UNIVERSITY OF CALIFORNIA
Los Angeles

The Effects of Numerical Scheme Resolvability for Large-Eddy Simulations

A dissertation submitted in partial satisfaction
of the requirements for the degree
Doctor of Philosophy in Aerospace Engineering

by

Ayaboe Edoh

2017

ABSTRACT OF THE DISSERTATION

The Effects of Numerical Scheme Resolvability for Large-Eddy Simulations

by

Ayaboe Edoh
Doctor of Philosophy in Aerospace Engineering
University of California, Los Angeles, 2017
Professor Ann R. Karagozian, Chair

The dissertation addresses the formulation of Large-Eddy Simulations (LES) with direct consideration of a base finite difference scheme, and with the intent of reducing numerical error influences on the closure model and ultimately the solution. As such, spectral characteristics of the explicitly-defined LES filter are considered with respect to the discretization method's spectral accuracy (i.e., resolvability). Analysis and development of discrete filtering stencils is undertaken, placing emphasis on the ability to specify desired scale-separation (e.g., cut-off wavenumber and scale-discriminant attenuation) relative to the computational grid. Assessment of the LES procedure is preceded by the establishment of a suitable base scheme, comprised of high-order discretizations and the addition of stabilization presented in a filter-based artificial dissipation form. Subsequent robustness and preservation of the overall solution accuracy is achieved by tuning the dissipation according to the dispersion characteristics of the underlying numerical method and seeking to deliberately remove the effects of discretization error. Extension to LES is then established by properly defining the explicit filter in relation to these numerical characteristics. Effectiveness of the procedure is evaluated by means of *a priori* and *a posteriori* inspection of turbulence calculations for the Burgers and Navier-Stokes equations, wherein the impacts of discretization and filter cut-off are assessed in light of scale-similarity and "perfect" (i.e., a DNS-

based closure) modeling. Results demonstrate the benefits of employing mutually tuned high-order discretizations and filters in the limit of the idealized "perfect" model, yet highlight the likely possibility of modeling error overshadowing such gains when actual closures, such as scale-similarity models, are used. In an attempt to enhance the scale-similarity models considered herein, the filter-based artificial dissipation is employed in order to enforce the prescribed LES field, and is shown to reduce overall model error. Meanwhile, its use is shown to be beneficial even in the presence of "perfect" modeling, wherein the dissipation can be tuned to specifically target discretization errors.

The dissertation of Ayaboe Edoh is approved.

Venkateswaran Sankaran

Christopher R. Anderson

Xiaolin Zhong

John J. Kim

Ann R. Karagozian, Committee Chair

University of California, Los Angeles

2017

Like many things, perseverance, too, must be practiced.

Table of Contents

List of Figures

List of Tables

Nomenclature

E_j	Inviscid flux vector
$G(x)$	Physical space representation of the filter kernel
P	Pressure
Q	Solution vector
R	Gas constant
R_o	Original residual to the governing equation
R_{AD}	Residual component corresponding to artificial dissipation
T	Temperature
V_j	Viscous flux vector
Δt	Time-step size
Δx	Mesh size
Δ	Filter width
δ_x^{2r}	Narrow dissipative stencil operator approximating ∂_x^{2r} to second order
γ	Ratio of specific heat coefficients, c_p/c_v
$\hat{\mathcal{G}}(k)$	Spectral response of an operator $\mathcal{G}$
λ	System eigenvalue
ϕ	Phase angle
ρ	Density
θ	Normalized wavenumber ($\theta = k\Delta x$)

c_p Specific heat coefficient at constant pressure

c_v Specific heat coefficient at constant volume

e_o Specific total energy

k Spatial wavenumber

k_c Cut-off wavenumber

$k_{1/2}$ Cut-off wavenumber such that $\hat{\mathcal{G}}(k_{1/2}) = 0.5$

k_{mod} Modified wavenumber of spatial scheme

u_j Velocity component

AD Artificial Dissipation

ADM Approximate Deconvolution Method

CD Central Difference discretization scheme

CFL Courant-Friedrichs-Lewy condition number

DNS Direct Numerical Simulations

DOF Degrees of freedom

DRP Dispersion Relation Preserving

fDNS Filtered DNS

FTGR Filter-to-grid ratio

LES Large-Eddy Simulations

NSE Navier-Stokes Equations

PPW Points-per-wave

RANS Reynolds-Averaged Navier-Stokes

Re Reynolds number

SF Solution-Filtering

SSM Scale-Similarity Model

TGV Taylor-Green Vortex

VNA Von Neuman Analysis

Acknowledgments

The culmination of this doctoral journey cannot forgo paying proper homage to the many friendships, relationships, and mentorships that I have been fortunate enough to hold through the years. Theses have served as a continuous source of support and motivation: whether by means of every-day interactions (no matter how seemingly menial) or by way of long-distance. In this regard, I am immensely fortunate and attribute my present circumstances to all who have nurtured me. Here, I wish to briefly offer acknowledgements and express a fraction of my full gratitude to some of these individuals (although in no way an exhaustive list!).

First, I would like to thank my family, who never let me forget that I was the youngest (not that I mind the perks) yet never seemed to doubt that I could fit into the large boots laid before me. À Maman (Emilienne Noellie Folly), tu as tout sacrifié pour nous tes enfants, et je te remercie infiniment. À travers toi, j'ai appris ce que c'est que de travailler avec conviction et détermination. Tu m'as toujours dit que c'est en travaillant avec joie–même quand la tâche est difficile–qu'on cultive un sens de fierté et de sagesse. J'espère pouvoir un jour atteindre une fraction de ta générosité et de ton humilité. À Papa (Koffi Otto Edoh), tu as toujours souhaité notre réussite scolaire et j'aurais tellement aimé que tu sois en vie pour que nous puissions partager cette expérience. Sachant que tu as fait ce même chemin dans ta jeunesse, j'espère que tu es fier du travail que j'ai accompli. To my younger-older sister, Melissa, you have served as my role model and compass. You are brilliant both intellectually and empathically. And while you were my reference for academic achievement, I thank you for always pushing me to nurture my other interests and appreciate the simple humanities of life. To my eldest-older sister, Nadia, thank you for your positivity and cheer through these years and for letting me be un tonton to your kids; I hope to set a good example for them. To my older brother, Henri, I appreciate the timely

words of encouragement. To my beautiful and soon-to-be partner, Stephanie, I am eternally grateful for your unending support and sacrifices, especially in these final months of writing where you had to deal with what must have seemed like a shell or shadow of my normal self. Thank you for weathering ground zero with selfless-ness and for continually pushing me forward. To RPK (Enrique, Toni, Ahmed, Alex), Croom, Lynne, Kamo, Jason, Bass, Silverio, Joel, members of Nu Delta and the many others not explicitly listed: I hope my sincerest gratitude needs no words.

Fortunately, I have also been blessed to have mirroring support structures in my academic realm. Thank you to Mr. Healey (Walter Johnson HS), to Professors Marzouk and Professor Balakrishnan (MIT), to Veronica Garcia and Ron Neumeyer (Rocketdyne) – champions at different chapters in this academic journey that each played a role in my eventual deliverance into the watchful care of my current academic parents: Professor Karagozian and Venke Sankaran. To Professor Karagozian, I am deeply thankful for your sincere kindness and the faith you have shown in my abilities. Your words of encouragement have always served as a needed boost and as fuel to work harder and to be worthy of being your student. To Venke, who entered my life amidst a time of uncertainty and who is unknowingly responsible for having kept me on this path, I thank you for your openness. I hold strong to your emphasis in finding the lingering profundities in often-overlooked nuances, and I hope to have it guide me in my quest to produce meaningful impact. To Nathan, who served as an academic older brother of sorts and now a friend: as tempting as it was to ask you for the answers, thank you for allowing me to step into myself. To members of my cohort (Chris, Hai, Ali, Richard, and Kyle), who I view as cousins on this track: thank you for providing aid in times of need (including reading and providing input for much of this dissertation!) and I hope we reconvene and collaborate to do something great one day. I would also like to extend my gratitude to the greater Air Force Research

Lab (AFRL) team (Josiah, Eder, Jun, Matt, Zoli, Michael, Chris, Clancy, Bill and others) and to the UCLA staff (Abel, Angie, Annie, and Evgenia).

Graciously, financial trepidation over the years has been minimal, and for this, I would like to acknowledge the various funding sources over the course of my career: the Gates Millennium Scholarship (GMS); the Graduate Opportunity and Cota-Robles Fellowships from UCLA; the Air Force Office of Scientific Research (AFOSR) under the directives of Drs. Fariba Fahroo, Chiping Li and Michael Kendra; as well as ERC Inc. (subcontracts RS130039, RS131108, PS150006). By these various means, I have been offered the precious luxuries of choice and financial peace-of-mind, essential to my studies.

I suppose, in short, one can simply state that "it takes a village". The impact that one's environment can have on the life course should never be underestimated. And in this way, I am fortunate to have been surrounded by positivity and encouragement. I hope that society can move to acknowledge the potential in all of its members and to embrace the immense richness in its diversity. I thank those before me who fought both the silent and historic battles, and who made this track a possibility. In the same token, my thoughts are with those that still remain locked out and overlooked. May we strive for true equity.

Vita

2012	Masters (Aerospace Engineering), UCLA, Los Angeles, California.
2010	B.S. (Course XVI: Aerospace Engineering), M.I.T. , Cambridge, Massachusetts.
2012-2016	Cotal-Robles Fellowship (UCLA)
2011-2012	Graduate Opportunity Fellowship (UCLA)
2006-2017	Gates Millennium Scholarship (The Gates Foundation)
2010-2011	Associate Engineer, Systems Performance, Pratt and Whitney Rocketdyne, Canoga Park, California

Publications and Presentations

A. K. Edoh and A. R. Karagozian, *Inspecting Interactions of Discretization, Filter Formulation, and Stabilization in LES: Lessons from the Taylor-Green Vortex* (AIAA 2017-3952), 23rd AIAA Computational Fluid Dynamics Conference. Denver, CO. June 2017.

A. K. Edoh and A. R. Karagozian, *Stabilized Scale-Similarity Modeling for Explicitly-Filtered LES* (AIAA 2017-1227), 55th AIAA Aerospace Sciences Meeting Grapevine, TX, January 2017.

A. K. Edoh, N. L. Mundis, A. R. Karagozian, V. Sankaran. *Discrete Filtering Formulations for Large-Eddy Simulations* (AIAA 2016-3794), 46th AIAA Fluid Dynamics Conference. Washington, DC. June 2016.

A. K. Edoh, N. L. Mundis, A. R. Karagozian, V. Sankaran. *The Role of Dispersion and Dissipation on Stabilization Strategies for Time Accurate Simulations* (AIAA 2016-0071), 54th AIAA Aerospace Sciences Meeting. San Diego, CA. January 2016.

N. L. Mundis, A. K. Edoh, V. Sankaran. *Highly-Accurate Filter-Based Artificial-Dissipation Schemes for Stiff Unsteady Fluid Systems* (AIAA 2016-0072), 54th AIAA Aerospace Sciences Meeting. San Diego, CA. January 2016.

N. L. Mundis, A. K. Edoh, V. Sankaran. *Optimal Runge-Kutta Schemes for High-order Spatial and Temporal Discretizations* (AIAA 2015-2752), 22nd AIAA Computational Fluid Dynamics Conference. Dallas, Texas. June 2015.

A. K. Edoh, A. R. Karagozian, C. L. Merkle, V. Sankaran. *Comparison of Artificial Dissipation and Filtering Schemes for Time-Accurate Simulations* (AIAA 2015-0284), 53rd AIAA Aerospace Sciences Meeting. Kissimmee, Florida. January 2015.

A. K. Edoh, A. R. Karagozian, C. L. Merkle, V. Sankaran. *Optimal Numerical Schemes for Time Accurate Compressible Large Eddy Simulations: Comparison of Artificial Dissipation and Filtering Schemes*, 67th Annual Meeting of the American Physical Society Division of Fluid Dynamics, San Francisco, CA, November, 2014.

CHAPTER 1

Large-Eddy Simulations

The need to study physical phenomena in engineering systems has seen the increased use of numerical simulations as a complement to experimental efforts. The ability to specify test conditions and interrogate the flow-field permits one to bypass diagnostic and facility challenges that would otherwise hinder investigating some of the taxing regimes considered in engineering efforts. These include areas such as turbulence and reactive flow, along with the strong instabilities that can arise from their coupling (Harvazinski et al., 2015).

