the global system for the linear advection case. Results are shown for a 20×20 grid

	RAS Preconditioner			No Preconditioning		
	c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}
FR	5.00	4.00	4.00	6.15	6.09	6.07
HFR	5.00	4.00	3.19	9.00	7.00	6.00
EFR	3.34	3.00	3.00	15.00	12.00	9.00

5.2.2.2 Euler Equations

The discretization considered in Section 3.2.2 can be extended to a system of conservation laws by applying the approach to every component of the vector of conserved variables \mathbf{u} in (2.5). Note here, \mathbf{u} and $\hat{\mathbf{u}}$ refer to the vector of solution and trace variables, respectively. For these problems, the Riemann flux can be written such that

$$\tilde{\mathfrak{F}}_{k,f}(\mathbf{u}_{k,f}^h, \hat{\mathbf{u}}_f^h) = \mathbf{F}(\hat{\mathbf{u}}_f^h) + \mathbf{S}(\mathbf{u}_{k,f}^h - \hat{\mathbf{u}}_f^h)\mathbf{n}_{k,f}, \quad (5.11)$$

where $\mathbf{S} = \mathbf{S}(\hat{\mathbf{u}})$ is the stabilization matrix. In this work, we make use of a local Lax-Friedrichs type solver with $\mathbf{S} = s\bar{\mathbf{I}}$ [135], where $s = \hat{c} + |\hat{\mathbf{v}} \cdot \mathbf{n}|$, $\bar{\mathbf{I}}$ is the identity matrix, and $\hat{\mathbf{v}}$ is the trace velocity. The speed of sound \hat{c} can be computed at the interface in terms of the trace variable via

$$\hat{c} = \sqrt{\frac{\bar{\gamma}\hat{P}}{\hat{\rho}}}. \quad (5.12)$$

The normal boundary flux $\tilde{\mathfrak{F}}^{\text{BC}}$ term can be obtained by first computing the values of $\hat{\mathbf{u}}^{\text{BC}}$ at the boundary. Then, it can be calculated using Equation (5.11) by setting $\hat{\mathbf{u}}^h = \hat{\mathbf{u}}^{\text{BC}}$. For the airfoil-vortex interaction problem, we make use of the weak-Riemann inflow, outflow and slip-wall boundary conditions to obtain the values of $\hat{\mathbf{u}}^{\text{BC}}$ as described in [136].

5.2.2.2.1 Isentropic Vortex For our second numerical example, consider the advection of an isentropic vortex, which is an exact solution to the Euler equations. This is a commonly used numerical simulation for the verification of high-order CFD codes [19]. The computational domain consists of a square of side length $L = 20$, initialized with

$$\rho = \left[1 - \frac{\beta^2 M^2 (\bar{\gamma} - 1) e^{2f}}{8\pi^2} \right]^{\frac{1}{\bar{\gamma}-1}}, \quad (5.13)$$

$$v_x = \frac{\beta y e^f}{2\pi R}, \quad (5.14)$$

$$v_y = 1 - \frac{\beta x e^f}{2\pi R}, \quad (5.15)$$

$$P = \frac{\rho^{\bar{\gamma}}}{\bar{\gamma} M^2}, \quad (5.16)$$

where ρ is the density, β is the vortex strength, M is the Mach number, v_x and v_y are the x and y components of the velocity field, respectively, $f = (1 - x^2 - y^2)/2R^2$ and R is the radius of the vortex. We set $R = 1.5$ and $\beta = 13.5$ for this problem [33] and apply periodic boundary conditions in all directions. The initial profile is advanced through one cycle using the second-order two-stage SDIRK_{2,2} method. We verify the problem on quadrilateral

elements with second to sixth-order spatial discretizations. Four levels of refinement are considered with 5×5 , 10×10 , 20×20 and 40×40 quadrilateral elements.

Table A.8 shows the L_2 norm of the density error for all considered levels of refinement using c_{DG} , c_{SD} and c_{HU} . Results with the c_{DG} parameter are consistently more accurate than all other considered values for FR, HFR, and EFR schemes. HFR methods achieve errors that are very close to that of the FR schemes. On the other hand, EFR methods have a larger error than their FR and HFR counterparts, but when used with c_{DG} , they achieve similar or smaller errors than the former methods with c_{SD} and c_{HU} . For the finer simulations, the error difference between EFR and FR methods is, however, in the order of 10^{-8} . All schemes achieved the expected $p + 1$ order of accuracy.

Figure 5.6 shows the sparsity of the LHS matrix for FR, HFR, and EFR methods for a problem involving the Euler equations on a grid with 25×25 elements and a polynomial of degree 3, which results in 124912, 89590 and 44740 nonzero entries, respectively. Due to the structured character of the grid, FR has 40% fewer nonzero elements than expected, as per Table 3.1.






Figure 5.6. LHS matrix views for $p = 3$ FR, HFR and EFR methods for a 25×25 -element grid for the isentropic vortex case


We investigate the effects of the c -parameter on the implicit solver. We consider $p = 3$ schemes with $\Delta t = 0.025$. The Jacobian matrix was considered constant for 10 iterations to reduce assembly time. Figure 5.7 contains several features of the implicit solver. In Figure 5.7a, we observe the average number of iterations required for the linear solver to arrive at a tolerance of 10^{-12} for values of c ranging from $c_-/2$ to $3c_{HU}/2$. Specific values of c have been identified with a triangle marker for c_{DG} , a square marker for c_{SD} , and a circle marker for c_{HU} . We observe that, similar to the linear advection case, increasing c can reduce the number of GMRES iterations. For all schemes, c_{SD} and c_{HU} required fewer GMRES iterations than c_{DG} . This is also the case for the total number of Newton iterations or linear solves, as shown in Figure 5.7b, where c_{DG} requires $\sim 2.5\%$ more linear solves than c_{HU} for



(a) Average number of GMRES iterations per linear solve



(b) Total number of linear solves



(c) L_2 norm of the density error against wall-clock time

Figure 5.7. Effects of the c -parameter on the performance of FR schemes. Results are shown for $p = 3$ for the isentropic vortex case. Values of c begin at $c_-/2$ and increase in the direction of the arrows. Markers for c_{DG} (▲), c_{SD} (■), c_{HU} (●) have been added

EFR and HFR and $\sim 5\%$ in the case of FR. However, we see from Figure 5.7c, that the L_2 norm of the density error is the smallest for c_{DG} . The additional GMRES iterations and linear solves were balanced with an error of between one-third to one-half of that obtained with c_{HU} for all schemes. In addition, c_{DG} methods only took 13% more time for FR and less than 10% more time for HFR and EFR compared to c_{HU} , suggesting that c_{DG} is more cost-effective than the other two parameters.

When comparing hybridized methods, HFR was able to obtain L_2 error values that agreed with FR up to three significant digits with speedups between 1.22 to 1.3 depending on c . EFR obtained 25% more numerical error for c_{DG} , 12% more for c_{SD} and 8% more for c_{HU} than the respective FR counterparts. We note that EFR has similar and smaller values of error for specific configurations, such as the coarsest levels for $p = 3$ using c_{DG} , as shown in the verification results.




Figure 5.8. Wall-clock time against number of elements for FR, HFR, and EFR for the isentropic vortex case considering c_{DG} (—), c_{SD} (---) and c_{HU} (-.-.) parameters. Results correspond to $p = 3$ schemes

We also consider the performance of HFR and EFR methods for different numbers of elements. Figure 5.8 shows the wall-clock time against the number of elements for all considered refinement levels in this section, considering $\Delta t = 0.025$ and $p = 3$. This value of Δt corresponds to CFL numbers $CFL \approx 0.002\sqrt{N}$. Note that this conservative CFL number was used to isolate the spatial discretization errors. Results are shown for c_{DG} with a continuous line, c_{SD} with a dashed line, and c_{HU} with a dash-dotted line. FR simulation times increment more quickly than hybridized schemes. Considering the c_{DG} results, the wall-clock time curve grows as $1.65N$ compared to $0.893N$ for HFR and $0.67N$ for EFR. With this value of c and considering the 40×40 grid, HFR and EFR were faster than FR by a factor of 1.77 and 2.3, respectively. Accordingly, hybridized methods are a more efficient approach to conventional FR implicit schemes, even for a large number of elements. In this plot, we also observe that changing c to c_{HU} reduces the computational cost of FR by a factor

of 1.3, but this is less significant for the hybridized methods, where the speedup obtained was only 1.12, and 1.05 for HFR and EFR, respectively.

5.2.2.2 Airfoil-Vortex Interaction The final numerical example is the interaction between a vortex and a NACA0012 airfoil at zero angle of attack. This case allows us to demonstrate an application of hybridized FR methods in a more complex flow setting. Starting from an initial steady subsonic flow, a vortex is released five chords upstream of the airfoil’s leading edge. Over time, the vortex convects toward the airfoil with a freestream velocity u_∞ . The disturbance added to the steady flow field as a consequence of the vortex depends on the tangential velocity given by [74, 137]

$$\bar{u}_t = \frac{\Gamma}{2\pi r} \left(\frac{r^2}{r^2 + r_v^2} \right), \quad (5.17)$$

where \bar{u}_t is nondimensionalized by the freestream velocity, r is the distance from the vortex center to any point in the domain, and r_v is the vortex radius, which we take to be $r_v = 0.018$, and we set $\Gamma = -0.283$. For this example, we considered the compressible Euler equations




Figure 5.9. Unstructured mesh of 6830 quadrilateral elements for the airfoil-vortex interaction case

with a Mach number $M = 0.25$ and a convective time $t_c = 1/u_\infty$. A mesh containing 6830 quadrilateral elements is used in this simulation, with refined elements along the vortex-free convection path as well as in areas of expected complex flow in proximity to the airfoil as shown in Figure 5.9. We use coarse elements downstream of the airfoil as we will only consider the effects before and during the direct interaction. FR, HFR, and EFR methods with c_{DG} , c_{SD} and c_{HU} parameters are considered in this case for schemes ranging from orders $p = 1$ to $p = 4$ using a second-order singly-diagonal implicit Runge-Kutta method and a nondimensional time step size $\Delta t_c = 0.001$, which corresponds to a CFL number of about 1.8. For reference, the maximum stable CFL number with explicit $RK_{4,4}$ is approximately 0.48, 0.23, 0.12, and 0.08 for $p = 1, 2, 3,$ and 4 , respectively. The density residual was converged to a tolerance of 10^{-10} . To create a balance between the assembly and the global solve time, the Jacobian matrix was updated every 10 time steps. Ideally, the vortex core pressure remains

unaltered, but generally, dissipative numerical schemes will cause the vortex to decay [137]. As a reference, we use a sixth-order FR discretization with c_{DG} on a mesh refined by splitting every element into four.




Figure 5.10. Pressure readings at $y = 0$ before the interaction in the airfoil-vortex interaction case with c_{DG}

Figure 5.10 shows the evolution of the pressure along $y = 0$ at four different snapshots during the free convective phase of the vortex, starting with the initial unaltered condition at $t_c = 0$. Clearly, the $p = 1$ simulation is overly dissipative, causing significant vortex decay. Here, EFR introduces significantly more dissipation than HFR and FR, which is consistent with the wave propagation properties defined earlier in this chapter. As the polynomial degree increases, the vortex strength is better preserved throughout the simulation, and the pressure readings become closer to the reference data. At $p = 4$, we notice the EFR results are very close to HFR, with some slight additional dissipation. Table 5.3 shows the vortex core pressure readings at $t_c = 4$. For $p \geq 3$ for all considered values of c , the pressure

Table 5.3. Vortex pressure core readings at $t_c = 4$ for the airfoil-vortex interaction case

p	1			2			3			4		
	c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}
FR	0.9700	0.9778	0.9920	0.8835	0.8987	0.9083	0.8299	0.8468	0.8532	0.7898	0.8012	0.8026
HFR	0.9700	0.9778	0.9919	0.8836	0.8987	0.9083	0.8299	0.8468	0.8532	0.7898	0.8012	0.8026
EFR	0.9800	0.9844	0.9931	0.8807	0.8985	0.9099	0.8323	0.8473	0.8533	0.7905	0.8006	0.8017

readings were almost identical between FR and HFR, with higher values of c introducing more dissipation. For the higher-order simulations, EFR shows slightly more damped vortex cores, with a relative error in the orders of 10^{-3} with respect to FR.

Figures 5.11 and 5.12 show the evolution of the lift and moment coefficients, respectively. Due to the vortex decay, both the lift and moment curves show a significant difference from the reference data for $p = 1$. This behaviour improves for higher polynomial degrees, where at $p = 4$, results are close to the reference data at least until $t \approx 6$. Based on these figures and consistent with the pressure readings, more dissipation for the EFR vortex occurs in comparison with HFR and FR, resulting in more damped moment coefficient curves, especially at low orders.

While results show less error for HFR at high orders, performance measures indicate that EFR methods may be more efficient at high orders. Figure 5.13 displays the density residual against the wall-clock time. It can be seen that HFR methods are slower than conventional FR for $p = 1$. This occurs since the DOF at this polynomial degree are equivalent for both methods, but the static condensation procedure introduces computation overhead. Moving toward higher orders shows that HFR methods were able to achieve 4.7 times faster simulations and EFR 6.7 times faster for $p = 4$ compared to FR with c_{DG} .

Table 5.4. Average time (in seconds) per linear solve for the airfoil-vortex interaction case

p	1			2			3			4		
	c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}
FR	0.082	0.081	0.075	0.385	0.348	0.285	1.266	1.053	0.995	4.862	3.118	2.454
HFR	0.078	0.078	0.067	0.212	0.194	0.186	0.442	0.406	0.399	0.796	0.739	0.726
EFR	0.009	0.008	0.008	0.051	0.051	0.053	0.140	0.141	0.144	0.317	0.324	0.327

Figure 5.14 shows the resulting wall-clock times for all considered methods. FR with c_{DG} has a larger wall-clock time in comparison to the other two correction functions by up to 2 times for $p = 4$, when compared to c_{HU} . This is associated with slower convergence of the GMRES solver. As shown in Table 5.4, the average time per linear solve was between 1.1 and 1.98 times longer for c_{DG} than for c_{HU} with FR, consistently 1.1 times longer in the

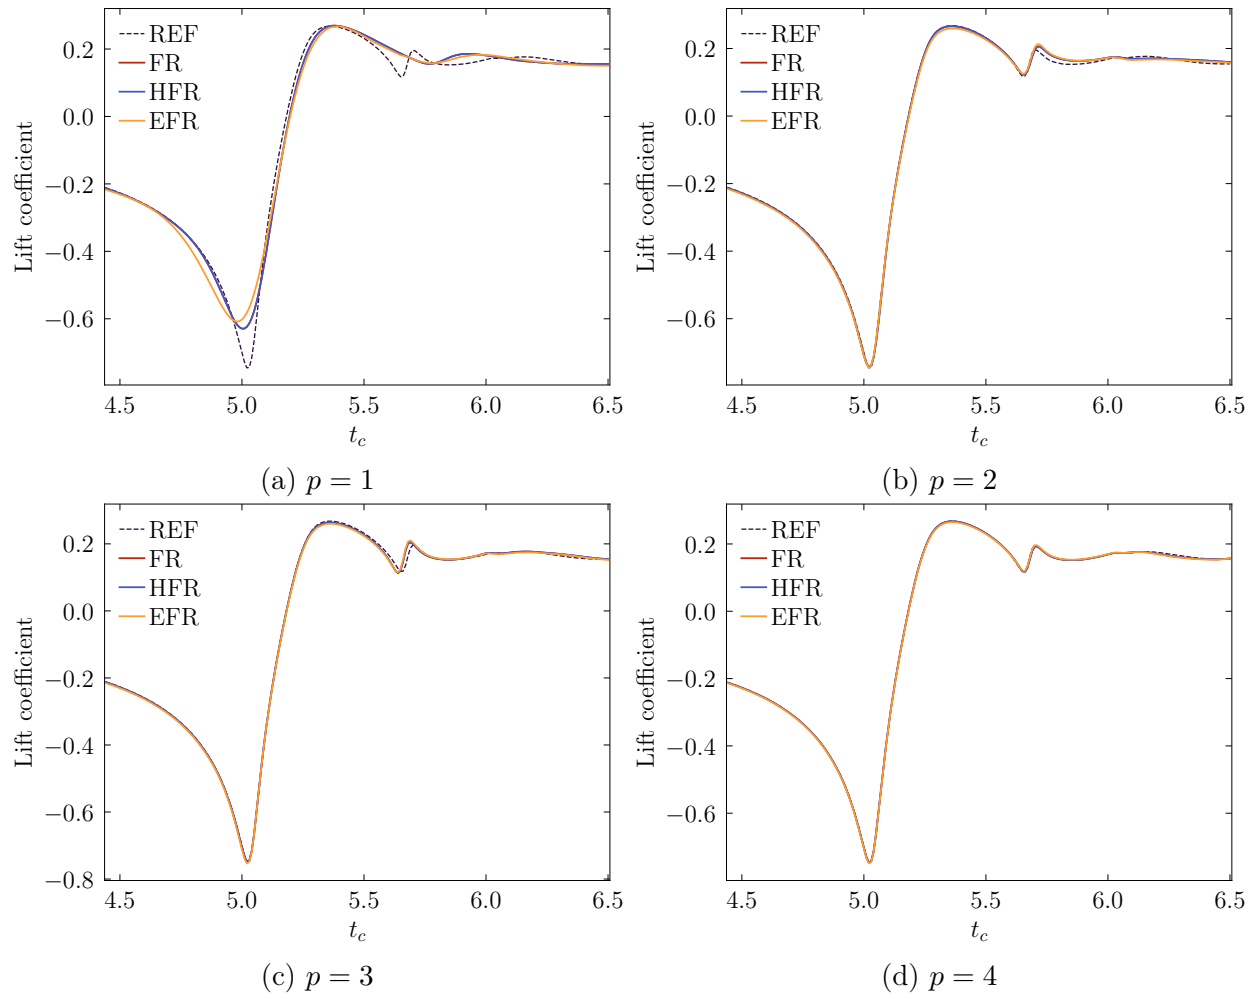


Figure 5.11. Evolution of the lift coefficient for the airfoil-vortex interaction case with c_{DG}




Figure 5.12. Evolution of the moment coefficient for the airfoil-vortex interaction case with c_{DG}




Figure 5.13. Density residual against wall-clock time for the airfoil-vortex interaction case with c_{DG}

case of HFR, and generally the same or even smaller for EFR. Hence, we observe that the effect of the c parameter in the wall-clock time is more prominent for the conventional FR method than for HFR and EFR. We also note that FR methods spend most of their time on linear solves when ignoring the Jacobian assembly time. Hybridized schemes, however, require computation of the local solves. It was observed that they can spend between 40 to 50% on these computations, depending on the polynomial degree. Despite this, they are still more efficient than conventional FR methods. From Table 5.5, we see that the values of α_{DOF} for this unstructured grid are close to those in Table 3.1. In addition, while the ratios of α_{NNZ} can provide insights into the general benefits of using hybridized schemes, the computational time is more sensitive to the performance of the GMRES solver.

We have analyzed the accuracy and performance of hybridized FR methods in advection problems. We now extend the approach to the full advection-diffusion regime in the next




Figure 5.14. Wall-clock times for the airfoil-vortex interaction case for polynomial degrees 1 to 4 using all considered parameters c

section.

Table 5.5. Coefficients α_{DOF} and α_{NNZ} resulting from the airfoil-vortex interaction problems at different polynomial degrees

p	α_{DOF}				α_{NNZ}			
	1	2	3	4	1	2	3	4
FR	3.904	8.784	15.616	24.400	41.496	144.399	346.055	680.973
HFR	3.856	5.784	7.712	9.640	53.412	120.177	213.649	333.826
EFR	0.999	2.928	4.855	6.784	8.852	45.339	108.531	198.430

5.3 Advection-Diffusion-Type Problems

5.3.1 Stability Analysis

The HFR methods considered in the previous section were shown to recover conventional FR formulations for linear advection problems, and EFR methods introduced additional dissipation via spectral analysis. In this section, we analyze the behaviour of hybridized FR methods for linear-diffusion problems that make use of the Vincent-Castonguay-Jameson-Huynh (VCJH) correction functions [29] and discuss connections to HDG methods for which linear stability proofs have been obtained. For this purpose, we consider the linear

advection-diffusion equation

$$\frac{\partial u}{\partial t} + \boldsymbol{\alpha} \cdot \nabla u - \beta \nabla^2 u = 0, \quad (5.18)$$

where $\boldsymbol{\alpha}$ is the advection speed, and β is the diffusion coefficient. We can rewrite this second-order problem as a system of first-order equations to be consistent with (3.2)

$$\frac{\partial u}{\partial t} + \nabla \cdot (\boldsymbol{\alpha} u - \beta \mathbf{q}) = 0, \quad (5.19)$$

$$\mathbf{q} - \nabla u = 0, \quad (5.20)$$

subjected to periodic boundary conditions.

5.3.1.1 Explicit Forms of the Numerical Trace

As previously stated in the implementation section, we consider hybridized forms of FR methods where the trace variable may belong to one of the finite-element spaces in (3.43), which lead to the so-called hybridized and embedded flux reconstruction schemes. At a given flux point, the Riemann flux for advection-diffusion is given by

$$\hat{\mathfrak{F}}_{k,f} = \boldsymbol{\alpha} \hat{u}_f^h - \beta \mathbf{q}_{k,f}^h + s_{k,f} (u_{k,f}^h - \hat{u}_f^h) \mathbf{n}_{k,f}, \quad (5.21)$$

where $s_{k,f} = s_{k,f}^{(c)} + s_{k,f}^{(v)}$. Because of the discontinuous nature of the trace polynomials in HFR with space \mathbb{M}^h (Equation (3.43a)), the conservation condition is applicable pointwise. In the case of the embedded schemes, it can be shown that they are conservative at the trace points [62], containing contributions from all elements in its vicinity. Application of the transmission conditions on an interior trace point to solve for \hat{u} yields the following explicit expressions for the trace variable for the HFR method

$$\hat{u}^{\mathbb{M}_p^h} = \frac{\{\{su\}\}}{\{\{s\}\}} - \frac{\beta \llbracket \mathbf{q} \rrbracket}{2 \{\{s\}\}}. \quad (5.22)$$

In the case of the EFR method, a simple explicit expression cannot be obtained solely in terms of u due to the global coupling resulting from the reduced space of the trace. If, for instance, we choose to under-integrate the transmission conditions by employing a quadrature such as GLL, the interior solution coupling is reduced, and a simplified expression can be written as

$$\hat{u}^{\tilde{\mathbb{M}}_p^{h,GLL}} = \frac{\sum_{\bar{F}} \{\{swu\}\}_{\bar{F}}}{\sum_{\bar{F}} \{\{s\}\}_{\bar{F}}} - \frac{\beta \sum_{\bar{F}} \llbracket w\mathbf{q} \rrbracket_{\bar{F}}}{2 \sum_{\bar{F}} \{\{s\}\}_{\bar{F}}}, \quad (5.23)$$

where \bar{F} refers to the faces intersecting the trace point. This form was applied in the spectral analysis of Section 5.2.1. Here w is a quadrature weight arising from the diagonal local mass matrix of the GLL quadrature. However, in the rest of the work, we employ exact

integration to mitigate possible aliasing issues of the GLL quadrature choice. From these expressions, assuming a homogeneous definition of the stabilization on the $(-)$ and $(+)$ sides of the interface for all faces, we note that the problem is undefined for $\{\{s\}\} = 0$, and then the following statement is a constraint for hybridized advection and advection-diffusion problems

$$\{\{s\}\} \neq 0. \quad (5.24)$$

5.3.1.2 Proof of Stability

We now devise a methodology to show energy stability of hybridized flux reconstruction methods. This will provide insights into the stabilization mechanisms of HFR methods alongside ESFR correction functions described in Chapter 3. Specifically, we consider the analysis of HFR methods on quadrilateral Cartesian grids, which have transformation Jacobian matrices of the form

$$\mathbf{J}_k = \begin{bmatrix} J_x & 0 \\ 0 & J_y \end{bmatrix}, \quad (5.25)$$

which, for these problems, $\hat{J}_{\bar{f}} = J_x$ and $\hat{J}_{\bar{f}} = J_y$ at a horizontal and vertical face, respectively, since the cross-terms are zero. This analysis has been widely studied for conventional FR schemes with several advective and diffusive Riemann solvers for one [138, 29, 139] and higher-dimensional [140, 141] problems. We seek to study the time evolution of the solution using a suitable Sobolev norm. We perform different algebraic manipulations to determine the constraints that will yield well-defined and linearly stable schemes. We make direct use of the proofs in the work of Sheshadri et al. [141, 142, 143] for conventional FR and augment them with algebraic manipulated forms of the transmission conditions.

First, we introduce two important equations that will enable the study via the following lemmas.

Lemma 5.3.1. *For hybridizable FR methods on Cartesian grids, the following holds*

$$\sum_{k=1}^N \int_{\partial\Omega_k} \hat{\mathfrak{F}}_k \cdot \mathbf{n}_k \hat{u}^h ds = 0. \quad (5.26)$$

Proof. On multiplying the strong form of the transmission conditions

$$[[\hat{\mathfrak{F}}]]_{\varepsilon_0^h} = 0, \quad [[\cdot]]_{\varepsilon_0^h} = \sum_{\bar{f} \in \varepsilon_0^h} [[\cdot]]_{\bar{f}}, \quad (5.27)$$

by an arbitrary test function ϕ from a space in (3.43) and integrating over each \bar{f} , we have

$$\sum_{\bar{f} \in \varepsilon^h} \int_{\bar{f}} \left(\llbracket \hat{\mathfrak{F}} \rrbracket_{\bar{f}} \right) \phi d\bar{f} = 0, \quad (5.28)$$

which can be rewritten using the identity

$$\sum_{\bar{f} \in \varepsilon^h} \int_{\bar{f}} \left(\llbracket \hat{\mathfrak{F}} \rrbracket_{\bar{f}} \right) \phi ds = \sum_{k=1}^N \int_{\partial\Omega_k} \hat{\mathfrak{F}}_k \cdot \mathbf{n}_k \phi ds = 0. \quad (5.29)$$

Recall that the problem has been defined to be periodic and that \hat{u}^h also belongs to one of the spaces in (3.43). Since $\int_f \phi df = \int_f \phi|_f df$ and both ϕ , \hat{u}^h belong to the same space, we substitute ϕ by the trace variable \hat{u}^h , and the proof is complete. \square

Lemma 5.3.2. *For all hybridizable FR methods, the conservativity condition holds, and therefore, the following holds as well*

$$\sum_{k=1}^N \sum_{f=1}^{N_f} \left[\frac{\partial^p (\hat{\mathfrak{F}}_k \cdot \mathbf{n}_k)}{\partial \psi^p} \frac{\partial^p \hat{u}^h}{\partial \psi^p} \right]_f = 0. \quad (5.30)$$

Proof. The procedure is similar to the previous lemma but uses differentiation. On differentiating the transmission conditions p times along the ψ -direction corresponding to each face, multiplying by $\frac{\partial^p \phi}{\partial \psi^p}$ and integrating over ε_0^h

$$\sum_{\bar{f} \in \varepsilon^h} \int_{\bar{f}} \frac{\partial^p (\llbracket \hat{\mathfrak{F}} \rrbracket_{\bar{f}})}{\partial \psi^p} \frac{\partial^p \phi}{\partial \psi^p} d\bar{f} = 0, \quad (5.31)$$

since both $\hat{\mathfrak{F}} \cdot \mathbf{n}$ and $\phi|_{\bar{f}} \in \mathbb{P}_p$, the integrand is a constant and hence

$$\sum_{\bar{f} \in \varepsilon^h} \left[\frac{\partial^p (\llbracket \hat{\mathfrak{F}} \rrbracket_{\bar{f}})}{\partial \psi^p} \frac{\partial^p \phi}{\partial \psi^p} d\bar{f} \right]_{\bar{f}} = 0, \quad (5.32)$$

where we have omitted the integration limits since we are considering a Cartesian grid with constant nonzero face Jacobians. Substituting ϕ by the trace variable and applying identity (5.29) completes the above lemma. \square

Lemma 5.3.3. *For the tensor-product FR formulation with VCJH correction functions, the following holds*

$$\frac{1}{2} \frac{d}{dt} \|u^h\|_{p,2}^2 = -\beta \|\mathbf{q}\|^2 + \Theta^{\text{FR}} + \Theta^{\text{HFR}}, \quad (5.33)$$

where

$$\begin{aligned}
\Theta^{\text{FR}} = & - \sum_{k=1}^N \int_{\partial\Omega_k} u_k^h(\hat{\mathfrak{F}}_k \cdot \mathbf{n}_k) ds \\
& + \sum_{k=1}^N \frac{1}{2} \int_{\partial\Omega_k} u_k^h(\mathbf{F}_k^{h,(c)} \cdot \mathbf{n}_k) ds \\
& - \sum_{k=1}^N \int_{\partial\Omega_k} (\hat{u}_k^h - u_k^h)(\mathbf{F}_k^{h,(v)} \cdot \mathbf{n}_k) ds \\
& - c \sum_{k=1}^N \sum_{f=1}^{N_f} \left[J_{\psi_k}^{2p+1} \frac{\partial^p u_k^h}{\partial \psi^p} \frac{\partial^p}{\partial \psi^p} (\hat{\mathfrak{F}}_k \cdot \mathbf{n}_k) \right]_f \\
& + c \sum_{k=1}^N \sum_{f=1}^{N_f} \left[J_{\psi_k}^{2p+1} \frac{1}{2} \frac{\partial^p u_k^h}{\partial \psi^p} \frac{\partial^p}{\partial \psi^p} (\mathbf{F}_k^{h,(c)} \cdot \mathbf{n}_k) \right]_f \\
& - c \sum_{k=1}^N \sum_{f=1}^{N_f} \left[J_{\psi_k}^{2p+1} \frac{\partial^p (\hat{u}_k^h - u_k^h)}{\partial \psi^p} \frac{\partial^p}{\partial \psi^p} (\mathbf{F}_k^{h,(v)} \cdot \mathbf{n}_k) \right]_f, \tag{5.34}
\end{aligned}$$

and

$$\Theta^{\text{HFR}} = \sum_{k=1}^N \int_{\partial\Omega_k} \hat{u}_k^h(\hat{\mathfrak{F}}_k \cdot \mathbf{n}_k) ds + c \sum_{k=1}^N \sum_{f=1}^{N_f} \left[J_{\psi_k}^{2p+1} \frac{\partial^p \hat{u}_k^h}{\partial \psi^p} \frac{\partial^p (\hat{\mathfrak{F}}_k \cdot \mathbf{n}_k)}{\partial \psi^p} \right]_f. \tag{5.35}$$

Proof. Here we have directly introduced the expressions obtained from the proof of stability of the FR method by Sheshadri et al. [141, 142]. The reader can refer to it for the proof of this expression. After algebraic manipulations and changes in the notation for the sake of consistency, a general expression for the stability of the FR method on Cartesian quadrilateral elements is given by

$$\frac{1}{2} \frac{d}{dt} \|u^h\|^2 = -\beta \|\mathbf{q}\|^2 + \Theta^{\text{FR}}, \tag{5.36}$$

where Θ^{FR} reads as in Equation (5.34). Here, ψ is a dummy coordinate variable such that $\psi = x$ for horizontal faces and $\psi = y$ for vertical faces. In addition, $\mathbf{F}^{(c)}$ and $\mathbf{F}^{(v)}$ refer to the advective and diffusive fluxes, and $\hat{\mathfrak{F}}_k \cdot \mathbf{n}$ is the total normal Riemann flux involving both advection and diffusion. Typically, explicit forms of the numerical trace \hat{u}^h are used to derive these stability proofs. However, since HFR methods implicitly define it, we leave it as a variable for this analysis. This also allows different function spaces for the trace to be considered. In this sense, we augment the above expression with Lemmas 5.3.1 and 5.3.2. Multiplying Equation (5.30) by c and adding it to (5.26) completes the proof. \square

With these tools, we are ready to state our theorem on the stability of HFR methods.