Advancements in computing power and resources have allowed researchers to make fundamental insights into these crucial areas by means of fully-resolved computations, (i.e.: Direct Numerical Simulations, DNS) (Kim et al., 1987; Poludnenko and Oran, 2010; Lagha et al., 2011). However, the large disparity in scales that exist even in moderate- to high-Reynolds Number (Re) flows typically relegates these efforts to simplified canonical configurations. The need to provide both fundamental and predictive information for the design of practical systems at high Reynolds numbers – which introduce complexities pertaining to both geometric and physical elements – has motivated the development of under-resolved formulations that are capable of approximating the full flow field. Part of this class of methods is the Large-Eddy Simulations (LES) methodology, in which one designates the scales to be computed and seeks to model influences of the neglected portion of the field through so-called closure terms to the governing equations. In most cases, one chooses to remove the small-scales from the calculation, believ-

ing these to exhibit universal characteristics, as in the case of turbulence. The premise then suggests the existence of tractable and relatively general model closure forms. In essence, the intention of LES constitutes an indirect calculation to the full problem, designed to render the computational workload manageable.

While the formulation for LES and its subsequent governing equations is typically associated with a low-pass filtering of spatial scales, it may also include temporal filtering. In this way, one may recover the Reynolds-Averaged Navier-Stokes (RANS) representation by considering a sufficiently large temporal filter width (Pruett, 2000; Pruett et al., 2003). RANS has proven to be a computationally efficient tool for engineering use (Hirsch and Tartinville, 2009; Corson et al., 2009); however, the fact that the equations implicitly solve for mean quantities makes it such that one may not be able to capture important time-sensitive phenomena. These shortcomings are generally explained by the fact that RANS implicitly models all scales, which may obfuscate important dynamics. As a response to this, one may thus consider reducing the temporal filter width, eventually recovering an Unsteady-RANS formulation (Speziale, 1998). Again, the utility of such a formulation needs to be judged relative to the dynamics one hopes to capture. In classic notions of turbulence where the spatial stiffness may dominate its temporal counterpart, the more traditional tendency of filtering relative to the spatial dynamics in LES may play the implicit role of employing the model contributions conservatively[1]. Nevertheless, the ability to specify the resolved versus modeled components – either spatially, temporally, or both – renders the LES formalism an attractive option for investigating a broad range of problems.

The various performances achievable by LES are largely anticipated by the as-

[1]The spatio-temporal nature of the fluid equations results in frequency and wavenumber cut-offs being related; in this way, filtering in space implies a temporal filtering (Garnier et al., 2009). In general, numerical restrictions on time-step sizes make it such that one is resolved well-enough relative to physical time-scales, thus corresponding to filter-to-grid ratios above unity in time.

sumptions employed in deriving the governing equations (e.g., temporal filtering with long-time averaging in order to recover RANS-like behavior). The consequences imposed by the underlying assumptions can then be seen to reside within the influence of the model terms. In this way, proper representation of the closures is paramount to achieving the intended predictive capabilities of the method. This has naturally spurred a deep interest and a need to develop suitable models, either by phenomenological (i.e.: functional) approaches or from a symmetry-preserving (i.e.: structural) standpoint (Sagaut, 2006). Also increasingly recognized, however, is the interplay of the different sources of error within LES, namely: 1) projection error as it pertains to representing the continuous field in a discrete manner (e.g., nodal schemes); 2) model error as it relates to estimating the behavior of the closure terms; and 3) discretization error as it affects calculation of the resolved scales. As a result, the continual advancement of LES for accurate and predictive calculations calls for a somewhat holistic perspective in the formulation of the overall approach. The current research effort thus seeks to investigate some of the inter-related consequence of these elements, notably focusing on the role that numerical discretization error plays in affecting the accuracy of the overall LES formulation.

1.1 Consequences of Numerical Error

As noted above, the suitability of the closure model in LES is fundamental to retrieving the expected dynamical predictions. However, another significant consideration is the influence of numerical error on the computational results. This is particularly true when considering "coarse-grid" LES, in which case even the smallest resolved scales may exhibit non-negligible dynamical energy. These high, grid-relative wavenumbers can become contaminated by discretization error. When considered in the absence of sufficiently well-resolved diffusive effects,

the lingering inaccuracies may then affect the model performance, which typically relies on resolved-scale information.

The work of Ghosal (1996) performs *a priori* type analysis in order to estimate the dynamical contributions of numerical error (i.e.: truncation and aliasing) relative to model terms derived from statistical closure theory. The study shows the latter to be overwhelmed in typical LES implementations that presume the filter width to be identical to the mesh grid (i.e.: grid- or implicit-filtering). Ghosal furthermore notes that low-order methods generally produce large truncation error but low aliasing error; meanwhile, the converse is highlighted to be the case for high-order discretizations. Kravchenko and Moin (1997) perform *a posteriori* computations that confirm the analysis of Ghosal, furthermore noting the impact of different discrete derivative representations (e.g, divergence versus skew-symmetric, rotational, or advection forms of the convective operator) on the non-linear stability of the methods (Blaisdell et al., 1996). Moreover, the linear analysis of Geurts and van der Bos (2005a) interprets the influence of discretization error as inducing high-pass effects on the equations, which alludes to the well-established notion that the numerical method inherently filters the LES solution.

Notable, particularly with respect to efforts in simulating turbulent reactive flow, is evidence of heightened sensitivity of the LES solution to the numerical scheme. Cocks et al. (2015) compare the performance of different algorithms, each employing identical closure models for the simulation of a bluff-body stabilized flame. In the case of non-reactive flow, given sufficient resolution, the codes are shown to be in agreement with each other and with the experimental data. When reactions are enabled, however, each implementation is seen to be qualitatively different. Furthermore, none are seen to match the reference data. This perceived increase in complexity is somewhat surprising since the span of turbulent scales is typically presumed to shrink due to the lower Reynolds number of reactive flow.

Nevertheless, the test case underscores the potential for numerical errors to be increasingly consequential in reacting LES.

A proposed response to the issues of discretization error has been the adoption of explicitly-filtered LES formulations, which consider the filter width to be larger than the mesh grid support. Analyses confirm that the dynamical influence of numerical errors is successfully diminished relative to model contributions. Such an approach, however, is not popularly used due to perceived inefficiencies when considering the computational degrees of freedom (DOF) required to resolve a given filter width. For example, a filter-to-grid ratio (FTGR) of two corresponds to a factor of sixteen in computational overhead (e.g., 2^4, considering a three-dimensional domain and the time-step being proportionally scaled according to a CFL ~ 1 condition). In some instances, explicit-filtering has been shown to be outperformed by implicit-filtering, likely owing to a fortuitous cancellation of numerical and modeling errors in the latter implementation (Lund, 2003). Grid-filtering formulations do not, however, provide a reliable means by which one can assess the overall algorithm accuracy, both in terms of the numerics and the modeling. This is in contrast with the explicitly-filtered context which allows for grid-converged LES (Bose et al., 2010; Radhakrishnan and Bellan, 2012). In such a case, one chooses to keep the filter-width constant while decreasing the grid resolution, thus making the numerically-error-prone modes less relevant[2].

While advancement of LES is directly tied to model development, the impact of taking the numerical treatment into consideration is two-fold: 1) minimizing the influence of discretization error on the resolved scale, which would naturally improve the LES solution; and 2) from the standpoint of model development,

[2]Fixing the LES filter while reducing the grid leads to grid-converged LES; the anticipated model contribution does not change yet succumbs to less numerical artifacts. On the other hand, if the filter width is reduced in tandem with the grid (as for implicitly-filtered LES), one eventually approaches the DNS limit; in this case, while the numerical errors are reduced, the model contribution has also decreased and it becomes difficult to ascertain its fundamental properties and limitations.

removal of numerical artifacts in order to present a more clear interpretation of the model performance. As previously mentioned, employing filter-to-grid ratios above unity can enforce the separation of poorly-resolved scales from the model contribution; yet, this can come with non-negligible computational overhead. Utilizing high-order or spectrally-optimized schemes such as Dispersion Relation Preserving (DRP) stencils (Tam and Webb, 1993) allows a larger portion of the resolved scales to be properly represented, thus reducing the need for large filter-to-grid ratios[3]. This increased resolvability should be particularly useful in the case of LES (Fauconnier et al., 2009) – which is inherently under-resolved. But while one may target a particular cut-off or filter width, these definitions are somewhat ambiguous. As a result, it becomes important to specify the spectral attenuation properties of the LES filter, which requires an understanding of its construction and manipulation. By employing an appropriately scale-discriminant cut-off, one may further ensure maximization of the resolved LES field given a desired filter-to-grid ratio. Selection of the filter is known to be reflected in the modeling but is also important for taking full advantage of a high-accuracy discretization. And so, this consideration also plays a role in the perceived efficiency of the explicit LES implementation. All in all, the aim resides in the idea of tuning the LES formulation (e.g., filter type, filter-to-grid ratio, etc ...) to the spectral capability or resolvability of the base scheme.

1.2 Focus and Outline of Current Studies

The present work aligns itself with efforts towards the implementation of high-order numerical schemes within the Cartesian component of a dual-mesh frame-

[3]Additional care should also be taken with respect to the potential of aliasing error; considerations such as the "3/2" rule (Orszag, 1971), and analogous concepts for the cubic non-linearities present in the compressible fluid equations, thus constitute a second guideline by which to select an appropriate filter-to-grid ratio.

work, intended for internal reactive flow LES calculations. As a result, the current work seeks to understand how best to leverage high-accuracy discretizations for LES and the eventual development and validation of compressible model closures geared towards turbulent reactive flows. Focusing within this Cartesian realm, uniform and periodic domains are considered for a finite difference implementation, the simplicity of which provides a clean background for understanding the respective consequences of the numerical scheme, the LES model, and the selected filter formulation.

In an incremental approach towards these motives, the thesis first focuses on the characterization of numerical schemes; it includes the construction and spectral analysis of discrete filtering stencils in Chapter 2, the formulation of filter-consistent stabilization methods in Chapter 3, and a study of the pairing of the stabilization and the high-order numerical schemes in Chapter 4 for long-time, temporally-accurate evolution of non-linear flow solutions. Next, the thesis inspects the impact of the numerical base scheme on LES calculations in relation to the filter formulation and model choice. Notably, Chapter 5 considers LES of the 1D Burgers equation as a prototype of the fluid equations and provides *a priori* and *a posteriori* analysis, assuming "perfect" and scale-similarity type closures. Chapter 6 then extends these lessons to the compressible Navier-Stokes equations (NSE) and uses the simulation of the Taylor-Green vortex within an explicitly-filtered LES framework as the testbed. Finally, Chapter 7 provides summarizing remarks and future work directions.

CHAPTER 2

Discrete Filtering Schemes

This chapter was taken with some modifications from the AIAA conference proceedings article "Discrete Filtering Formulations for Large-Eddy Simulation" (Edoh et al., June 2016).

2.1 Background

Fundamental to the concept of LES is the notion of scale-separation, which is responsible for decomposing the full solution field into the "large eddies" (to be resolved) and the small scales (to be modeled). This delimitation of content may be expressed in different ways; however, the idea of low-pass filtering is perhaps most common and intuitive with respect to the act of suppressing high-wavenumber or high-frequency modes.

The manner in which the small-scales are removed can largely be understood from inspecting the spectral response of the associated filter. The presumption of a spectrally-sharp Fourier cut-off is abstractly desirable as it clearly truncates the effective resolution without altering lower wavenumber content. As a consequence, it furthermore relieves the LES model from needing to account for large-scale effects, which are influenced by the geometry; the tractability of modeling such problem-dependent phenomena is improbable and would otherwise be contrary to the idea of universal modeling in LES. An *a priori* type investigation by Stefano and Vasilyev (2002), for instance, compare a model built from a spectrally-sharp

filter to one that is built from a smooth top-hat filter. The study highlights a reduced susceptibility to model error for the former, along with more accurate representations of the energy cascade process. Such characteristics emphasize the utility of employing a sharp filter (in the spectral sense). Unfortunately, practical implementation of the Fourier filter reveals some difficulties. For example, *a priori* studies have shown low correlations between Smagorinsky and Scale-Similarity models and true closures when employing a spectrally-sharp filter (Liu et al., 1995). In addition, *a posteriori* model implementations meant to represent the associated cross-scale interactions incur difficulties. For example, dynamic model procedures utilizing spectral-sharp test-filtering have demonstrated poor performance (Liu et al., 1994). A possible explanation for the lack of success – particularly with respect to models of the eddy-viscosity type – in being able to represent the effects of the sharp-spectral filter may be related to their inability to properly emulate the exponential order of the spectral filter (Pruett and Sochacki, 2001). An even more practical hinderance to the use of Fourier spectral filtering relates to geometric limitations, namely the need for homogeneous (i.e., periodic) boundaries and a uniform grid. Such requirements largely limit general use. Discrete filtering is therefore presented as a simple and thus attractive alternative that also provides tractable design of the spectral response characteristics.

Discrete filtering operates in physical space by using linear combinations of nodal solution values. This allows one to achieve localized spectral attenuation that also accommodates implementation in complex geometries. Nodal-based element methods for instance (e.g., Spectral Difference, Discontinuous-Galerkin etc) are globally unstructured; however, one can utilize the structured inter-element node arrangement in order to systematically derive filter stencils – the respective weights of which can be pre-calculated based on the available support and desired spectral response traits (Mirzaee et al., 2012; Lodato et al., 2013). Discrete filtering on fully unstructured meshes is also possible, albeit more algorithmi-

cally expensive and challenging with respect to manipulating spectral response characteristics. Marsden et al. (2002) leverage local interpolation functions, while Haselbacher and Vasilyev (2003) provide a more efficient procedure based on least-squares gradient reconstruction. Regardless, in most cases, specification of the spectral response is determined by a shape optimization procedure that yields the proper filter stencil coefficients (von Storch, 1978; Lele, 1992; Vasilyev et al., 1998; Bogey and Bailly, 2004). Even on uniform structured grids, such optimization techniques have been used to derive suitable high-order boundary closure schemes (Gaitonde and Visbal, 2000) or to generate self-similar schemes for consistent multi-resolution LES studies (Radhakrishnan and Bellan, 2012).

In the special case of uniform grids, one may avoid further complications associated with such shape optimization procedures by employing particular "families" of filter stencils that are easily tunable. The predetermined manipulation of their properties provides greater flexibility, which then simplifies their use in applications such as dynamic adaptation of spectral response characteristics, either in the spatial or temporal sense. With regards to such filter schemes, of note is the class of Purser filters (Purser, 1987), from which the Shapiro (Shapiro, 1971, 1975) and binomial stencils (Jahne) emerge for use as optimal low- and high-pass discrete filters, respectively. Also useful are compact Padé-type implementations such as the scale-discriminant Tangent stencils and the naturally invertible Sine stencils (Raymond and Garder, 1991). The following sections detail the respective formulations of these and other stencils, relating them to the general construction of discrete filters, and furthermore evaluates their spectral response characteristics with the intended goal of understanding how to apply such forms in LES.

2.2 Construction

First considering the one-dimensional (1D) case for simplicity, a general explicit discrete stencil can be written as,

$$\bar{u}_i = \sum_{r=J}^{K} b_r u_{i+r} \ . \tag{2.1}$$

Such a weighted-sum perspective thus draws natural links between discrete filtering and notions such as averaging or interpolation. The above can then be summarized as a general operator,

$$\bar{u} = \mathcal{G}\{u\} \ , \tag{2.2}$$

where u is the input signal and $\bar{u}$ is the filtered result. In order to understand the spectral consequences of these stencils, it is helpful to consider the transfer function of the operator which may be derived by substituting the following Fourier representation of the solution,

$$u(x_i) = \sum_{k} \hat{u}(k) e^{\imath k x_i} \ , \tag{2.3}$$

into Equation 2.1[1]. In this way, the effect of the filtering operation in Equation 2.2 on the solution modes can be understood as a simple multiplication of the filter's transfer function with the original signal in Fourier space, $\mathcal{F}\{\mathcal{G}\{u\}\} \sim \hat{\mathcal{G}}(k) \cdot \hat{u}(k)$. Depending on the stencil coefficients b_r, the filter transfer function may be complex. Therefore, the dissipative and dispersive effects on the input function are understood by re-rewriting the transfer function in complex exponential form,

$$\begin{aligned} \hat{\mathcal{G}} &= A e^{\imath \phi} \ , \\ &\text{where} \quad \begin{cases} A = |\mathrm{Re}\{\hat{\mathcal{G}}\}| + |\mathrm{Im}\{\hat{\mathcal{G}}\}| \ , \\ \phi = \tan^{-1}\{\mathrm{Im}\{\hat{\mathcal{G}}\}/\mathrm{Re}\{\hat{\mathcal{G}}\}\} \end{cases} . \end{aligned} \tag{2.4}$$

[1]Note: The imaginary number, $\imath = \sqrt{-1}$ is used in order to avoid confusion with the index i.

As such, A corresponds to growth effects and ϕ corresponds to phase effects. From this perspective, it then becomes evident that successive applications of the filter (e.g., $\bar{u} = \mathcal{G}\{\mathcal{G}\{...\{u\}\}\}$) can be summarized as $\hat{\bar{u}} = [\hat{\mathcal{G}}^n(k) \cdot \hat{u}]$ for n applications of the filter. Consequently, the overall impact on amplitude is geometric (i.e., $A_{\text{total}} = A^n$) while that on the phase is additive (i.e., $\phi_{\text{total}} = n \cdot \phi$). Implicit in this type of analysis of the resulting transfer functions is the assumption that the stencil in question is applied globally – which furthermore supposes a constant mesh size. Proper interpretations of hybrid stencils may require alternate methods. Jordan (2007), for example, performs a composite modified wavenumber analysis of boundary and interior point stencils on a uniform mesh and demonstrates different predictions on spectral performance for the overall coupled hybrid system compared to predictions derived from inspection of the schemes individually.

As mentioned, customization of the spectral characteristics can be achieved via shape optimization, where error in the spectral transfer function is minimized relative to a target response. Equation 2.5 shows how this is done relative to the L_2 energy norm, L_2

$$\min_{b_r} \quad \int_{k_{r,\min}}^{k_{r,\max}} \left[\Re\{\hat{\mathcal{G}}(k) - \hat{\mathcal{G}}_t(k)\}\right]^2 dk + \int_{k_{i,\min}}^{k_{i,\max}} \left[\Im\{\hat{\mathcal{G}}(k) - \hat{\mathcal{G}}_t(k)\}\right]^2 dk \ , \quad (2.5)$$

where $\hat{\mathcal{G}}_t(k)$ is the target spectral response shape. The optimization is then taken with respect to the stencil coefficients such as to minimize the functional (see Appendix C of Vasilyev et al. (1998) for details). Further control in the design process is achievable by specifying the bounds of interest $[k_{\min}, k_{\max}]$ or by increasing error sensitivity through a logarithmic change of variables (Bogey and Bailly, 2004). As an alternate to the energy norm, one may consider an infinity norm and employ optimization strategies based on the Remez multiple exchange algorithm, as done by Linders and Nordstrom (2015) with respect to the efficient derivation of new optimized dispersion-relation preserving (DRP) finite difference schemes.

A more traditional expression for filtering in the context of LES is the convo-

lution of the input signal with a particular kernel, $G(x)$,

$$\bar{u}(x) = \int [G(x-y) \cdot u(y)]\, dy, \tag{2.6}$$

$$\text{with} \quad \int G(y)dy = 1\ . \tag{2.7}$$

Equation 2.7 is a normalization that calls for the filtering operation to preserve constant modes. Some well-used functions include the top-hat, Gaussian, and spectral-sharp filter kernels (see Table 6.1 for a list of their physical and spectral space operators in the instance of symmetric applications). Written in the convolution form above, the filtering operation is then executed by evaluating the corresponding integral. In this respect, using kernels such as the top-hat (or other compact generalizations such as the B-splines used in the Smoothness-Increasing Accuracy Conserving methods (Mirzaee et al., 2012)) limits the need to introduce possible errors associated with forcibly needing to truncate the kernel. In the case where the kernel is non-compact, one typically chooses to truncate the evaluation, as is typically done with the Gaussian function, which features sufficient physical space decay for such a suitable approximation to be made.

Figure 2.1 plots the respective physical and spectral functions for the symmetric forms of the Gaussian, top-hat and spectral-sharp kernels. Figure 2.1(a) highlights their spatial compactness (or lack thereof), while Figure 2.1(b) demonstrates their performances with respect to achieving the spectral scale separation abstracted in the idea of LES. As evident, the top-hat filter is convenient to use in physical space but is highly oscillatory and slowly-decaying in spectral space because of its sinc function representation. The direct opposite is true of the spectral sharp filter which is a top-hat function in wavenumber space yet is a sinc function in physical space. On the other hand, the Gaussian kernel maintains its functional structure in both physical and spectral space and shows suitable decay characteristics, although it is nowhere as sharp as the spectral function in wavenumber space.

filter	$G(x)$	$\hat{\mathcal{G}}(k) = \mathcal{F}\{G(x,\Delta)\}$
Gaussian	$\left(\frac{6}{\pi\Delta^2}\right)^{\frac{1}{2}} \exp\left(-\frac{6x^2}{\Delta^2}\right)$	$\exp\left(-\frac{k^2\Delta^2}{24}\right)$
top-hat	$\begin{cases} 1/\Delta & \text{if } \lvert x\rvert \leq \Delta/2 \\ 0 & \text{otherwise} \end{cases}$	$\frac{\sin(\frac{1}{2}k\Delta)}{\frac{1}{2}k\Delta} = \text{sinc}\left(\frac{1}{2}k\Delta\right)$
spectral cut-off	$\frac{1}{\pi x}\sin\left(\frac{\pi x}{\Delta}\right) = \frac{1}{\Delta}\text{sinc}\left(\frac{\pi x}{\Delta}\right)$	$\begin{cases} 1 & \text{if } \lvert k\rvert \leq \pi/\Delta \\ 0 & \text{otherwise} \end{cases}$

Table 2.1: Physical kernel and spectral function representations of symmetric filter kernels.

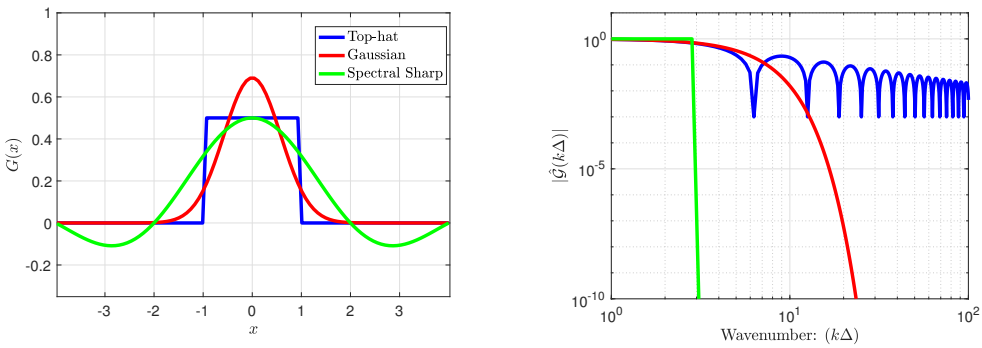

(a) (b)

Figure 2.1: Gaussian, Top-hat, and Spectral Sharp symmetric filters with filter-width $\Delta = (2\Delta x)$ from Table 6.1: a) Physical space kernel $G(x)$, Spectral space function $|\hat{\mathcal{G}}(k)|$.

For cases where the input signal u is a known function (e.g., finite element type numerical schemes express the solution variable in terms of known closed-form bases), the convolution integral in Equation 2.6 can be derived analytically (i.e., quadrature-free). Usually, however, the convolution operation is estimated

by a numerical quadrature rule which directly reverts the integral expression to the discrete form introduced in Equation 2.1. Important to note are the possible oversights that may stem from focusing on the design and characteristics of the filter kernel without taking into account the quadrature rule for numerical integration of the convolution integral. For instance, Figure 2.2 plots two interpretations of a symmetric top hat kernel with filter width $\Delta = (2\Delta x)$. The plot compares the exact spectral representation with the transfer functions recovered when utilizing either Simpson's rule ($[b_{-1}, b_0, b_1] = [1/6, 2/3, 1/6]$) or the Trapezoidal rule ($[b_{-1}, b_0, b_1] = [1/4, 1/2, 1/4]$) to evaluate the convolution integral discretely. Notably, the quadrature choice can have a strong impact on the spectral characteristics of the effective filter operation, particularly towards high wavenumbers. It can thus be advantageous to consider the discrete filter representation of Equation 2.1 directly in order to anticipate such deviations in performance.