Theorem 5.3.4. *Using a tensor-product formulation of the hybridized FR methods with VCJH correction functions, it can be shown that if*

- *The stabilization parameters $s_{\pm} = s_{\pm}^{(c)} + s_{\pm}^{(v)}$ are chosen such that $s_{\pm} > \frac{\boldsymbol{\alpha} \cdot \mathbf{n}_{\pm}}{2}$ and*
- *The correction parameter satisfies $c \geq 0$,*

then the following expression holds for the two-dimensional linear advection-diffusion equation with periodic boundary conditions on Cartesian quadrilateral elements

$$\frac{1}{2} \frac{d}{dt} \|u^h\|_{p,2}^2 \leq 0, \quad (5.37)$$

for a broken Sobolev norm of the solution given by

$$\|u^h\|_{p,2}^2 = \sum_{k=1}^N \int_{\Omega_k} \left[(u_k^h)^2 + \frac{c}{2} \left(\left(\frac{\partial^p u_k^h}{\partial \tilde{x}^p} \right)^2 + \left(\frac{\partial^p u_k^h}{\partial \tilde{y}^p} \right)^2 \right) + \frac{c^2}{4} \left(\frac{\partial^{2p} u_k^h}{\partial \tilde{x}^p \partial \tilde{y}^p} \right)^2 \right] d\Omega_k. \quad (5.38)$$

Proof. To state this proof, we can rewrite the equations in Lemma 5.3.3 as a summation over all faces in the computational domain. Note that this is valid since all integrations and derivatives in the previous equations are performed over the borders of the elements. Hence, we consider one of these faces with either horizontal or vertical direction with left and right elements Ω_- and Ω_+ and with outward unit normal vectors \mathbf{n}_- and \mathbf{n}_+ , respectively. Let us now expand each of the terms in these equations at a given face. The first term in Equation (5.34) can be written as follows

$$\begin{aligned} & - \int_{\bar{f}} \left(u_-^h \left[\mathbf{F}(\hat{u}_{\bar{f}}^h, u_-^h, \mathbf{q}_-^h) \cdot \mathbf{n}_- + s_-(u_-^h - \hat{u}_{\bar{f}}^h) \right] + u_+^h \left[\mathbf{F}(\hat{u}_{\bar{f}}^h, u_+^h, \mathbf{q}_+^h) \cdot \mathbf{n}_+ + s_+(u_+^h - \hat{u}_{\bar{f}}^h) \right] \right) d\bar{f} \\ & = - \int_{\bar{f}} \left(u_-^h [\boldsymbol{\alpha} \cdot \mathbf{n}_- \hat{u}_{\bar{f}}^h - \beta \mathbf{q}_-^h \cdot \mathbf{n}_- + s_-(u_-^h - \hat{u}_{\bar{f}}^h)] \right. \\ & \quad \left. + u_+^h [\boldsymbol{\alpha} \cdot \mathbf{n}_+ \hat{u}_{\bar{f}}^h - \beta \mathbf{q}_+^h \cdot \mathbf{n}_+ + s_+(u_+^h - \hat{u}_{\bar{f}}^h)] \right) d\bar{f}, \end{aligned} \quad (5.39)$$

where we have expanded the definitions of the Riemann solver according to Equation (5.21) and considered a total stabilization parameter $s = s^{(c)} + s^{(v)}$. Similarly, the second term involving the convective flux can be expanded for this face

$$\begin{aligned} & \frac{1}{2} \int_{\bar{f}} (u_-^h \mathbf{F}^{(c)}(u_-^h) \cdot \mathbf{n}_- + u_+^h \mathbf{F}^{(c)}(u_+^h) \cdot \mathbf{n}_+) d\bar{f} \\ & = \frac{1}{2} \int_{\bar{f}} (\boldsymbol{\alpha} \cdot \mathbf{n}_- (u_-^h)^2 + \boldsymbol{\alpha} \cdot \mathbf{n}_+ (u_+^h)^2) d\bar{f}, \end{aligned} \quad (5.40)$$

and the third term involving the diffusion component of the flux can be written

$$\begin{aligned}
& - \int_{\bar{f}} \left((\hat{u}_{\bar{f}}^h - u_-^h) \mathbf{F}^{(v)}(\mathbf{q}_-^h) \cdot \mathbf{n}_- + (\hat{u}_{\bar{f}}^h - u_+^h) \mathbf{F}^{(v)}(\mathbf{q}_+^h) \cdot \mathbf{n}_+ \right) d\bar{f} \\
& = - \int_{\bar{f}} \left((\hat{u}_{\bar{f}}^h - u_-^h)(-\beta \mathbf{q}_-^h \cdot \mathbf{n}_-) + (\hat{u}_{\bar{f}}^h - u_+^h)(-\beta \mathbf{q}_+^h \cdot \mathbf{n}_+) \right) d\bar{f}. \tag{5.41}
\end{aligned}$$

Finally, we consider the first term of Equation (5.35)

$$\begin{aligned}
& \int_{\bar{f}} \left(\hat{u}_{\bar{f}}^h [\boldsymbol{\alpha} \cdot \mathbf{n}_- \hat{u}_{\bar{f}}^h - \beta \mathbf{q}_-^h \cdot \mathbf{n}_- + s_-(u_-^h - \hat{u}_{\bar{f}}^h)] + \hat{u}_{\bar{f}}^h [\boldsymbol{\alpha} \cdot \mathbf{n}_+ \hat{u}_{\bar{f}}^h - \beta \mathbf{q}_+^h \cdot \mathbf{n}_+ + s_+(u_+^h - \hat{u}_{\bar{f}}^h)] \right) d\bar{f} \\
& = \int_{\bar{f}} \left(\hat{u}_{\bar{f}}^h [-\beta \mathbf{q}_-^h \cdot \mathbf{n}_- + s_-(u_-^h - \hat{u}_{\bar{f}}^h)] + \hat{u}_{\bar{f}}^h [-\beta \mathbf{q}_+^h \cdot \mathbf{n}_+ + s_+(u_+^h - \hat{u}_{\bar{f}}^h)] \right) d\bar{f}, \tag{5.42}
\end{aligned}$$

where we have used $\mathbf{n}_- = -\mathbf{n}_+$ to cancel out the advective flux on the trace variable. After adding all of the above contributions, we write

$$\Theta_{\bar{f}}^A = \int_{\bar{f}} \left[\bar{s}_-(u_-^h - \hat{u}_{\bar{f}}^h)^2 + \bar{s}_+(u_+^h - \hat{u}_{\bar{f}}^h)^2 \right] d\bar{f}, \tag{5.43}$$

where we have introduced $\bar{s}_{\pm} = s_{\pm} - \frac{\boldsymbol{\alpha} \cdot \mathbf{n}_{\pm}}{2}$. Note the exchange of energy between the two adjacent elements is implicitly done via the trace variable. In a similar manner, we can obtain the contributions from the derivative terms and write

$$\Theta_{\bar{f}}^B = J_{\psi_k}^{2p+1} \left[\bar{s}_- \left(\frac{\partial u_-^h}{\partial \psi^p} - \frac{\partial \hat{u}_{\bar{f}}^h}{\partial \psi^p} \right)^2 + \bar{s}_+ \left(\frac{\partial u_+^h}{\partial \psi^p} - \frac{\partial \hat{u}_{\bar{f}}^h}{\partial \psi^p} \right)^2 \right], \tag{5.44}$$

for which we omit the derivation since it follows a similar procedure. Considering a periodic domain, the sum over all faces $\bar{f} \in \varepsilon^h$ results in the following stability statement for HFR schemes on quadrilateral elements

$$\frac{1}{2} \frac{d}{dt} \|u^h\|^2 = -\beta \|\mathbf{q}\|^2 - \sum_{\bar{f} \in \varepsilon^h} \left(\Theta_{\bar{f}}^A + c \Theta_{\bar{f}}^B \right). \tag{5.45}$$

From the above statements, we observe that for $c \geq 0$ and

$$\bar{s}_{\pm} \geq 0 \quad \Rightarrow \quad s_{\pm} \geq \frac{\boldsymbol{\alpha} \cdot \mathbf{n}_{\pm}}{2}, \tag{5.46}$$

the hybridized form of FR for advection-diffusion satisfies

$$\frac{d}{dt} \|u^h\|^2 \leq 0. \tag{5.47}$$

However, from the explicit forms of the numerical trace defined in Equations (5.22) and (5.23), we see that the method is undefined for $s_- = -s_+$ and hence the inequality becomes strict

$$s_{\pm} > \frac{\boldsymbol{\alpha} \cdot \mathbf{n}_{\pm}}{2}, \quad (5.48)$$

or equally

$$s_{k,f} > \frac{\boldsymbol{\alpha} \cdot \mathbf{n}_k}{2}, \quad (5.49)$$

where the strict inequality has to be satisfied on at least one face side [54]. Note that this proof of stability recovers that of the HDG method [59] for $c = 0$, and hence we show via this analysis that the stability proof of linear advection-diffusion HDG can be recovered from HFR methods with this particular value of c . \square

5.3.1.3 Connection to Standard FR Schemes

In this section, we show the connection of hybridized methods with conventional FR formulations for a typical choice of the stabilization parameter. It is important to note that only discontinuous trace polynomials may recover existing FR formulations for purely convective problems. To establish a connection with standard FR formulations, we consider pure advection and pure diffusion scenarios. A typical choice of stabilization for problems involving advection and diffusion is

$$s_{k,f} = s_{k,f}^{(c)} + s_{k,f}^{(v)} = \bar{\lambda} |\boldsymbol{\alpha} \cdot \mathbf{n}| + \tau_v, \quad (5.50)$$

where $\bar{\lambda}$ is an upwinding constant and τ_v is the so-called diffusion stabilization parameter. Generally, one can take it to be

$$\tau_v = \frac{\beta}{\ell}, \quad (5.51)$$

with ℓ a diffusive-length scale.

5.3.1.4 Advection Regime

First, we consider the case of pure advection ($\beta = 0$). For the above choice of stabilization with $\tau_v = 0$, the energy statement reads

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \|u^h\|^2 = & - \sum_{\bar{f} \in \mathcal{E}^h} \left[\int_{\bar{f}} \frac{|\boldsymbol{\alpha} \cdot \mathbf{n}|}{2} \left[\bar{\zeta}_- (u_-^h - \hat{u}_{\bar{f}}^h)^2 + \bar{\zeta}_+ (u_+^h - \hat{u}_{\bar{f}}^h)^2 \right] d\bar{f} \right. \\ & \left. + J_{\psi_k}^{2p+1} \frac{|\boldsymbol{\alpha} \cdot \mathbf{n}|}{2} c \left[\bar{\zeta}_- \left(\frac{\partial u_-^h}{\partial \psi^p} - \frac{\partial \hat{u}_{\bar{f}}^h}{\partial \psi^p} \right)^2 + \bar{\zeta}_+ \left(\frac{\partial u_+^h}{\partial \psi^p} - \frac{\partial \hat{u}_{\bar{f}}^h}{\partial \psi^p} \right)^2 \right] \right]_{\bar{f}}, \end{aligned} \quad (5.52)$$

which implies that for the general case where the neighbouring interface solutions can take any arbitrary value, nonpositivity can be guaranteed for $\bar{\lambda} \geq \frac{1}{2}$. Here we have defined

$$\bar{\zeta}_{\pm} = 2\bar{\lambda} \pm \text{sign}(\boldsymbol{\alpha} \cdot \mathbf{n}_-). \quad (5.53)$$

While this suggests that central-like approaches ($\bar{\lambda} \rightarrow 0$) are not suitable choices for hybridization of pure advection problems, this can be mitigated by choice of discontinuous trace polynomials. As previously discussed in [32, 59], this finite-dimensional function space leads to the exact formulation of standard FR schemes, where the relationship with the trace variable and the interface solution value is equal in magnitude when the same stabilization parameter is used on both sides, i.e.,

$$(u_-^h - \hat{u}_{\bar{f}}^h) = (\hat{u}_{\bar{f}}^h - u_+^h), \quad \text{for } s_- = s_+, \quad (5.54)$$

which is a consequence of the local conservation for HFR, $[\hat{\mathfrak{F}}]_{\bar{f}} = 0$. This results in a less strict range of stable $\bar{\lambda}$ parameters, as the stability statement becomes

$$\frac{1}{2} \frac{d}{dt} \|u^h\|^2 = \sum_{\bar{f} \in \epsilon^h} \left[-\frac{\bar{\lambda}}{2} \int_{\bar{f}} |\boldsymbol{\alpha} \cdot \mathbf{n}| (u_- - u_+)^2 d\bar{f} - c \frac{\bar{\lambda}}{2} \left[J_{\psi_k}^{2p+1} |\boldsymbol{\alpha} \cdot \mathbf{n}| \left(\frac{\partial u_-^h}{\partial \psi^p} - \frac{\partial u_+^h}{\partial \psi^p} \right)^2 \right]_{\bar{f}} \right], \quad (5.55)$$

consistent with the analysis of the FR methods in [143, 141]. However, the implicit characteristic of \hat{u} still requires $\bar{\lambda} > 0$ for the problem to be well-defined. This means that the exact central FR scheme for advection cannot be recovered with this type of stabilization. Note that we can find an explicit form of the linear-advection common flux with discontinuous trace polynomials and show that it takes the following form [59]

$$\hat{\mathfrak{F}} \cdot \mathbf{n}_- = \frac{\boldsymbol{\alpha} \cdot \mathbf{n}_- s_- + s_+ s_-}{s_+ + s_-} u_- + \frac{\boldsymbol{\alpha} \cdot \mathbf{n}_- s_+ - s_+ s_-}{s_+ + s_-} u_+. \quad (5.56)$$

A possible way to define a central HFR method can be shown if the stabilization parameters are taken to be different on each side of the interface and are defined as follows

$$s_- = \gamma |\boldsymbol{\alpha} \cdot \mathbf{n}| + \boldsymbol{\alpha} \cdot \mathbf{n}_-, \quad (5.57)$$

$$s_+ = f(\gamma, \boldsymbol{\alpha}) |\boldsymbol{\alpha} \cdot \mathbf{n}| - \boldsymbol{\alpha} \cdot \mathbf{n}_-, \quad f(\gamma, \boldsymbol{\alpha}) = \frac{\gamma \text{sign}(\boldsymbol{\alpha} \cdot \mathbf{n}_-)}{2\gamma + \text{sign}(\boldsymbol{\alpha} \cdot \mathbf{n}_-)}, \quad (5.58)$$

for $\gamma > 0$. This recovers the exact central FR scheme for pure advection and no dissipative mechanism. The trace is thereby defined by

$$\hat{u} = \left(\frac{\text{sign}(\boldsymbol{\alpha} \cdot \mathbf{n}_-)}{2\gamma} + 1 \right) u_- - \frac{\text{sign}(\boldsymbol{\alpha} \cdot \mathbf{n}_-)}{2\gamma} u_+. \quad (5.59)$$

The proof of (5.57)-(5.59) can be shown by seeking the forms of the stabilization parameters in (5.56) that yield equal terms multiplying u_- and u_+ . However, it is well-known that fully-central methods are inconvenient for applications of physical interest due to their lack of dissipation.

5.3.1.5 Diffusive Regime

In the case of pure diffusion ($\boldsymbol{\alpha} = \mathbf{0}$), we set $\bar{\lambda} = 0$ in Equation (5.50) and obtain that the evolution of the L_2 energy is governed by

$$\begin{aligned} \frac{1}{2} \frac{d}{dt} \|u^h\|^2 = & -\beta \|\mathbf{q}\|^2 \\ & - \tau_v \sum_{\bar{f} \in \varepsilon^h} \left[\int_{\bar{f}} ((u_-^h - \hat{u}_{\bar{f}}^h)^2 + (u_+^h - \hat{u}_{\bar{f}}^h)^2) d\bar{f} \right. \\ & \left. + J_{\psi_k}^{2p+1} c \left[\left(\frac{\partial u_-^h}{\partial \psi^p} - \frac{\partial \hat{u}_{\bar{f}}^h}{\partial \psi^p} \right)^2 + \left(\frac{\partial u_+^h}{\partial \psi^p} - \frac{\partial \hat{u}_{\bar{f}}^h}{\partial \psi^p} \right)^2 \right] \right], \end{aligned} \quad (5.60)$$

which shows that for this type of problem with an arbitrary positive diffusion coefficient β , stability is observed for any value of the viscous stabilization $\tau_v > 0$. Contrary to the pure advection regime, hybridized methods for diffusion have a particular form of the numerical trace that cannot recover existing FR-LDG schemes for any finite value of τ_v . LDG approaches make use of interface solution values that take the form

$$\hat{u}^{\text{FR}} = \{\{u\}\} - \zeta \llbracket u \rrbracket, \quad (5.61)$$

where ζ is a directional parameter. However, hybridized LDG (LDG-H) methods result in numerical traces defined by

$$\hat{u}^{\text{HFR}} = \{\{u\}\} - \zeta \llbracket u \rrbracket - \theta \llbracket \mathbf{q} \rrbracket. \quad (5.62)$$

Comparing the explicit definitions in Section 5.3.1.1, we see that

$$\theta = \frac{1}{2\{\{s\}\}}, \quad (5.63)$$

and $\theta \neq 0$ for any finite value of the stabilization. We also note that (5.60) is consistent with the HDG method in [59] if we set $c = 0$.

5.3.2 Local Post-Processing

Post-processing techniques have been widely used to improve the accuracy of numerical solutions. This is possible due to the optimal convergence rates of the solution and flux in locally-conservative methods involving diffusion operators [54, 144, 57]. Thus, this approach will only be applied to our problems with discontinuous trace variables and not the EFR method, whose DG equivalent has been shown to display suboptimal flux convergence in [61] since the flux is not single-valued at the flux points. This approach has been applied to numerous types of problems involving steady-state and time-dependent problems. The procedure generally consists of using a Raviart-Thomas projection of the flux to obtain a better approximation in $H(\text{div}; \Omega)$ and solving a local problem for the solution. Recall that in the FR approach, we use correction functions to create a C_0 -continuous flux function. With these correction functions, the discontinuous flux can be reconstructed to take the values of the Riemann fluxes, upgrading it to \mathbb{P}_{p+1} . In this section, we present a modified version of the post-processing method presented in [59], which leverages the operators already defined in the FR framework without the need to create additional RT formulations.

The first step in obtaining an elementwise superconvergent solution u_k^{h*} is to post-process the flux. While this first step can be done on the total advective and diffusive fluxes for linear problems [59], we choose to work with the viscous flux only for simplicity. On each element, we reconstruct the diffusive flux by computing

$$\tilde{\mathbf{F}}_k^{*(v)} = \tilde{\mathbf{F}}_k^{hD(v)}(\tilde{\mathbf{x}})|_{\tilde{\mathbf{x}}_s^*} + \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}^*} \mathbf{g}_f^m(\tilde{\mathbf{x}})[\tilde{H}(\tilde{\mathbf{x}})]_{\tilde{\mathbf{x}}=\tilde{\mathbf{x}}_f^{*m}}, \quad (5.64)$$

where $\tilde{\mathbf{x}}_f^{*m}$ are N_s^* post-processing points that define polynomials of degree $p^* = p + 1$ and \mathbf{g} is the correction vector function of the same degree. From this post-processed flux, we can

now obtain a more accurate solution by solving

$$-\beta \left[\sum_{i=1}^{N_s^*} \tilde{\nabla} \varphi_i^* \cdot \nabla u_{k,i}^* - \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}^*} \tilde{\nabla} \cdot \mathbf{g}_f^{*m} (\nabla u_k^* \cdot \tilde{\mathbf{n}}_f^m) \Big|_{\mathbf{x}_{f,m}^r} \right] \quad (5.65a)$$

$$= \sum_{i=1}^{N_s^*} \tilde{\mathbf{F}}_{k,i}^{*(v)} \cdot \tilde{\nabla} \varphi_i^* - \sum_{f=1}^{N_f} \sum_{m=1}^{N_{r,f}^*} \tilde{\nabla} \cdot \mathbf{g}_f^{*m} (\tilde{\mathbf{F}}_{k,f}^{*(v)} \cdot \tilde{\mathbf{n}}_f^m) \Big|_{\mathbf{x}_{f,m}^r}, \quad (5.65b)$$

$$\int_{\Omega_k} (u_k^h - u_k^{h*}) d\mathbf{x} = 0.$$

where

$$\nabla u^* = \mathbf{J}^{-T} \tilde{\nabla} u^*(\tilde{\mathbf{x}}) = \mathbf{J}^{-T} \sum_{j=1}^{N_s^*} U_{k,j}^* \tilde{\nabla} \varphi_j^*, \quad (5.66)$$

which can be shown to be the FR discretization of the following problem at the element level

$$\nabla \cdot (-\beta \nabla u) = \nabla \cdot \mathbf{F}_k^{*(v)}, \quad (5.67a)$$

$$-\beta \nabla u \cdot \mathbf{n} = \mathbf{F}_k^{*(v)} \cdot \mathbf{n}, \quad (5.67b)$$

$$\int_{\Omega_k} (u - u^*) d\mathbf{x} = 0. \quad (5.67c)$$

The last statement ensures elementwise conservation of the solution. This post-processing leverages the existing FR correction functions to enable superior convergence of the methods and extends the post-processing schemes to the full family of FR schemes. This post-processing can be applied to schemes involving any of the VCJH correction functions and recovers a scaled form of the conventional HDG post-processing when \mathbf{g} is constructed with c_{DG} . Otherwise, the post-processing schemes seem to be new. Later in this work, we perform numerical examples to showcase the superconvergent characteristics of hybridized FR methods. In all cases, we use the same correction function for both the solution and the post-processing steps for the sake of consistency.

5.3.3 Numerical Examples

In this section, we perform numerical experiments to discuss the stability, performance, and accuracy of hybridized FR methods in advection-diffusion problems. We mainly consider three values of the c parameter that recover existing high-order formulations. These include c_{DG} , c_{SD} and c_{HU} for $p = 1$ to $p = 4$ schemes. For the linear problems, the post-processing scheme is applied with the same values of these correction parameters. The L_2 -norm of the

solution error is measured via

$$E_{L_2}(\Omega^h) = \sqrt{\frac{1}{|\Omega|} \sum_{k=1}^N \int_{\Omega_k} (u_k^h - u^e)^2 d\Omega_k}, \quad (5.68)$$

where u^e is the analytical solution. Furthermore, we first consider a linear steady-state problem and two unsteady cases. Then, we present a nonlinear problem involving the compressible Navier-Stokes equations. All simulations are carried out serially on a 3.2 GHz Intel Core i5-5600 processor with 16Gb of RAM. The implicit system uses an exact Jacobian and is solved via the RAS preconditioner in the PETSc framework [112] with. The order convergence tables obtained in this section have been placed in Appendix A for the sake of brevity.

5.3.3.1 Steady-State Linear Advection-Diffusion

Consider the linear advection-diffusion equation with a source term chosen such that the exact solution is given by


$$u(\mathbf{x}) = xy \frac{(1 - e^{(x-1)\alpha_x})(1 - e^{(y-1)\alpha_y})}{(1 - e^{(1-\alpha_x)})(1 - e^{(1-\alpha_y)})}, \quad (5.69)$$

defined on $\Omega = [0, 1]^2$. Dirichlet boundary conditions are obtained directly from the exact solution. This case has been used to analyze the accuracy and post-processing of steady-state HDG methods for the weakly convection-dominated regime in [59]. The advection velocity is set to $\boldsymbol{\alpha} = [\alpha_x, \alpha_y] = [25, 25]$ and the diffusion coefficient to $\beta = 1$. Due to the relatively large advection velocity, a boundary layer is expected to form toward the right and top ends of the domain. We consider the L_2 -norm of the error in a reduced space $\Omega^{L_2} = [0.1, 0.9]^2$ to exclude the resolution of the boundary layer. The grid was generated using the following stretching function

$$\mathbf{x} = \frac{1}{a} \tanh \left(\frac{[i, j]}{\sqrt{N}} \operatorname{arctanh} a \right), \quad 0 \leq i, j \leq \sqrt{N} - 1, \quad (5.70)$$

where N is the total number of elements, and we set $a = 0.995$.

We make use of four levels of refinement with 5×5 , 10×10 , 20×20 , and 40×40 quadrilateral elements. Table A.2 shows the L_2 -norm of the error for $p = 1$ to $p = 4$ standard and hybridized FR schemes with correction parameters c_{DG} , c_{SD} and c_{HU} . From the stability section, we observed that c acts as an added dissipation mechanism to the c_{DG} schemes, for which $c = 0$. In this table, it can be observed that all schemes achieved the expected $p + 1$ order of accuracy in logarithmic scale and that c_{DG} has the smallest L_2 error for each of the



(a) Error without post-processing (b) Error after post-processing

Figure 5.15. Contours of the solution and post-processed solution for the steady-state linear advection-diffusion problem for a $p = 2$ solution on a 10×10 grid and $c = c_{SD}$. Linear colour scaling adjusted to $[0, 6 \times 10^{-3}]$ from white to black

considered types of discretization. Among these, the considered hybridized formulations in this work are more accurate than the standard FR-LDG formulation at the coarsest to finest levels. Specifically, the EFR method displayed the smallest error levels compared to HFR and FR for the finer levels of refinement. For instance, $p = 3$ schemes with c_{DG} displayed errors of 6.16×10^{-8} , 5.99×10^{-8} and 5.01×10^{-8} for FR, HFR and EFR, respectively. For HFR, where discontinuous traces are used, and for conventional FR methods, we applied the post-processing scheme in Section 5.3.2 and show the results obtained in Table A.3. As established in the literature, LDG methods are defined as in Equation (5.61), and they do not possess the superconvergence property in the general case [144], but can be shown to superconverge for Cartesian grids with special choices of the common fluxes [145]. For the sake of completeness, we show the results for this specific configuration in Table A.4 of this work but do not consider it anymore as it is not a feature of the arbitrary case. Interestingly, all c_{DG} methods for HFR were able to achieve the expected $p + 2$ -order of accuracy after post-processing. However, when $c \neq c_{DG}$, only methods with $p > 1$ were able to achieve the superconvergent behaviour. An example contour with the error in the computational domain is displayed in Figure 5.15, where the error levels can be seen to decrease by an order of magnitude.

To further visualize the impact of the c -parameter on this superconvergent behaviour, we perform an additional set of simulations and compute the order of accuracy for values of $c \in [0, 2c_{HU}]$ and display the results in Figure 5.16. It is known that the expected $p + 1$ convergence of FR methods is lost at large values of c [31]. For the considered relatively small range of c -parameters, the order of accuracy of the solution before post-processing slowly reduces as c increases, but consistent with the results in the convergence study, rapid decay of

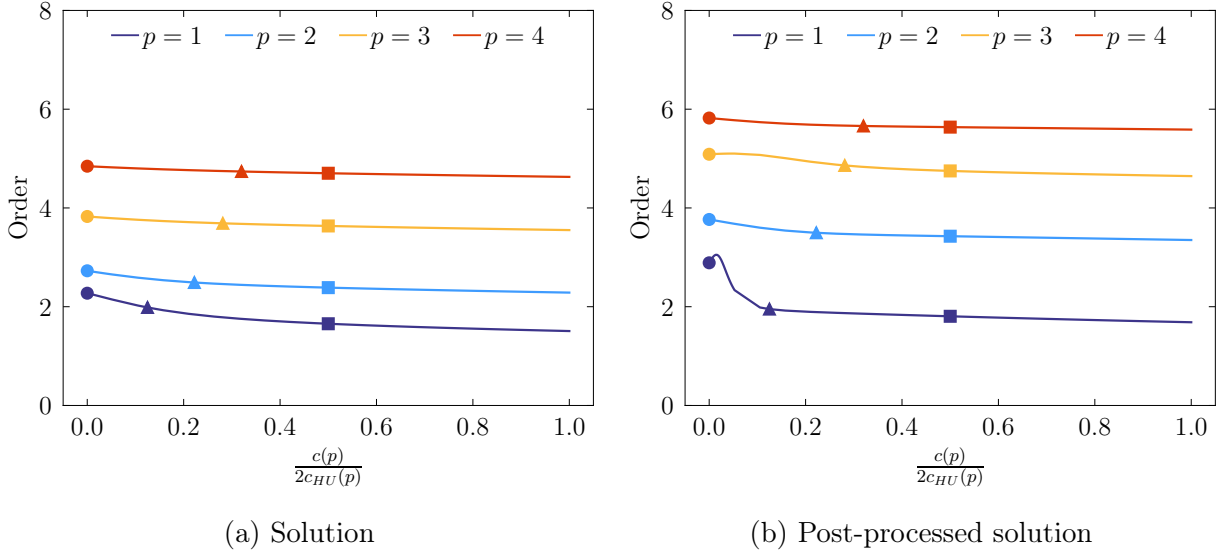




Figure 5.16. Effects of the correction function on the grid convergence rate of hybridized FR methods with discontinuous traces for the steady-state linear advection-diffusion problem. Markers for c_{DG} (\circ), c_{SD} (Δ), c_{HU} (\square) highlight these corrections at polynomial degrees $p = 1$ to $p = 4$

the superconvergent order is seen at $p = 1$ from third order to second in the vicinity of $c \rightarrow 0$. Furthermore, we analyze the performance of using hybridized methods as opposed to implicit FR schemes in terms of the number of nonzeros, time spent per time step size, and the number of GMRES iterations. We denote them NNZ , t_w and N_{GMRES} , respectively. A bar is also placed on top of these quantities when we have considered their average per linear solve. For hybridized methods, t_w accounts for both the solution of the system and the recovery of the internal solution via the local problems in Equation (3.54). Results are shown in Figure 5.17 for each of the considered schemes and polynomial degrees on the finest 40×40 grid to reduce timing errors. It can be seen in Figure 5.17a that, for the highest polynomial degree considered ($p = 4$), the number of nonzeros in the system for HFR is reduced by half and by about four times for EFR in comparison with FR schemes. Despite these metrics, it can be seen that the time spent on solving these systems before final convergence was achieved, with a small fraction of that used in FR. Specifically, considering the $p = 4$ simulations, we observe a reduction of 20-22 times for the HFR method and 40-48 times for the EFR method, depending on the value of c . This large difference can be attributed to the number of GMRES iterations shown in Figure 5.17c. Unlike the hybridized schemes, FR schemes required a significantly larger number of implicit iterations to reach convergence for the Krylov solver. Hence, hybridized FR methods have a significant benefit over standard implicit FR for all values of c in terms of performance and accuracy for this steady-state problem.



(a) Relative number of nonzeros

(b) Relative solver time



(c) Number of implicit iterations

Figure 5.17. Performance metrics for the steady-state linear advection-diffusion problem. Markers for c_{DG} (\circ), c_{SD} (Δ), c_{HU} (\square) have been added at polynomial degrees $p = 1$ to $p = 4$. For reference $NNZ_{\max} = 1336250$, $\bar{t}_{w,\max} = 2.00s$, $\bar{N}_{\text{GMRES},\max} = 420$

5.3.3.2 Advection-Diffusion of a Sine Wave

Next, verification of unsteady linear advection-diffusion is performed in this section. Consider the unsteady linear diffusion equation with initial condition

$$u(\mathbf{x}) = \sin\left(\frac{2\pi}{L}x\right) \sin\left(\frac{2\pi}{L}y\right), \quad (5.71)$$

in a square domain of side length L with periodic boundary conditions. The viscous stabilization was chosen to be $\tau_v = \beta = 0.1$, and the simulation was run for one cycle on grids of 5×5 , 10×10 , 20×20 and 40×40 . Then, the L_2 norm of the error was computed


considering the analytical solution

$$u^e(\mathbf{x}, t) = \exp(-2\beta\pi^2 t) \sin[\pi(x - \alpha_x t)] \sin[\pi(y - \alpha_y t)], \quad (5.72)$$

after one cycle. The time discretization is done with a third-order single-diagonally implicit Runge Kutta method SDIRK₃ using a consistent time step size $\Delta t = 1 \times 10^{-4}$ in all grids to reduce the temporal error. This corresponds to a CFL number of 0.1 on the finest grids. Table A.5 shows the orders of accuracy for the considered correction functions and polynomial degrees 1 to 4, where $p + 1$ convergence was observed for all considered schemes. Similar to the previous problem, hybridized methods displayed smaller error levels than the standard FR discretization for all considered values of c , with EFR being the most accurate at the finer grid levels. Specifically, $p = 3$ schemes show errors in the L_2 norm of 2.07×10^{-8} , 1.25×10^{-8} and 1.21×10^{-8} for the FR, HFR and FR methods with c_{DG} , respectively. A reduction of about half of the error was obtained with the hybridized schemes. Note that, as before, c_{DG} showed the best accuracy out of the three correction functions.