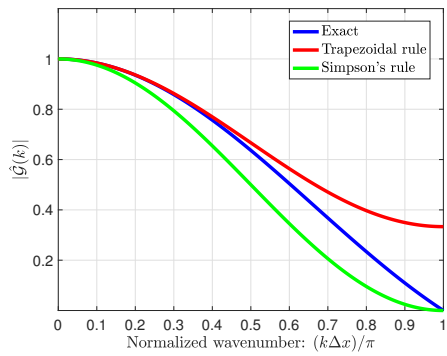

Figure 2.2: Spectral response for the top-hat filter assuming $\Delta = (2\Delta x)$, comparing the exact representation versus discrete approximations to the convolution integral utilizing Simpson and Trapezoidal quadrature rules.

Reconciling the discrete filtering stencils with the convolution form can be accommodated by defining the associated kernel as a sum of dirac delta functions, $G_{\text{disc}}(x) = \sum_r b_r \cdot \delta(x + \Delta x_r)$. The coefficients b_r can thus be seen to holistically contain quadrature, kernel, and basis function information (in the case of finite

element schemes) relative to a discrete approximation of the filter convolution equation. Consequently, understanding how to prescribe the discrete filter weights is central to the construction and manipulation of transfer function characteristics.

A preliminary approach in seeking to define these coefficients is to require the discrete filtering operation to preserve a specific set of monomials:

$$x^m = \sum_{r=J}^{K} b_r \int \left[\delta(x_r - y) \cdot y^m\right] dy, \quad \text{where} \quad x_r = x + \Delta x_r \ . \tag{2.8}$$

Substituting a useful change of variables ($z = x_r - y$) and evaluating about $x = 0$ (Mirzaee et al., 2012), one recovers the constraints for the coefficients:

$$x^m = \sum_{r=J}^{K} b_r \int \left[\delta(z) \cdot (x + \Delta x_r - z)^m (-1)\right] dz \ , \tag{2.9}$$

$$\rightarrow \quad 0^m = \sum_{r=J}^{K} b_r (\Delta x_r)^m \quad (\text{taking } x = 0) \ . \tag{2.10}$$

Closer inspection of these conditions reveals them to correspond to eliminating high-order terms ($\sim \partial^m u$) in the Taylor series expansion of $\bar{u}$ with respect to the unfiltered variable, u. The first condition (x^0) guarantees that constant modes are preserved and also communicates that $\bar{u}$ is a zeroth-order derivative of u – and thus essentially an interpolated variable of the input. The preservation of higher order monomials by the filter convolution can then be seen as demanding a higher order interpolation of the input signal, such that $\bar{u} = u + O(\Delta x^M)$. An equivalent derivation of the filter order is available by considering the Taylor-series expansion of the stencil (Equation 2.1) and utilizing the spectral response of the filter $\hat{\mathcal{G}}(k)$ to determine the associated series coefficients (Pruett and Sochacki, 2001):

$$\bar{u} = \sum_{m=0}^{\infty} d_m (\Delta x)^m \partial^m u \ , \quad d_m \equiv \frac{\partial_k^m \hat{\mathcal{G}}(k)\big|_{k=0}}{m! \left(\sqrt{-1}\right)^m} \ . \tag{2.11}$$

In other words, the number of vanishing derivatives of the spectral response at $k = 0$ corresponds to the filter order.

Also interesting to note is that Equation 2.10 can be viewed as a set of discrete moment conditions. This is consequential for such things as commutation error

in LES, which results from ignoring the effects of differentiating spatially-varying filters (e.g., $\mathcal{E}_{\text{com}} = \overline{\partial u} - \partial \bar{u} \neq 0$ for non-constant filter widths, $\Delta(x)$). Such effects are typically excluded in LES modeling but can be substantial depending on the spatial variation of the filter width. As an attempted relief to such concerns, it has been shown that the magnitude of commutation error decreases in accordance with the number of vanishing moments of the filter kernel (Vasilyev et al., 1998). Recalling that the monomial conditions of Equations 2.9-2.10 simultaneously express discrete moment relations as well as Taylor series coefficients thus leads one to conclude that commutation error decreases with the use of high-order filters[2]. The top-hat and Gaussian filters can both be shown to be second-order; meanwhile, the spectral sharp filter is essentially of infinite order. Through the design of the discrete filter coefficients, one is able to control the desired order of the kernel and also control the possible influence of commutation errors on the solution – typically judged relative to numerical errors inherited from the base scheme.

Further specification of the filter stencil coefficients is then typically made relative to desired spectral characteristics such as the elimination of odd-even modes (i.e., $\hat{\mathcal{G}}(k_{\text{Nyquist}} = \pi/\Delta x) = 0$) or the adherence to an effective cut-off frequency (i.e., $\hat{\mathcal{G}}(k_{1/2}) = 0.5$) (Lund, 1997). These requirements appear as additional constraints to the optimization procedure of Equation 2.5. Figure 2.3 shows the transfer function of a five-point, fourth-oder biased discrete filter stencil ($J = -1$,$K = 3$) with coefficients listed in Table 2.2. It is apparent that the resulting response has an induced phase component due to the stencil asymmetry. Fortunately, these misgivings are relegated to high wavenumbers where there is also sufficient damping from the scheme. Another observation is the pres-

[2]The observation that commutation error decreases with filter order is typically used to rationalize the subsequent neglect of commutation modeling for high-order filters; however, (Geurts and van der Bos, 2005b) importantly shows that the magnitude of the typical non-linear closures (i.e., $\overline{\bar{u}\bar{u}} - \bar{u}\bar{u}$) is also reduced with respect to filter order and that both contributions (non-linear and commutation) are of similar magnitude. Therefore, the idea of modeling one without the other loses validity.

ence of overshoots ($|\hat{\mathcal{G}}| > 1$), which suggests modal growth and thus can threaten solution stability[3]. These may be reduced by establishing a suitable optimization constraint that would enforce monotonically decreasing responses such that $\partial_k \hat{\mathcal{G}}(k) \leq 0$. Nevertheless, such biased stencils, are typically used as boundary closure schemes; therefore, the unwanted phase and growth effects are expected to be localized. Sensitivity of the interior solution accuracy to boundary information, however, may nevertheless be consequential. For such considerations, methods such as eigenvalue matrix analysis (Kim, 2010) or energy estimates (e.g., Kreiss Theory (Larsson and Gustafsson, 2008) may provide additional insights on global stability properties relating to the coupling of different stencils. Furthermore, the optimization procedure provided in Equation 2.5 may be used to reduce phase effects while maintaing the desired spectral damping properties, as shown in Vasilyev et al. (1998).

$$\hat{\mathcal{G}}(0) = 1;\ \hat{\mathcal{G}}(\pi/\Delta x) = 0;\ \partial_k^m \hat{\mathcal{G}}(0) = 0 \text{ (for } m = [1, 2, 3]) \tag{2.12}$$

b_{-1}	b_0	b_1	b_2	b_3
1/16	3/4	3/8	-1/4	1/16

Table 2.2: Biased stencil coefficients of fourth-order biased stencil from Table 5 (with $\alpha_f = 0$) of Gaitonde and Visbal (2000).

While the development so far has remained general, the subsequent analysis will focus on uniform grid arrangements as this relates to the intended Cartesian portion of the dual-mesh code in development. Conclusions may also be valid in a transformed computational space. The fact that the filtering at hand is meant to achieve scale-separation for LES furthermore requires the convolution operator to be dissipative to leading order (rather than dispersive) – and thus the

[3]Of note are the biased "de-centered" explicit Shapiro stencils of Falissard (2015) which can achieve $2R + 1$ order and have a monotonic modulus.

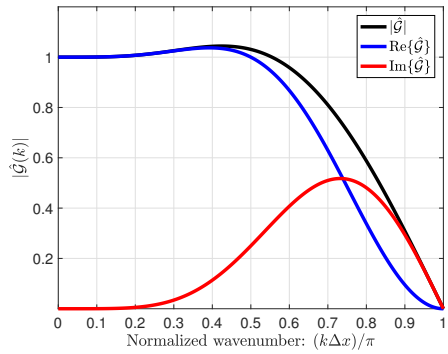

Figure 2.3: Fourth-order biased stencil from Table 5 (with $\alpha_f = 0$) of Gaitonde and Visbal (2000).

corresponding stencils are required to be at least of order even$\{M\}$. Most ideally, the filter operation is expected to be purely dissipative with no phase error($^*\phi = 0$) which suggests the use of symmetric and unbiased stencils, where $b_j = b_{-j}$. In such circumstances, it is worthwhile to consider established families of discrete filter stencils for the design of spectral characteristics. These provide sufficient flexibility in prescribing such things as cut-off and high-wavenumber damping, in addition to controlling the level of scale-separation. The following sections thus explore explicit and implicit (i.e., Padé) discrete filter stencil families, offering useful generalizations and quantifying their spectral performance in light of their anticipated use in explicitly-filtered LES formulations.

2.3 Explicit Stencil Families

Explicit discrete filtering schemes present a simple way of calculating filtered quantities, $\bar{u}$, based on surrounding input values, u. Ascribing to the call for a purely dissipative operator allows Equation 2.1 to be re-written as a symmetric stencil,

$$\bar{u}_i = b_0 u_i + \sum_{r=1}^{R} b_r \cdot (u_{i+r} + u_{i-r}) \ , \tag{2.13}$$

with a transfer function,

$$\hat{\mathcal{G}}(k) \;=\; b_0 + \sum_{r=1}^{R} 2b_r \cdot \cos(rk\Delta x) \;, \tag{2.14}$$

$$\text{where} \quad \begin{cases} \cos\theta \;=\; 1 - 2\sin^2(\theta/2) \;, \\ \cos(r\theta) \;=\; 2\cos[(r-1)\theta] - \cos[(r-2)\theta] \end{cases} \;. \tag{2.15}$$

This representation can then be re-expressed in terms of dissipative difference stencils by employing the trigonometric identities in Equation 2.15. By first recognizing that the second-order, narrow dissipative stencils satisfying,

$$\delta_x^{2r} = \{\delta_x^2\}^r \quad \text{and} \quad \delta_x^{2r} = \frac{\partial^{2r}}{\partial x^{2r}} + O(\Delta x^2) \;, \tag{2.16}$$

have the following spectral function,

$$\mathcal{F}\{\delta_x^{2r} u\} = \frac{1}{(\Delta x)^{2r}}[(-4)^r \cdot \sin^{2r}(k\Delta x/2)] \cdot \hat{u}(k) \;, \tag{2.17}$$

then the original stencil of Equation 2.13 can be re-written solely in terms of dissipative differences:

$$\bar{u}_i = \left[\epsilon_{EF,0} + \sum_{r=1}^{R} \epsilon_{EF,2r} (\Delta x)^{2r} \delta_x^{2r}\right] \{u_i\} \;. \tag{2.18}$$

This highlights the damping or attenuating effect of filtering and further provides insight into how such discrete filtering stencils are relatable to artificial dissipation methods – useful as scheme stabilization agents (see Chapter 3). Evoking differencing stencils also hints at differential filtering techniques (Germano, 1986; Sagaut, 2006) which offer additional flexibility in filter implementations on unstructured meshes through the solution of auxiliary equations (Mullen and Fischer, 1999; Bull and Jameson, 2016). Referring to the filtering in term of operators, one can also re-write Equation 2.18 as,

$$\bar{u} = \mathcal{G}\{u\} = [\mathcal{I} + \mathcal{D}]\{u\} \;, \tag{2.19}$$

which makes further delimitation between the actions of scaling constant modes by $\mathcal{I}$ and of differencing gradients by $\mathcal{D}$. Modification of the latter is responsible for tuning the dissipation characteristics.

As previously noted, one would like to tailor the spectral characteristics of the discrete filter scheme through selection of the stencil coefficients. Remembering that constant modes should be preserved, it becomes evident that $\epsilon_{EF,0} = 1$. The following section explores the specification of the remaining $\epsilon_{EF,2r}$ coefficients, largely from a spectrally-motivated perspective. This is shown with respect to the class of explicit Purser filter stencils; some useful generalizations are furthermore offered for the manipulation of spectral properties.