The solution was post-processed using correction functions of the same c -parameter. Results and convergence orders are shown in Table A.6 for the post-processed solution. FR-LDG did not show super accuracy for the case in which the switch was taken to be in an arbitrary fashion, but HFR achieved the expected order $p + 2$ for the post-processed solution. Similar to the previous case, methods with $c \neq c_{DG}$ do not exhibit the $p + 2$ order of accuracy at $p = 1$.

In terms of performance, Figure 5.18 shows the error in the L_2 norm resulting from the 40×40 grids with all polynomial degrees against the wall-clock time. We observe the wall-clock time required to obtain a certain level of the L_2 error. Two values of the time step size were chosen to perform the comparison, namely $\Delta t = 5 \times 10^{-4}$ and $\Delta t = 1 \times 10^{-4}$, for which results appear in Figures 5.18a and 5.18b, respectively. These represent CFL numbers of 0.5 and 0.1, respectively. Clearly, the use of a large time step size is detrimental to the accuracy of the solution for the finest simulations in all runs. However, hybridized methods were able to achieve smaller levels of the L_2 norm in both cases at a fraction of the cost. Specifically, at $p = 4$, simulations were between 14.06 and 17.52 times faster for HFR and between 22.5 to 30.1 times faster for EFR, compared to standard implicit FR schemes for the considered values of c . The speedups for all other runs on the finest grids are shown in Table 5.6. Reducing the time step size from $\Delta t = 5 \times 10^{-4}$ to $\Delta t = 1 \times 10^{-4}$ improved the performance of the standard FR formulation by ~ 1.6 times per linear solve, but the improvements for the hybridized methods were not significant. To further investigate the influence of the time step size on the performance results, we carry out an additional set of simulations on the 40×40 grid for a larger range of time step sizes at $p = 3$. We show the results in Figure 5.19. Consistently, reducing the time step size improves the performance of



(a) Relative solver time with $\Delta t = 5 \times 10^{-4}$ (b) Relative solver time with $\Delta t = 1 \times 10^{-4}$

Figure 5.18. Performance metrics for the unsteady linear advection-diffusion problem. Markers for c_{DG} (\circ), c_{SD} (Δ), c_{HU} (\square) have been added at polynomial degrees $p = 1$ to $p = 4$. For reference, $\bar{t}_{w,\max}(\Delta t = 1 \times 10^{-4}) = 0.1522$ and $\bar{t}_{w,\max}(\Delta t = 5 \times 10^{-4}) = 0.2409$

the standard FR formulation per linear solve, where a reduction by a factor of ten showed an effect of 2.5 times better performance per step. However, the impact of the time step size was less significant for the hybridized methods, whose performance per linear solve remains almost constant with the considered Δt increase, especially for the EFR schemes.




Figure 5.19. Normalized time to solve the linear system as a function of the time step size. The performance of FR methods has a larger dependence on the size of the time step size as opposed to hybridized methods. Markers for c_{DG} (\circ), c_{SD} (Δ), c_{HU} (\square) have been added at polynomial degrees $p = 1$ to $p = 4$

Table 5.6. Speedup factors for the unsteady linear advection-diffusion problem with 1600 elements. Values are shown for two time step sizes and compared against the corresponding FR formulation with the same c -parameter

Δt	p	HFR			EFR		
		c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}
1×10^{-4}	1	0.76	0.89	0.84	2.00	2.07	1.81
	2	2.71	2.75	2.59	4.46	4.35	4.73
	3	7.72	7.49	6.49	10.84	11.53	9.92
	4	15.22	14.37	12.81	22.60	20.42	17.29
5×10^{-4}	1	1.08	1.08	1.03	2.95	2.58	2.24
	2	4.41	4.25	3.52	7.78	6.78	6.00
	3	10.43	9.59	9.62	17.42	16.18	16.26
	4	17.52	16.51	14.06	30.12	26.35	22.56

5.3.3.3 Advection-Diffusion of a Gaussian Profile

We now consider unsteady advection-diffusion of a Gaussian profile to analyze the stabilization mechanism of hybrid FR methods. To this end, we make use of a $[-5, 5]^2$ domain with periodic boundary conditions and an initial condition

$$u(\mathbf{x}, 0) = e^{-(x^2+y^2)}. \quad (5.73)$$

We consider a pure advection problem with $\beta = 0$, $\boldsymbol{\alpha} = [1, 1]$, and an advection-diffusion problem with $\beta = 0.01$ and the same advection velocity. We run a set of simulations using a third-order SDIRK method with a small time step size $\Delta t = 0.005$ on a 20×20 grid, representing a CFL number of 0.01. This small value helps mitigate temporal errors. All runs are performed with $p = 3$ spatial discretizations for five convective times $t_c = 5$ for the advection problem and for one convective time $t_c = 1$ for the advection-diffusion case.

First, we discuss the purely advective case. Recall that $\bar{\lambda} = 0$ results in undefined hybridized methods. For HFR methods, we have previously shown their equivalence in the stability analysis section when the conservation law includes only the advection operator. For the HFR method, we consider schemes ranging from the central approach of Equations (5.57)-(5.58), to increasing values of the upwinding parameter $\bar{\lambda} \in \{0.3, 0.5, 0.7, 1.0\}$, where $\bar{\lambda} = 1$ represents the upwind scheme. For the EFR method, we consider $\bar{\lambda} \in \{0.001, 0.3, 0.5, 0.7, 1.0\}$. For all methods, we compute the evolution of the solution energy for correction parameters c_{DG} , c_{SD} , and c_{HU} . Results are shown in Figure 5.20 with a zoomed-in version in Figure 5.21. The behaviour in these figures is consistent with that previously observed in [29, 143], where larger values of c introduce additional dissipation and $\bar{\lambda} \rightarrow 0$ approaches a central FR method,

for which the expected oscillatory behaviour of $c \neq 0$ is observed. As opposed to HFR, EFR methods require a larger value of $\bar{\lambda}$ for stability. As we previously discussed in the analysis of the stability section, we cannot easily guarantee stable EFR methods for $\bar{\lambda} < \frac{1}{2}$, and our experiments reveal that the minimum value is problem-dependent. Clearly, results for $\bar{\lambda} = 0.3$ are unstable for $c = 0$, blowing up after about 35 convective times, but the added dissipation of c_{SD} and c_{HU} kept the simulations stable for this $\bar{\lambda}$ parameter for the duration of these simulations. This is consistent with the analysis, where larger values of c result in more negative $\frac{d\|u\|_2}{dt}$. While EFR methods may appear to be stable for values $\bar{\lambda} < \frac{1}{2}$ for some problems, they may be only mildly stable and may blow up later in long-time integration simulations.




Figure 5.20. Time-evolution of the solution energy for hybridized FR methods for the Gaussian Pulse case with $\beta = 0$ for different values of $\bar{\lambda}$. FR and HFR methods exhibit exactly the same curves, so only one is shown on the left. The value $\bar{\lambda}$ refers to the central FR method recovered via HFR with stabilization in Equations (5.57)-(5.58) and $\bar{\lambda} \rightarrow 0$ for EFR. Line strokes define methods with c_{DG} (—), c_{SD} (---), c_{HU} (⋯⋯)

Next, we consider the advection-diffusion case on the same computational grid. Similarly, we plot the energy of the solution against the convective time in Figure 5.22. In this case, we show results for FR, HFR, and EFR, none of which are equivalent. In the case of FR, larger $\bar{\lambda}$ means larger numerical dissipation, consistent with the purely advective case. Due to the physical diffusion, the energy of the solution is seen to decay for the linear advection-diffusion problem with $\beta = 0.01$. A larger range of $\bar{\lambda}$ parameters is stable for the EFR method compared to the purely advective case due to physical dissipation. However, it can be seen that for the smallest considered value, an increase in energy is seen close to the end of the simulation for all considered values of c , with larger c introducing additional dissipation. Results look very similar among FR, HFR, and EFR from this view, and we show a zoomed-in




Figure 5.21. Zoom of the time-evolution of the solution energy for hybridized FR methods for the Gaussian Pulse case with $\beta = 0$ for different values of $\bar{\lambda}$. The value $\bar{\lambda}$ refers to the central FR method recovered via HFR with stabilization in Equations (5.57)-(5.58) and $\bar{\lambda} \rightarrow 0$ for EFR. Line strokes define methods with c_{DG} (—), c_{SD} (---), c_{HU} (⋯)

version for a time range $0.4 < t_c < 0.45$ in Figure 5.23, moments before the peak of the Gaussian pulse reaches the periodic boundary. Smaller values of $\bar{\lambda} \rightarrow 0$ for HFR and EFR make the gradient jump term grow in the trace equation. See Equations (5.22), and (5.23). Then, setting $\bar{\lambda} \rightarrow 0$ does not recover the behaviour of a central scheme when diffusion operators appear. Hence, for advection-diffusion problems, the HFR central scheme cannot be recovered with the formulation in (5.57), (5.58) for the advection-diffusion case. With this example, we observe that results are consistent with the analytical findings of the stability section and demonstrate a range of stable HFR and EFR methods for advection-diffusion.

5.3.3.4 Planar Couette flow

Finally, we present a problem involving the compressible Navier-Stokes equations. Planar Couette flow is a well-known two-dimensional case to perform verification of the viscous fluxes given its simplification of the Navier-Stokes equations. This problem consists of viscous flow between two plates separated by a distance ℓ . A moving wall is located at $y = \ell$ with temperature T_e and constant velocity v_e , which drives the flow in the positive x -direction. At $y = 0$, a fixed wall ($v_w = 0$) with temperature T_w is placed. Due to the no-slip condition, the flow variables are equal to those of the walls at $y = 0$ and $y = \ell$, respectively. Consequently, the flow experiences a temperature gradient due to viscous dissipation. The exact temperature profile can be computed from

$$T = T_w + \left[T_e - T_w + \frac{\text{Pr}}{2c_p} v_e^2 \right] \frac{y}{\ell} - \frac{\text{Pr}}{2c_p} v_e^2 \left(\frac{y}{\ell} \right)^2, \quad (5.74)$$

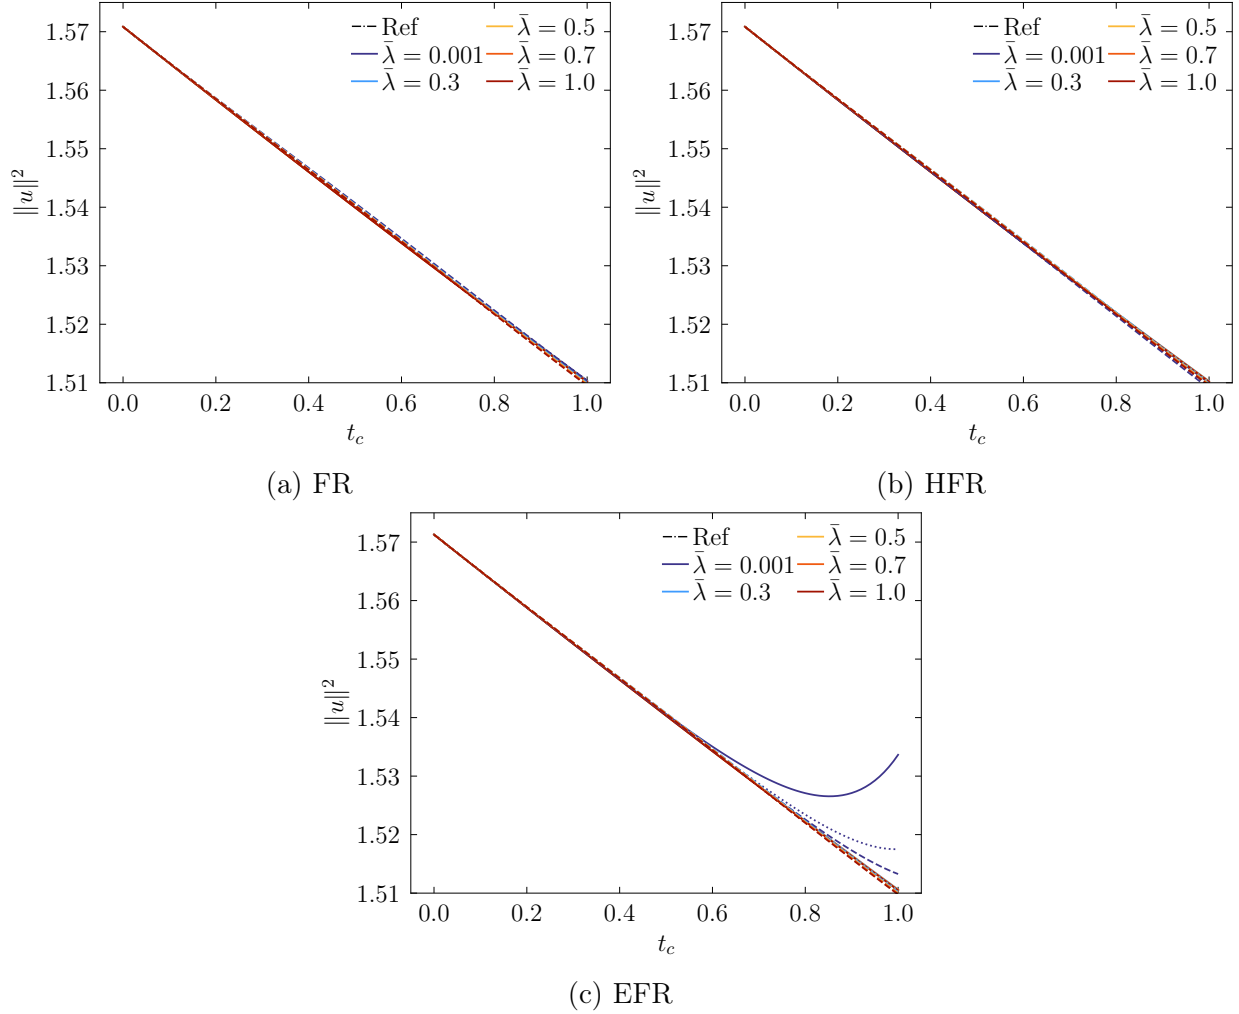


Figure 5.22. Time-evolution of the solution energy for hybridized FR methods for the Gaussian Pulse case with $\beta = 0.01$ for different values of $\bar{\lambda}$. Line strokes define methods with c_{DG} (—), c_{SD} (---), c_{HU} (⋯)

where $\text{Pr} = 0.71$. The Mach number is set to 0.1 and $\text{Re} = 5$. We consider four levels of refinement using grids composed of 4×2 , 8×4 , 16×8 and 32×16 quadrilateral elements with $p = 1 - 4$ schemes and the same c -parameters considered in the previous section. For stabilization, we consider a Lax-Friedrichs type matrix for the inviscid fluxes such as that in [32] for the Euler equations, and constant viscous stabilization of the form

$$s_{f,k}^{(v)} = \frac{1}{\text{Re}}, \quad (5.75)$$

for the hybridized formulations. For standard FR, we also make use of a Lax-Friedrichs Riemann solver for the inviscid fluxes and the LDG method with an arbitrary directional switch for the viscous component. The simulations were initialized with a stationary problem

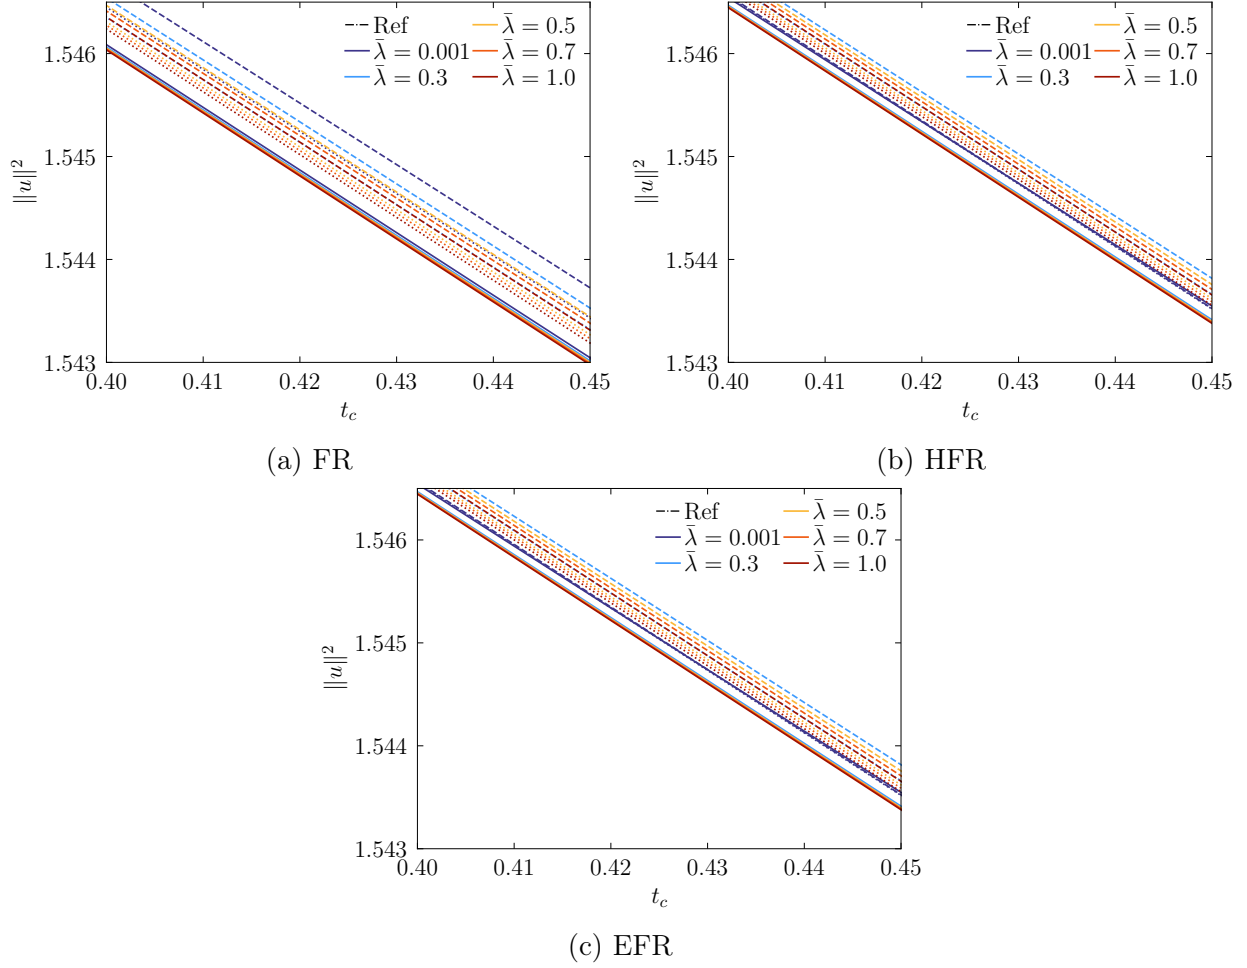


Figure 5.23. Zoom of the time-evolution of the solution energy for hybridized FR methods for the Gaussian Pulse case with $\beta = 0.01$ for different values of $\bar{\lambda}$. Line strokes define methods with c_{DG} (—), c_{SD} (---), c_{HU} (.....)

($\mathbf{v} = \mathbf{0}$) and allowed to run until the density residual converged to a tolerance of 10^{-10} in the L_∞ norm, which proved to be sufficient for the temperature error in the L_2 norm to converge. For convergence acceleration, a relaxation factor was used with an implicit Euler scheme to converge the nonlinear residuals via a time-step ramp function. Specifically, the following function at the i -th iteration was used to update the time step size

$$\Delta t_i = \begin{cases} \Delta t_0 & i < 20, \\ 2^{1/16} \Delta t_{i-1} & \text{mod}(i, n) = 0 \text{ and } \Delta t_{i-1} < 10^4 \Delta t_0, \\ \Delta t_{i-1} & \text{otherwise,} \end{cases} \quad (5.76)$$

which was chosen empirically to accelerate the convergence of the FR simulation consistent with the number of Jacobian updates every n iterations. This was used for all grids and




Figure 5.24. Performance metrics for the planar Couette problem

Table 5.7. Speedup factors for the Couette problem with 32×16 elements for the HFR and EFR methods. These factors take into account the time to solve the global and local problems but exclude the assembly time of the Jacobian matrices

p	HFR			EFR		
	c_{DG}	c_{SD}	c_{HU}	c_{DG}	c_{SD}	c_{HU}
1	2.31	1.72	1.53	14.84	11.11	9.33
2	6.91	5.46	3.57	21.95	16.66	10.74
3	18.10	11.73	10.77	49.45	30.31	27.26
4	40.38	25.12	19.86	90.7	54.24	41.73

polynomial degrees. The value of the base time step size was set to $\Delta t_0 = 1 \times 10^{-5}$. The Jacobian matrix was computed exactly and updated every five time steps to reduce the computational cost associated with its assembly. Verification is presented in Table A.7, where the L_2 norm of the error is computed for the aforementioned levels of refinement. It is interesting to see that for this nonlinear problem, HFR displayed the lowest L_2 -norm, followed by EFR and then FR. Specifically, L_2 errors of 6.17×10^{-10} , 4.75×10^{-10} and 4.86×10^{-10} were obtained for FR, HFR, and EFR, respectively. The expected orders of accuracy were obtained in all cases and considered values of c . Consistent with our previous experiments, we analyze the performance based on the time spent on the solution of the linear system, and for the hybridized method, this accounts for the solution of the local problems. The number of nonzeros in the implicit system reduces between 4 to 7 times for the HFR and EFR methods compared to FR for the finest problems, as shown in Figure 5.24. Interestingly, the effect of the c -parameter on the FR simulations can be significant. For $c = c_{DG}$, simulations are twice as expensive as those using c_{HU} . This can be attributed to fewer nonzero entries for this

value of c as well as a stiffer problem resulting from setting $c = 0$. This is consistent with our observations in [32] for the advection regime. We observed significant speedup values, which are shown in Table 5.7. For example, between 19.86 and 40.38 times faster simulations were observed for the HFR discretizations and between 41.73 to 90.7 times faster solutions for the EFR discretizations. Plots of the residual and temperature error against the wall-clock time




Figure 5.25. Temperature and density residual against wall-clock time in seconds for a $p = 4$ scheme with different correction functions on the 16×8 grid. Line strokes represent methods with c_{DG} (—), c_{SD} (---), c_{HU} (⋯)

are shown in Figure 5.25 for the finest grids, and $p = 4$ runs as an example of the evolution of the convergence. The speedups obtained with this problem are significantly larger than the ratios of nonzeros. In Figure 5.26, we show the time spent to solve the linear system at every iteration. Here, we observe that the increasing time step size was detrimental to the FR time spent on solving these systems much more so than for the hybridized methods. Hence, we can significantly reduce the cost of implicit FR simulations for nonlinear viscous problems via hybridization.




Figure 5.26. Average time to solve the linear system per iteration ($p = 4$ finest grids). Markers represent methods with c_{DG} (\bullet), c_{SD} (\blacktriangle), c_{HU} (\blacksquare)

Chapter 6

Implicit Polynomial-Adaptive Hybridized Formulations

6.1 Overview

Practical applications of engineering interest require large computational domains. When using high-order methods, increasing the polynomial degree impacts the computation time significantly due to a large increase in degrees of freedom. Polynomial adaptation algorithms allow us to modify the degree of the polynomial representing the solution within each element, such that high resolution is maintained only in regions of interest. In this chapter, we develop a formulation that generalizes polynomial adaptation to both discontinuous and continuous hybridized FR methods, which can be applied at intervals during a simulation. Specifically, we present p -adaptive hybridizable flux reconstruction schemes to solve vortex-dominated flows involving the compressible Euler and Navier-Stokes equations. We make use of a feature-based vorticity indicator to assign solution polynomial degrees and a global projection to determine the new values of the trace after the adaptation procedure. Finally, we demonstrate the application of these methods with a series of numerical examples.

6.2 Implementation

In practice, only regions of flow complexity, such as unsteady vortex shedding, require high resolution. The local nature of flux reconstruction schemes allows a natural implementation of adaptivity algorithms. In order/ p -adaptation, we increase or decrease the degree of solution polynomials within elements depending on a form of error estimation. Feature-based indicators have been shown to be useful for unsteady flows. In particular, the nondimensional vorticity indicator of [83] has shown potential for vortex-dominated flows. The indicator

determines the maximum nondimensional vorticity magnitude within each element Ω_k , i.e.

$$\Phi_k = \frac{\bar{Z}_k l_k}{v_\infty}, \quad (6.1)$$

where \bar{Z}_k is the maximum vorticity magnitude

$$\bar{Z}_k = \max_{1 < i < N_s} |Z_{k,i}|, \quad Z_{k,i} = \nabla \times \mathbf{v}_{k,i}, \quad (6.2)$$

l_k is the maximum length between two mapping points in the element, and v_∞ is the freestream velocity. Then, we follow Algorithm 2, where the new polynomial degree is decided inside each element based on a maximum polynomial degree bound p_{\max} , a threshold vector $\mathbf{c} = [c_1, \dots, c_{p_{\max}}]$ [83], and a tolerance factor ϵ . This threshold vector determines the ranges of the indicator for which a given polynomial degree will be chosen. Depending on whether the new polynomial is higher or lower, an element-wise interpolation or L_2 projection of the interior solution is respectively applied to the interior elements. For hybridized methods, the updated trace can be found in line with the transmission conditions by enforcing

$$\sum_{\bar{f} \in \varepsilon^h} \int_{\bar{f}} \llbracket \hat{\mathfrak{F}}(\bar{\mathbb{P}}u^h, \hat{u}^h) \rrbracket_{\bar{f}} \phi_{\bar{f}}^p d\bar{f} = 0, \quad (6.3)$$

where $\bar{\mathbb{P}}u^h$ is the locally projected interior solution, readily available from Algorithm 2, and $\phi_{\bar{f}}^p$ is the new trace basis function at face \bar{f} . The degree of the trace nodal basis function after element adaptation is determined as the maximum polynomial degree between the two neighbouring elements at a given interface. The procedure is briefly explained in Algorithm 3. We note that for the hybridized method with discontinuous traces, Equation (6.3) yields a problem local to each face, which can be solved efficiently. Hence, for general problems, a nonlinear system of equations must be solved at every adaptation call. However, this does not need to be performed at every time step but instead at a reasonable fraction of the characteristic time.

To visualize the underlying projection and its relation to a p -uniform formulation, consider its application to linear advection with constant velocity $\boldsymbol{\alpha}$ using the HFR method. Equation (6.3) becomes at an interface \bar{f}

$$\int_{\bar{f}} \llbracket \boldsymbol{\alpha} \cdot \mathbf{n} \hat{u}^h + s(\bar{\mathbb{P}}u^h - \hat{u}^h) \rrbracket_{\bar{f}} \phi_{\bar{f}}^p d\bar{f} = 0, \quad \forall \bar{f} \in \varepsilon_0^h, \quad (6.4)$$

Algorithm 2: Vorticity-based indication of element-wise polynomial degree

```

for  $\Omega_k \in \mathcal{T}^h$  do
  Compute  $\Phi_k$  using Equation (6.1)
  if  $\Phi_k < c_1\epsilon$  then
    |  $p_k = 1$ 
  end
  else
    | for  $i \in 2, \dots, p_{\max}$  do
      | | if  $\Phi_k < c_i\epsilon$  then
        | | |  $p_k = i$ 
      | | end
    | end
  end
end

```

Algorithm 3: Projection Algorithm

```

for  $\Omega_k \in \mathcal{T}^h$  do
  | Compute  $u^{\bar{\mathbb{P}}} = \bar{\mathbb{P}}u_k^h$ 
end
for  $\bar{f} \in \varepsilon_0^h$  do
  |  $p_{\bar{f}} = \max(p_+, p_-)$ 
end
  Solve for  $\hat{u}$  using Equation (6.3)

```

or equivalently for homogeneous face stabilization

$$\int_{\bar{f}} \llbracket \bar{\mathbb{P}}u^h - \hat{u}^h \rrbracket_{\bar{f}} \phi_{\bar{f}}^p d\bar{f} = 0, \quad \forall \bar{f} \in \varepsilon_0^h, \quad (6.5)$$

noting that $\llbracket \boldsymbol{\alpha} \cdot \mathbf{n} \rrbracket = 0$. Since $\hat{\mathfrak{F}}_{\bar{f}}$ is a polynomial of the same degree as the new basis function $\phi_{\bar{f}}^p$, $p_{\bar{f}} = \max(p_-, p_+)$, the problem can be shown to become

$$\sum_{j=1}^{N_{r,\bar{f}_p}} \left[u_j^{\bar{\mathbb{P}},+} + u_j^{\bar{\mathbb{P}},-} - 2\hat{u}_j^h \right] \int_{\bar{f}} \phi_j \phi_i d\bar{f} = 0, \quad (6.6)$$

where $u_j^{\bar{\mathbb{P}},+}$ is the projected solution at the + side of the interface. N_{r,\bar{f}_p} is the number of trace points on a face after the projection. Hence, it can be shown from here that the relation in brackets is satisfied pointwise for Gauss-Legendre points, typically used in HFR. In this case, after application of the adaptation projection, the form of the trace in linear advection

becomes

$$\hat{u}_j^h = \frac{u_j^{\mathbb{P},+} + u_j^{\mathbb{P},-}}{2}, \quad (6.7)$$

which is consistent with the standard flux reconstruction form of the fluxes in the adaptive setting [74]. This consistency is also observed between p -uniform HFR and FR formulations for linear advection as outlined in Chapter 5.

6.3 Numerical Examples

In this section, we present a series of numerical examples to showcase the benefits of p -adaptation in hybridized FR schemes. We will discuss results for an isentropic vortex problem, flow over a cylinder at $\text{Re} = 150$ and flow over a NACA0012 airfoil at $\text{Re} = 10,000$. All of the performance computations correspond to serial runs.

6.3.1 Isentropic Vortex

First, we consider the isentropic vortex problem governed by the compressible Euler equations. This problem consists of a vortex advecting through a periodic square domain of side length $L = 40$. We choose this length so that the vortex occupies a smaller portion of the domain, similar to typical aerodynamics of external flows. With this problem, we will showcase the benefits of the p -adaptive algorithm. The initial conditions are specified from the following primitives

$$\rho = \left[1 - \frac{\beta^2 M^2 (\bar{\gamma} - 1) e^{2f}}{8\pi^2} \right]^{\frac{1}{\bar{\gamma}-1}}, \quad (6.8)$$

$$v_x = \frac{\beta y e^f}{2\pi R}, \quad (6.9)$$

$$v_y = 1 - \frac{\beta x e^f}{2\pi R}, \quad (6.10)$$

$$P = \frac{\rho^{\bar{\gamma}}}{\bar{\gamma} M^2}, \quad (6.11)$$

where ρ is the density, β is the vortex strength, v_i is the velocity in the i -th direction, $\bar{\gamma} = 1.4$ is the ratio of specific heats, M is the Mach number, which we set to 0.4. The radius of the vortex R is set to 1.5 and $f = (1 - x^2 - y^2)/2R^2$. We consider a vortex strength $\beta = 13.5$ and make use of the second-order, two-stage SDIRK method with a small time step size $\Delta t = 5 \times 10^{-3}$ to allow spatial error to dominate. The corresponding CFL number is approximately 2×10^{-3} . The implicit residual was converged to a tolerance of 10^{-10} .


We run this problem for one advective cycle and repeatedly decrease the value of ϵ while

maintaining the number of elements constant (10×10 quadrilaterals). This causes the sensitivity of the indicator to increase gradually, thereby resulting in more elements with higher polynomial degrees. For a given problem, lower ϵ signifies a simulation with more degrees of freedom. Figure 6.1 compares the p -distribution for two values of ϵ . Smaller values of ϵ yield a larger number of elements with higher polynomial degrees. Figure 6.2 shows curves containing different metrics to demonstrate the benefits of polynomial adaptation in both hybridized and standard FR formulations. Figure 6.2a shows the L_2 norm of the density error against the number of internal degrees of freedom. For the same grid, increasing the polynomial degree can significantly increase the number of degrees of freedom. The decrease in numerical error occurs at a slower rate than for the adaptive simulations. When adaptation is used, only elements with a given value of vorticity are assigned a higher polynomial degree, thereby decreasing the L_2 norm of the error while maintaining the degrees of freedom at a lower value than a p -uniform simulation at the same error level. These values are tabulated in 6.1. Elements where the local velocity gradients vanish do not benefit from using high order since vorticity is approximately zero. In addition, we see that for the same number of internal degrees of freedom DOF_v , both HFR and EFR methods result in solutions with similar values of numerical error, with EFR having slightly higher error than HFR and FR in most of the runs.


In Figure 6.2b, we show the total wall-clock time against error for the same runs. For the lowest error attained by these simulations, adaptive FR is about twice as fast, HFR close to three times, and EFR about four times faster than p -uniform FR schemes. When solving the nonlinear system of equations involved in this problem, the number of internal degrees of freedom only affects the size of the FR problem directly, but for hybridized methods, the trace DOF (DOF_t) dictates it. This translates into a decrease in the computational cost for the same number of internal DOF. It is therefore observed that hybridized p -adaptive simulations are twice as fast as their FR counterparts.

Table 6.1. Resulting average number of degrees of freedom for the p -adaptive runs for the isentropic vortex problem


		$\epsilon^0/1$	$\epsilon^0/2$	$\epsilon^0/4$	$\epsilon^0/8$	$\epsilon^0/16$
FR	$\overline{\text{DOF}}_v$	7349.33	8169.62	9241.91	10139.50	10956.09
HFR	$\overline{\text{DOF}}_v$	7349.32	8169.62	9241.91	10139.50	10956.09
	$\overline{\text{DOF}}_t$	6761.73	6991.86	7265.55	7492.49	7701.76
EFR	$\overline{\text{DOF}}_v$	7353.12	8169.24	9241.00	10136.63	10957.13
	$\overline{\text{DOF}}_t$	1963.55	2191.70	2465.02	2691.13	2902.22



(a) Vorticity magnitude



(b) $\epsilon = 0.2$



(c) $\epsilon = 0.01$

Figure 6.1. Portion of the isentropic vortex domain showing contours of (a) vorticity, and polynomial distribution for (b) $\epsilon = 0.2$ and (c) $\epsilon = 0.01$ after one time step. Smaller ϵ yields more degrees of freedom. Contours correspond to an EFR simulation with 10×10 quadrilateral elements

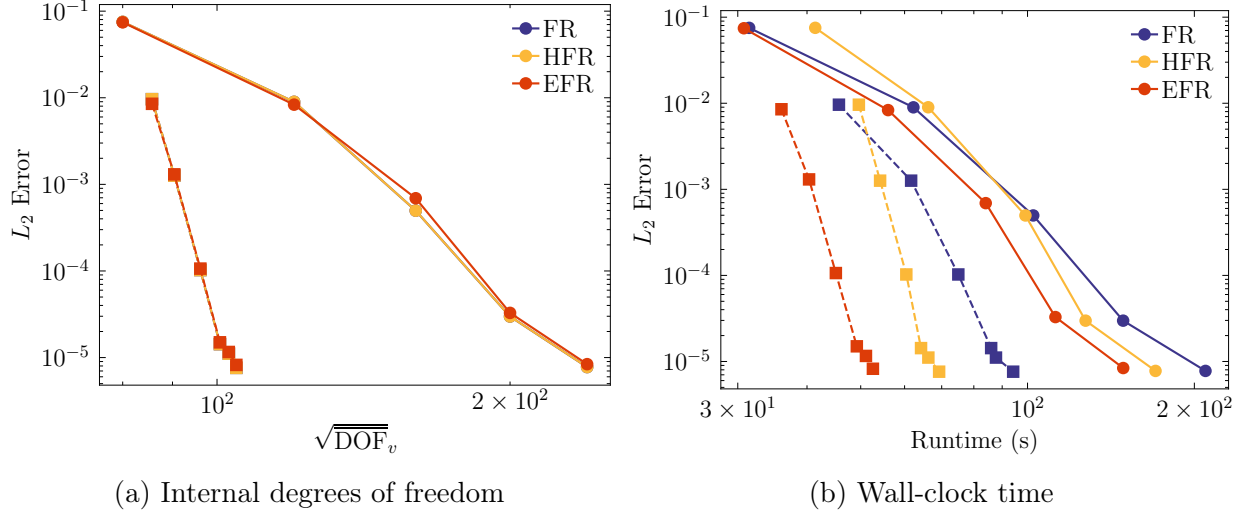


Figure 6.2. Performance metrics for the isentropic vortex problem in terms of internal degrees of freedom DOF_v and runtime in seconds. Results are shown for p -uniform runs (\bullet) and p -adaptive runs (\blacksquare)

6.3.2 Cylinder at $\text{Re} = 150$




Figure 6.3. Portion of the cylinder mesh composed of 3564 quadrilateral elements refined toward the cylinder walls to resolve the boundary layer

In our second numerical example, we simulate flow over a cylinder at $\text{Re} = 150$. This case has been previously used to demonstrate the suitability of p -adaptation in [74]. We consider the flow with $M = 0.1$. We make use of a mesh of 3564 quadrangles refined at the wall and growing away from it, as shown in Figure 6.3. A second-order two-stage SDIRK scheme was used to advance the solution in time, with a $\Delta t/t_c = 6.25 \times 10^{-3}$, which corresponds to about 100 times the maximum stable value of an explicit $\text{RK}_{4,4}$ method using $p = 5$ for this problem. This yields a CFL number ≈ 8 . Here, $t_c = tu_\infty/D$ is the convective time. The exact Jacobian matrices were computed every 100 iterations to match the interval for the adaptation algorithm, which was also applied every 100 iterations. This interval was chosen such that the Jacobian computation only takes roughly 10% of the total runtime. The

vorticity indicator tolerance was set to $\epsilon = 0.1$. The convergence tolerance was set to 10^{-5} for all cases. After the transient effects of the simulation passed, the simulation was continued for an additional 200 convective times, and averaged statistics were computed, which are compared against the works of [74, 146]. For this problem, antialiasing [147, 148, 149] was used for the EFR method to stabilize the simulation.

A comparison of vorticity contours for the same adaptive cases at the same instant is observed in Figure 6.4. Close to the cylinder, the unsteady vortex interaction pattern shows similar behaviour for the three cases, with good agreement up to 15 convective lengths D . At this point, the physics is already heavily underresolved due to the large element sizes. Here, the EFR method qualitatively shows less smooth vorticity contours.

Figure 6.5 shows the polynomial distribution at an instantaneous snapshot corresponding to a moment of minimum lift for the adaptive FR, HFR, and EFR discretizations. It is expected that regions with higher vorticity magnitude, such as the boundary layers, will require higher resolution. The overall distribution of the adapted polynomials is similar between FR, HFR, and EFR. $p = 4$ elements can be seen in the boundary layer and cores of the unsteady vortices. The polynomial degree is gradually reduced to $p = 3$ in the vicinity of $p = 4$ and so on until $p = 1$. The latter is used mostly in the far-field regions. Right from the vicinity of the back of the cylinder to about 15 characteristic lengths downstream, the three considered schemes agree well in the distribution of polynomial degrees for most elements. For the downstream cells beyond this point and in some regions upstream of the cylinder, EFR flagged higher vorticity values and hence introduced higher polynomial degrees. This can be associated with higher solution jumps, resulting in spurious local vorticity values. This results in EFR having slightly more degrees of freedom than HFR and FR, as shown in Table 6.2.

We compare our numerical results against the reference data of [74] and the references therein. Specifically, we compare the averaged drag coefficient \bar{c}_d , the amplitudes of the c_l and c_d curves, Δc_l , Δc_d , respectively, and the Strouhal number St , the latter three have been obtained via an FFT signal decomposition. From Table 6.2, it is clear that $p = 1$ is significantly less accurate than any higher polynomial degree. Overall, for the p -uniform simulations, results start to converge beyond $p = 3$ for FR, HFR, and EFR, showing little change after this value. Similar values are obtained among the three schemes. The adaptive simulations also agree well with the reference data with a relative error smaller than 0.5% in all quantities with respect to the $p = 5$ simulations. This was done at a fraction of the degrees of freedom, which we show in the last two columns of this table. The resulting internal degrees of freedom $\overline{\text{DOF}}_v$ and the averaged trace degrees of freedom $\overline{\text{DOF}}_v$ are presented. HFR and EFR formulations required only about a 1/2 and 1/3 of the degrees of freedom of the p -adaptive FR scheme in the implicit system, respectively.




(a) FR



(b) FR




(c) HFR




(d) EFR


Figure 6.4. Instantaneous snapshots of the cylinder problem showing contours of vorticity for FR, HFR, and EFR at moments of minimum lift. Similar contours are observed between the p -adaptive and p -uniform simulations



(a) FR



(b) HFR



(c) EFR

Figure 6.5. Instant snapshots of the cylinder problem showing the adaptation algorithm for FR, HFR and EFR at moments of minimum lift. Contour lines of vorticity are superimposed

Table 6.2. Averaged results for the cylinder problem for FR, HFR, and EFR schemes including both p -uniform and p -adaptive runs

	Degree	\bar{c}_d	Δc_d	Δc_l	St	$\overline{\text{DOF}}_v$	$\overline{\text{DOF}}_t$
FR	1	1.2907	0.0179	0.4659	0.1801	14256	-
	2	1.3234	0.0256	0.5155	0.1842	32076	-
	3	1.3246	0.0257	0.5163	0.1843	57024	-
	4	1.3245	0.0256	0.5161	0.1843	89100	-
	5	1.3245	0.0257	0.5163	0.1843	128304	-
	adaptive	1.3253	0.0257	0.5163	0.1836	39845	-
HFR	1	1.3042	0.0213	0.4884	0.1818	14256	14026
	2	1.3241	0.0256	0.5157	0.1843	32076	21039
	3	1.3245	0.0257	0.5163	0.1843	57024	28052
	4	1.3246	0.0257	0.5161	0.1843	89100	35065
	5	1.3244	0.0257	0.5161	0.1843	128304	42078
	adaptive	1.3274	0.0257	0.5183	0.1840	39713	22165
EFR	1	1.3190	0.0228	0.4997	0.1824	14256	3677
	2	1.3285	0.0259	0.5182	0.1846	32076	10690
	3	1.3258	0.0257	0.5169	0.1844	57024	17703
	4	1.3249	0.0256	0.5162	0.1843	89100	24716
	5	1.3243	0.0257	0.5158	0.1843	128304	31729
	adaptive	1.3277	0.0256	0.5186	0.1841	40669	12089
Cagnone et al. [74]	1.3246	0.0258	0.5166	0.1836			
Inoue et al. [146]	1.3200	0.0260	0.5200	0.1830			

A table with speedup factors is shown in Table 6.3, where we have computed $t_{p_{\text{FR}}}/t_{p_{\text{HFR}}}$ and $t_{p_{\text{FR}}}/t_{p_{\text{EFR}}}$, where p_{FR} , p_{HFR} and p_{EFR} are the polynomial degrees of a given FR, HFR, and EFR scheme. We first note that, consistent with previous observations [32], hybridized methods show increased performance benefits at higher polynomial degrees. However, by looking at the diagonals of these tables, where we compare equal polynomial degrees for the hybridized and standard formulations, we see that at least 1.26 faster simulations were achieved at $p = 1$. As the order is increased, speedups of up to 14.5 are observed for the EFR method at $p = 5$ for p -uniform discretizations of the same degree. Hybridized adaptive methods achieved speedups of 5.34 and 6.42 for the HFR and EFR methods with respect to the adaptive standard FR scheme. Furthermore, we see the benefits of using p -adaptive hybridization, which allows us to obtain results comparable to $p = 4$ to $p = 5$ simulations by reducing the computational costs by between 15 to 45 times. Finally, we show the runtime in seconds of 100 time steps with each of the considered schemes in Figure 6.6. These runtimes were computed serially on a 3.2 GHz Intel Core i5-5600 processor with 16Gb of RAM. The cost of increasing the polynomial degree by one for a fixed mesh in FR scales by a factor of 3 on average. For the hybridized methods, the computational cost only increases by a factor of 1.5 on average.

Table 6.3. Speedup factors for the cylinder problem. Results compare runtime ratios between the FR method and the corresponding scheme for all considered p -adaptive and p -uniform runs

$p_{\text{HFR}} \backslash p_{\text{FR}}$	1	2	3	4	5	adaptive
1	1.26	5.93	20.42	51.88	129.42	18.84
2	0.54	2.55	8.79	22.33	55.70	8.11
3	0.28	1.34	4.62	11.73	29.27	4.26
4	0.16	0.76	2.63	6.67	16.64	2.42
5	0.11	0.50	1.72	4.37	10.89	1.59
adaptive	0.36	1.68	5.79	14.71	36.69	5.34

$p_{\text{EFR}} \backslash p_{\text{FR}}$	1	2	3	4	5	adaptive
1	1.67	7.89	27.17	69.03	172.20	25.07
2	0.79	3.73	12.84	32.61	81.36	11.85
3	0.42	2.00	6.87	17.46	43.56	6.34
4	0.23	1.10	3.77	9.58	23.91	3.48
5	0.14	0.66	2.28	5.79	14.45	2.10
adaptive	0.43	2.02	6.95	17.66	44.06	6.42




Figure 6.6. Runtime in seconds for the cylinder problem. Results show the time spent to solve 100 iterations for p -uniform simulations (\bullet) and p -adaptive (\blacksquare) with $\epsilon = 0.1$

6.3.3 NACA 0012 Airfoil at $\text{Re} = 10,000$




Figure 6.7. NACA 0012 airfoil mesh consisting of 8658 quadrilateral elements refined at the wall boundaries

Finally, we present flow over a NACA0012 airfoil at $\text{Re} = 10,000$ with an angle of attack $\alpha = 2$ deg and a chord c . This simulation compares unsteady vortex shedding from the airfoil between hybridized and standard FR formulations. We run this problem at a Mach number $M = 0.2$. The computational grid is composed of 8658 quadrilateral elements, with refinement toward the airfoil walls to capture the formation of the boundary layer. Downstream of the airfoil, the element size uses a 1.13 growth rate to maintain relatively good resolution for trailing vortex within two to three convective lengths. The downstream length of the domain is $10c$. A portion of the mesh is shown in Figure 6.7. We employ the second-order two-stage SDIRK method for time stepping with a $\Delta t = 1 \times 10^{-3} t_c$ in this problem, which gives a CFL number of approximately 2.5 based on the freestream conditions and the minimum element


length. This numerical example is run with uniform solution polynomial degrees $p = 1$ to $p = 5$ as well as with p -adaptation using FR, HFR, and EFR. For the adaptation algorithm, we use a tolerance for the vorticity indicator of $\epsilon = 0.1$. A convergence tolerance of 10^{-5} is used for the implicit unsteady density residual.

Vorticity contours are shown in Figure 6.8 up to three chord lengths, where we compare runs between a $p = 5$ -uniform standard FR scheme against the p -adaptive runs. The overall periodic Karman shedding pattern corresponds to that observed in similar runs of up to $\alpha = 4$ degrees [150], which is well captured by the p -adaptive simulations as well. From these contours, we observe that both the hybridized and standard FR p -adaptive runs can capture the unsteady vortex shedding similarly. A plot of the averaged C_p coefficient is shown in Figure 6.10, where results are shown for the p -adaptive runs, using a uniform $p = 5$ standard FR as reference. Note that results are overlapping at most locations, indicating comparable averaged pressure distributions over the airfoil among the considered cases.

Figure 6.9 displays the distribution of polynomial degrees over the domain for each of the considered schemes. Here we show instantaneous snapshots at moments of minimum lift. The algorithm set $p = 5$ elements in proximity to the airfoil leading edge at around $x/c \lesssim 0.15$ and gradually decreases the polynomial degree to $p = 4$ and then to $p = 3$ over the rest of the airfoil. A small recirculation region due to flow separation is observed starting at $x/c \approx 0.65$ on the suction side of the airfoil, where $p = 2$ elements are used since the vorticity magnitude in this region is small compared to the boundary layer. Vortices shedding off the airfoil are tracked by the adaptation algorithm with $p = 5$ elements close to their core, which also gradually decreases to $p = 1$ cells in the far field. The three considered schemes have comparable polynomial distributions, which result in similar numbers of degrees of freedom, as observed in Table 6.4.

Quantitative results are shown in Table 6.4. These include averaged values of lift, drag, and moment coefficient at the quarter-chord, as well as the Strouhal number and the resulting number of degrees of freedom. These quantities were obtained by performing a simple time-averaging operation over a sufficiently long window of time. We compare our results against the numerical experiments of Ikeda [150] and experimental data of Ohtake [151] in c_l . p -adaptive results showed close agreement with the p -uniform data with error below 0.5%. Results are in good agreement with the reference data, with a relative difference of approximately 6% in all numerical results and less than 20% compared to digitized experiments of [151].


Adaptive FR simulations required only about 1/4 the number of internal degrees of freedom of a $p = 5$ simulation. Hybridized methods were able to obtain comparable results for this problem by solving systems of equations that depend on DOF_t . This number of degrees of freedom is 2 to 3.5 times smaller than that used in a standard p -adaptive simulation.



(a) FR p -uniform



(b) FR



(c) HFR



(d) EFR

Figure 6.8. Instant snapshots of the airfoil problem showing contours of vorticity for FR, HFR, and EFR at moments of minimum lift




Figure 6.9. Instant snapshots of the airfoil problem with superimposed contours of vorticity for FR, HFR, and EFR at moments of minimum lift




Figure 6.10. Averaged C_p coefficient for the adaptive schemes on the upper (●) and lower (▲) walls of the airfoil. Reference is a $p = 5$ FR result

This results in significant speedup factors, which we summarize in Table 6.5. Similar to the cylinder problem, the benefit of using hybridization becomes more important as we go to higher orders. For instance, using HFR at $p = 1$ was not beneficial and, in fact, resulted in increased computational cost. For $p > 1$, speedups higher than unity can be observed in the diagonal of these tables, where each polynomial degree for FR is compared with its respective hybridized formulation. Below this diagonal, it can also be seen that using hybridization at higher polynomial degrees with FR can still result in speedups for implicit formulations. We obtained 4 and 5.73 times faster simulations than a typical p -adaptive FR method using p -adaptive HFR and EFR, respectively. Furthermore, speedups in excess of 22 and 32 times were obtained with p -adaptive hybridized methods with respect to a $p = 5$ uniform standard FR formulation. Finally, we show the time spent on the implicit global system (solid) and the solution of the local equations (diagonal hatch). These results were computed serially on 2.40 GHz AMD Rome 7532 processors. As shown in Figure 6.11, hybridized methods take only a fraction of the time required by FR, and increased benefits are observed for higher polynomial degrees. However, they spend some additional time on the local solves in addition to the global system. We note that the local solves scale linearly with the number of processors as they are element-wise operations. Hence, their implementation in high-performance systems is natural. With this example, we have shown that we can reduce the cost of p -adaptive FR by hybridization, including the globally-coupled EFR method.




Figure 6.11. Overview of the runtime spent on the global (solid) and local (diagonal hatch) solves for the airfoil problem in 100 iterations

Table 6.4. Results for the airfoil problem for FR, HFR, and EFR methods with p -uniform and p -adaptive discretizations

	Degree	\bar{c}_l	\bar{c}_d	$\bar{c}_{m,c/4}$	St	$\overline{\text{DOF}}_v$	$\overline{\text{DOF}}_t$
FR	1	0.0402	0.0430	0.0179	2.4187	34632	-
	2	0.0409	0.0431	0.0171	2.4691	77922	-
	3	0.0409	0.0436	0.0170	2.4677	138528	-
	4	0.0409	0.0439	0.0170	2.4691	216450	-
	5	0.0409	0.0438	0.0170	2.4667	311688	-
	adaptive	0.0409	0.0438	0.0170	2.4679	80368	-
HFR	1	0.0411	0.0387	0.0181	2.4084	34632	34226
	2	0.0410	0.0430	0.0172	2.4735	77922	51339
	3	0.0409	0.0431	0.0171	2.4660	138528	68452
	4	0.0409	0.0441	0.0170	2.4691	216450	85565
	5	0.0409	0.0436	0.0170	2.4691	311688	102678
	adaptive	0.0409	0.0435	0.0170	2.4670	79946	47833
EFR	1	0.0444	0.0383	0.0184	2.4390	34632	8859
	2	0.0410	0.0431	0.0172	2.4754	77922	25972
	3	0.0409	0.0431	0.0171	2.4672	138528	43085
	4	0.0409	0.0441	0.0170	2.4685	216450	60198
	5	0.0409	0.0437	0.0170	2.4691	311688	77311
	adaptive	0.0409	0.0435	0.0170	2.4680	80415	22517
Ikeda et al. [150]		0.0435	-	-	2.4556		
Ohtake et al. [151]		0.0500	0.0430	0.0160	-		

Table 6.5. Speedup factors for the airfoil problem. Results compare runtime ratios between the FR method and the corresponding scheme for all considered p -adaptive and p -uniform runs

$p_{\text{HFR}} \backslash p_{\text{FR}}$	1	2	3	4	5	adaptive
1	0.88	4.05	11.74	30.92	71.71	12.68
2	0.33	1.53	4.43	11.66	27.03	4.78
3	0.19	0.86	2.51	6.60	15.32	2.71
4	0.11	0.50	1.45	3.82	8.85	1.56
5	0.07	0.35	1.00	2.64	6.12	1.08
adaptive	0.28	1.28	3.70	9.75	22.60	4.00

$p_{\text{EFR}} \backslash p_{\text{FR}}$	1	2	3	4	5	adaptive
1	1.45	6.67	19.35	50.96	118.19	20.89
2	0.58	2.69	7.82	20.58	47.72	8.44
3	0.30	1.38	4.02	10.58	24.53	4.34
4	0.17	0.79	2.29	6.02	13.96	2.47
5	0.11	0.52	1.50	3.95	9.16	1.62
adaptive	0.40	1.83	5.31	13.98	32.42	5.73

Chapter 7

Implicit-Explicit HFR-FR Formulations

7.1 Overview

A largely disparate range of element sizes typically appears in simulations at high Reynolds numbers, with large elements in the far field and elements in orders of magnitude smaller in proximity to walls. This introduces what is referred to as geometry-induced stiffness since the maximum time-step size for explicit simulations is dictated by the smallest element size. For high-order schemes, using explicit methods can be very restrictive in terms of the allowable time-step size to maintain stability. On the other hand, fully implicit methods can become prohibitively expensive for these types of problems regarding computation time per step and memory requirements. As previously stated, the cost of implicit methods scales with $O(p^d)$ for standard FR and $O(p^{d-1})$ for hybridized FR, where p is the polynomial degree, and d is the dimension of the problem. Hence, the feasibility of employing implicit time-stepping for large-scale computations is limited at high orders.

A more efficient approach involves both explicit and implicit IMEX time-stepping methods. These schemes are able to leverage the stability of implicit schemes for stiff terms while mitigating their cost by using explicit methods for nonstiff terms. To demonstrate their application, consider an ordinary differential equation of the form

$$u' = f(u) + g(u), \tag{7.1}$$

where $f(u)$ is the nonstiff part of the problem and $g(u)$ is the stiff portion. Consider an implicit s -stage diagonally-implicit Runge-Kutta (DIRK) scheme for the stiff region associated with a matrix and vector of coefficients of a Butcher tableau [152] given by $\mathbf{A} \in \mathbb{R}^{s \times s}$, $\mathbf{b} \in \mathbb{R}^s$, $\mathbf{c} \in \mathbb{R}^s$. For the nonstiff part, consider an explicit $\sigma = s + 1$ -stage RK method with respective coefficients given by $\bar{\mathbf{A}} \in \mathbb{R}^{\sigma \times \sigma}$, $\bar{\mathbf{b}} \in \mathbb{R}^\sigma$, $\bar{\mathbf{c}} \in \mathbb{R}^\sigma$. To compensate for the difference in the size of the matrices of coefficients, a first row and a first column of zeros

are padded into the implicit tableau. In order to be paired, the implicit and explicit schemes must satisfy $\bar{\mathbf{c}} = [0 \ \mathbf{c}]^T$. The resulting form of a general IMEX Butcher tableau can be seen in Table 7.1.

0	0	0	\dots	0	\bar{c}_1	0	0	0	\dots	0	
c_1	0	a_{11}	0	\dots	0	\bar{c}_2	\bar{a}_{21}	0	0	\dots	0
c_2	0	a_{21}	a_{22}	\dots	0	\bar{c}_3	\bar{a}_{31}	\bar{a}_{32}	0	\dots	0
\vdots	\vdots	\vdots	\vdots	\ddots	\vdots	\vdots	\vdots	\vdots	\vdots	\ddots	\vdots
c_s	0	a_{s1}	a_{s2}	\dots	a_{ss}	\bar{c}_σ	$\bar{a}_{\sigma 1}$	$\bar{a}_{\sigma 2}$	$\bar{a}_{\sigma 3}$	\dots	0
	0	b_1	b_2	\dots	b_s		\bar{b}_1	\bar{b}_2	\bar{b}_3	\dots	\bar{b}_σ

(a) Implicit part
(b) Explicit part

Table 7.1. General form of Butcher tableaus for IMEX schemes

To advance the solution from time level n to $n + 1$ by a time-step Δt , the first stage is always explicit. Then implicit and explicit solves are alternated, as shown in Algorithm 4 [153, 44]. Using IMEX schemes can yield significantly smaller implicit systems to solve, as it is dedicated to a portion of the problem when considering geometry-induced stiffness. In addition, these schemes are linearly stable, maintain the expected orders of accuracy [44], and can have superior performance compared to purely explicit and purely implicit methods for LES simulations [153]. Furthermore, by employing a procedure similar to our work in Chapter 4, optimized IMEX methods [48] can be obtained, resulting in additional speedups. The IMEX approach can be further leveraged by introducing hybridization to solve the implicit portion. Previous works on IMEX methods with hybridization have been developed to tackle stiffness associated with shallow water systems [154]. This was done to separate the faster gravity wave from that of the nonlinear advection operator. However, time-splitting applications typically struggle to compete with purely explicit methods. Hence, this section develops an efficient IMEX formulation for geometry-induced stiffness.

Algorithm 4: Time integration using an IMEX scheme for one time-step

Set $\bar{R}_1 = f(u_n)$.
for $i \leftarrow 1$ **to** s **do**
 Solve for R_i in $R_i = g(u_i)$, where

$$u_i = u_n + \Delta t \sum_{j=1}^i a_{i,j} R_j + \Delta t \sum_{j=1}^i \bar{a}_{i+1,j} \bar{R}_j.$$

 Evaluate
 $\bar{R}_{i+1} = f(u_i)$.
end
Compute the value at the next time step
$$u_{n+1} = u_n + \Delta t \sum_{j=1}^s b_j R_j + \Delta t \sum_{j=1}^{\sigma} \bar{b}_j \bar{R}_j.$$

7.2 Formulation

This section proposes an IMEX formulation by pairing conventional FR and HFR methods to tackle geometry-induced stiffness. While hybridized methods have also been developed in explicit formulations, they require explicit trace definitions [155, 156] and nonlinear solvers [157] and hence the benefits over a standard FR formulation in a general nonlinear problem are still not clear. The explicit form of the FR method is suitable for nonstiff problems. FR methods are locally conservative and have demonstrated potential for modern parallel computer architectures. Hence, we employ it to solve the moderate to large elements associated with lower stiffness in the domain. We introduce hybridization, which is expected to reduce the size of the implicit solver by employing HFR or EFR formulations for the smallest elements associated with the stiff portions of the domain. We refer to the proposed approach as hybridized IMEX methods. These schemes are expected to reduce the cost of a purely implicit method while increasing the allowable time-step size and improving the constrained stability posed by the explicit formulation.

We are interested in integrating the equation

$$\frac{d\mathbf{u}}{dt} = \mathbf{R}(\mathbf{u}(t)), \quad (7.2)$$

subject to an appropriate initial condition, where \mathbf{R} typically contains the divergence of the flux after applying a spatial discretization such as the FR method. In order to integrate this equation for geometry-induced stiffness, an s -stage IMEX method with order q to advance a solution from time level n to $n + 1$ will be employed according to a modified version of Algorithm 4.

To this end, consider a computational domain Ω subdivided into two regions, as shown

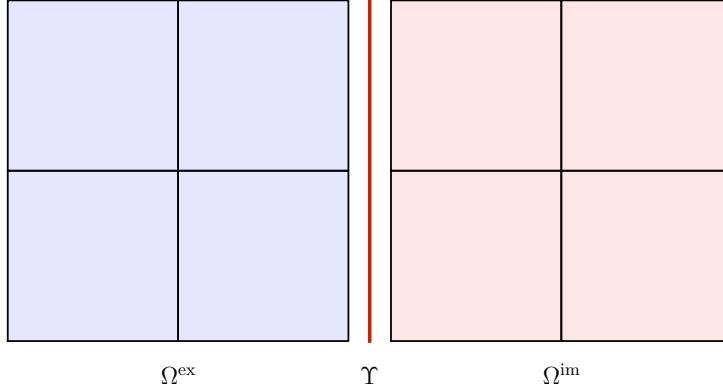


Figure 7.1. Reference domain partitioning for IMEX schemes

in Figure 7.1. The implicit region is denoted Ω^{im} and the explicit region Ω^{ex} . The interface between these two regions is denoted $\Upsilon = \Omega^{\text{ex}} \cap \Omega^{\text{im}}$. The solution and trace vectors can be defined


$$\mathbf{u} = \begin{bmatrix} \mathbf{u}^{\text{ex}} \\ \mathbf{u}^{\text{im}} \end{bmatrix}, \quad \hat{\mathbf{u}} = \hat{\mathbf{u}}^{\text{im}}, \quad (7.3)$$

respectively, where $\mathbf{u}^{\text{im}} \in \mathbb{R}^{N^{\text{im}}N_s}$, $\mathbf{u}^{\text{ex}} \in \mathbb{R}^{N^{\text{ex}}N_s}$, and $\hat{\mathbf{u}}^{\text{im}} \in \mathbb{R}^{\hat{N}N_s}$. N^{im} , N^{ex} , \hat{N} are the number of implicit elements, explicit elements, and trace points. The interior solution is found in both explicit and implicit subdomains, but the trace is only defined at $\varepsilon^{h,\text{im}} \setminus \Upsilon$ since $\hat{\mathbf{u}}^{\text{ex}} = \emptyset$. Hence, we will refer to the trace in the implicit side as $\hat{\mathbf{u}}$. After hybridization of the implicit portion and applying the above definitions, the ODE in Equation (7.2) can be generalized to a system of the form


$$\frac{d\mathbf{u}}{dt} = \begin{bmatrix} \mathbf{R}^{\text{ex}}(\mathbf{u}) \\ \mathbf{R}^{\text{im}}(\mathbf{u}, \hat{\mathbf{u}}) \end{bmatrix} \quad \text{in } \mathcal{T}^h, \quad (7.4)$$

$$\mathbf{G}(\mathbf{u}, \hat{\mathbf{u}}) = \mathbf{0} \quad \text{in } \varepsilon^{h,\text{im}}, \quad (7.5)$$

where $\varepsilon^{h,\text{im}} = \varepsilon^h \cap \Omega^{\text{im}}$. The resulting distribution of trace and solution points, $\hat{\mathbf{x}}$ and \mathbf{x}^s , respectively, can be seen in Figure 7.2 for HFR and EFR methods. Furthermore, $\mathbf{R}^{\text{ex}}(\mathbf{u})$ is associated with the standard FR equations in (3.35), $\mathbf{R}^{\text{im}}(\mathbf{u}, \hat{\mathbf{u}})$ with the hybridized Equations in (3.42a)-(3.42b) and $\mathbf{G}(\mathbf{u}, \hat{\mathbf{u}})$ with the transmission conditions in (3.42c). Due to the temporally non-overlapping nature of IMEX schemes, the explicit equations can be solved separately from the implicit portion. Applying the IMEX-RK method, the solution at



(a) HFR



(b) EFR

Figure 7.2. Distribution of trace and solution points in this configuration for HFR and EFR IMEX methods for a $p = 2$ discretization

each stage can be found by first computing the value of the explicit solution via

$$\mathbf{u}_i^{\text{ex}} = \mathbf{u}_n^{\text{ex}} + \Delta t \sum_{j=1}^i \bar{a}_{i+1,j} \mathbf{R}_j^{\text{ex}}, \quad (7.6)$$

where $\mathbf{R}_j^{\text{ex}} = \mathbf{R}(\mathbf{u}_{j-1})$ for $j > 1$ and $\mathbf{R}_j^{\text{ex}} = \mathbf{R}(\mathbf{u}_n)$ for $j = 1$. The explicit residual only depends on known values of the solution at stages $i < s$ since $\bar{a}_{i,j} = 0$ for $j \geq i$. Here, the residual is a function of the solution in the explicit subdomain and in the elements in direct contact with the IMEX interface of the implicit portion, which is always known for the required indices. After computing (7.6), the implicit solution can be readily obtained via

$$\mathbf{u}_i^{\text{im}} = \mathbf{u}_n^{\text{im}} + \Delta t \sum_{j=1}^i a_{i,j} \mathbf{R}_j^{\text{im}}(\mathbf{u}_j, \hat{\mathbf{u}}_j), \quad (7.7a)$$

$$\mathbf{G}(\mathbf{u}_i, \hat{\mathbf{u}}_i) = 0, \quad (7.7b)$$

which employs hybridization and Equation (7.7b) represents discrete transmission conditions with the form

$$\sum_{\bar{f} \in \varepsilon_0^{h,\text{im}} \setminus \Upsilon} \int_{\bar{f}} \llbracket \hat{\mathfrak{F}}(u_i, \hat{u}_i) \rrbracket_{\bar{f}} \phi ds + \sum_{\bar{f} \in \Upsilon} \int_{\bar{f}} \llbracket \hat{\mathfrak{F}}^{\text{FR}}(u_i^{\text{im}}, u_i^{\text{ex}}) \rrbracket_{\bar{f}} \phi ds + \sum_{\bar{f} \in \varepsilon_\theta^{h,\text{im}}} \int_{\bar{f}} \mathfrak{F}_{\bar{f}}^{\text{BC}} \phi ds = 0, \quad (7.8)$$

where the typical transmission equations of Equation (3.42c) have been augmented with an interface condition to weakly enforce conservation along the IMEX interface and hence globally in the domain. Then, the solution at the next step can be found by

$$\mathbf{u}_{n+1} = \mathbf{u}_n + \Delta t \sum_{j=1}^s b_j \mathbf{R}_j^{\text{im}} + \Delta t \sum_{j=1}^{\sigma} \bar{b}_j \mathbf{R}_j^{\text{ex}}. \quad (7.9)$$

At the IMEX interface, we apply the standard FR fluxes and introduce them as boundary conditions for the hybridized portion. Consequently, at the interface, there is no trace definition. The local conservation property of the FR method [26] enables the use of this approach to both HFR and EFR methods. Recall that our EFR implementation uses discontinuous traces at the boundaries throughout this work. In the proposed hybridized formulation, discontinuous traces are also used at the IMEX interface for both methods. This allows for using paired HFR-FR and EFR-FR methods to tackle geometry-induced stiffness. The nonstiff portion will retain its local conservation properties, and the stiff portion will be globally conservative for the EFR method and locally conservative for the HFR method. Specifically, for the HFR methods with discontinuous traces, the transmission conditions reduce to a pointwise conservation statement, which can be shown to yield the same definition

of the traces in linear advection throughout the domain and recovers a standard FR IMEX method in these cases only. The proof is trivial and is omitted for brevity.