2.3.1 Explicit Purser Filters

The explicit stencils of Purser (1987) draw inspiration from the fact that the filter transfer function $\hat{\mathcal{G}}(k)$ resembles the complement of a probability cumulative distribution function (CDF) defined by $F_{CDF}(z)$ for the random variable z. The shape of the CDF, in turn, is determined by its derivative: the probability density (PDF) function $f_{PDF}(z)$. By selecting a polynomial-based PDF with finite or semi-finite support (i.e., a one-sided distribution), one can then integrate and write the corresponding CDF as a rational function. Subsequent derivation of the associated spectral transfer function is achieved by a change of variables on the integral limits. By using a suitable basis for these limits (e.g., $\mathcal{F}\{\delta^2\}$), the filter response will then be expressed in polynomials of the δ^2 difference operator and one can then easily identify the related coefficients $\epsilon_{EF,2r}$.

Purser originally chooses the Beta PDF distribution (Johnson et al., 1994) as a candidate, which has finite support over an interval $[0, 1]$. Employing proper normalization, the Purser filter's transfer function is written as,

$$\hat{\mathcal{G}}(k) \;=\; 1 - \int_0^{\hat{g}_\ell(k)} f_{PDF}(z) dz \;, \tag{2.20}$$

with,

$$f_{PDF}(z) = \frac{(C_R + C_S + 1)!}{C_R! C_S!} z^{C_R} (1 - z)^{C_S} \quad \text{and} \quad \hat{g}_\ell(k) = \sin^2\left(\frac{k\Delta x}{2}\right) \;. \tag{2.21}$$

Substituting in the Beta distribution PDF as well as the change of variable into the integral limits then gives the Purser filter as,

$$\hat{\mathcal{G}}_{\text{Purser}}(k) = 1 - \left[\frac{(C_R+C_S+1)!}{C_R!C_S!}\right]\sum_{k=0}^{C_S}\begin{pmatrix} C_S \\ k \end{pmatrix}\frac{(-1)^k}{(k+C_R+1)}\sin^{2(k+C_R+1)}\left(\frac{k\Delta x}{2}\right). \tag{2.22}$$

This can then be put in a difference form of as,

$$\bar{u}_i = \left[1+\sum_{r=C_R+1}^{R}\epsilon_{EF,2r}(\Delta x)^{2r}\delta_x^{2r}\right]\{u_i\}\ , \tag{2.23}$$

$$\text{with} \quad \begin{cases} \epsilon_{EF,2r} = \left[\frac{(C_R+C_S+1)!}{C_R!C_S!}\right]\begin{pmatrix} C_S \\ r-C_R-1 \end{pmatrix}\left(\frac{(-1)^{r-C_R}}{r}\right)\left(\frac{-1}{4}\right)^r \\ C_R \geq 0\ ,\ C_s \geq 0\ , C_R+C_S+1=R \end{cases}. \tag{2.24}$$

Several benefits arise from deriving transfer functions from a PDF distribution. Firstly, the transfer function is guaranteed to be monotonic because the CDF is strictly increasing; this is a reasonable requirement for achieving clear scale-separation in LES. Secondly, distribution properties such as the median, mean, variance and higher-order moments of the PDF are typically well understood in terms of the available distribution parameters and provide natural insight into characterizing and specifying such things as the effective cut-off frequency (k_c) or sharpness of the filter response. Assuming the effective cut-off frequency to be $k_c = k_{1/2}$ such that $|\hat{\mathcal{G}}(k_{1/2})| = 0.5$, the median is sufficient to characterize the filter width $\Delta_{1/2} = \pi/k_{1/2}$. The median for the Beta distribution can then be approximated based on the PDF parameters:

$$\sin^2(k_{1/2}\Delta x/2) \approx \frac{C_R+2/3}{C_R+C_S+4/3}\ . \tag{2.25}$$

This estimate of the median is valid for $(C_R, C_S) \geq 1$, for which there is a 4% relative error; when $(C_R, C_S) \geq 2$, the relative error of the expression further decreases

to a 1% (Kerman, 2011). Alternatively, Purser (1987) provides an asymptotic measure based on the mode and mean of a lightly skewed unimodal distribution that requires the assumption that both (C_R, C_S) are large.

Considering a fixed polynomial order $R = C_R + C_S + 1$ of the Beta distribution, one is able to generate a basis of response functions according to the pair of $[C_R, C_S]$ parameters. These schemes correspond to a set stencil size of $(2R + 1)$. The resulting bases can then be interpolated in order to produce intermediate response shapes as required. In general, the C_R parameter controls roll-off, while higher values increases the range of low-frequency preservation as well as sharpness. Meanwhile the C_S parameter controls smoothness towards the odd-even modes $(k\Delta x) \rightarrow \pi$, and thus it works as a complement to C_R. This behavior is intuitively understandable by considering the PDF and noting that larger values of C_R shift the mean towards the upper limit of the bounds (i.e., high wavenumers) while larger values of C_S shift the mean towards the lower limit. With respect to the difference form of Equation 2.23, it becomes evident that increasing C_S for a fixed R introduces lower derivatives $\delta^{2r<2R}$ and results in reducing the formal order of the filter according to the stencil's Taylor-series expansion. In other words, a Purser filter that has $2 \cdot C_R$ vanishing discrete moments is of order $2 \cdot (C_R + 1) = 2 \cdot (R - Cs)$. The case of $[C_R = R - 1, C_S = 0]$ thus corresponds to the maximal order filter given a set stencil size. Such maximal order stencils are classically referred to as Shapiro stencils. Shapiro derived the associated b_r coefficient analytically (Shapiro, 1971, 1975), and these have been re-discovered and validated through the optimization efforts of Bogey and Bailly (2004). The class of Shapiro filters approach an ideal low-pass response towards infinite order (i.e., as $R \rightarrow \infty$), wherein all modes are preserved with the exception of an eliminated Nyquist frequency, $k_{\text{Nyquist}} = \pi/\Delta$. The associated Shapiro spectral response is,

$$\hat{\mathcal{G}}_{\text{Shapiro}}(k) = 1 - \sin^{2R}(k\Delta x/2) \ . \tag{2.26}$$

These filters are of order $2R$ and inherently build on the derivative operator

$\mathcal{F}\{\delta_x^{2r}\} \sim \sin^{2r}(k\Delta x/2)$. Also of note are the odd-numbered Binomial stencils which correspond to the special case of $[C_R = 0, C_S = R - 1]$ for the Purser schemes. These stencils have spectral response,

$$\hat{\mathcal{G}}_{\text{Binomial,odd}}(k) = \frac{1}{2^R}[1 + \cos(k\Delta x)]^R \,, \tag{2.27}$$

$$= \frac{1}{2^{R-1}}[1 - \sin^2(k\Delta x/2)]^R \,, \tag{2.28}$$

and are inherently second-order. They are built from averaging operators $\mathcal{F}\{\mu^{2r}\} \sim [1+\cos(k\Delta x)]^r$. The Binomial stencils can furthermore be used to approximate an ideal high-pass response as $R \to \infty$, thus serving as complements to the Shapiro stencils (e.g., $\hat{\mathcal{G}}_{\text{high-pass}} = 1 - \hat{\mathcal{G}}_{\text{low-pass}}$). Figure 2.4 demonstrates the response characteristics of the Shapiro filters, noting how an ideal low-pass filter is approached as the filter order $2 \cdot (C_R + 1)$ increases. Meanwhile Figure 2.5 plots the transfer function of odd-stenciled Binomial filters, confirming these to be second-order and suggesting their complement $(1 - \hat{\mathcal{G}})$ to approach and ideal high-pass response.

The Shapiro and Binomial filters are special cases of the Purser stencil; however, it is useful to understand the general impact of the parameters $[C_R, C_S]$ on the transfer function. Figure 2.6 plots response functions for the basis $R = 3$, corresponding to a stencil width of 7. Figure 2.6(a) shows how the effective cutoff is affected: as C_R decreases, C_S increases accordingly and producers lower cut-offs. But while plotting the response in terms of $(k\Delta x)$ suggests that the filter sharpness main remain unchanged, the filter width is inversely proportional to wavenumber which means that as the cut-off decreases, more scales are attenuated for a given change in $\theta = (k\Delta x)$. Therefore in order to truly preserve scale-discriminance, the spectral response would need to become more sharp as $k_c \to 0$. Figure 2.6(b) plots the transfer function error $|1 - \hat{\mathcal{G}}(k)|$ on a logarithmic scale and demonstrates how including lower derivates in the difference expression of Equation 2.23 (i.e., higher values of C_S) reduces the formal filter order and thus reduces preservation properties. Figure 2.7 plots the same information for the basis $R = 6$. Intuitively,

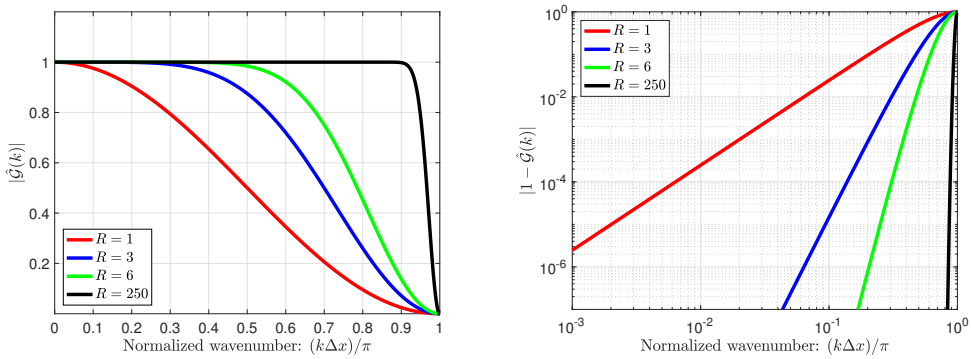

(a) (b)

Figure 2.4: Shapiro filters (see Equation 2.26) of order $2 \cdot (C_R + 1)$ corresponding to Purser filters with parameters $[C_R, C_S] = [R - 1, 0]$: a) Growth factor $|\hat{\mathcal{G}}|$, and b) Growth factor error $|1 - \hat{\mathcal{G}}|$.

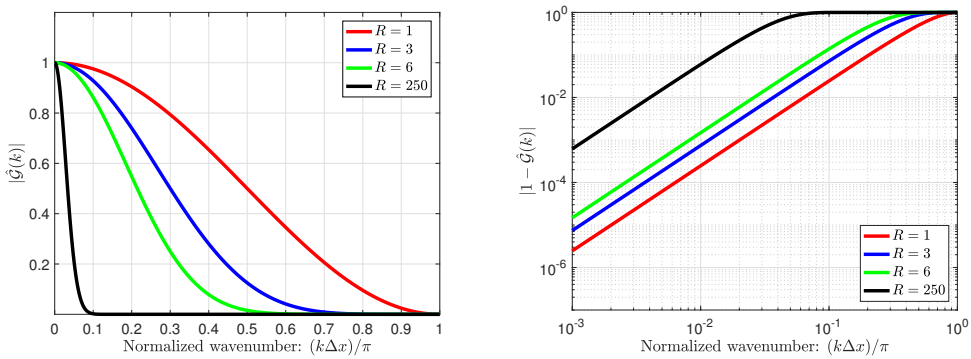

(a) (b)

Figure 2.5: Binomial filters (see Equation 2.27) of second order corresponding to Purser filters with parameters $[C_R, C_S] = [0, R - 1]$: a) Growth factor $|\hat{\mathcal{G}}|$, and b) Growth factor error $|1 - \hat{\mathcal{G}}|$.

it shows how larger stencil widths can generate a larger set of basis response functions that also cover a wider range of cut-offs. The increase in the density of the bases furthermore reveals that one can interpolate and yield intermediate transfer functions that better retain the sharpness qualities of the individual schemes.

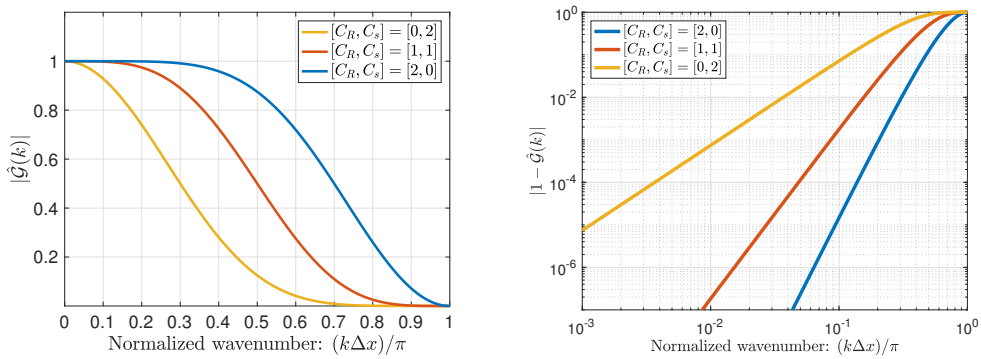

(a) (b)

Figure 2.6: Purser filter response basis for $R = (C_R + C_S + 1) = 3$: a) Growth factor $|\hat{\mathcal{G}}|$, and b) Growth factor error $|1 - \hat{\mathcal{G}}|$.