7.3 Numerical Examples

This section presents a series of numerical examples to showcase the benefits of using our proposed hybridized IMEX formulations to tackle geometrical stiffness in two and three dimensions. We will first present verification of our approach via linear advection and then demonstrate capabilities via nonlinear numerical examples. To evaluate geometry-induced stiffness, the following elementwise stiffness indicator is used

$$E_s = \frac{|\Omega_k|}{|\partial\Omega_k|}, \quad (7.10)$$

where $|\Omega_k|$ is the volume of the element and $|\partial\Omega_k|$ is the sum of the face areas. Hence, the indicator has dimensions of length. For elements with high aspect ratios as well as for very small elements, the sum of the face areas is significantly larger than its volume, which means that elements with high numerical stiffness will have a high value of E_s . The use of a user-provided cutoff E_s determines the implicit ratio, which we define to be the number of implicit elements N^{im} in relation to the total number of elements N in the domain

$$\text{IF} = \frac{N^{\text{im}}}{N}, \quad (7.11)$$

so that a higher cutoff value of E_s yields more implicit elements. For simplicity, this chapter will refer to FR, HFR, and EFR methods as IMEX discretizations with implicit portions solved using the FR, HFR, and EFR methods, respectively. For all runs, the explicit portion uses a standard FR discretization. In the following problems, IMEX_{*s,q*} methods with *s* stages and *q* order are considered.

7.3.1 Verification

We perform verification of linear advection by considering a periodic square domain. The domain is split into quadrilateral elements. Along the center, a band of seven layers of stretched elements is placed with a stretching ratio of 2, as shown in Figure 7.3. Within this band, elements are flagged as implicit, whereas uniform elements away from this section remain explicit. The implicit portion of the domain is solved using a hybridized form, and the explicit portion uses a standard FR method. Both implicit and explicit regions use a solution polynomial degree of seven to reduce contamination arising from spatial error. The

initial condition is a Gaussian profile

$$u(\mathbf{x}, t) = \exp\left(-\frac{1}{20}[(x - x_c)^2 + (y - y_c)^2]\right), \quad (7.12)$$

where $(x_c, y_c) = (10, 10)$ is the center coordinate of the domain of size 20×20 . After one convective time $t_c = 20$, the L_2 -norm of the error is computed. Results of the L_2 norm of the error against the exact solution are shown in Table 7.2 for the IMEX_{3,2} [40], IMEX_{5,3} [43] and AIMEX_{10,2} method [48] with Butcher tableaus included in the appendix for the IMEX methods and are available for AIMEX as supplementary material in [48]. Due to the temporal error dominating the L_2 norm, results for the EFR and HFR methods differ only beyond single precision, and no difference can be observed in the tabulated values. In addition, results for the approach with HFR are equivalent to those of the FR single scheme, as expected for linear advection. Hence, only the HFR approach is shown. The second and third orders of temporal accuracy are recovered for the three considered IMEX schemes. Of the three considered IMEX methods, AIMEX_{10,2} provides relatively lower error than IMEX_{3,2} for a given time-step size Δt , allowing for a larger time step due to the optimized explicit stability polynomial with only two implicit solves. While the IMEX_{5,3} method is more accurate than the other two, it requires five implicit solves per time step. Thus, we use the optimized AIMEX_{10,2} method for the rest of this chapter.




Figure 7.3. Distribution of explicit and implicit elements for the IMEX verification of linear advection

Table 7.2. Convergence table showing the L_2 -norm of the solution error and the order of accuracy for linear advection using multiple IMEX schemes

Scheme	Δt	HFR	Order	EFR	Order
IMEX _{3,2}	0.02	6.78×10^{-6}	-	6.78×10^{-6}	-
	0.01	1.67×10^{-6}	2.02	1.67×10^{-6}	2.02
	0.005	4.17×10^{-7}	2.00	4.17×10^{-7}	2.00
	0.0025	1.04×10^{-7}	2.00	1.04×10^{-7}	2.00
IMEX _{5,3}	0.02	2.19×10^{-6}	-	2.19×10^{-6}	-
	0.01	2.73×10^{-7}	3.00	2.73×10^{-7}	3.00
	0.005	3.41×10^{-8}	3.00	3.41×10^{-8}	3.00
	0.0025	4.28×10^{-9}	3.00	4.28×10^{-9}	3.00
AIMEX _{10,2}	0.02	4.28×10^{-6}	-	1.83×10^{-6}	-
	0.01	4.60×10^{-7}	1.99	4.60×10^{-7}	1.99
	0.005	1.15×10^{-7}	2.00	1.15×10^{-7}	2.00
	0.0025	2.89×10^{-8}	2.00	2.89×10^{-8}	2.00

7.3.2 Laminar Flow over a Circular Cylinder

In this section, we study flow over a cylinder at $\text{Re} = 150$, based on the cylinder diameter D , to simulate unsteady vortex shedding. This problem belongs to the laminar regime. Hence, a two-dimensional approach is suitable for this simulation. Here, we want to validate our implementation of hybridized IMEX discretizations. We make use of a computational domain divided into 3090 quadrilateral elements refined toward the cylinder walls to capture the gradients due to the boundary layer, as shown in Figure 7.4. We use this case as a baseline problem to analyze the performance of IMEX HFR-FR schemes. We choose to run this problem at Mach number $M = 0.1$, similar to our problem in Chapter 6, to compare against the reference data of Cagnone [74]. In this case, the resolution within the boundary layer is increased to test the benefit of these methods in stiff regions. In Figure 7.5, a histogram containing the distribution of element sizes is provided. We observe disparate element sizes with differences of up to three orders of magnitude between the largest and the smallest elements.

By increasing the cutoff element size, the maximum allowable time-step size is expected to increase. The maximum stable Δt are computed for a range of implicit factors using a simple bisection algorithm and shown in Figure 7.6. We capped the maximum time-step size to $t_c/200$ to maintain accuracy. This value is typically achieved at $\text{IF} > 0.5$, a relatively large implicit factor that requires significant computational resources in industrial-scale problems, as we will discuss in the last numerical example of this chapter. Our main interest is focused




Figure 7.4. Mesh for the laminar cylinder case at $Re = 150$ consisting of 3090 quadrilateral elements




Figure 7.5. Distribution of element sizes for the laminar cylinder case

on regions with moderately low implicit factors. Note from this plot that the value of Δt_{\max} is similar for the FR, HFR, and EFR methods, with a few differences due to the shape of the stability polynomial in AIMEX methods optimized for the FR method only. We compute speedups against FR explicit formulations with $\text{IF} = 0$. Speedups near two orders of magnitudes for $p = 4$ simulations are observed at higher IF values in Figure 7.7. These speedups were computed using serial simulations. We observed hybridized methods to yield up to 2.5 times faster results than an FR method at $\text{IF} = 0.2$. This translates to 10 times faster results compared to an explicit approach while utilizing significantly less memory than a fully implicit method. Furthermore, an implicit factor $\text{IF} = 0.2$ is chosen, which is expected to provide optimal speedup factors without using significant memory.

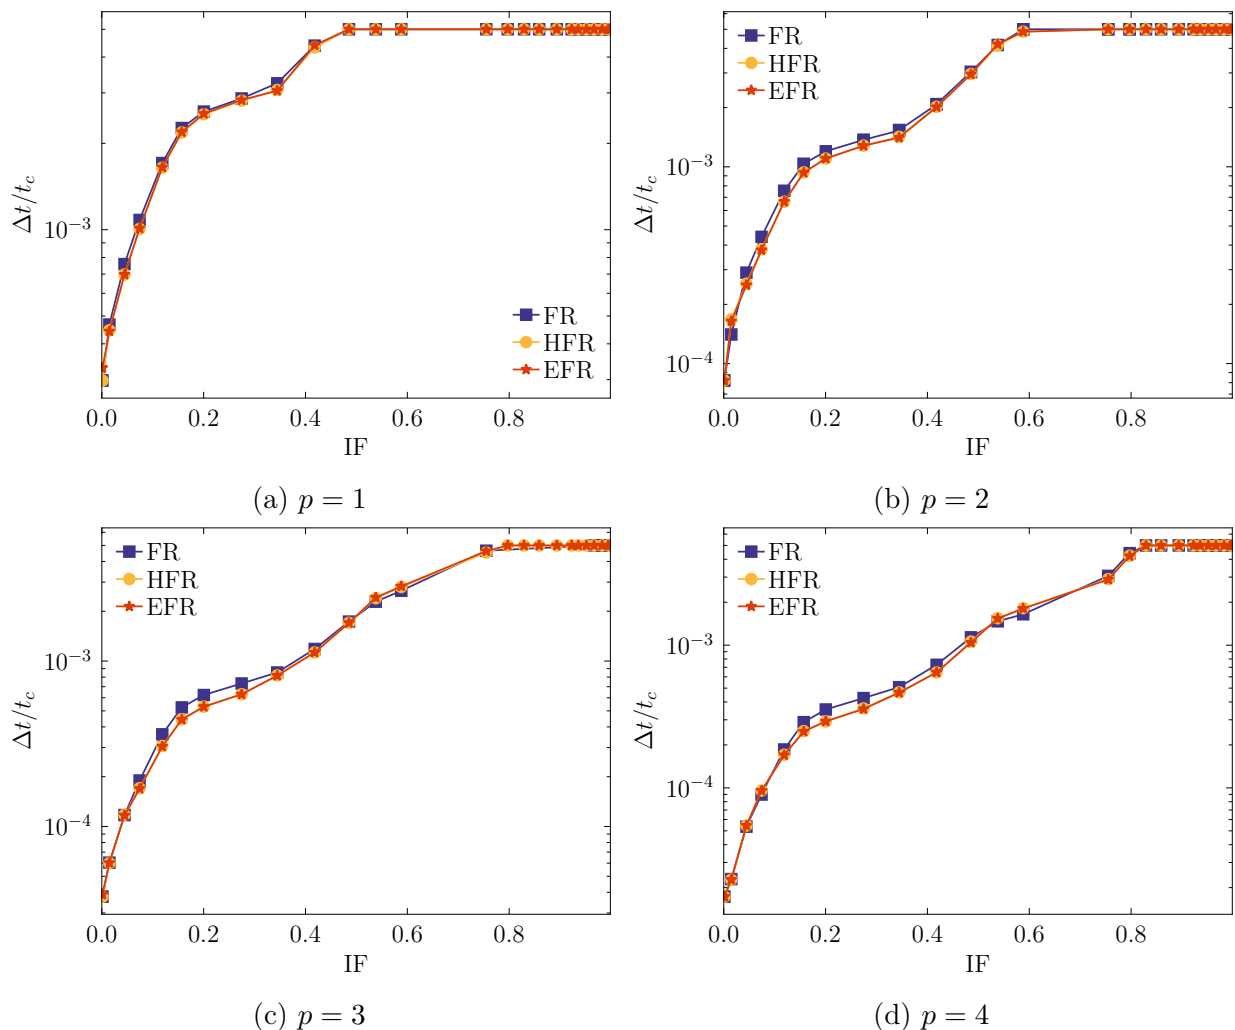


Figure 7.6. Maximum stable time-step size for multiple implicit fractions for the AIMEX_{10,2} scheme

We perform a series of simulations with polynomial degrees $p = 1$ to $p = 4$ for 200 convective times $t_c = U_\infty/D$. The evolution of the lift and drag coefficients for the $p = 1$ and

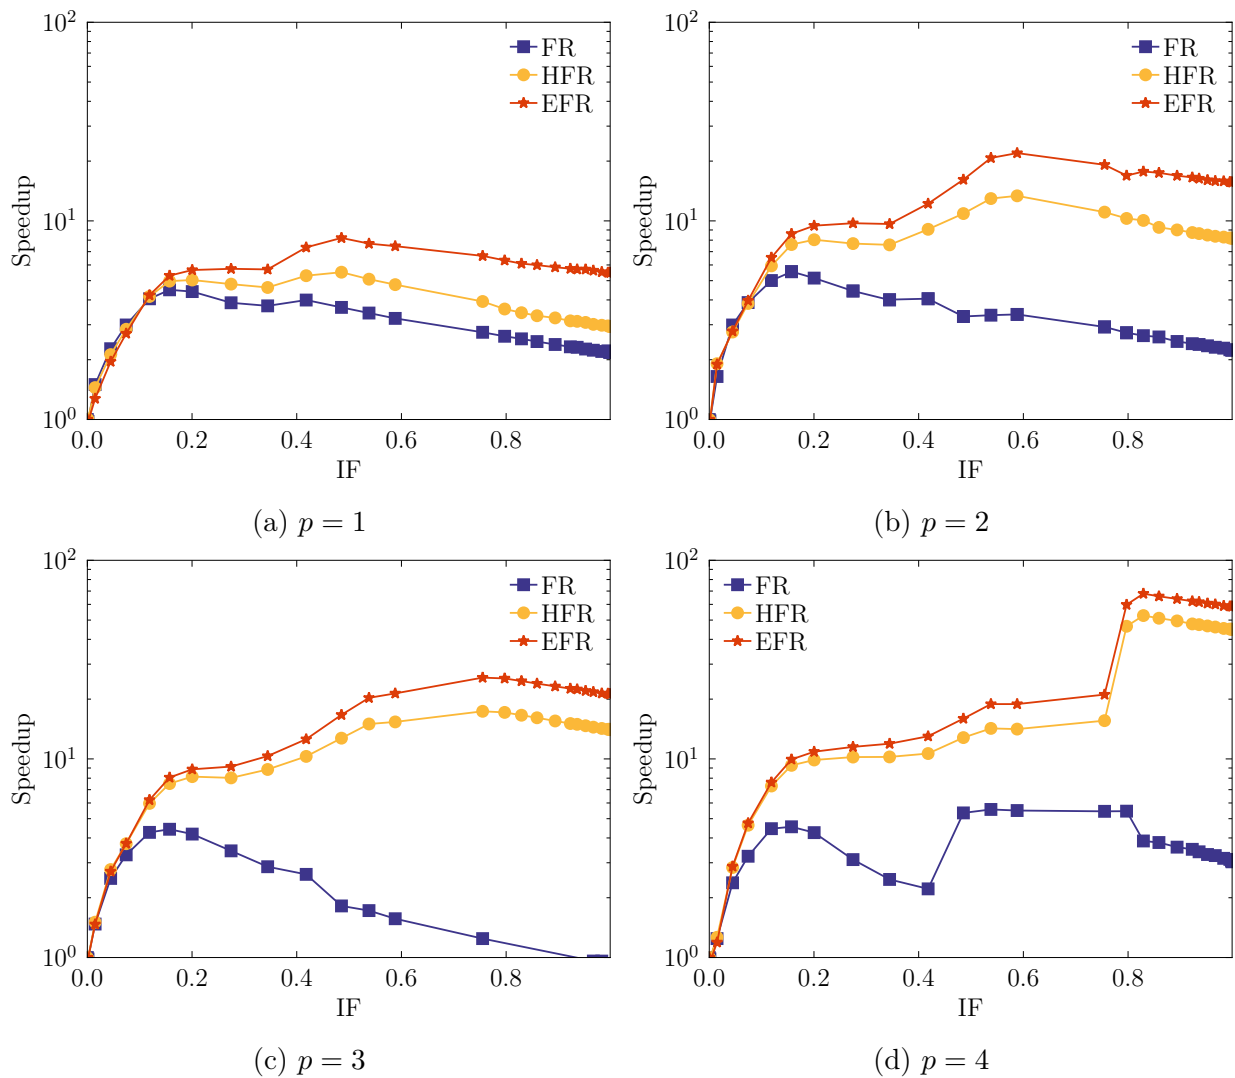


Figure 7.7. Speedup factors for multiple implicit fractions for the AIMEX_{10,2} scheme

$p = 4$ HFR and EFR schemes is shown in Figure 7.8. We observe a periodic sinusoidal wave associated with the characteristic vortex shedding of this problem. By looking at this plot, it is clear that the frequency of the low-order simulation is different from the $p = 4$ results. However, no significant difference is observed between the hybridized and standard IMEX results. More specifically, the Strouhal number for this problem converges to 0.1844 after $p = 3$. $p = 1$ is heavily underresolved and underpredicts this result, as shown in Table 7.3. The converged results represent less than 0.5% relative error compared to the numerical results of Cagnone [74] and less than 1% compared to the experimental data. Results from the hybridized methods agree with the FR results, with differences of less than 0.2% in all quantities. Hence, the proposed hybridized IUMEX formulations behave similarly to FR, especially at higher orders. The performance benefit of this approach is significant against explicit FR methods. We now evaluate them in a three-dimensional cylinder case in the next section.

Table 7.3. Summary of results for the cylinder at $\text{Re} = 150$ for dual and single scheme IMEX methods

Implicit Scheme	p	\bar{c}_d	Δc_d	Δc_l	St
FR	1	1.2962	0.0205	0.4879	0.1785
	2	1.3271	0.0257	0.5178	0.1842
	3	1.3289	0.0258	0.5189	0.1844
	4	1.3295	0.0258	0.5192	0.1844
HFR	1	1.3020	0.0206	0.4901	0.1787
	2	1.3277	0.0258	0.5187	0.1842
	3	1.3267	0.0256	0.5161	0.1844
	4	1.3312	0.0259	0.5211	0.1844
EFR	1	1.4207	0.0212	0.5099	0.1787
	2	1.3280	0.0258	0.5188	0.1842
	3	1.3285	0.0258	0.5178	0.1844
	4	1.3309	0.0259	0.5207	0.1844
Cagnone [158], $p = 4$		1.3246	0.0258	0.5166	0.1836
Inoue [146]		1.3200	0.0260	0.5200	0.1830

7.3.3 Turbulent Flow over a Circular Cylinder

Three-dimensional flow over a cylinder at $\text{Re} = 1000$ is simulated in this section. The computational grid is composed of 37080 hexahedral elements. Along the spanwise direction, a length of $L_z = 2\pi$ is used, a few units over the minimum length required to resolve the

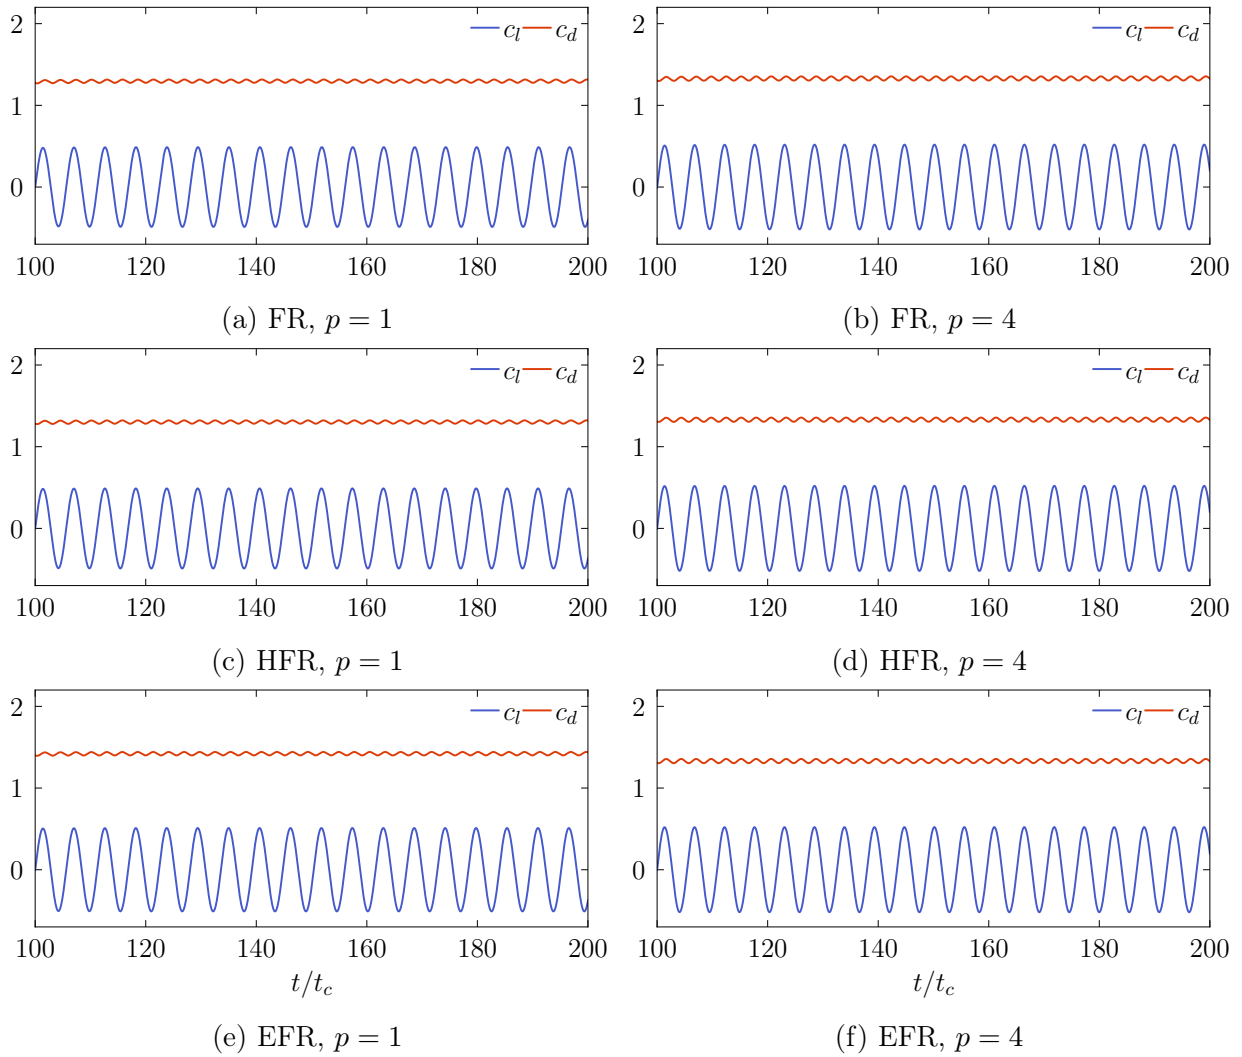


Figure 7.8. Evolution of drag and lift coefficients for the cylinder at $Re = 150$

three-dimensionality, which is $4D$ [159]. This length is divided using a grid spacing of $\Delta z = \pi/6D$ in the streamwise direction, which results in 40 layers of elements. Beyond the cylinder, the boundaries were placed at a distance of $40D$ downstream to minimize the effects of boundary conditions. In the previous section, we observed that an implicit factor close to 0.2 performs well in the 2D case and will also be used in this problem. In reality, large-scale problems demand an overwhelming amount of memory, so high implicit factors or fully implicit methods require availability of a vast amount of resources at high orders. Similarly, the AIMEX_{10,2} method is employed here for time integration with $\Delta t/tc = 2.6 \times 10^{-2}$, which is decreased by half per unit increase in polynomial degree. We distribute implicit and explicit elements as shown in Figure 7.9, and apply hybridized methods on the implicit portions of the domain. We converge our implicit residuals to a tolerance of 10^{-6} . The smaller elements in the vicinity of the cylinder walls are flagged as implicit, and larger elements away from the walls are flagged as explicit.




Figure 7.9. Distribution of implicit (red) and explicit (blue) elements in the computational domain for the turbulent cylinder case

We ran this simulation for 200 convective times and averaged the statistics for the last $100t_c$ to remove the initial transient effects. Results are shown in Table 7.4, where the drag coefficient C_D , lift coefficient fluctuations $C_{L,rms}$, and the Strouhal number St are shown. We compare our results against the data of [160], which provided reference values for a coarse and a fine problem. Relative convergence of the C_D can be observed as the order is increased, which is within 1% of the coarse results and 15% of the DNS data. The mean lift fluctuations are closer for the HFR method than for the EFR method, which is known to introduce additional error. Overall, results converge to the reference data. The Strouhal number was captured well for the EFR method at $p = 4$. In the case of $p = 1$, the EFR method did not transition, which caused a significant discrepancy with the reference St as opposed to the

other values. The spectra of the C_L signal are shown in 7.10, where the convergence to 0.21 is seen as the order is increased.

Table 7.4. Summary of results for the turbulent cylinder case

Scheme	p	\bar{C}_D	$C_{L,rms}$	St
HFR	1	2.152	1.046	0.210
	2	1.156	0.640	0.204
	3	1.009	0.361	0.207
	4	0.998	0.333	0.209
EFR	1	2.907	0.052	0.225
	2	1.151	0.656	0.205
	3	1.012	0.379	0.209
	4	1.000	0.340	0.210
Zhao et al. (coarse) [160]		1.092	0.310	0.210
Zhao et al. (fine) [160]		1.170	0.335	0.210




Figure 7.10. Spectra of the unsteady lift coefficient fluctuations for the turbulent cylinder problem

Performance results for this problem demonstrate that hybridized methods are also suitable for three-dimensional problems. Results are shown for the time spent on the global solves t_G , local solves t_L , block Jacobian computations and implicit matrix assembly t_J , and right-hand-side residual computations in the implicit and explicit portions (t_{Rim} and t_{Rex} , respectively), which together add to the overall wall-clock time t_w . We show speedups against standard FR IMEX schemes t_w/t_w^{FR} and against explicit runs at IF = 0, t_w/t_w^{ex} with the AIMEX tableaus. Results are tabulated in Table 7.5. The timing results are computed for

100 time steps. The time to assemble the Jacobian matrix is associated with a single call, as it was updated every 100 time steps to reduce overhead. These results are computed on 2.4GHz AMD Rome 7532 CPUs using 64 cores. The time spent on the global solutions takes a significantly higher proportion of the total FR computations due to the large size of these systems. For hybridized methods, the solution of the local problems adds overhead, which has linear scalability and can be done efficiently. We observed speedup factors around 2.5 against IMEX and explicit FR using the EFR method, which consistently provided faster runs than FR for $p > 1$. The HFR method, however, is only a benefit at $p = 4$. This is consistent with the increased number of trace unknowns appearing in hexahedral elements. Contours of Q-criterion are shown for simulations using $p = 2$ and $p = 4$ in Figure 7.11, where


Table 7.5. Summary of performance metrics for the turbulent cylinder case for 100 time steps

Scheme	p	t_G	t_L	t_J	$t_{R^{im}}$	$t_{R^{ex}}$	t_w	t_w/t_w^{FR}	t_w/t_w^{ex}
FR	1	187.18	-	3.10	3.72	26.39	217.33	1.00	0.76
	2	232.34	-	16.04	4.46	51.55	288.44	1.00	1.43
	3	599.79	-	85.16	8.65	116.29	724.92	1.00	1.50
	4	1522.91	-	363.96	16.38	248.18	1787.85	1.00	1.15
HFR	1	200.99	0.56	2.71	3.14	25.45	230.15	0.94	0.72
	2	208.59	3.32	12.50	3.74	46.32	261.97	1.10	1.58
	3	433.31	16.89	54.76	8.17	111.03	569.40	1.27	1.91
	4	922.10	60.34	195.94	16.71	248.13	1247.28	1.43	1.65
EFR	1	201.28	0.59	2.61	3.19	25.86	230.93	0.94	0.72
	2	153.48	2.79	12.32	4.50	53.55	214.32	1.35	1.93
	3	225.92	16.51	52.08	8.67	114.95	366.05	1.98	2.97
	4	452.98	56.68	189.40	16.52	247.93	774.11	2.31	2.66


the behaviour of these vortical structures can be observed. They result from instabilities caused by the complex shedding phenomena associated with this value of Re. At higher polynomial degrees, finer turbulent structures are observed, which is expected due to the increased resolution of the $p = 4$ method against $p = 2$. Overall, contour results from both methods are in good agreement with each other and with the reference [160]. Results from this problem demonstrate the suitability of IMEX methods for problems of moderate stiffness.

7.3.4 Turbulent Flow over a Multi-Element Airfoil


Finally, this section presents a wall-resolved large-eddy simulation of a multi-element 30P30N airfoil at $Re = 1.7 \times 10^6$ and an angle of attack 5.5 degrees. The reference chord is denoted c and represents the stowed airfoil. This problem is commonly used in the




(a) HFR, $p = 2$



(b) EFR, $p = 2$



(c) HFR, $p = 4$



(d) EFR, $p = 4$

Figure 7.11. Side view of Q-criterion contours for the turbulent cylinder problem

aeroacoustics community, particularly within the AIAA Benchmark Problems in Airframe Noise Computations workshops [161]. This airfoil makes use of a slat at the leading edge to increase maximum lift, allowing the main wing to operate at higher angles of attack before stalling. These components introduce complex flow behaviour and significantly contribute to the acoustic field during the landing phase of an aircraft. Multiple researchers have experimentally [162] and numerically [163, 164, 165] produced reference data, and an overview of results from the aforementioned workshop is available in [161]. This problem has been previously demonstrated in the context of wall-modeled LES and wall-resolved LES [165, 166], the latter with generally dense refinement within the slat cove region only and coarse everywhere else. The relatively high Reynolds number makes this problem computationally challenging. We generate a computational grid of 549280 hexahedral elements, shown in Figure 7.12, with zoomed-in views for the slat and the flap. The spanwise length is $c/9$, with 40 elements uniformly refined, which is slightly above the coarse grid with 30 layers used in [165].