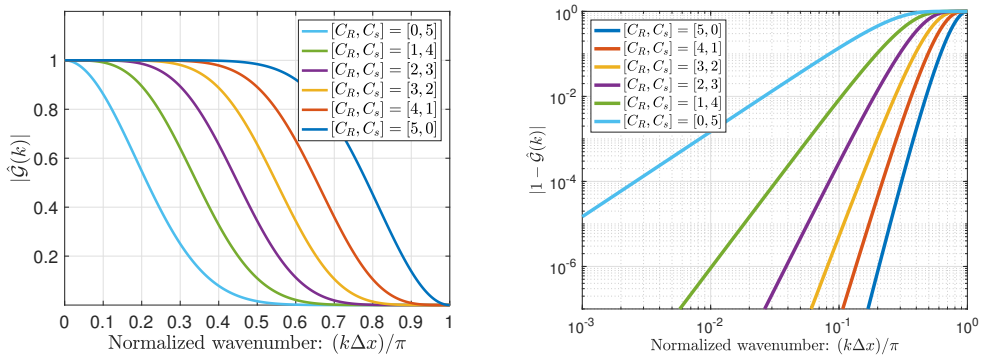

(a) (b)

Figure 2.7: Purser filter response basis for $R = (C_R + C_S + 1) = 6$: a) Growth factor $|\hat{\mathcal{G}}|$, and b) Growth factor error $|1 - \hat{\mathcal{G}}|$.

Of further note is the way in which all transfer functions derived thus far

inherently eliminate odd-even modes, $\hat{\mathcal{G}}(k_{\text{Nyquist}}) = 0$. This is a consequence of the Beta distribution having finite support and the fact that Equation 2.20 is integrating across the full interval during the derivation of the discrete stencils. In order to alter the high-wavenumber damping of the scheme, one can choose to either re-scale the CDF (i.e., the integral in Equation 2.20) or to re-scale the integral limits. Assuming a re-scaling factor $\mu \in [0, 1]$, the latter approach results in modified coefficients $\epsilon'_{EF,2r\geq 2} = \mu^r \cdot \epsilon_{EF,2r\geq 2}$, while the former technique is more simple and yields $\epsilon'_{EF,2r\geq 2} = \mu \cdot \epsilon_{EF,2r\geq 2}$. Re-scaling the CDF is chosen as the preferred method; the impact of such modifications is more intuitive and insensitive to the underlying PDF distribution. In this way, the re-scaled operator can simply be expressed as,

$$\bar{u} = [\mathcal{I} + \mu \cdot \mathcal{D}]\{u\}. \tag{2.29}$$

2.3.1.1 Generalized Explicit Purser Filters

Central to the derivation of Purser filters are two things: 1) the selection of an appropriate PDF function, and 2) properly transforming the integral limits of the CDF in order to parameterize the response in terms of spectral functions with known difference representations. As a result, manipulating either of these components can lead to new generalizations of the Purser filters.

For example, one could derive a PDF shape – which implicitly controls the derivative of the filter's spectral response – and use Lagrange interpolation in order to re-write the distribution as a polynomial function. This potentially provides an alternate way of designing custom filter transfer functions. In this way, cut-off and sharpness properties can be manipulated. Alternatively, one could use such a procedure in order to derive self-similar stencils that maintain the proper damping characteristics between grids of different resolution, useful for conducting studies of grid-converged LES (Radhakrishnan and Bellan, 2012).

Regarding the second suggestion provided above, one can alternatively seek to manipulate the integral limits. For instance, a general high-order term such as $\hat{g}_\ell(k) = [\sum_{i\geq 1} -\epsilon_{2i} \sin^{2i}(k\Delta x/2)]$ could be used, permitting that $\hat{g}_\ell \in [0, 1]$. Such transformations could be used to cleverly modulate how the CDF behaves with respect to wavenumbers k, thus allowing for additional customizability of the filter response shape (Edoh et al., June 2016). A natural inclination is to use the traditional Purser filters as the integral limit. Defining the integral limit function as a traditional Purser filter of parameters $R_\ell = C_{R_\ell} + C_{S_\ell} + 1$, the resulting generalized Purser stencil then reads,

$$\hat{\mathcal{G}}_{\text{Purser, gen}}(k) = 1 - \left[\frac{(C_R + C_S + 1)!}{C_R! C_S!}\right] \sum_{k=0}^{C_S} \begin{pmatrix} C_S \\ k \end{pmatrix} \frac{(-1)^k}{(k + C_R + 1)} z^{(k+C_R+1)} \Big|^{z=\hat{g}_\ell} , \tag{2.30}$$

with the a difference form of the stencil being,

$$\bar{u}_i = \left[1 + \sum_{r=C_R+1}^{R} \epsilon_{EF,2r} \left\{ \sum_{r_\ell=C_{R_\ell}+1}^{R_\ell} (-1) \cdot \epsilon_{EF,2r_\ell} (\Delta x)^{2r_\ell} \delta^{2r_\ell} \right\}^r \right] \{u_i\} , \tag{2.31}$$

where coefficients $\epsilon_{EF,2r_\ell}$ are based on applying $[C_{R_\ell}, C_{S_\ell}]$ to Equation 2.24. In this way, one naturally creates stencils based on powers of $(\sum_{r_\ell}^{R_\ell} \delta^{2r_\ell})$.

The overall filter stencil is of length $2R_\ell \cdot R + 1$, while the formal filter order is $2 \cdot (C_{R_\ell} + 1) \cdot (C_R + 1)$. The median, can furthermore be estimated using a generalization of 2.25,

$$\hat{g}_\ell(k_{1/2}) \approx \frac{C_R + 2/3}{C_R + C_S + 4/3} . \tag{2.32}$$

In order to demonstrate the effects of such a modification, Figure 2.8 looks at the effects of $[C_R, C_S]$ on a $R = 3$ Beta distribution that utilizes integral limits based on $[C_{R_\ell}, C_{S_\ell}] = [0, 2]$ and $[C_{R_\ell}, C_{S_\ell}] = [2, 0]$ stencils. In other words, the integral limits correspond to the $\{-\mathcal{D}\}$ operators of $2R = 6$ Binomial and Shapiro filters, respectively. These are then compared to the basis set of traditional Purser filters (e.g., $[C_{R_\ell}, C_{S_\ell}] = [0, 0]$) of equivalent stencil width, $R = 9$. The higher R

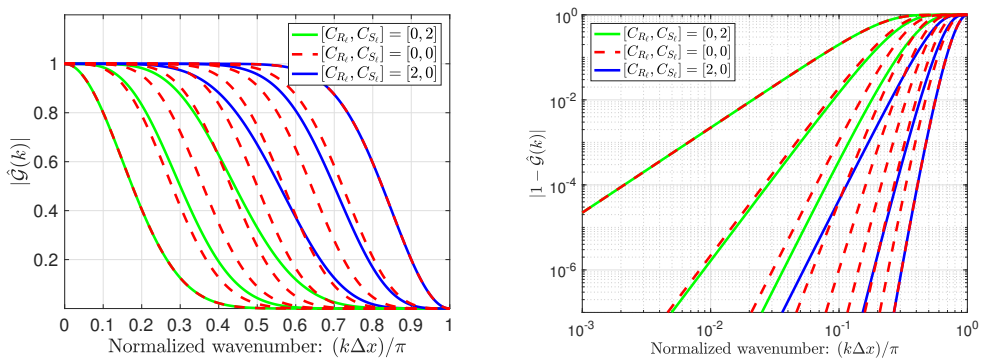

(a) (b)

Figure 2.8: Comparison of traditional Purser filter with $R = (C_R + C_S + 1) = 9$ (in red) versus generalized Purser filter response basis (see Equation 2.31) with $R = (C_R + C_S + 1) = 3$ for $[C_{R_\ell}, C_{S_\ell}] = [0, 2]$ (in green) and $[C_{R_\ell}, C_{S_\ell}] = [2, 0]$ (in blue) : a) Growth factor $|\hat{\mathcal{G}}|$, and b) Growth factor error $|1 - \hat{\mathcal{G}}|$.

value leads to a much richer set of responses for the given stencil size, which translates to more precise interpolation and manipulation of the filter characteristics. Comparing the classic Purser filter to the proposed generalizations, it is evident that the set $[C_{R_\ell}, C_{S_\ell}, C_R, C_S] = [2, 0|2, 0]$ is equivalent to $[0, 0|8, 0]$ (an eighteenth-order, nineteen point Shapiro filter) and that the set $[C_{R_\ell}, C_{S_\ell}, C_R, C_S] = [0, 2|0, 2]$ is equivalent to $[0, 0|0, 8]$ stencil (a second-order, nineteen point Binomial filter). Besides these, the generalized Purser formulations are shown to yield a new distinct set of response functions. New interpolations can furthermore be constructed by mixing the bases generated from different integral functions. In general, the traditional Purser schemes are preferable; although, utilizing the generalized forms in "cascade" filtering techniques (i.e., repeated applications of a base operator as done by Jeanmart and Winckelmans (2007)) may at times be algorithmically more efficient.

2.4 Implicit Stencil Families

Implicit filtering stencils – otherwise referred to as Compact or Padé schemes – utilize the surrounding filtered variables, $\bar{u}$, in addition to the surrounding input variables, u. Applying the interpretation of filtering as a convolution operation, this amounts to the following general form:

$$\int \left[G_\ell(x-y) \cdot \bar{u}(x) \right] dy = \int \left[G_r(x-y) \cdot u(x) \right] dy \ . \tag{2.33}$$

These implicit schemes require the solution of a linear system and thus incur additional computational effort. Nevertheless, they remain attractive due to their ability to match the capabilities of high-order explicit schemes on a compact stencil width. For instance, Gaitonde and Visbal (2000) employ biased tri-diagonal type Padé filters in order to maintain high-order boundary closures. Also, ideal low-pass representations can be well-approximated despite short stencil widths (Alpert, 1981). The benefits from this can be understood in another way: implicit stencils can achieve high performance without needing very high discrete derivative operators that become increasingly dubious to use on non-smooth, under-resolved features. This furthermore highlights the potential benefits for the use of implicit filtering stencils in LES computations.

In the case of a uniform grid, symmetric implicit stencils are written as,

$$\bar{u}_i + \sum_{\ell=1}^{L} a_\ell \left(\bar{u}_{i+\ell} + \bar{u}_{i-\ell} \right) = b_0 u_i + \sum_{r=1}^{R} b_r \left(u_{i+r} + u_{i-r} \right) \ , \tag{2.34}$$

with spectral function,

$$\hat{\mathcal{G}}(k) = \frac{b_0 + \sum_{r=1}^{R} 2b_r \cdot \cos(rk\Delta x)}{1 + \sum_{\ell=1}^{L} 2a_\ell \cdot \cos(\ell k \Delta x)} \ , \tag{2.35}$$

and a difference operator representation,

$$\left[\epsilon_{IF,0} + \sum_{\ell=1}^{L} \epsilon_{EF,2\ell} (\Delta x)^{2\ell} \delta_x^{2\ell} \right] \{\bar{u}_i\} = \left[\epsilon_{EF,0} + \sum_{r=1}^{R} \epsilon_{EF,2r} (\Delta x)^{2r} \delta_x^{2r} \right] \{u_i\} \ , \tag{2.36}$$

such that,

$$[\mathcal{I} + \mathcal{D}_\ell]\{\bar{u}\} = [\mathcal{I} + \mathcal{D}_r]\{u\} \ . \tag{2.37}$$

Again, consistency with the preservation of constant modes is achieved by setting $\epsilon_{IF,0} = \epsilon_{EF,0}$. In the case where one looks to scale the damping properties of the stencil in similar fashion to Equation 2.29, manipulation of the spectral response in Equation 2.35 to be of the form $\hat{\mathcal{G}} = (\hat{\mathcal{I}} + \mu\hat{\mathcal{D}})$ suggests the following operator arrangement for implicit stencil re-scaling:

$$[\mathcal{I} + \mathcal{D}_\ell]\{\bar{u}\} = [\mathcal{I} + \mu \cdot \mathcal{D}_r + (1-\mu) \cdot \mathcal{D}_\ell]\{u\} \ . \tag{2.38}$$

The formal order of the stencil can furthermore be determined by Equation 2.11 or can be estimated from,

$$\bar{u} \approx \frac{1}{\epsilon'_{IF,0}} \cdot \left[\epsilon'_{EF,0} u + \sum_{r=0}^{\infty} \epsilon'_{EF,2r} (\Delta x)^{2r} \partial^{2r} u - \sum_{\ell=0}^{\infty} \epsilon'_{IF,2\ell} (\Delta x)^{2\ell} \partial^{2\ell} u \right] \ , \tag{2.39}$$

where the prime notation (ϵ') takes into account the Taylor-series contributions of the dissipative differencing stencils, δ_x^{2r}.