The entropically-damped artificial compressibility method is used in this problem, which resolved stability issues encountered with the compressible Navier-Stokes equations. This is an appropriate choice since the baseline Mach number for this problem is 0.1, which is within the incompressible range. The EDAC equations use an incompressibility factor, which we set to $\Theta = 100$. This value was chosen to maintain a sensible time-step size and reduce the effects of artificial compressibility [95], as discussed in Chapter 2. This problem typically uses a Rusanov-type Riemann solver, with stabilization computed from Davis estimates of the maximum eigenvalues of the EDAC equations [95]. For the hybridized EDAC equations, we employ the following convective stabilization parameter

$$s = \frac{3}{2}|\hat{\boldsymbol{v}}| + \hat{d}, \quad \hat{d}^2 = \frac{\hat{\boldsymbol{v}}^2}{4} + \hat{P} + \Theta, \quad (7.13)$$

which leads to an isotropic stabilization operator $s\mathbf{I}$. This stabilization leads to a new hybridized formulation of the EDAC equations, first used in this work. Simulations are run for 20 convective times, and the statistics are averaged for the last $10t_c$. The time-step sizes considered here are $\Delta t/t_c = 8 \times 10^{-5}$ for $p = 1$ and $\Delta t/t_c = 3.125 \times 10^{-5}$ for $p = 2$. For $10t_c$, our computations made use of 0.0091 core years for $p = 1$, and 0.097 core years for $p = 2$.

Following the promising results of the EFR method in the turbulent cylinder problem, we employ this method in the implicit portion of our IMEX domain but will also include performance results against the HFR and FR methods. The implicit factor is chosen to solve $\sim 20\%$ of elements with the EFR method at $\text{IF} = 0.23$, consistent with our previous analyses in the cylinder problems. A histogram with the element size in the x -axis is provided in Figure 7.13, showing element sizes with over four orders of magnitude in difference. The

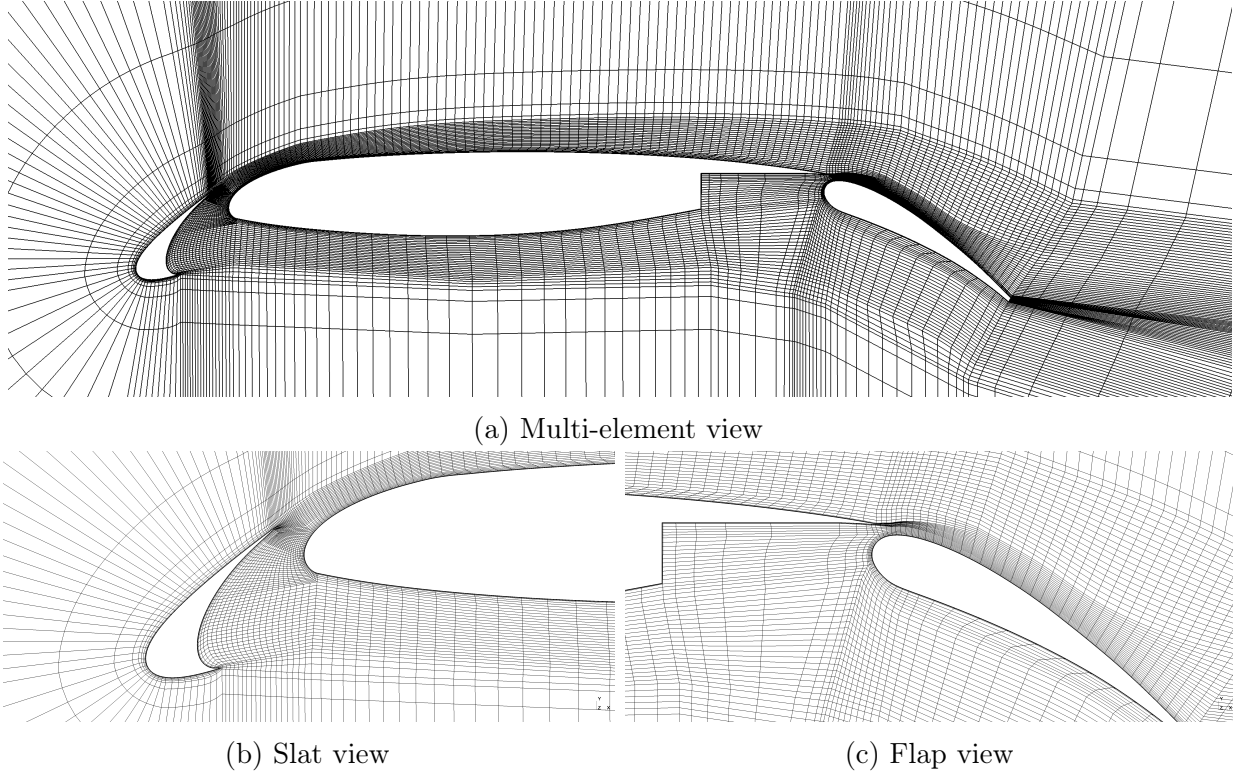


Figure 7.12. Computational grid for the multi-element airfoil

final distribution of implicit (red) and explicit (blue) elements in the domain partition is shown in Figure 7.14. Elements near the walls will be resolved implicitly, and medium to large elements will be solved explicitly. A tolerance of 10^{-4} is used to converge the unsteady implicit residuals.

Contours of averaged vorticity are shown in Figure 7.15 for $p = 1$ and $p = 2$ simulations. Results from the $p = 1$ simulations are highly dissipative and display smeared-out regions of vorticity. Simulations at $p = 2$ are already in good agreement with the PIV visualizations from Pascioni et al. [167], where the expected detached shear layer emerging from the slat is observed with increased definition. In the instantaneous plots, similar behaviour is observed between the $p = 1$ and $p = 2$ simulations, resulting in the latter being a more accurate representation of the complex vortex interaction within the slat cove.

A similar conclusion can be drawn from the contours of Q-criterion, which are shown for a full view in Figure 7.16. The second-order ($p = 1$) simulations resolve a smaller range of scales compared to third order ($p = 2$). At the shown value of Q-criterion, many more structures are shown in the higher-order simulation emanating from the gap between the cove and airfoil all the way to the upper side of the flap downstream. A zoomed-in portion within the slat cove is shown in Figure 7.17, clearly showing the shear layer and reattachment region in the slat. For $p = 1$, the detail in the turbulent structures is small, whereas much


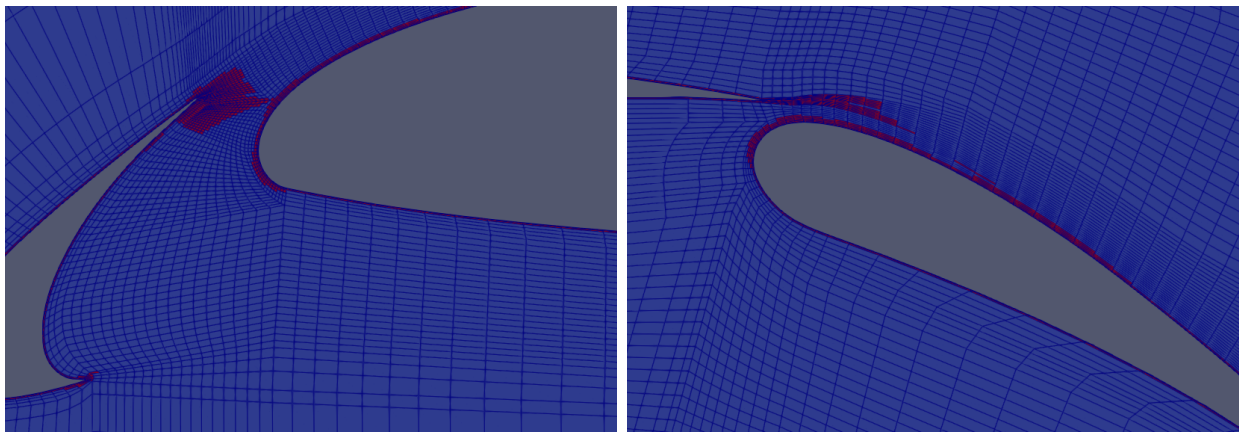


Figure 7.13. Distribution of element sizes for the multi-element airfoil grid



(a) Slat cove zoom-in


(b) Flap zoom-in

Figure 7.14. Distribution of implicit and explicit elements after IMEX partitioning. Blue elements are solved explicitly, and red elements are solved implicitly



(a) Averaged, $p = 1$

(b) Averaged, $p = 2$



(c) Instantaneous, $p = 1$

(d) Instantaneous, $p = 2$

Figure 7.15. Vorticity contours for the multi-element airfoil problem

more detail can be observed for the $p = 2$ results.






Figure 7.16. Q-criterion contours for the multi-element airfoil coloured by velocity magnitude

Time-averaged plots of the pressure coefficient are shown in Figure 7.18. Results are compared against experimental data by Florida State University [162] and from Muyarama et al. [168]. We note that these experimental values were originally performed in wind tunnel facilities at different angles of attack. They were carried out in closed-wall wind tunnels with significant end-wall effects, altering the effective attack angle. However, they were later compared to numerical simulations and were deemed appropriate as reference data at an angle of attack 5.5 degrees [162]. Increasing the order to $p = 2$ improves the agreement between the current results and the reference data, particularly for the main airfoil and the flap. The C_p plots are close to the reference data for $p = 2$. The slat in our problem is still underresolved and may require a higher degree in the computations or refinement of the grid in this area. We also compare velocity profiles along lines normal to the shear layer, shown in Figure 7.19. Results for the two computations considered in this section are shown in Figure 7.20, namely $p = 1$ and $p = 2$. Overall, an improved agreement can be observed for the third-order results, especially in proximity to the shear layer. In this region, the higher-order results follow the increases in velocity to the freestream conditions. The low-order method shows more dissipated behaviour in these regions, which is expected and consistent with



(a) HFR, $p = 2$



(b) EFR, $p = 2$

Figure 7.17. Q-criterion contours zoomed in the slat for the multi-element airfoil coloured by streamwise vorticity

the numerical error of these schemes. The $p = 2$ results show oscillatory behaviour for the higher-order results at the crossing of the shear layer in L1. In general, there is still some discrepancy in areas close to the mean quantities, especially for $p = 1$.




Figure 7.18. Pressure coefficient C_p for the multi-element airfoil problem

Finally, we measure computational performance. Specifically, the time spent in 100 time steps for FR, HFR, and EFR methods is computed. These timing metrics are shown and compared with the standard FR IMEX approach and with their explicit counterparts when the implicit factor is zero, which results in an optimal explicit-Runge Kutta method. Results are shown in Table 7.6. These runtimes were computed on 2.4GHz AMD Rome 7532 CPUS using 1024 cores. Similar to previous results in the cylinder, significant time is spent on the computation of the global problems, especially at higher order, accounting for 80% of the time in FR, over 70% in HFR, and over 50% in EFR, with most of the remaining time in the explicit portion. Compared to the standard FR IMEX approach, EFR IMEX methods achieved 1.5 to 4.1 times faster simulations, representing at least 15 times faster than an explicit FR simulation of this problem. Hence, we have demonstrated that hybridized IMEX methods can significantly speed up computations of moderately high Reynolds numbers.






Figure 7.19. Line plot locations in the vicinity of the slat

Table 7.6. Summary of performance metrics for the multi-element airfoil problem for 100 time steps


Scheme	p	t_G	t_L	t_J	$t_{R^{im}}$	$t_{R^{ex}}$	t_w	t_w/t_w^{FR}	t_w/t_w^{ex}
FR	1	14.66	-	4.72	0.65	7.79	27.81	1.00	35.83
	2	101.41	-	19.48	2.04	28.04	150.97	1.00	24.56
	3	983.59	-	134.93	6.28	55.52	1180.32	1.00	4.03
HFR	1	54.61	1.08	1.29	0.63	7.51	65.13	0.43	15.33
	2	218.56	6.46	8.21	2.33	28.49	264.04	0.57	14.06
	3	370.49	15.90	60.97	6.42	52.03	505.81	2.33	9.40
EFR	1	7.72	1.25	1.13	0.67	7.79	18.56	1.50	53.80
	2	62.86	6.50	7.50	2.40	27.97	107.21	1.41	34.62
	3	153.08	15.61	59.02	6.27	51.18	285.17	4.14	16.67




(a) L1




(b) L2




(c) L3




(d) L4



(e) L5



(f) L6



(g) L7

Figure 7.20. Plots of sampled normalized velocity magnitude for the multi-element airfoil across slat shear layer

Chapter 8

Conclusions

This thesis investigated and developed solution-accelerating strategies for high-order flux reconstruction methods. A framework was proposed to advance high-order methods toward complex three-dimensional problems of industrial interest, considering explicit and implicit implementations. To this end, optimal Runge-Kutta methods were developed for the explicit portion of this thesis, and hybridization of the FR approach was introduced for the implicit part of this work. These approaches were combined with polynomial adaptation and with hybridized IMEX methods to further reduce computational cost. In the next paragraphs, we provide detailed observations of the methods developed in this thesis.

In the context of explicit simulations, optimal stability polynomials were generated for the flux reconstruction approach in one, two, and three dimensions. The element types considered were lines, quadrilaterals, triangles, prisms, tetrahedra, and hexahedra. Larger optimized time steps were achieved by increasing the number of Runge-Kutta stages and adding terms to the stability polynomial. The relative efficiency increased with the number of stages for all element types. While quadrilateral and hexahedral elements did not benefit significantly compared to line elements, triangular, prismatic, and tetrahedral elements demonstrated modest performance improvement. The optimal stability polynomials yielded speedup factors up to 1.97 compared to classical Runge-Kutta methods. These schemes maintained their designed order of accuracy in linear advection. They allowed substantial reductions in computational cost with minimal impact on accuracy, making them a feasible solution for accelerating high-order methods. In addition, the spectral properties of these methods were analyzed for the fully-discrete schemes in a one-dimensional setting. It was shown that the spectral properties of optimal RK methods vary by decreasing the dissipation of the highest frequencies, an unwanted effect for ILES simulations. However, these properties improve as we move slightly from the stability limit, providing the expected speedup factors close to 2.

To advance implicit FR discretizations, hybridized flux reconstruction methods with a range of correction functions were introduced, namely the HFR and EFR schemes. The HFR

method allows discontinuous traces, and the EFR method enforces continuity of the traces on the internal faces. In this work, the trace in all methods was considered discontinuous at the boundaries. Through numerical analysis, EFR methods were compared with HFR and FR methods. It was found that an under-integrated form of the EFR method introduces more numerical error at low orders but behaves similarly to FR at higher polynomial degrees. Verification of the implementation used in this work via linear advection and advection of an isentropic vortex confirmed the expected order of accuracy for all hybridized schemes. In the isentropic vortex case, HFR achieved speedup factors of 1.4, while EFR achieved 2.3 compared to implicit FR. Dissipative correction functions slightly reduced the computational burden of the linear solver. EFR with c_{DG} proved to be a cost-effective choice over FR with c_{SD} and c_{HU} . EFR outperformed HFR in a complex airfoil vortex interaction problem, and both were significantly faster than FR with speedup factors of 4.3 for HFR and 6.7 for EFR, considering a fifth-order scheme. With these studies, it was determined that hybridization can be effectively used for flux reconstruction methods. In the context of advection-diffusion problems, we demonstrated linear stability constraints with the HFR method and the relationship with standard FR schemes. Significant speedups of linear diffusion and advection-diffusion problems were observed by almost two orders of magnitude for a Couette problem. In addition, the superconvergence properties of HFR methods in the diffusion-dominated regime for the range of energy-stable correction functions were evaluated, empirically showing that it is only a property of the c_{DG} method at $p = 1$ and within its vicinity for $p > 1$.

Next, polynomial adaptation was introduced to the hybridized FR framework. We presented a procedure for locally adapting the polynomial degrees in each cell by solving a global projection problem, which determines the new value of the trace in the functionally nonconforming setting. Due to the discontinuous nature of the HFR method, this problem is, in fact, local but remains global for EFR. Since adaptation is not performed at every iteration, the cost of this adaptation procedure for EFR is akin to that of a time step. Hence, the interval for adaptation can be chosen considering this overhead. A series of numerical examples demonstrated the applicability, performance, and benefits of this approach. Namely, an isentropic vortex, flow over a cylinder at $Re = 150$ and flow over a NACA 0012 airfoil at $Re = 10,000$ were dynamically adapted using a vorticity indicator. Adaptive hybridized methods achieved comparable errors to p -adaptive standard formulations at a fraction of the cost in the vortex problem. In the cylinder case, speedup factors of 5.34 and 6.78 times using p -adaptive HFR and EFR were respectively obtained against p -adaptive FR formulations. The speedups between p -adaptive and standard FR formulations are significantly higher, especially at higher polynomial degrees. In this sense, p -adaptive HFR and EFR formulations were about 40 times faster than a p -uniform FR formulation at $p = 5$. Similarly, results

for an airfoil problem demonstrated comparable results with p -adaptive hybridized methods up to 5.7 times faster than p -adaptive FR runs and up to 33 times faster simulations than p -uniform FR at $p = 5$.

Finally, novel implicit-explicit methods were developed for larger-scale simulations of increased geometry-induced stiffness. Combining hybridized and standard FR formulations can be done conservatively using the FR formulation's fluxes at the interface. Via a performance analysis in two and three dimensions, it was shown that for problems with moderately low stiffness, the obtained speedups are significant in two dimensions. For a laminar cylinder problem, speedups in excess of 6 were obtained at the highest polynomial degrees considering EFR against FR IMEX methods. In the three-dimensional moderately stiff setting via a turbulent cylinder at $\text{Re} = 1,000$, these speedups are only observed at higher polynomial degrees $p > 3$ for HFR and $p \geq 2$ for EFR. However, results for a multi-element airfoil at $\text{Re} = 1.7 \times 10^6$ demonstrated that in problems where geometry-induced stiffness is a significant contributor, such as in moderately-high Reynolds numbers, the speedups of IMEX methods are significant. Performance speedups of EFR methods against FR-IMEX schemes were in excess of four, resulting in simulations at least 15 times faster than explicit counterpart formulations.

8.1 Practical Summary

Some general findings in the form of recommendations summarize based on the above observations in a practical sense

- Optimal Runge-Kutta methods enable faster simulations where explicit methods are the method of choice for time integration. Specifically, they are suitable for problems with moderately low stiffness or when computational resources are not sufficient for implicit time stepping.
- Hybridization is a suitable solution-acceleration strategy for implicit flux reconstruction schemes. It enables reduced computational cost while maintaining accuracy. It showed significant potential in two-dimensional formulations at low and high orders. In three dimensions, however, the EFR method enables speedups even at low orders, unlike the HFR method, which showed suitability in most cases for ($p \geq 3$) only.
- Incorporating polynomial adaptation into the hybridized FR framework reduces the overall computational cost of simulations and yields good agreement with p -adaptive FR. This is true for both the HFR and EFR methods, which make use of a global projection to obtain the value of the trace after each adaptation call. This call should

be made in sensible intervals, which can be a fraction of the characteristic time of the problem.

- While hybridization enables significant reductions in computational cost over standard FR, industrial-scale problems require IMEX methods in order to alleviate restrictive stability constraints of explicit methods. Combining hybridized and standard formulations allows the computation of moderately high Re, wall-resolved LES problems using the EFR method with faster computations than standard IMEX schemes in low and high orders.

8.2 Future Work

Possible research avenues are listed below

- Polynomial adaptation was applied to problems in two dimensions in this thesis to demonstrate solution-acceleration of vortex-dominated flows. The natural next step is the extension to three dimensions alongside the development of dynamical load-balancing algorithms that consider the distribution of the trace over parallel ranks. In addition to this, exploring additional indicators can improve polynomial adaptive results in general. A yet unexplored indicator can be the use of solution jumps and the postprocessing of hybridized methods to flag elements that require additional resolution.
- IMEX with hybridization demonstrated significant potential in flow computations using the FR approach. The properties of this method can be expanded to simulate additional applications that benefit from domain subdivisions with localized stiffness. This can be, for instance, fluid-solid interaction problems, where hybridization can be used in the solid portion and FR in the fluid portion. This will allow a seamless extension of flow solvers to more complex applications.
- While hybridization via IMEX methods accelerated a wall-resolved LES simulation in this work, introducing additional regularization is necessary to maintain stability of underresolved simulations. Hence, incorporating entropy-stable formulations and tools is an important factor in advancing high-order methods toward industrial-scale computations.
- Each of the methodologies considered in this work advanced high-order methods to more complex applications. Combining polynomial adaptation, hybridization, and IMEX methods is yet to be explored in practical underresolved applications in performance and robustness. The main challenge for this idea will be stability, which is expected to improve with the previous suggestion introducing nonlinearly stable formulations.

References

- [1] IATA. *Climate Change and CORSIA Fact Sheet*. 2017.
- [2] A. Vedantham. *Aviation and the Global Atmosphere: a Special Report of IPCC Working Groups I and III*. 1999.
- [3] A. Krein and G. Williams. “Flightpath 2050: Europe’s Vision for Aeronautics”. In: *Innovation for Sustainable Aviation in a Global Environment: Proceedings of the Sixth European Aeronautics Days, Madrid 30* (2012).
- [4] *Canada’s Aviation Climate Action Plan 2022-2030*.
- [5] Kolmogorov Andrei Nikolaevich et al. “The Local Structure of Turbulence in Incompressible Viscous Fluid for Very Large Reynolds Numbers”. In: *Proceedings of the Royal Society of London. Series A: Mathematical and Physical Sciences* 434.1890 (1991), pp. 9–13.
- [6] P. Davidson. *Turbulence: An Introduction for Scientists and Engineers*. Oxford University Press, 2015. 647 pp. ISBN: 978-0-19-872259-5.
- [7] P. M. Sforza. *Commercial Airplane Design Principles*. Elsevier, 2014.
- [8] J. P. Slotnick et al. *CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences*. NF1676L-18332. 2014.
- [9] J. S. Hesthaven and T. Warburton. *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications*. Springer Science & Business Media, 2007.
- [10] Z. J. Wang. “High-Order Computational Fluid Dynamics Tools for Aircraft Design”. In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 372.2022 (2014), p. 20130318.
- [11] S. A. Orszag and G. S. Patterson Jr. “Numerical Simulation of Three-Dimensional Homogeneous Isotropic Turbulence”. In: *Physical Review Letters* 28.2 (1972), p. 76.
- [12] P. R. Spalart. “Strategies for Turbulence Modelling and Simulations”. In: *International Journal of Heat and Fluid Flow* 21.3 (2000), pp. 252–263.

- [13] M. Germano et al. “A Dynamic Subgrid-Scale Eddy Viscosity Model”. In: *Physics of Fluids A: Fluid Dynamics* 3.7 (1991), pp. 1760–1765.
- [14] J. P. Boris. “On Large Eddy Simulation Using Subgrid Turbulence Models”. In: *Whither Turbulence? Turbulence at the Crossroads: Proceedings of a Workshop Held at Cornell University, Ithaca, NY, March 22–24, 1989*. Springer, 2005, pp. 344–353.
- [15] B. C. Vermeire, S. Nadarajah, and P. G. Tucker. “Implicit Large Eddy Simulation Using the High-Order Correction Procedure via Reconstruction Scheme”. In: *International Journal for Numerical Methods in Fluids* 82.5 (2016), pp. 231–260.
- [16] H. Zhu et al. “Implicit Large-Eddy Simulation for the High-Order Flux Reconstruction Method”. In: *AIAA Journal* 54.9 (2016), pp. 2721–2733.
- [17] Y. Li and Z. J. Wang. “A Priori and a Posteriori Evaluations of Subgrid Stress Models with the Burgers’ Equation”. In: *53rd AIAA Aerospace Sciences Meeting*. American Institute of Aeronautics and Astronautics, 2015.
- [18] P. Durbin and G. Medic. *Fluid Dynamics with a Computational Perspective*. Fluid Dynamics with a Computational Perspective v. 10. Cambridge University Press, 2007. ISBN: 978-0-521-85017-9.
- [19] Z. J. Wang et al. “High-Order CFD Methods: Current Status and Perspective”. In: *International Journal for Numerical Methods in Fluids* 72.8 (2013), pp. 811–845.
- [20] C. W. Shu. “High Order ENO and WENO Schemes for Computational Fluid Dynamics”. In: *High-Order Methods for Computational Physics* (1999), pp. 439–582.
- [21] A. Harten et al. “Uniformly High Order Accurate Essentially Non-Oscillatory Schemes, III”. In: *Journal of Computational Physics* 131.1 (1997), pp. 3–47.
- [22] B. Cockburn and C. W. Shu. “TVB Runge-Kutta Local Projection Discontinuous Galerkin Finite Element Method for Conservation Laws. II. General Framework”. In: *Mathematics of Computation* 52.186 (1989), pp. 411–435.
- [23] B. Cockburn, S. Y. Lin, and C. W. Shu. “TBV Runge-Kutta Local Projection Discontinuous Galerkin Finite Element Method for Conservation Laws III: One-Dimensional Systems”. In: *Journal of Computational Physics* 84.1 (1989), pp. 90–113.
- [24] Z. J. Wang. “Spectral (Finite) Volume Method for Conservation Laws on Unstructured Grids. Basic Formulation: Basic Formulation”. In: *Journal of Computational Physics* 178.1 (2002), pp. 210–251.

- [25] Y. Liu, M. Vinokur, and Z. J. Wang. “Discontinuous Spectral Difference Method for Conservation Laws on Unstructured Grids”. In: *Computational Fluid Dynamics 2004: Proceedings of the Third International Conference on Computational Fluid Dynamics, ICCFD3, Toronto, 12–16 July 2004*. Springer, 2006, pp. 449–454.
- [26] H. T. Huynh. “A Flux Reconstruction Approach to High-Order Schemes Including Discontinuous Galerkin Methods”. In: *18th AIAA Computational Fluid Dynamics Conference*. 2007, p. 4079.
- [27] H. T. Huynh. “A Reconstruction Approach to High-Order Schemes Including Discontinuous Galerkin for Diffusion”. In: *47th AIAA Aerospace Sciences Meeting*. American Institute of Aeronautics and Astronautics.
- [28] Z. J. Wang and H. Gao. “A Unifying Lifting Collocation Penalty Formulation Including the Discontinuous Galerkin, Spectral Volume/Difference Methods for Conservation Laws on Mixed Grids”. In: *Journal of Computational Physics* 228.21 (2009), pp. 8161–8186.
- [29] P. E. Vincent, P. Castonguay, and A. Jameson. “A New Class of High-Order Energy Stable Flux Reconstruction Schemes”. In: *Journal of Scientific Computing* 47.1 (2011), pp. 50–72. ISSN: 1573-7691.
- [30] P. Castonguay, P. E. Vincent, and A. Jameson. “A New Class of High-Order Energy Stable Flux Reconstruction Schemes for Triangular Elements”. In: *Journal of Scientific Computing* 51.1 (2012), pp. 224–256.
- [31] P. E. Vincent, P. Castonguay, and A. Jameson. “Insights from von Neumann Analysis of High-Order Flux Reconstruction Schemes”. In: *Journal of Computational Physics* 230.22 (2011), pp. 8134–8154.
- [32] C. A. Pereira and B. C. Vermeire. “Performance and Accuracy of Hybridized Flux Reconstruction Schemes”. In: *Journal of Computational Physics* 457 (2022), p. 111039.
- [33] B. C. Vermeire, F. D. Witherden, and P. E. Vincent. “On the Utility of GPU Accelerated High-Order Methods for Unsteady Flow Simulations: A Comparison with Industry-Standard Tools”. In: *Journal of Computational Physics* 334 (2017), pp. 497–521.
- [34] R. C. Moura, S. J. Sherwin, and J. Peiró. “Eigensolution Analysis of Spectral/High-Order Continuous Galerkin Approximations to Advection–Diffusion Problems: Insights into Spectral Vanishing Viscosity”. In: *Journal of Computational Physics* 307 (2016), pp. 401–422.
- [35] G. Mengaldo et al. “Spatial Eigensolution Analysis of Discontinuous Galerkin Schemes with Practical Insights for Under-Resolved Computations and Implicit LES”. In: *Computers & Fluids* 169 (2018), pp. 349–364.

- [36] C. A. Pereira and B. C. Vermeire. “Spectral Properties of High-Order Element Types for Implicit Large Eddy Simulation”. In: *Journal of Scientific Computing* 85.2 (2020), pp. 1–38.
- [37] B. C. Vermeire, J. S. Cagnone, and S. Nadarajah. “ILES Using the Correction Procedure via Reconstruction Scheme”. In: *51st AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. 2013, p. 1001.
- [38] P. Fernández. “High-Order Implicit Large-Eddy Simulation for Transitional Aerodynamics Flows”. PhD thesis. Massachusetts Institute of Technology, 2016.
- [39] R. C. Moura, J. Peiró, and S. J. Sherwin. “Implicit LES Approaches via Discontinuous Galerkin Methods at Very Large Reynolds”. In: *Direct and Large-Eddy Simulation XI*. Springer, 2019, pp. 53–59.
- [40] U. M. Ascher, S. J. Ruuth, and R. J. Spiteri. “Implicit-Explicit Runge-Kutta Methods for Time-Dependent Partial Differential Equations”. In: *Applied Numerical Mathematics*. Special Issue on Time Integration 25.2 (1997), pp. 151–167. ISSN: 0168-9274.
- [41] M. Shoeybi et al. “An Adaptive Implicit–Explicit Scheme for the DNS and LES of Compressible Flows on Unstructured Grids”. In: *Journal of Computational Physics* 229.17 (2010), pp. 5944–5965.
- [42] W. Hundsdorfer and J. Jaffré. “Implicit–Explicit Time Stepping with Spatial Discontinuous Finite Elements”. In: *Applied Numerical Mathematics* 45.2-3 (2003), pp. 231–254.
- [43] A. Kanevsky et al. “Application of Implicit–Explicit High Order Runge–Kutta Methods to Discontinuous-Galerkin Schemes”. In: *Journal of Computational Physics* 225.2 (2007), pp. 1753–1781. ISSN: 0021-9991.
- [44] B. C. Vermeire and S. Nadarajah. “Adaptive IMEX Schemes for High-Order Unstructured Methods”. In: *Journal of Computational Physics* 280 (2015), pp. 261–286.
- [45] D. I. Ketcheson and A. J. Ahmadi. “Optimal Stability Polynomials for Numerical Integration of Initial Value Problems”. In: *Communications in Applied Mathematics and Computational Science* 7.2 (2012), pp. 247–271. ISSN: 2157-5452, 1559-3940.
- [46] E. J. Kubatko, B. A. Yeager, and D. I. Ketcheson. “Optimal Strong-Stability-Preserving Runge–Kutta Time Discretizations for Discontinuous Galerkin Methods”. In: *Journal of Scientific Computing* 60.2 (2014), pp. 313–344. ISSN: 1573-7691.
- [47] B. C. Vermeire, N. A. Loppi, and P. E. Vincent. “Optimal Runge–Kutta Schemes for Pseudo Time-Stepping with High-Order Unstructured Methods”. In: *Journal of Computational Physics* 383 (2019), pp. 55–71.