Perhaps the most widely used implicit filter stencils are the tri-diagonal and penta-digaonal schemes of Lele (1992), which take advantage of tractable matrix inversion techniques (e.g., the Thomas algorithm) and furthermore offer simple tuning of the filter cut-off k_c. But while these stencils are largely motivated by Taylor-series analysis, it is possible to derive more general implicit formulas based spectral analysis. Most simply, one can consider the transfer function of Equation 2.35 as a rational function,

$$\hat{\mathcal{G}}(k) = \frac{\hat{\alpha}(k)}{\hat{\beta}(k)} \ . \tag{2.40}$$

In order to guarantee a stable scheme, one then needs to ensure:

$$|\hat{\alpha}(k)| \leq |\hat{\beta}(k)| \ \forall \ k \quad \text{s.t.} \quad |\hat{\mathcal{G}}(k) \leq 1| \ . \tag{2.41}$$

This observation can be used to greatly simplify and generalize the derivation of new implicit filter stencils.

The following section thus builds on explicit schemes from the previous section by introducing the class of implicit Purser filters, which includes the sub-class of implicit Shapiro stencils (also referred to as Long filters (Long et al., 1978)). In addition to this, the family of Tangent stencils (Raymond, 1988) is presented and shown to be attractive due to its enhanced sharpness capabilities.

2.4.1 Implicit Purser Filters

The proposed extension of Purser filters to implicit stencils is first motivated by considering the Long filter (Long et al., 1978) (also referred to as an implicit Shapiro filter) which takes the spectral form,

$$\hat{\mathcal{G}}_{\text{Long}}(k) = \frac{1 - \sin^{2R}(k\Delta x/2)}{1 - (1 - \delta_\Delta)\sin^{2R}(k\Delta x/2)} \ , \quad \text{with} \quad \delta_\Delta \in \langle 0, \infty \rangle \ . \tag{2.42}$$

The tuning parameter δ_Δ is used to prescribe the effective cut-off of the filter. The response approaches an ideal low-pass as $\delta_\Delta \to 0$, which corresponds to a spectral-sharp filter with $\Delta = (2\Delta x)$. This behavior is attainable regardless of the formal order of the filter. For instance, Figure 2.9(a) plots the response of a second-order Long filter for various values of the tuning parameter and compares these to a high-order 10th-order explicit Shapiro filter. While the low-wavenumber performance is dominated by the formal order of the scheme (see Figure 2.9(b)), the low-order implicit method is able to surpass the tenth-order explicit Shapiro filter at high wavenumbers, in terms of estimating an ideal low pass. This points to the strong potential utility of these implicit schemes.

Since the Long filter is built on explicit Shapiro stencils (which are a sub-class of the Purser filters), it is reasonable to think that general implicit Purser filters can be built in a similar fashion. These are written to take the operator form,

$$[\mathcal{I} + (1 - \delta_\Delta) \cdot \mathcal{D}_{\text{Purser,exp}}]\{\bar{u}\} = [\mathcal{I} + \mathcal{D}_{\text{Purser,exp}}]\{u\} \ , \tag{2.43}$$

with the spectral representation:

$$\hat{\mathcal{G}}_{\text{Purser,imp}}(k) = \frac{\hat{\mathcal{G}}_{\text{Purser,exp}}(k)}{1 + (1 - \delta_\Delta) \cdot [\hat{\mathcal{G}}_{\text{Purser,exp}}(k) - 1]} \ , \quad \text{with} \quad \delta_\Delta \in \langle 0, \infty \rangle \ . \quad (2.44)$$

According to the requirement of Equation 2.41, it is apparent that this form is naturally stable and that it recovers the equivalent explicit Purser stencil when $\delta_\Delta = 1$, a fact that will prove consequential with respect to assessing the invertibility of the stencil (see Section 2.6.1). The resulting implicit Purser filter can be shown to be of equivalent order and stencil width as the underlying explicit Purser stencil that it is built from. Figure 2.10(a) plots the response of an implicit Purser filter built from $[C_R, C_S] = [0, 1]$ (i.e., a Binomial filter), which has a five-point stencil width. By inspecting the response error plot of Figure 2.10(b), it is evident that the original order of the explicit filter is $2 \cdot (C_R + 1)$, and that this is maintained for the implicit scheme. By looking at the response itself, we notice that smaller values of the tuning parameter δ_Δ are required to approximate an ideal low-pass filter. This is related to the underlying explicit Purser stencil which is sub-optimal in terms of low-wavenumber preservation – this is in contrast to the Shapiro stencils utilized in Long filter formulation.

In keeping with the spectrally-motivated derivation of implicit stencils, it is also possible to think of Padé stencils that are made from different explicit Purser schemes. This would allow the construction of tri-diagonal type schemes that could be used for efficient inversion. For example, the following is a general expression for Purser-based Padé filters of different implicit and explicit stencil widths:

$$[\mathcal{I} + \mathcal{D}^{(C_{R_1}, C_{S_1})}_{\text{Purser,exp}}] \ \{\bar{u}\} = [\mathcal{I} + \mathcal{D}^{(C_{R_2}, C_{S_2})}_{\text{Purser,exp}}] \ \{u\} \ . \quad (2.45)$$

In this case, neither stability nor formal order are naturally guaranteed; they depend on the judicious selection of $[C_{R_1}, C_{S_1}]$ and $[C_{R_2}, C_{S_2}]$. These parameters would also be responsible for tuning the filter cut-off, although this is much less evident than the simple way in which the schemes of Equation 2.44 are tuned (via

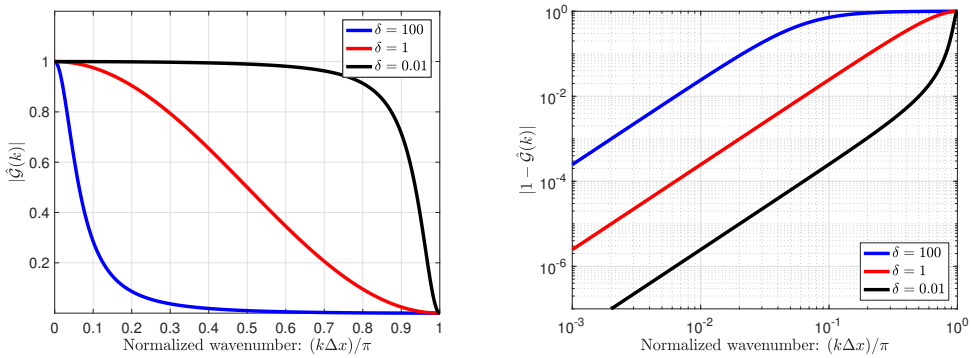

(a) (b)

Figure 2.9: Second-order Long filter response as a function of tuning parameter δ_Δ (see Equation 2.42), also equivalent to a implicit Purser filter response built from explicit stencil $[C_R, C_S] = [0, 0]$: a) Growth factor $|\hat{\mathcal{G}}|$, and b) Growth factor error $|1 - \hat{\mathcal{G}}|$.

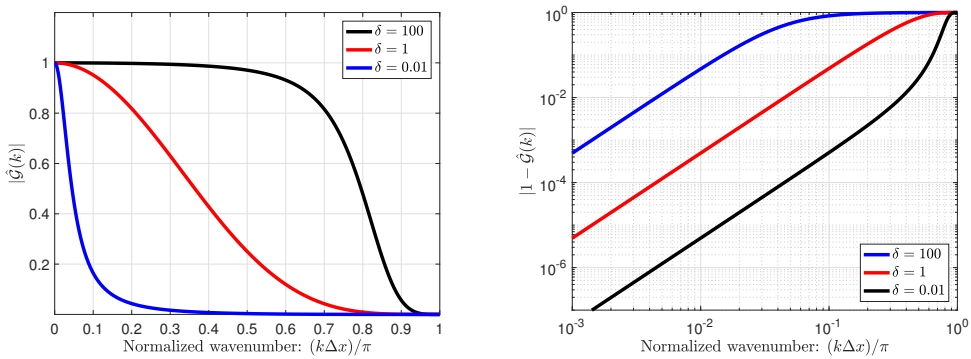

(a) (b)

Figure 2.10: Second-order, five-point implicit Purser filter response built from explicit stencil $[C_R, C_S] = [0, 1]$ as a function of turning parameter δ_Δ (see Equation 2.44): a) Growth factor $|\hat{\mathcal{G}}|$, and b) Growth factor error $|1 - \hat{\mathcal{G}}|$.

simple coefficient scaling of the implicit stencil by δ_Δ). As before, Equation 2.41 would need to be satisfied in order to guarantee stability. In addition, considering Equation 2.39 suggests the formal order will be determined by $\min(C_{R_1}, C_{R_2})$, which is impractical since changing these is necessary for manipulating the spectral shape. This differs from the Lele schemes which are constructed around Taylor-series considerations and so are able to maintain their formal order despite changes to the tuning parameter. A similar feature may be achievable with the scheme proposed in Equation 2.45, but this is not easily evident. Nevertheless, it is important to acknowledge the special case of Equation 2.45 in which $\mathcal{D}_{\text{Purser,exp}}^{(C_{R_2},C_{S_2})} = 0$, as it relates to a discrete generalization of the inverse Helmholtz differential filter of Germano (1986)and its hyper-viscous forms (i.e., Butterworth filters (Butterworth, 1930)). These schemes have the benefit of being naturally invertible (i.e., $|\hat{\mathcal{G}}(k)| > 0$), which facilitates deconvolution procedures (e.g., retrieval of the unfiltered solution from filtered data), necessary as part of certain LES modeling techniques (Bull and Jameson, 2016).

Briefly retreating to the fundamental idea that Purser filters are derived from CDF distributions under a change of variable, it also becomes evident that implicit Purser schemes arise from rational polynomial representations of PDF functions, $f_{PDF}(z)$. While these may be derived from an optimization procedure, one may also consider the one-sided distributions (only valid for $z \geq 0$) listed in Table 6.2. As previously stated, using such established distributions is advantageous as their shape characteristics are well characterized by function parameters and can then be used to describe the derived filter transfer function. Inspecting the polynomial form of the respective CDFs, one can surmise that substituting $z = \hat{g}_\ell(k)$ would result in implicit representations. However, because these distributions are one-sided, the CDF does not reach unity on a finite domain which means that the resulting transfer function of the filter will naturally be invertible. Increasing the high-wavenumber damping is then tantamount to scaling up the integral limit of

Distirbution	$f_{PDF}(z)$	$F_{CDF}(z)$
Log-logistic	$\frac{(\beta/\alpha)(z/\alpha)^{\beta-1}}{(1+(z/\alpha)^{\beta})^2}$	$\frac{z^{\beta}}{\alpha^{\beta}+z^{\beta}}$
Dagum	$\frac{ap}{z}\left(\frac{(z/b)^{ap}}{[(z/b)^{a}+1]^{p+1}}\right)$	$[1+(z/b)^{-a}]^{-p}$

Table 2.3: One-sided probability and cumulative distribution functions (Tadikamalla, 1980).

Equation 2.20 by a factor $\mu \geq 0$. These proposed schemes, along with the generalizations introduced above in Equations 2.43 and 2.45 provide great flexibility in designing the spectral characteristics of implicit filters.