- [48] B. C. Vermeire and S. H. Nasab. “Accelerated Implicit-Explicit Runge-Kutta Schemes for Locally Stiff Systems”. In: *Journal of Computational Physics* 429 (2021), p. 110022.
- [49] B. C. Vermeire, N. A. Loppi, and P. E. Vincent. “Optimal Embedded Pair Runge-Kutta Schemes for Pseudo-Time Stepping”. In: *Journal of Computational Physics* 415 (2020), p. 109499.
- [50] A. Jameson. “Solution of the Euler Equations for Two Dimensional Transonic Flow by a Multigrid Method”. In: *Applied mathematics and computation* 13.3-4 (1983), pp. 327–355.
- [51] K. J. Fidkowski et al. “P-Multigrid Solution of High-Order Discontinuous Galerkin Discretizations of the Compressible Navier–Stokes Equations”. In: *Journal of Computational Physics* 207.1 (2005), pp. 92–113.
- [52] H. Yang, F. Li, and J. Qiu. “Dispersion and Dissipation Errors of Two Fully Discrete Discontinuous Galerkin Methods”. In: *Journal of Scientific Computing* 55 (2013).
- [53] B. C. Vermeire and P. E. Vincent. “On the Behaviour of Fully-Discrete Flux Reconstruction Schemes”. In: *Computer Methods in Applied Mechanics and Engineering* 315 (2017), pp. 1053–1079.
- [54] B. Cockburn, J. Gopalakrishnan, and R. Lazarov. “Unified Hybridization of Discontinuous Galerkin, Mixed, and Continuous Galerkin Methods for Second Order Elliptic Problems”. In: *SIAM Journal on Numerical Analysis* 47.2 (2009), pp. 1319–1365. ISSN: 0036-1429.
- [55] B. F. De Veubeke and O. C. Zienkiewicz. “Displacement and Equilibrium Models in the Finite Element Method”. In: *International Journal for Numerical Methods in Engineering* 52.3 (2001), pp. 287–342.
- [56] B. Cockburn. “Static Condensation, Hybridization, and the Devising of the HDG Methods”. In: *Building Bridges: Connections and Challenges in Modern Approaches to Numerical Partial Differential Equations*. Ed. by G. R. Barrenechea et al. Cham: Springer International Publishing, 2016, pp. 129–177. ISBN: 978-3-319-41640-3.
- [57] B. Cockburn, B. Dong, and J. Guzmán. “A Superconvergent LDG-hybridizable Galerkin Method for Second-Order Elliptic Problems”. In: *Mathematics of Computation* 77.264 (2008), pp. 1887–1916.
- [58] B. Cockburn, W. Qiu, and K. Shi. “Conditions for Superconvergence of HDG Methods for Second-Order Elliptic Problems”. In: *Mathematics of Computation* 81.279 (2012), pp. 1327–1353.

- [59] N. C. Nguyen, J. Peraire, and B. Cockburn. “An Implicit High-Order Hybridizable Discontinuous Galerkin Method for Linear Convection–Diffusion Equations”. In: *Journal of Computational Physics* 228.9 (2009), pp. 3232–3254.
- [60] S. Güzey, B. Cockburn, and H. K. Stolarski. “The Embedded Discontinuous Galerkin Method: Application to Linear Shell Problems”. In: *International Journal for Numerical Methods in Engineering* 70.7 (2007), pp. 757–790.
- [61] B. Cockburn et al. “An Analysis of the Embedded Discontinuous Galerkin Method for Second-Order Elliptic Problems”. In: *SIAM Journal on Numerical Analysis* 47.4 (2009), pp. 2686–2707.
- [62] D. S. Kamenetskiy. “On the Relation of the Embedded Discontinuous Galerkin Method to the Stabilized Residual-Based Finite Element Methods”. In: *Applied Numerical Mathematics* 108 (2016), pp. 271–285.
- [63] N. C. Nguyen, J. Peraire, and B. Cockburn. “A Class of Embedded Discontinuous Galerkin Methods for Computational Fluid Dynamics”. In: *Journal of Computational Physics* 302 (2015), pp. 674–692. ISSN: 0021-9991.
- [64] A. Jaust et al. “FESTUNG: A MATLAB/GNU Octave Toolbox for the Discontinuous Galerkin Method. Part III: Hybridized Discontinuous Galerkin (HDG) Formulation”. In: *Computers & Mathematics with Applications* 75.12 (2018), pp. 4505–4533.
- [65] B. Cockburn et al. “A Hybridizable Discontinuous Galerkin Method for Steady-State Convection-Diffusion-Reaction Problems”. In: *SIAM Journal on Scientific Computing* 31.5 (2009), pp. 3827–3846.
- [66] Y. Chen and B. Cockburn. “Analysis of Variable-Degree HDG Methods for Convection–Diffusion Equations. Part I: General Nonconforming Meshes”. In: *IMA Journal of Numerical Analysis* 32.4 (2012), pp. 1267–1293.
- [67] M. Giacomini, R. Sevilla, and A. Huerta. “Tutorial on Hybridizable Discontinuous Galerkin (HDG) Formulation for Incompressible Flow Problems”. In: *Modeling in Engineering Using Innovative Numerical Methods for Solids and Fluids*. Springer, 2020, pp. 163–201.
- [68] A. Cesmelioglu, B. Cockburn, and W. Qiu. “Analysis of a Hybridizable Discontinuous Galerkin Method for the Steady-State Incompressible Navier-Stokes Equations”. In: *Mathematics of Computation* 86.306 (2017), pp. 1643–1670.
- [69] N. C. Nguyen, J. Peraire, and B. Cockburn. “An Implicit High-Order Hybridizable Discontinuous Galerkin Method for Nonlinear Convection–Diffusion Equations”. In: *Journal of Computational Physics* 228.23 (2009), pp. 8841–8855.

- [70] P. Fernandez et al. “Implicit Large-Eddy Simulation of Compressible Flows Using the Interior Embedded Discontinuous Galerkin Method”. In: *54th AIAA Aerospace Sciences Meeting*. 2016, p. 1332.
- [71] J. Vila-Pérez et al. “Hybridisable Discontinuous Galerkin Formulation of Compressible Flows”. In: *Archives of Computational Methods in Engineering* 28.2 (2021), pp. 753–784.
- [72] K. J. Fidkowski. “Comparison of Hybrid and Standard Discontinuous Galerkin Methods in a Mesh-Optimisation Setting”. In: *International Journal of Computational Fluid Dynamics* 33.1-2 (2019), pp. 34–42.
- [73] P. Fernandez, N. C. Nguyen, and J. Peraire. “The Hybridized Discontinuous Galerkin Method for Implicit Large-Eddy Simulation of Transitional Turbulent Flows”. In: *Journal of Computational Physics* 336 (2017), pp. 308–329.
- [74] J. S. Cagnone and S. K. Nadarajah. “A Stable Interface Element Scheme for the P-Adaptive Lifting Collocation Penalty Formulation”. In: *Journal of Computational Physics* 231.4 (2012), pp. 1615–1634.
- [75] H. Gao and Z. J. Wang. “A Residual-Based Procedure for Hp-Adaptation on 2-d Hybrid Meshes”. In: *49th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. 2011, p. 492.
- [76] I. Babuska, B. A. Szabo, and I. N. Katz. “The P-Version of the Finite Element Method”. In: *SIAM Journal on Numerical Analysis* 18.3 (1981), pp. 515–545.
- [77] M. Kompenhans et al. “Comparisons of P-Adaptation Strategies Based on Truncation- and Discretisation-Errors for High Order Discontinuous Galerkin Methods”. In: *Computers & Fluids* 139 (2016), pp. 36–46.
- [78] L. Wang and D. J. Mavriplis. “Adjoint-Based h–p Adaptive Discontinuous Galerkin Methods for the 2D Compressible Euler Equations”. In: *Journal of Computational Physics* 228.20 (2009), pp. 7643–7661.
- [79] K. J. Fidkowski and D. L. Darmofal. “Review of Output-Based Error Estimation and Mesh Adaptation in Computational Fluid Dynamics”. In: *AIAA Journal* 49.4 (2011), pp. 673–694.
- [80] K. J. Fidkowski and Y. Luo. “Output-Based Space–Time Mesh Adaptation for the Compressible Navier–Stokes Equations”. In: *Journal of Computational Physics* 230.14 (2011), pp. 5753–5773.
- [81] L. Shi and Z. J. Wang. “Adjoint-Based Error Estimation and Mesh Adaptation for the Correction Procedure via Reconstruction Method”. In: *Journal of Computational Physics* 295 (2015), pp. 261–284.

- [82] E. Ching, Y. Lv, and M. Ihme. “Entropy Residual as a Feature-Based Adaptation Indicator for Simulations of Unsteady Flow”. In: *54th AIAA Aerospace Sciences Meeting*. 2016, p. 0837.
- [83] R. Ghoreishi and B. C. Vermeire. “Vorticity-Based Polynomial Adaptation for Moving and Deforming Domains”. In: *Computers & Fluids* 231 (2021), p. 105160.
- [84] F. Naddei, M. de la Llave Plata, and V. Couaillier. “A Comparison of Refinement Indicators for P-Adaptive Discontinuous Galerkin Methods for the Euler and Navier-Stokes Equations”. In: *2018 AIAA Aerospace Sciences Meeting*. 2018, p. 0368.
- [85] C. Roy. “Review of Discretization Error Estimators in Scientific Computing”. In: *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. 2010, p. 126.
- [86] S. J. Kamkar et al. “Combined Feature-Driven Richardson-based Adaptive Mesh Refinement for Unsteady Vortical Flows”. In: *AIAA Journal* 50.12 (2012), pp. 2834–2847.
- [87] G. Giorgiani, S. Fernández-Méndez, and A. Huerta. “Hybridizable Discontinuous Galerkin P-Adaptivity for Wave Propagation Problems”. In: *International Journal for Numerical Methods in Fluids* 72.12 (2013), pp. 1244–1262. ISSN: 1097-0363.
- [88] G. Giorgiani, S. Fernández-Méndez, and A. Huerta. “Hybridizable Discontinuous Galerkin with Degree Adaptivity for the Incompressible Navier–Stokes Equations”. In: *Computers & Fluids* 98 (2014), pp. 196–208.
- [89] J. M. Hoermann et al. “An Adaptive Hybridizable Discontinuous Galerkin Approach for Cardiac Electrophysiology”. In: *International Journal for Numerical Methods in Biomedical Engineering* 34.5 (2018), e2959.
- [90] K. Fidkowski. “Output-Based Mesh Optimization for the Embedded Discontinuous Galerkin Method”. In: *AIAA Aviation 2019 Forum*. 2019, p. 2950.
- [91] A. J. Chorin. “A Numerical Method for Solving Incompressible Viscous Flow Problems”. In: *Journal of Computational Physics* 135.2 (1997), pp. 118–125.
- [92] A. Jameson. “Time Dependent Calculations Using Multigrid, with Applications to Unsteady Flows Past Airfoils and Wings”. In: *10th Computational Fluid Dynamics Conference*. 1991, p. 1596.
- [93] R. L. Panton. *Incompressible Flow*. John Wiley & Sons, 2013.
- [94] J. R. Clausen. “Entropically Damped Form of Artificial Compressibility for Explicit Simulation of Incompressible Flow”. In: *Physical Review E* 87.1 (2013), p. 013309.

- [95] W. Trojak et al. “Artificial Compressibility Approaches in Flux Reconstruction for Incompressible Viscous Flow Simulations”. In: *Computers & Fluids* 247 (2022), p. 105634.
- [96] P. Zwanenburg and S. Nadarajah. “Equivalence between the Energy Stable Flux Reconstruction and Filtered Discontinuous Galerkin Schemes: Numerical Validation”. In: *54th AIAA Aerospace Sciences Meeting*. 2016, p. 1824.
- [97] D. M. Williams and A. Jameson. “Energy Stable Flux Reconstruction Schemes for Advection–Diffusion Problems on Tetrahedra”. In: *Journal of Scientific Computing* 59.3 (2014), pp. 721–759.
- [98] D. N. Arnold. “An Interior Penalty Finite Element Method with Discontinuous Elements”. In: *SIAM Journal on Numerical Analysis* 19.4 (1982), pp. 742–760.
- [99] F. Bassi and S. Rebay. “A High-Order Accurate Discontinuous Finite Element Method for the Numerical Solution of the Compressible Navier–Stokes Equations”. In: *Journal of Computational Physics* 131.2 (1997), pp. 267–279.
- [100] F. Bassi and S. Rebay. “A High Order Discontinuous Galerkin Method for Compressible Turbulent Flows”. In: *Discontinuous Galerkin Methods*. Ed. by B. Cockburn, G. E. Karniadakis, and C. W. Shu. Lecture Notes in Computational Science and Engineering. Springer Berlin Heidelberg, 2000, pp. 77–88. ISBN: 978-3-642-59721-3.
- [101] B. Cockburn and C. W. Shu. “The Local Discontinuous Galerkin Method for Time-Dependent Convection-Diffusion Systems”. In: *SIAM Journal on Numerical Analysis* 35.6 (1998), pp. 2440–2463.
- [102] V. V. Rusanov. *Calculation of Interaction of Non-Steady Shock Waves with Obstacles*. NRC, Division of Mechanical Engineering, 1962.
- [103] P. D. Lax. “Weak Solutions of Nonlinear Hyperbolic Equations and Their Numerical Computation”. In: *Communications on Pure and Applied Mathematics* 7.1 (1954), pp. 159–193.
- [104] P. L. Roe. “Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes”. In: *Journal of Computational Physics* 43.2 (1981), pp. 357–372.
- [105] E. F. Toro, M. Spruce, and W. Speares. “Restoration of the Contact Surface in the HLL-Riemann Solver”. In: *Shock waves* 4 (1994), pp. 25–34.
- [106] J. Peraire and P. O. Persson. “The Compact Discontinuous Galerkin (CDG) Method for Elliptic Problems”. In: *SIAM Journal on Scientific Computing* 30.4 (2008), pp. 1806–1824.

- [107] P. Fernández. “Entropy-Stable Hybridized Discontinuous Galerkin Methods for Large-Eddy Simulation of Transitional and Turbulent Flows”. PhD thesis. Massachusetts Institute of Technology, 2019.
- [108] J. C. Butcher. “Practical Runge–Kutta Methods for Scientific Computation”. In: *The ANZIAM Journal* 50.3 (2009), pp. 333–342.
- [109] O. C. Zienkiewicz and R. L. Taylor. *The Finite Element Method for Solid and Structural Mechanics*. Elsevier, 2005.
- [110] G. Karniadakis and S. Sherwin. *Spectral/ Hp Element Methods for Computational Fluid Dynamics*. Oxford University Press, 2013.
- [111] X. C. Cai and M. Sarkis. “A Restricted Additive Schwarz Preconditioner for General Sparse Linear Systems”. In: *SIAM Journal on Scientific Computing* 21.2 (1999), pp. 792–797.
- [112] S. Balay et al. *PETSc/TAO Users Manual*. ANL-21/39 - Revision 3.19. Argonne National Laboratory, 2023.
- [113] X. Roca, C. Nguyen, and J. Peraire. “Scalable Parallelization of the Hybridized Discontinuous Galerkin Method for Compressible Flow”. In: *21st AIAA Computational Fluid Dynamics Conference*. 2013, p. 2939.
- [114] A. Huerta et al. “Efficiency of High-Order Elements for Continuous and Discontinuous Galerkin Methods”. In: *International Journal for Numerical Methods in Engineering* 96.9 (2013), pp. 529–560.
- [115] H. J. Stetter. *Analysis of Discretization Methods for Ordinary Differential Equations*. Springer Tracts in Natural Philosophy. Berlin Heidelberg: Springer-Verlag, 1973. ISBN: 978-3-642-65473-2.
- [116] E. Hairer, S. P. Nørsett, and G. Wanner. *Solving Ordinary Differential Equations I: Nonstiff Problems*. Berlin, Heidelberg: Springer-Verlag, 1993. ISBN: 978-0-387-56670-2.
- [117] S. Hedayati Nasab, C. A. Pereira, and B. C. Vermeire. “Optimal Runge-Kutta Stability Polynomials for Multidimensional High-Order Methods”. In: *Journal of Scientific Computing* 89.1 (2021), p. 11.
- [118] C. Pereira and B. Vermeire. “Fully-Discrete Analysis of High-Order Spatial Discretizations with Optimal Explicit Runge–Kutta Methods”. In: *Journal of Scientific Computing* 83.3 (2020), pp. 1–35.
- [119] B. C. Vermeire. “Paired Explicit Runge-Kutta Schemes for Stiff Systems of Equations”. In: *Journal of Computational Physics* 393 (2019), pp. 465–483. ISSN: 0021-9991.

- [120] S. Diamond and S. Boyd. “CVXPY: A Python-Embedded Modeling Language for Convex Optimization”. In: *Journal of Machine Learning Research* 17.83 (2016), pp. 1–5.
- [121] A. Agrawal et al. “A Rewriting System for Convex Optimization Problems”. In: *Journal of Control and Decision* 5.1 (2018), pp. 42–60.
- [122] R. C. Moura, S. J. Sherwin, and J. Peiró. “Linear Dispersion–Diffusion Analysis and Its Application to under-Resolved Turbulence Simulations Using Discontinuous Galerkin Spectral/ Hp Methods”. In: *Journal of Computational Physics* 298 (2015), pp. 695–710.
- [123] K. Van den Abeele. “Development of High-Order Accurate Schemes for Unstructured Grids”. In: *Phd Thesis in Vrije Universiteit Brussel* (2009).
- [124] B. C. Vermeire and P. E. Vincent. “On the Properties of Energy Stable Flux Reconstruction Schemes for Implicit Large Eddy Simulation”. In: *Journal of Computational Physics* 327 (2016), pp. 368–388.
- [125] P. Castonguay. “High-Order Energy Stable Flux Reconstruction Schemes for Fluid Flow Simulations on Unstructured Grids”. Stanford University, 2012.
- [126] B. Cockburn and C. W. Shu. “Runge–Kutta Discontinuous Galerkin Methods for Convection-Dominated Problems”. In: *Journal of Scientific Computing* 16.3 (2001), pp. 173–261. ISSN: 1573-7691.
- [127] W. M. Van Rees et al. “A Comparison of Vortex and Pseudo-Spectral Methods for the Simulation of Periodic Vortical Flows at High Reynolds Numbers”. In: *Journal of Computational Physics* 230.8 (2011), pp. 2794–2805.
- [128] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge university press, 2012.
- [129] W. Guo, X. Zhong, and J. M. Qiu. “Superconvergence of Discontinuous Galerkin and Local Discontinuous Galerkin Methods: Eigen-Structure Analysis Based on Fourier Approach”. In: *Journal of Computational Physics* 235 (2013), pp. 458–485.
- [130] F. Q. Hu and H. L. Atkins. “Eigensolution Analysis of the Discontinuous Galerkin Method with Nonuniform Grids: I. One Space Dimension”. In: *Journal of Computational Physics* 182.2 (2002), pp. 516–545. ISSN: 0021-9991.
- [131] N. A. Adams. “Implicit Subgrid-Scale Modeling by Adaptive Local Deconvolution”. In: *Direct and Large-Eddy Simulation V*. Ed. by R. Friedrich, B. J. Geurts, and O. Métais. ERCOFTAC Series. Springer Netherlands, 2004, pp. 15–22. ISBN: 978-1-4020-2313-2.
- [132] Z. J. Wang and Y. Li. “A Mathematical Analysis of Scale Similarity”. In: *Communications in Computational Physics* 21.1 (2017), pp. 149–161. ISSN: 1815-2406, 1991-7120.

- [133] R. Moura, P. Fernandez, and G. Mengaldo. “Diffusion and Dispersion Characteristics of Hybridized Discontinuous Galerkin Methods for Under-Resolved Turbulence Simulations”. In: *APS Division of Fluid Dynamics Meeting Abstracts*. 2017, F31–007.
- [134] P. Fernandez et al. “Non-Modal Analysis of Spectral Element Methods: Towards Accurate and Robust Large-Eddy Simulations”. In: *Computer Methods in Applied Mechanics and Engineering* 346 (2019), pp. 43–62. ISSN: 0045-7825.
- [135] J. Peraire, N. Nguyen, and B. Cockburn. “A Hybridizable Discontinuous Galerkin Method for the Compressible Euler and Navier-Stokes Equations”. In: *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. 2010, p. 363.
- [136] G. Mengaldo et al. “A Guide to the Implementation of Boundary Conditions in Compact High-Order Methods for Compressible Aerodynamics”. In: *7th AIAA Theoretical Fluid Mechanics Conference*. 2014, p. 2923.
- [137] M. Scully. “Computation of Helicopter Rotor Wake Geometry and Its Influence on Rotor Harmonic Airloads”. Thesis. Massachusetts Institute of Technology, 1975.
- [138] P. Castonguay et al. “Energy Stable Flux Reconstruction Schemes for Advection–Diffusion Problems”. In: *Computer Methods in Applied Mechanics and Engineering* 267 (2013), pp. 400–417.
- [139] S. Quaegebeur et al. “Stability of Energy Stable Flux Reconstruction for the Diffusion Problem Using Compact Numerical Fluxes”. In: *SIAM Journal on Scientific Computing* 41.1 (2019), A643–A667.
- [140] S. Quaegebeur, A. Cicchino, and S. Nadarajah. “Stability of Energy Stable Flux Reconstruction for the Diffusion Problem Using Compact Numerical Fluxes on Quadrilateral Elements”. In: *AIAA Scitech 2020 Forum*. 2020, p. 1048.
- [141] A. Sheshadri and A. Jameson. “Erratum to: On the Stability of the Flux Reconstruction Schemes on Quadrilateral Elements for the Linear Advection Equation”. In: *Journal of Scientific Computing* 67.2 (2016), pp. 791–794.
- [142] A. Sheshadri and A. Jameson. “An Analysis of Stability of the Flux Reconstruction Formulation on Quadrilateral Elements for the Linear Advection–Diffusion Equation”. In: *Journal of Scientific Computing* 74.3 (2018), pp. 1757–1785.
- [143] A. Sheshadri and A. Jameson. “On the Stability of the Flux Reconstruction Schemes on Quadrilateral Elements for the Linear Advection Equation”. In: *Journal of Scientific Computing* 67.2 (2016), pp. 769–790.

- [144] B. Cockburn, J. Guzmán, and H. Wang. “Superconvergent Discontinuous Galerkin Methods for Second-Order Elliptic Problems”. In: *Mathematics of Computation* 78.265 (2009), pp. 1–24.
- [145] B. Cockburn et al. “Superconvergence of the Local Discontinuous Galerkin Method for Elliptic Problems on Cartesian Grids”. In: *SIAM Journal on Numerical Analysis* 39.1 (2001), pp. 264–285.
- [146] O. Inoue and N. Hatakeyama. “Sound Generation by a Two-Dimensional Circular Cylinder in a Uniform Flow”. In: *Journal of Fluid Mechanics* 471 (2002), pp. 285–314.
- [147] F. D. Witherden, J. S. Park, and P. E. Vincent. “An Analysis of Solution Point Coordinates for Flux Reconstruction Schemes on Tetrahedral Elements”. In: *Journal of Scientific Computing* 69 (2016), pp. 905–920.
- [148] R. M. Kirby and G. E. Karniadakis. “De-Aliasing on Non-Uniform Grids: Algorithms and Applications”. In: *Journal of Computational Physics* 191.1 (2003), pp. 249–264.
- [149] A. D. Beck et al. “On the Influence of Polynomial De-Aliasing on Subgrid Scale Models”. In: *Flow, Turbulence and Combustion* 97 (2016), pp. 475–511.
- [150] T. Ikeda and T. Atobe. “Numerical Studies of Acoustic Effects on 2D Airfoil Aerodynamics at a Low Reynolds Number”. In: *50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*. 2012, p. 700.
- [151] T. Ohtake, Y. Nakae, and T. Motohashi. “Nonlinearity of the Aerodynamic Characteristics of NACA0012 Aerofoil at Low Reynolds Numbers”. In: *Japan Society of Aeronautical Space Sciences* 55.644 (2007), pp. 439–445.
- [152] J. Butcher. *Numerical Methods for Ordinary Differential Equations*. John Willy and Sons, Ltd, 2002.
- [153] P. O. Persson. “High-Order LES Simulations Using Implicit-Explicit Runge-Kutta Schemes”. In: *49th AIAA Aerospace Sciences Meeting*. 2011, p. 684.
- [154] S. Kang, F. X. Giraldo, and T. Bui-Thanh. “IMEX HDG-DG: A Coupled Implicit Hybridized Discontinuous Galerkin and Explicit Discontinuous Galerkin Approach for Shallow Water Systems”. In: *Journal of Computational Physics* 401 (2020), p. 109010. ISSN: 0021-9991.
- [155] M. Stanglmeier et al. “An Explicit Hybridizable Discontinuous Galerkin Method for the Acoustic Wave Equation”. In: *Computer Methods in Applied Mechanics and Engineering* 300 (2016), pp. 748–769.

- [156] G. Nehmetallah et al. “An Explicit Hybridizable Discontinuous Galerkin Method for the 3D Time-Domain Maxwell Equations”. In: *Spectral and High Order Methods for Partial Differential Equations* (2020), p. 513.
- [157] A. Samii et al. “A Comparison of the Explicit and Implicit Hybridizable Discontinuous Galerkin Methods for Nonlinear Shallow Water Equations”. In: *Journal of Scientific Computing* 80 (2019), pp. 1936–1956.
- [158] J. S. Cagnone, B. C. Vermeire, and S. Nadarajah. “A P-Adaptive LCP Formulation for the Compressible Navier–Stokes Equations”. In: *Journal of Computational Physics* 233 (2013), pp. 324–338.
- [159] C. Lei, L. Cheng, and K. Kavanagh. “Spanwise Length Effects on Three-Dimensional Modelling of Flow over a Circular Cylinder”. In: *Computer Methods in Applied Mechanics and Engineering* 190.22-23 (2001), pp. 2909–2923.
- [160] M. Zhao, L. Cheng, and T. Zhou. “Direct Numerical Simulation of Three-Dimensional Flow Past a Yawed Circular Cylinder of Infinite Length”. In: *Journal of Fluids and Structures* 25.5 (2009), pp. 831–847. ISSN: 08899746.
- [161] M. M. Choudhari and D. P. Lockard. “Assessment of Slat Noise Predictions for 30P30N High-Lift Configuration from BANC-III Workshop”. In: *21st AIAA/CEAS Aeroacoustics Conference*. 2015, p. 2844.
- [162] K. A. Pascioni and L. N. Cattafesta. “Aeroacoustic Measurements of Leading-Edge Slat Noise”. In: *22nd AIAA/CEAS Aeroacoustics Conference*. 2016, p. 2960.
- [163] N. Ashton, A. West, and F. Mendonça. “Flow Dynamics Past a 30P30N Three-Element Airfoil Using Improved Delayed Detached-Eddy Simulation”. In: *AIAA Journal* 54.11 (2016), pp. 3657–3667.
- [164] M. Terracol and E. Manoha. “Wall-Resolved Large Eddy Simulation of a Highlift Airfoil: Detailed Flow Analysis and Noise Generation Study”. In: *20th AIAA/CEAS Aeroacoustics Conference*. 2014, p. 3050.
- [165] M. M. Choudhari and M. R. Khorrami. “Effect of Three-Dimensional Shear-Layer Structures on Slat Cove Unsteadiness”. In: *AIAA Journal* 45.9 (2007), pp. 2174–2186.
- [166] J. Shi, H. Yan, and Z. J. Wang. “Towards Direct Computation of Aeroacoustic Noise with the High-Order FR/CPR Method”. In: *2018 AIAA/CEAS Aeroacoustics Conference*. 2018, p. 4095.
- [167] K. Pascioni, L. N. Cattafesta, and M. M. Choudhari. “An Experimental Investigation of the 30P30N Multi-Element High-Lift Airfoil”. In: *20th AIAA/CEAS Aeroacoustics Conference*. 2014, p. 3062.

- [168] M. Murayama et al. “Experimental Study on Slat Noise from 30P30N Three-Element High-Lift Airfoil at JAXA Hard-Wall Low-speed Wind Tunnel”. In: *20th AIAA/CEAS Aeroacoustics Conference*. 2014, p. 2080.
- [169] G. Gassner and D. A. Kopriva. “A Comparison of the Dispersion and Dissipation Errors of Gauss and Gauss–Lobatto Discontinuous Galerkin Spectral Element Methods”. In: *SIAM Journal on Scientific Computing* 33.5 (2011), pp. 2560–2579.
- [170] A. Chaudhuri et al. “Explicit Discontinuous Spectral Element Method with Entropy Generation Based Artificial Viscosity for Shocked Viscous Flows”. In: *Journal of Computational Physics* 332 (2017), pp. 99–117.
- [171] D. De Grazia et al. “Connections between the Discontinuous Galerkin Method and High-Order Flux Reconstruction Schemes”. In: *International Journal for Numerical Methods In Fluids* 75.12 (2014), pp. 860–877.
- [172] G. Mengaldo et al. “On the Connections Between Discontinuous Galerkin and Flux Reconstruction Schemes: Extension to Curvilinear Meshes”. In: *Journal of Scientific Computing* 67.3 (2016), pp. 1272–1292. ISSN: 0885-7474, 1573-7691.
- [173] S. Terrana, C. Nguyen, and J. Peraire. “GPU-accelerated Large Eddy Simulation of Hypersonic Flows”. In: *AIAA Scitech 2020 Forum*. 2020, p. 1062.
- [174] G. Karypis and V. Kumar. *METIS: A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices*. Manual. 1998.
- [175] J. Dahm. “Toward Accurate, Efficient, and Robust Hybridized Discontinuous Galerkin Methods”. PhD Thesis. University of Michigan. 2017.

Appendix A

HFR Grid Convergence Tables

Table A.1. L_2 norm of the error for the linear advection problem. Results are presented for HFR/FR and EFR methods

p	Mesh size	FR/HFR						EFR					
		c_{DG}	Order	c_{SD}	Order	c_{HU}	Order	c_{DG}	Order	c_{SD}	Order	c_{HU}	Order
1	5×5	9.71E-2	-	1.05E-1	-	1.25E-1	-	1.09E-1	-	1.18E-1	-	1.33E-1	-
	10×10	5.14E-2	0.916	7.31E-2	0.517	1.10E-1	0.105	7.88E-2	0.467	9.94E-2	0.247	1.23E-1	0.120
	20×20	1.58E-2	1.705	3.71E-2	0.979	7.87E-2	0.082	4.02E-2	0.971	6.19E-2	0.683	9.03E-2	0.440
	40×40	2.85E-3	2.468	1.29E-2	1.528	3.92E-2	0.980	1.35E-2	1.577	2.47E-2	1.326	4.39E-2	1.042
2	5×5	4.24E-2	-	5.68E-2	-	7.18E-2	-	4.05E-2	-	6.00E-2	-	7.22E-2	-
	10×10	7.51E-3	2.498	1.57E-2	1.856	2.53E-2	1.287	6.31E-3	2.682	1.59E-2	1.921	2.41E-2	1.586
	20×20	5.11E-4	3.876	1.83E-3	3.097	3.71E-3	2.318	4.48E-4	3.817	1.78E-3	3.157	3.34E-3	2.848
	40×40	4.40E-5	3.539	1.39E-4	3.718	2.92E-4	3.910	4.96E-5	3.174	1.43E-4	3.638	2.83E-4	3.562
3	5×5	1.38E-2	-	2.04E-2	-	2.57E-2	-	1.29E-2	-	1.96E-2	-	2.46E-2	-
	10×10	5.61E-4	4.619	1.62E-3	3.653	2.55E-3	2.935	6.18E-4	4.386	1.66E-3	3.562	2.56E-3	3.263
	20×20	2.47E-5	4.505	5.24E-5	4.954	8.11E-5	4.181	2.74E-5	4.496	5.61E-5	4.887	8.49E-5	4.914
	40×40	1.52E-6	4.023	2.52E-6	4.377	3.52E-6	3.444	1.67E-6	4.039	2.59E-6	4.440	3.51E-6	4.597
4	5×5	2.97E-3	-	5.45E-3	-	6.90E-3	-	3.12E-3	-	5.60E-3	-	7.04E-3	-
	10×10	6.85E-5	5.440	1.33E-4	5.353	1.78E-4	4.927	6.96E-5	5.486	1.36E-4	5.365	1.80E-4	5.289
	20×20	1.60E-6	5.424	2.66E-6	5.649	3.41E-6	4.217	1.75E-6	5.316	2.82E-6	5.592	3.55E-6	5.665
	40×40	5.03E-8	4.986	8.27E-8	5.006	1.05E-7	4.585	5.84E-8	4.901	9.33E-8	4.917	1.17E-7	4.924
5	5×5	6.17E-4	-	1.11E-3	-	1.37E-3	-	6.17E-4	-	1.14E-3	-	1.39E-3	-
	10×10	2.10E-6	8.198	5.86E-6	7.572	8.13E-6	6.218	2.80E-6	7.782	6.46E-6	7.459	8.68E-6	7.318
	20×20	9.75E-8	4.429	1.62E-7	5.180	1.96E-7	5.455	1.22E-7	4.526	1.79E-7	5.173	2.12E-7	5.353
	40×40	1.55E-9	5.976	2.56E-9	5.979	3.11E-9	5.166	2.27E-9	5.746	3.07E-9	5.864	3.65E-9	5.860

Table A.2. L_2 norm of the error for the steady-state linear advection-diffusion problem

p	Mesh	FR			HFR				EFR										
		c_{DG}	Order	c_{SD}	Order	c_{HU}	Order	c_{DG}	Order	c_{SD}	Order	c_{HU}	Order						
1	5×5	5.79E-03	-	6.39E-03	-	9.82E-03	-	4.02E-03	-	4.94E-03	-	7.92E-03	-	6.49E-03	-	7.15E-03	-	7.87E-03	-
	10×10	1.31E-03	2.14	1.57E-03	2.03	2.62E-03	1.91	1.09E-03	1.88	1.23E-03	2.01	2.07E-03	1.93	1.45E-03	2.16	1.56E-03	2.20	1.64E-03	2.26
	20×20	3.01E-04	2.12	3.83E-04	2.03	6.98E-04	1.91	2.25E-04	2.27	3.11E-04	1.98	6.58E-04	1.65	4.94E-04	1.55	4.99E-04	1.65	4.74E-04	1.79
	40×40	7.78E-05	1.95	9.58E-05	2.00	1.75E-04	2.00	6.80E-05	1.73	9.48E-05	1.71	2.03E-04	1.69	1.67E-04	1.57	1.59E-04	1.65	1.34E-04	1.82
2	5×5	1.37E-03	-	1.29E-03	-	1.33E-03	-	6.94E-04	-	5.95E-04	-	8.84E-04	-	7.78E-04	-	5.54E-04	-	5.73E-04	-
	10×10	1.54E-04	3.15	1.71E-04	2.92	2.06E-04	2.69	1.10E-04	2.66	1.29E-04	2.20	1.80E-04	2.30	1.09E-04	2.84	1.04E-04	2.41	1.20E-04	2.25
	20×20	1.98E-05	2.96	2.32E-05	2.88	2.98E-05	2.79	1.66E-05	2.73	2.30E-05	2.49	3.44E-05	2.39	1.55E-05	2.81	1.78E-05	2.55	2.36E-05	2.35
	40×40	2.42E-06	3.03	2.96E-06	2.97	4.03E-06	2.89	2.30E-06	2.85	3.42E-06	2.75	5.27E-06	2.70	2.09E-06	2.89	2.65E-06	2.75	3.70E-06	2.67
3	5×5	2.43E-04	-	2.48E-04	-	2.62E-04	-	1.69E-04	-	1.63E-04	-	1.90E-04	-	1.72E-04	-	1.50E-04	-	1.59E-04	-
	10×10	1.60E-05	3.92	1.93E-05	3.68	2.27E-05	3.53	1.29E-05	3.71	1.73E-05	3.24	2.20E-05	3.11	1.24E-05	3.79	1.47E-05	3.35	1.78E-05	3.16
	20×20	1.01E-06	3.99	1.28E-06	3.91	1.56E-06	3.86	9.09E-07	3.83	1.34E-06	3.69	1.77E-06	3.64	8.21E-07	3.92	1.06E-06	3.80	1.35E-06	3.72
	40×40	6.16E-08	4.03	8.10E-08	3.98	1.02E-07	3.94	5.99E-08	3.92	9.37E-08	3.84	1.26E-07	3.81	5.01E-08	4.03	6.77E-08	3.97	8.95E-08	3.91
4	5×5	3.69E-05	-	4.11E-05	-	4.46E-05	-	2.62E-05	-	3.26E-05	-	3.87E-05	-	2.59E-05	-	2.89E-05	-	3.30E-05	-
	10×10	1.37E-06	4.75	1.69E-06	4.61	1.92E-06	4.54	1.18E-06	4.48	1.68E-06	4.28	2.03E-06	4.25	1.14E-06	4.51	1.48E-06	4.28	1.73E-06	4.25
	20×20	4.33E-08	4.98	5.73E-08	4.88	6.74E-08	4.83	4.10E-08	4.85	6.29E-08	4.74	7.81E-08	4.70	3.90E-08	4.87	5.48E-08	4.76	6.60E-08	4.72
	40×40	1.34E-09	5.01	1.79E-09	5.00	2.14E-09	4.98	1.35E-09	4.93	2.12E-09	4.89	2.66E-09	4.88	1.27E-09	4.95	1.83E-09	4.90	2.24E-09	4.88

Table A.3. L_2 -norm of the post-processed solution error for the steady-state linear diffusion problem

p	Mesh size	FR			HFR								
		c_{DG}	Order	c_{SD}	Order	c_{HU}	Order	c_{DG}	Order	c_{SD}	Order	c_{HU}	Order
1	10×10	5.39E-03	-	7.46E-03	-	1.70E-02	-	2.22E-03	-	2.99E-03	-	7.37E-03	-
	20×20	9.84E-04	2.14	1.34E-03	2.03	3.06E-03	1.91	2.39E-04	3.22	5.50E-04	2.44	1.86E-03	1.98
	40×40	2.04E-04	2.12	3.08E-04	2.03	6.96E-04	1.91	3.23E-05	2.89	1.22E-04	2.18	4.71E-04	1.98
	80×80	5.22E-05	1.95	7.23E-05	2.00	1.60E-04	2.00	4.33E-06	2.90	3.47E-05	1.81	1.38E-04	1.77
2	10×10	1.84E-03	-	1.90E-03	-	1.99E-03	-	3.57E-04	-	3.73E-04	-	4.98E-04	-
	20×20	1.21E-04	3.15	1.37E-04	2.92	1.56E-04	2.69	2.46E-05	3.86	3.13E-05	3.58	4.35E-05	3.52
	40×40	1.19E-05	2.96	1.36E-05	2.88	1.57E-05	2.79	1.81E-06	3.77	2.61E-06	3.58	3.81E-06	3.52
	80×80	1.36E-06	3.03	1.59E-06	2.97	1.87E-06	2.89	1.25E-07	3.86	1.90E-07	3.78	2.84E-07	3.75
3	10×10	2.72E-04	-	3.52E-04	-	3.99E-04	-	4.80E-05	-	5.02E-05	-	6.40E-05	-
	20×20	1.04E-05	3.92	1.39E-05	3.68	1.60E-05	3.53	1.77E-06	4.76	1.86E-06	4.75	2.52E-06	4.67
	40×40	5.02E-07	3.99	6.75E-07	3.91	7.83E-07	3.86	5.20E-08	5.09	6.08E-08	4.94	9.66E-08	4.70
	80×80	2.83E-08	4.03	3.78E-08	3.98	4.41E-08	3.94	1.56E-09	5.06	2.07E-09	4.87	3.55E-09	4.77
4	10×10	3.53E-05	-	4.28E-05	-	4.61E-05	-	7.73E-06	-	9.81E-06	-	1.13E-05	-
	20×20	8.25E-07	4.75	1.03E-06	4.61	1.13E-06	4.54	1.49E-07	5.70	2.13E-07	5.53	2.53E-07	5.48
	40×40	2.05E-08	4.98	2.52E-08	4.88	2.76E-08	4.83	2.63E-09	5.82	3.99E-09	5.74	4.79E-09	5.72
	80×80	6.27E-10	5.01	7.70E-10	5.00	8.44E-10	4.98	5.00E-11	5.72	7.01E-11	5.83	8.34E-11	5.85

Table A.4. L_2 -norm of the solution error for the steady-state linear diffusion problem with consistent LDG switch

p	Mesh size	FR (no post-processing)						FR (post-processed)					
		c_{DG}	Order	c_{SD}	Order	c_{HU}	Order	c_{DG}	Order	c_{SD}	Order	c_{HU}	Order
1	5×5	8.56E-03	-	9.86E-03	-	1.33E-02	-	1.03E-01	-	1.24E-01	-	1.98E-01	-
	10×10	1.24E-03	2.79	1.75E-03	2.49	3.46E-03	1.95	2.86E-02	3.26	3.28E-02	2.22	5.47E-02	2.11
	20×20	3.13E-04	1.98	4.30E-04	2.03	8.45E-04	2.03	5.93E-03	2.91	8.07E-03	2.44	1.66E-02	2.29
	40×40	8.03E-05	1.96	1.12E-04	1.94	2.31E-04	1.87	1.78E-03	2.90	2.45E-03	1.98	5.09E-03	1.92
2	5×5	1.89E-03	-	2.14E-03	-	2.44E-03	-	1.74E-02	-	1.69E-02	-	2.64E-02	-
	10×10	1.64E-04	3.53	2.00E-04	3.42	2.55E-04	3.26	2.87E-03	4.14	3.47E-03	3.97	4.93E-03	3.83
	20×20	2.06E-05	2.99	2.82E-05	2.83	4.02E-05	2.67	4.35E-04	4.10	6.10E-04	4.02	9.20E-04	3.90
	40×40	2.57E-06	3.01	3.76E-06	2.91	5.65E-06	2.83	6.02E-05	4.07	9.00E-05	4.02	1.39E-04	3.94
3	5×5	3.04E-04	-	3.81E-04	-	4.41E-04	-	4.36E-03	-	4.38E-03	-	5.27E-03	-
	10×10	1.70E-05	4.16	2.25E-05	4.08	2.74E-05	4.00	3.37E-04	4.96	4.61E-04	4.92	5.93E-04	4.96
	20×20	1.04E-06	4.03	1.51E-06	3.89	1.96E-06	3.81	2.38E-05	5.11	3.54E-05	5.21	4.70E-05	5.30
	40×40	6.47E-08	4.01	9.93E-08	3.93	1.32E-07	3.89	1.57E-06	5.05	2.46E-06	5.11	3.33E-06	5.17
4	5×5	4.82E-05	-	6.22E-05	-	7.02E-05	-	6.87E-04	-	8.92E-04	-	1.07E-03	-
	10×10	1.43E-06	5.08	1.98E-06	4.98	2.33E-06	4.91	3.08E-05	5.97	4.47E-05	6.04	5.44E-05	6.09
	20×20	4.49E-08	4.99	6.72E-08	4.88	8.23E-08	4.82	1.07E-06	6.04	1.66E-06	6.12	2.06E-06	6.18
	40×40	1.45E-09	4.95	2.21E-09	4.93	2.74E-09	4.91	3.44E-08	5.35	5.53E-08	5.56	6.97E-08	5.79

Table A.5. L_2 norm of the error for the unsteady advection-diffusion problem

p	Mesh size	FR						EFR									
		c_{DG}	Order	c_{SD}	Order	c_{HU}	Order	c_{DG}	Order	c_{SD}	Order	c_{HU}	Order				
1	5×5	1.72E-02	-	2.28E-02	-	5.19E-02	-	1.28E-02	-	2.10E-02	-	5.50E-02	-	3.19E-02	-	6.02E-02	-
	10×10	3.68E-03	2.23	6.10E-03	1.90	2.02E-02	1.36	2.20E-03	1.85	5.83E-03	2.54	2.06E-02	1.42	8.33E-03	2.33	8.81E-03	2.04
	20×20	8.33E-04	2.14	1.52E-03	2.00	5.62E-03	1.84	4.34E-04	2.34	1.57E-03	1.89	5.93E-03	1.78	8.80E-04	2.85	1.95E-03	2.18
	40×40	1.90E-04	2.13	3.89E-04	1.97	1.49E-03	1.91	9.93E-05	2.13	4.12E-04	1.93	1.60E-03	1.91	1.31E-04	2.74	4.71E-04	2.05
2	5×5	1.02E-03	-	1.42E-03	-	2.36E-03	-	7.24E-04	-	1.18E-03	-	2.16E-03	-	8.58E-04	-	1.19E-03	-
	10×10	1.24E-04	3.03	1.64E-04	3.12	2.47E-04	3.26	8.43E-05	3.10	1.16E-04	3.34	1.99E-04	3.44	9.43E-05	3.19	1.14E-04	3.39
	20×20	1.52E-05	3.04	1.85E-05	3.15	2.56E-05	3.27	1.02E-05	3.04	1.34E-05	3.12	2.17E-05	3.19	1.09E-05	3.11	1.20E-05	3.25
	40×40	1.90E-06	3.00	2.22E-06	3.06	2.91E-06	3.14	1.27E-06	3.01	1.71E-06	2.97	2.79E-06	2.96	1.31E-06	3.05	1.40E-06	3.10
3	5×5	8.04E-05	-	1.22E-04	-	1.64E-04	-	5.34E-05	-	7.30E-05	-	1.02E-04	-	5.43E-05	-	7.29E-05	-
	10×10	5.21E-06	3.95	6.73E-06	4.18	8.39E-06	4.29	3.23E-06	4.05	4.16E-06	4.13	5.61E-06	4.18	3.22E-06	4.07	4.05E-06	4.17
	20×20	3.22E-07	4.01	4.14E-07	4.02	5.12E-07	4.03	2.00E-07	4.02	2.64E-07	3.98	3.58E-07	3.97	1.97E-07	4.03	2.52E-07	4.01
	40×40	2.07E-08	3.96	2.53E-08	4.03	3.05E-08	4.07	1.25E-08	4.00	1.71E-08	3.95	2.37E-08	3.92	1.21E-08	4.02	1.61E-08	3.97
4	5×5	5.37E-06	-	6.98E-06	-	8.28E-06	-	3.29E-06	-	4.33E-06	-	5.45E-06	-	3.34E-06	-	4.12E-06	-
	10×10	1.64E-07	5.04	2.21E-07	4.98	2.65E-07	4.97	1.01E-07	5.03	1.31E-07	5.05	1.62E-07	5.07	1.02E-07	5.03	1.20E-07	5.10
	20×20	5.16E-09	4.99	6.75E-09	5.04	7.96E-09	5.05	3.14E-09	5.01	4.22E-09	4.95	5.31E-09	4.93	3.14E-09	5.02	3.75E-09	5.00
	40×40	1.64E-10	4.98	2.07E-10	5.03	2.40E-10	5.05	9.81E-11	5.00	1.36E-10	4.95	1.74E-10	4.93	9.76E-11	5.01	1.19E-10	4.98

Table A.6. L_2 -norm of the post-processed solution error for the unsteady linear diffusion problem

p	Mesh size	FR						EFR					
		c_{DG}	Order	c_{SD}	Order	c_{HU}	Order	c_{DG}	Order	c_{SD}	Order	c_{HU}	Order
1	5×5	1.56E-02	-	1.96E-02	-	4.86E-02	-	1.14E-02	-	1.85E-02	-	5.08E-02	-
	10×10	3.31E-03	2.24	5.82E-03	1.75	1.99E-02	1.29	1.62E-03	2.82	5.33E-03	1.80	1.97E-02	1.36
	20×20	7.11E-04	2.22	1.45E-03	2.01	5.58E-03	1.83	2.08E-04	2.96	1.49E-03	1.84	5.88E-03	1.75
	40×40	1.58E-04	2.17	3.75E-04	1.95	1.49E-03	1.90	2.62E-05	2.99	3.96E-04	1.91	1.58E-03	1.89
2	5×5	9.25E-04	-	1.36E-03	-	2.17E-03	-	3.18E-04	-	7.83E-04	-	1.56E-03	-
	10×10	8.51E-05	3.44	1.23E-04	3.46	1.87E-04	3.53	2.42E-05	3.72	5.85E-05	3.74	1.14E-04	3.77
	20×20	9.44E-06	3.17	1.25E-05	3.30	1.71E-05	3.45	1.66E-06	3.86	3.97E-06	3.88	7.68E-06	3.90
	40×40	1.13E-06	3.06	1.43E-06	3.13	1.83E-06	3.23	1.08E-07	3.94	2.58E-07	3.95	4.96E-07	3.95
3	5×5	3.70E-05	-	6.32E-05	-	8.51E-05	-	1.03E-05	-	2.85E-05	-	4.13E-05	-
	10×10	2.70E-06	3.77	4.52E-06	3.80	5.88E-06	3.86	2.46E-07	5.39	7.15E-07	5.32	1.00E-06	5.36
	20×20	1.48E-07	4.19	2.26E-07	4.33	2.84E-07	4.37	5.96E-09	5.37	1.83E-08	5.29	2.40E-08	5.38
	40×40	1.10E-08	3.75	1.47E-08	3.94	1.75E-08	4.02	1.65E-10	5.18	5.28E-10	5.12	6.69E-10	5.17
4	5×5	3.50E-06	-	5.31E-06	-	6.25E-06	-	7.69E-07	-	1.58E-06	-	1.83E-06	-
	10×10	8.29E-08	5.40	1.18E-07	5.49	1.36E-07	5.53	1.57E-08	5.61	3.04E-08	5.70	3.54E-08	5.70
	20×20	2.52E-09	5.04	3.37E-09	5.13	3.79E-09	5.16	2.66E-10	5.89	5.20E-10	5.87	6.17E-10	5.84
	40×40	8.13E-11	4.96	1.05E-10	5.01	1.17E-10	5.01	4.26E-12	5.97	8.41E-12	5.95	1.01E-11	5.93

Table A.7. L_2 norm of the temperature error for the planar Couette flow case

p	Mesh size	FR			HFR			EFR						
		c_{DG}	Order	c_{SD}	c_{HU}	Order	c_{SD}	c_{DG}	Order	c_{HU}	Order			
1	4×2	2.02E-2	-	3.02E-2	7.43E-2	-	1.78E-2	1.49E-2	-	2.00E-2	1.51E-2	2.09E-2	-	
	8×4	5.22E-3	1.95	6.69E-3	1.36E-2	2.45	4.82E-3	1.91	5.79E-3	1.79	3.55E-3	2.07	4.00E-3	1.92
	16×8	1.26E-3	2.05	1.47E-3	2.71E-3	2.33	1.10E-3	2.14	1.06E-3	2.00	1.45E-3	1.97	1.04E-3	1.94
2	32×16	3.11E-4	2.02	3.59E-4	6.12E-4	2.15	2.42E-4	2.18	2.34E-4	1.91	3.86E-4	1.98	2.56E-4	2.03
	4×2	1.70E-3	-	2.56E-3	4.01E-3	-	1.50E-3	1.53E-3	-	1.76E-3	1.55E-3	1.81E-3	-	
	8×4	2.20E-4	2.95	2.95E-4	4.33E-4	3.21	1.65E-4	2.90	2.07E-4	2.84	2.00E-4	2.93	2.13E-4	2.86
3	16×8	2.77E-5	2.99	3.57E-5	5.01E-5	3.11	2.34E-5	3.03	2.54E-5	2.94	2.46E-5	3.03	2.65E-5	3.00
	32×16	3.55E-6	2.97	4.69E-6	6.56E-6	2.93	2.75E-6	3.09	3.09E-6	2.92	2.92E-6	3.07	3.14E-6	3.08
	4×2	1.55E-4	-	2.33E-4	3.14E-4	-	1.44E-4	1.35E-4	-	1.61E-4	1.45E-4	1.64E-4	-	
4	8×4	1.08E-5	3.84	1.56E-5	2.02E-5	3.96	9.12E-6	3.88	1.03E-5	3.73	9.26E-6	3.87	1.05E-5	3.78
	16×8	7.13E-7	3.92	9.92E-7	1.28E-6	3.98	5.74E-7	3.99	6.83E-7	3.92	8.39E-7	3.98	6.93E-7	3.92
	32×16	4.62E-8	3.95	6.55E-8	8.43E-8	3.92	3.48E-8	4.04	4.30E-8	3.99	5.50E-8	4.04	4.26E-8	4.02
4	4×2	1.47E-5	-	2.21E-5	2.74E-5	-	1.29E-5	1.29E-5	-	1.57E-5	1.44E-5	1.60E-5	-	
	8×4	5.54E-7	4.73	8.24E-7	1.01E-6	4.76	4.75E-7	4.76	5.74E-7	4.63	6.61E-7	4.75	5.82E-7	4.62
	16×8	1.89E-8	4.88	2.82E-8	3.46E-8	4.87	1.55E-8	4.94	1.98E-8	4.86	2.35E-8	4.93	2.02E-8	4.85
32×16	6.17E-10	4.93	9.34E-10	1.15E-9	4.91	4.75E-10	5.03	6.22E-10	4.99	7.58E-10	5.02	6.34E-10	4.99	

Table A.8. L_2 norm of the isentropic vortex for FR, HFR, and EFR with c_{DG} , c_{SD} , c_{HU} for degrees $p = 1$ to $p = 5$

p	Mesh size	FR			HFR			EFR			
		c_{DG}	Order	c_{SD}	c_{HU}	Order	c_{DG}	Order	c_{SD}	c_{HU}	Order
1	5×5	2.07E-1	-	2.24E-1	-	2.07E-1	-	2.12E-1	-	2.27E-1	-
	10×10	1.71E-1	0.280	1.95E-1	0.105	1.71E-1	0.278	1.89E-1	0.165	1.98E-1	0.100
	20×20	4.57E-2	1.901	1.13E-1	0.790	1.97E-1	0.789	1.97E-1	0.724	1.66E-1	0.255
	40×40	8.24E-3	2.472	3.49E-2	1.693	9.99E-2	1.693	9.99E-2	1.879	5.86E-2	1.505
2	5×5	1.41E-1	-	1.88E-1	-	1.41E-1	-	1.37E-1	-	1.96E-1	-
	10×10	2.43E-2	2.537	4.17E-2	1.287	2.39E-2	2.559	4.17E-2	2.174	8.51E-2	1.287
	20×20	5.20E-3	2.223	9.51E-3	2.318	5.20E-3	2.199	9.53E-3	2.129	1.71E-2	2.316
	40×40	3.89E-4	3.742	6.30E-4	3.917	1.14E-3	3.911	1.14E-3	3.905	3.92E-4	3.846
3	5×5	6.24E-2	-	9.15E-2	-	6.25E-2	-	9.20E-2	-	5.73E-2	-
	10×10	1.05E-2	2.567	1.34E-2	2.772	1.05E-2	2.569	1.34E-2	2.780	1.71E-2	2.950
	20×20	3.43E-4	4.943	6.80E-4	4.181	3.43E-4	4.944	6.81E-4	4.298	9.47E-4	4.178
	40×40	1.63E-5	4.394	5.91E-5	3.444	1.63E-5	4.395	5.90E-5	3.528	8.68E-5	3.448
4	5×5	1.80E-2	-	5.28E-2	-	1.80E-2	-	5.25E-2	-	7.22E-2	-
	10×10	2.23E-3	3.012	1.96E-3	4.751	2.22E-3	3.017	1.96E-3	4.746	2.39E-3	4.916
	20×20	2.31E-5	6.589	8.78E-5	4.481	2.31E-5	6.588	8.76E-5	4.482	1.28E-4	4.219
	40×40	8.59E-7	4.751	3.80E-6	4.529	9.13E-7	4.660	3.98E-6	4.460	5.47E-6	4.552
5	5×5	5.10E-2	-	4.49E-2	-	5.11E-2	-	4.50E-2	-	5.07E-2	-
	10×10	2.76E-4	7.527	5.28E-4	6.411	2.76E-4	7.531	5.28E-4	6.413	6.17E-4	6.220
	20×20	5.83E-6	5.567	1.17E-5	5.455	5.83E-6	5.567	1.17E-5	5.499	1.41E-5	5.455
	40×40	8.89E-8	6.036	3.04E-7	5.265	8.89E-8	6.036	3.04E-7	5.265	3.92E-7	5.166

Appendix B

Gauss-Lobatto Collocation in EFR

In Chapter 5, we considered spectral analysis of the HFR and EFR methods, where the latter uses a GLL collocation to reduce the element coupling in the semidiscrete operator. This common integration approach in some finite element methods leads to the so-called *mass lumping*. For example, the discontinuous Galerkin spectral-element method (DGSEM) is generally formulated using collocation-type formulations, which leads to diagonal local mass matrices of the form

$$\mathbf{M}_{ij} = \int_{\Omega} \phi_i \phi_j d\Omega = \delta_{ij} w_i, \quad (\text{B.1})$$

where w_i is an integration quadrature weight and δ_{ij} is the Kronecker delta. When used with GLL points, the schemes are shown to be less stiff, allowing for larger time step sizes than when Gauss-Legendre points are used. Because the integrand is of degree $2p$ and Gauss-Lobatto quadratures are accurate up to $2p - 1$, this results in under-integration. Previous analyses have shown that this approach leads to filtered forms of the nodal DG method with exact integration [169] and has been used to introduce artificial dissipation in the presence of shocks [170]. A similar connection to FR methods using the c_{HU} or correction parameter (also known as the g_2 correction function) has also been investigated [171, 172], showing an equivalence between this FR method and the DGSEM with GLL points. The dissipation and dispersion error can be affected by this choice since the filter acts by damping the highest frequency modes [169]. In our EFR method, choosing GLL points on the faces allows for better integration properties than equidistant points while maintaining nodes at the face ends. When collocated, this leads to under-integrated local mass matrices over the faces, which are eventually assembled into a global operator. Note that the standard FR method does not explicitly need quadratures for integration due to its divergence form, so this only affects the integration of the transmission conditions in the hybridized formulation. Furthermore, because the HFR method with discontinuous traces is used, conservation is pointwise. Hence, this discussion is only applicable to the EFR method.

Our spectral analysis can be directly extended to the EFR method with exact integration. In this case, the coupling between interior elements in the semidiscrete interior-solution-based operator is stronger, and additional blocks must be added to perform the analysis. In addition, we must employ large domains to eliminate the effects of the periodic boundaries due to the increased coupling. Note that the semidiscrete operator in the analysis in Chapter 5 is not computed when solving the hybridized problems but is only used to analyze the properties of the method. For completeness, we include additional dispersion and dissipation curves in this section. After doing this, the resulting matrix will also be block-circulant and can be obtained via Equation (5.9). However, the number of neighbouring blocks N_b is higher, but the influence of elements reduces away from the element in study. Similarly, we show the equivalent dispersion and dissipation properties of the EFR method with exact integration of the transmission conditions in Figure B.1 and B.2. When compared against the figures in Chapter 5, it can be seen that the EFR method with exact integration improves the numerical error at low wavenumbers, especially the dissipation at $p = 1$. However, it is still more dissipative than the HFR/FR method for the majority of considered correction parameters. As we go to higher orders, results with and without exact integration start to converge, with the former showing more agreement with the HFR methods.







Figure B.1. Dissipation curves for different c -parameters at $\theta = \pi/6$ for different polynomial degrees with exact integration of the transmission conditions. Results are highlighted for c_{DG} (—), c_{SD} (---) and c_{HU} (···)




(a) $p = 1$



(b) $p = 2$



(c) $p = 3$



(d) $p = 4$

Figure B.2. Dispersion curves for different c -parameters at $\theta = \pi/6$ for different polynomial degrees with exact integration of the transmission conditions. Results are highlighted for c_{DG} (—), c_{SD} (---) and c_{HU} (···)

Appendix C

HFR Parallelization

High-order methods are suitable for modern many-core computer architectures due to the locality of the approaches when using explicit time stepping. The FR approach has demonstrated performance of over 50% in CPUs and GPUs in these configurations. Implicit time-stepping does not typically achieve such high peak FLOPs due to increased communication costs involved in the linear solvers [173], which require significant memory bandwidth. To parallelize high-order methods, graph partitioning algorithms such as METIS [174] are typically employed to distribute the elements across ranks. Commonly, this partitioning seeks to reduce the communication between elements in different regions in the domain by minimizing the number of straddling edges and their associated weights. Partitioning based on element connectivity is the standard approach employed in explicit and implicit FR methods. In hybridization, the solution of the implicit system depends on the traces rather than on the interior values. However, operations typical of the FR method, such as right-hand-side evaluations, are still required, and communication between elements must be minimized. The conventional elementwise FR partitioning is also suitable for hybridization but requires additional steps to distribute the trace over the ranks as a preprocessing step [113, 175]. Hence, a simple algorithm to perform the distribution can be implemented, which is only required in the initial stage of the solver and when recomputations of the domain connectivity are required due to adaptation calls. In the next paragraph, we describe our algorithm.

As described in 5, we subdivide a domain Ω into multiple partitions Ω^i . Then, we loop over all trace variables and assign them based on a cost scalar value c^i , which measures the associated overhead with that rank. We seek to minimize this cost value by distributing the traces such that both requirements of off-trace components and the number of trace variables within each rank are balanced. Our algorithm assumes that the ranks are relatively balanced from the initial partitioning and starts by assuming a high cost where all trace variables lie on off-rank processors. The number of all unassigned trace variables is initially stored in

N_{offp} and N_{rem} at every rank. These variables are, respectively, the number of off-rank traces and a counter containing the local traces that have not been assigned yet. In addition, we keep the number of MPI interface traces assigned to a given rank in N_{MPI} .

After looping over every trace, the trace is distributed to reduce the need to send/receive calls with the trace information while maintaining a balanced number of trace points across the partitions. The algorithm's success depends on the proper partitioning of the elements. We then solve the implicit system and perform scatter trace operations between partitions using the Portable, Extensible Toolkit for Scientific Computation (PETSc) [112]. Profiling results from PETSc demonstrate balancing factors between partitions ranging from 1.0 and 1.05, which we deem acceptable to perform our parallel simulations. Note that this approach is suitable for both HFR and EFR implementations.

Algorithm 5: Trace distribution after element-based partitioning

Partition Ω into n subdomains Ω^i using a k -way edge-cut minimization algorithm

Set \hat{N}^i to the total number of traces in Ω^i

Store all \hat{N}_{int}^i trace points corresponding to interior faces in $\mathcal{D}^i = \{\hat{u}_0^i, \dots, \hat{u}_{\hat{N}_{\text{int}}^i}^i\}$

$N_{\text{offp}}^i \leftarrow$ Number of traces in Ω^i not yet assigned

$N_{\text{rem}}^i \leftarrow$ Number of traces in Ω^i not yet assigned

$N_{\text{MPI}}^i \leftarrow 0$

$c^i \leftarrow N_{\text{offp}}^i - N_{\text{rem}}^i - N_{\text{MPI}}^i - N_{\text{int}}^i$

for Every Trace u_j in Ω **do**

 Check which ranks $\hat{\Omega}$ have u_j

 Compare cost functions c in $\hat{\Omega}$

 Assign u_j to rank $\Omega^m \in \hat{\Omega}$, where $m \leftarrow \text{argmax}\{c^0, \dots, c^n\}$

for $\Omega^j \in \hat{\Omega}$ **do**

if $\Omega^j = \Omega^m$ **then**

$N_{\text{offp}}^j \leftarrow N_{\text{offp}}^j - 1$

$N_{\text{MPI}}^j \leftarrow N_{\text{MPI}}^j - 1$

end

$N_{\text{rem}}^j \leftarrow N_{\text{rem}}^j - 1$

end

 Recompute c^i for all ranks in $\hat{\Omega}$

end

Appendix D

IMEX Butcher Tableaus

We present coefficients of the Butcher tableaus corresponding to the implicit and explicit components of the IMEX_{3,2} and IMEX_{5,4} methods, as employed in Chapter 7. The AIMEX_{10,2} method is defined for each solution polynomial degree and omitted for brevity. The coefficients of the corresponding tableaus, however, can be found in their original publication as supplementary material [48].

0	0	0	0
0.2928932188	0	0.2928932188134520	0
1	0	0.7071067811865480	0.2928932188134520
	0	0.7071067811865480	0.2928932188134520

(a) Implicit part

0	0	0	0
0.2928932188	0.2928932188134520	0	0
1	-0.9428090415820630	1.9428090415820640	0
	0	0.7071067811865480	0.2928932188134520

(b) Explicit part

Table D.1. IMEX_{3,2} Butcher tableau

$\frac{1}{2}$	0	$\frac{1}{2}$	0	0	0
$\frac{2}{3}$	0	$\frac{1}{6}$	$\frac{1}{2}$	0	0
$\frac{1}{2}$	0	$-\frac{1}{2}$	$\frac{1}{2}$	$\frac{1}{2}$	0
1	0	$\frac{3}{2}$	$-\frac{3}{2}$	$\frac{1}{2}$	$\frac{1}{2}$
	0	$\frac{3}{2}$	$-\frac{3}{2}$	$\frac{1}{2}$	$\frac{1}{2}$

(a) Implicit part

0	0	0	0	0	0
$\frac{1}{2}$	$\frac{1}{2}$	0	0	0	0
$\frac{2}{3}$	$\frac{11}{18}$	$\frac{1}{18}$	0	0	0
$\frac{1}{2}$	$\frac{5}{6}$	$-\frac{5}{6}$	$\frac{1}{2}$	0	0
1	$\frac{1}{4}$	$\frac{7}{4}$	$\frac{3}{4}$	$-\frac{7}{4}$	0
	$\frac{1}{4}$	$\frac{7}{4}$	$\frac{4}{3}$	$-\frac{7}{4}$	0

(b) Explicit part

Table D.2. IMEX_{5,3} Butcher tableau