2.4.2 Tangent Filters

The family of implicit Tangent filters introduced by Raymond (1988) are defined by the transfer function,

$$\hat{\mathcal{G}}_{\text{Tangent}}(k) = \frac{\cos^{2R}(k\Delta x/2)}{\left[\cos^{2R}(k\Delta x/2)+\delta_{\Delta}\sin^{2R}(k\Delta x/2)\right]} \tag{2.46}$$

$$= \left[1+\delta_{\Delta}\tan^{2R}(k\Delta x/2)\right]^{-1} \quad , \quad \text{with} \quad \delta_{\Delta} \in [0,\infty\rangle \ , \tag{2.47}$$

with the difference form of Equation 2.36 being defined by the following coefficients:

$$\begin{aligned} \epsilon_{EF,2r\geq 2} &= \begin{pmatrix} R \\ R-r \end{pmatrix} \frac{1}{4^r} \ , \\ \epsilon_{IF,2\ell\neq 2L} &= \epsilon_{EF,2r} \ , \\ \epsilon_{IF,2L} &= [1+(-1)^{R}\delta_{\Delta}]\epsilon_{EF,2R} \quad \text{with} \quad L=R \ . \end{aligned} \tag{2.48}$$

Re-written in the form of Equation 2.47, the Tangent filter can be understood to be a combination of averaging $[\cos^{2R}(k\Delta x/2)]$ and derivative $[\sin^{2R}(k\Delta x/2)]$ operators – in other words, it cleverly combines Binomial and Shapiro stencils (Raymond and Garder, 1991).

The tuning parameter δ_Δ is used to control the effective filter cut-off frequency, and the spectral response has the special property of yielding $\Delta_{1/2} = (4\Delta x)$ for $\delta_\Delta = 1$, regardless of the filter order $2R$. This suggests the stencil to be naturally well-conditioned for mid-wavenumber damping (see Section 2.6.1 for more on stencil conditioning) and thus amenable to filtering for aliasing control, which requires a $\Delta_c = (3\Delta x)$ cut-off for quadratic non-linearities (Orszag, 1971) (e.g., incompressible flow) and a $\Delta_c = (4\Delta x)$ cut-off for cubic non-linearities (e.g., compressible flow in terms of primitive variables).

2.5 Spectral Response Characteristics

Sections 2.3 and 2.4 above have introduced explicit and implicit discrete filter families that offer the useful flexibility of specifying spectral response characteristics. Amongst the desired aspects to manipulate are the effective filter width, k_c, as well as the degree of scale-discriminant attenuation. In the context of scheme stabilization (see Chapter 3), this can play a role in terms of minimizing dissipative error. Meanwhile, with respect to LES, such metrics dictate the amount of scale-separation and thus inherently define the type of information embedded within the resolved scales.

Expanding on the relevance to LES, focus is traditionally placed on the filter width as this identifies the effective LES resolution. However, the definition is a vague characterization, and as a result, two filters of equivalent k_c may have very different spectral sharpness behavior. Such an example is the top-hat and spectral sharp filters shown in Figure 2.1(b). Evidently, the top-hat filter exhibits a stronger roll-off at low-wavenumbers which suggests suppressed activity in the large-scales compared to the spectral-sharp filter. This furthermore implies that a model corresponding to the former would need to supply dynamical information across a range of scales, rather than the classic perspective of only needing to

represent features below the cut-off. Also important to consider are the implications of response sharpness for implicit LES (ILES) formulations that utilize periodic filtering of the solution in order to represent the presumed dissipative effects of a closure model on the resolved field (also termed relaxation-filtering (Berland et al., 2008; Fauconnier, 2013; Aubard et al., 2013)). In this context, repeated applications of a filter operation can remove significant signal content depending on the scale-preservation properties of the scheme, thus lowering the effective cut-off wavenumber over time (i.e., systematically coarsening the effective LES resolution). As a result, assessing the filter cut-off along with the associated scale-discriminant characteristics can be important in evaluating the implications of applying different filter stencils.

The effective filter width may be defined in various ways. Most traditionally, one uses the second-moment of the filter kernel,

$$\Delta_c = \sqrt{\int x^2 G(x) dx} \ , \tag{2.49}$$

thus identifying the standard deviation of the function as the cut-off scale. This definition, however, is not valid for filters of order greater than two because of the vanishing moments related to filter order. As a result, alternate and accurate characterizations are important and necessary. Lund (1997) highlights how erroneously quantifying the effective cut-off metric can introduce significant levels of error in dynamic LES modeling, for example. Examining the evolution of resolved kinetic energy, Lund notes the sensitivity of results to improper specification of filter-to-grid ratio, $\alpha = (\Delta_c / 2\Delta x)$; he demonstrates noticeable error between even slight deviations of $\alpha = 2.0$ versus $\alpha = 1.96$

To aid in the proper description of the filter width for discrete filters, Lund proposes three alternatives to the traditional second-moment designation: 1) using the cut-off wavenumber k_c corresponding to the second moment of spectral response function $\hat{\mathcal{G}}(k)$; 2) choosing $k_c = k_{1/2}$ such that $\hat{\mathcal{G}}(k_{1/2}) = 0.5$; and 3)

fitting a trigonometric interpolation to the discrete filter weights in physical space and choosing the first zero-crossing as the effective cut-off, noting that the resulting interpolant is analogous to the spectral filter's physical-space sinc function kernel. Of these three choices, option 2 is most typically used due to its simplicity, although option 1 is deemed most accurate by Lund. The subsequent analyses assume this definition, unless noted otherwise.

2.5.1 Scale-Separation Performance

As previously noted, the nature of scale-separation can also influence requirements of LES closure model and implicitly defines the way information is preserved on the resolved scales, $k \gtrsim k_c$. In this respect, it is useful to parametrize the performance of the available discrete filter schemes. This is explored here as a function of cut-off resolution ($k_c = k_{1/2} = \pi/\Delta_{1/2}$) and of the filter order, which is directly related to preservation properties of the scheme. As part of this, sharpness is also quantified by inspecting a measure of the transition width:

$$\Delta\theta = \theta_2 - \theta_1 \quad \text{s.t.} \quad \begin{cases} |\hat{\mathcal{G}}(\theta_1 = k_1\Delta x)| = 0.95 \cdot |\hat{\mathcal{G}}(\theta = 0)| + 0.05 \cdot |\hat{\mathcal{G}}(\theta = \pi)| \\ |\hat{\mathcal{G}}(\theta_2 = k_2\Delta x)| = 0.95 \cdot |\hat{\mathcal{G}}(\theta = \pi)| + 0.05 \cdot |\hat{\mathcal{G}}(\theta = 0)| \end{cases} . \tag{2.50}$$

The spectral response comparisons highlight the utility and potential limitations of the different stencils. The explicit Purser, implicit Purser, and Tangent stencils are considered. In addition, the popular tri-diagonal scheme of Lele (1992) is considered and defined for arbitrary order $2R$ as the following,

$$\hat{\mathcal{G}}_{\text{Lele,tri}}(k) \;=\; \frac{b_0 + 2\sum_{r=1}^{R} b_r \cos\left(r\ k\Delta x\right)}{1 + 2\delta_\Delta \cos\left(k\Delta x\right)} \;, \tag{2.51}$$

where the stencil coefficients are constructed from Taylor-series considerations as,

$$\begin{array}{rcl} 1 & = & b_0 + 2\left[\left(\sum_{r=1}^{R} b_r\right) - \delta_\Delta\right] \ , \\ 0 & = & \left(\sum_{r=1}^{R} b_r \ r^m\right) - \delta_\Delta \ , \quad \text{with } m = 2 \cdot [1, ..., R-1] \ , \\ 0 & = & b_0 + 2\sum_{r=1}^{R}(-1)^r b_r \ , \\ 1/2 & = & b_0 + 2\left[\sum_{r=1}^{R} b_r \cos\left(r \ k_{1/2}\Delta x\right)\right] - \delta_\Delta \cos\left(k_{1/2}\Delta x\right) \end{array} \quad , \tag{2.52}$$

with $\delta_\Delta \in \langle 0.5, 0.5 \rangle$ for stable operation. Figures 2.11 and 2.12 show the performance maps for different discrete filtering schemes and communicates their respective spectral sharpness characteristics as a function of the filter width ($\Delta_{1/2}$) and the filter order ($2R$). Also included is information on the associated stencil width, speaking to the potential cost of the scheme.

Figure 2.11 shows the map of the classic explicit Purser filters. The class of Binomial stencils is found on the constant line for second-order filters; these are shown to have the lowest filter resolution $\Delta_{1/2}$ (i.e., lowest $\theta_{1/2} = k_{1/2}\Delta x$) for a given stencil length. On the other hand, one can identify the Shapiro filters as the schemes yielding the smallest filter cut-off (i.e., highest $\theta_{1/2} = k_{1/2}\Delta x$) for a given stencil length; they form the right boarder of the map. In general, it is apparent that achieving sharp responses requires very wide stencils. For example, a transition width of approximately $\Delta\theta = 0.2\pi$ is achieved with a forty-point explicit Purser scheme. Such hinderances serve to highlight the potential benefits of employing high-order implicit formulations instead.

Figure 2.12(a) shows the performance map for the class of Long filters, which can be interpreted as implicit Purser filters built from Shapiro stencils. As before, it it apparent that sharpness increases with the filter order. Furthermore, tracking the lines of constant order, it is apparent that increasing $2R$ eventually leads to strong asymmetry in the performance map relative to $\theta_{1/2}$ (at least based on the chosen definitions of θ_1 and θ_2). The cut-off wavenumber at which sharpness is sub-optimal can be seen to coincide with the underlying explicit stencil from which the Padé representation is built (i.e., the case when $\delta_\Delta = 1$). Therefore, as

the formal order of the scheme is increased, the Long stencil is seen to be largely under-performing with respect to sharpness for $\Delta_{1/2} \sim (3\Delta x)$.

In order to recover performance in this range, it is possible to consider a Padé representation built from another explicit Purser stencil such that the sub-optimal point is shifted away from the desired cut-off range. This is demonstrated in Figure 2.12(b), which showcases the use of explicit stencils tuned to $\Delta_{1/2} = (4\Delta x)$ to build the implicit method. In this way, the scheme is made sub-optimal at this mid-wavenumber and sharpness is regained for use at the higher and lower modes. A consequence of these manipulations is the fact that larger stencils are required for a given scheme order, and this stems from the fact that the base explicit Purser filters to be employed are sub-optimal in this respect. Therefore, further shifting the design point towards coarser resolutions would entail an additional price in the required stencil width (for a given order scheme).

Judiciously selecting the underlying explicit Purser stencil in the derivation of implicit schemes can also be applied to the naturally-invertible implicit Purser formulation of Equation 2.45 (taking $\mathcal{D}^{(C_{R_2},C_{S_2})} = 0$). In the case Shapiro type stencils are utilized, one recovers the implicit Sine filters (Raymond and Garder, 1991) (discrete analogues to the inverse Helmholtz filter of Germano (Germano, 1986) and hyper-viscous interpretations such as the Butterworth filters (Butterworth, 1930)). In the instance one uses arbitrary explicit Purser stencils, one can produce generalized Sine filters. In both cases, the stencils are inherently invertible ($|\hat{\mathcal{G}}| > 1$), which may be desirable for LES modeling requiring deconvolution procedures (Bull and Jameson, 2016). The performance maps of such filtering methods are respectively shown in Figure 2.12(c) and 2.12(d) for an implicit Sine filter and a generalized implicit Sine filter derived from Purser stencils with mid-wavenumber cut-offs tuned to $\Delta_{1/2} = (4\Delta x)$. As before, using non-Shapiro stencils incurs a cost in stencil width relative to the formal order of the scheme. Nevertheless, doing so can allow one to achieve sharper characteristics as

demonstrated in comparing Figures 2.12(c) and 2.12(d).

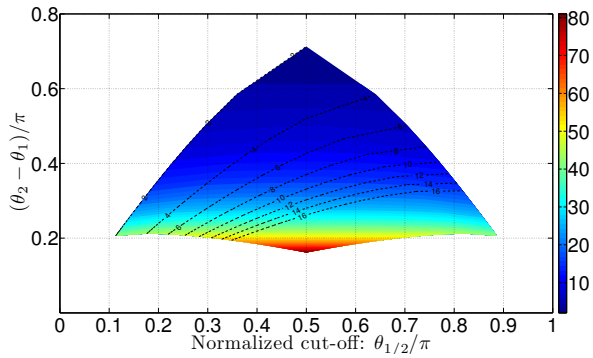

Figure 2.11: Map of filter transition width ($\Delta\theta/\pi$) (as defined by Equation 2.50) for traditional explicit Purser filters (see Equation 2.22) versus filter cut-off frequency ($k_{1/2}\Delta x/\pi$) (as defined by Equation 2.25) for $2R$-th order filters (denoted by dashed contour lines), colored by stencil width ($2R+1$).

The performance map of the Tangent stencil is provided in Figure 2.12(e). Most evident is the fact that its sharpness characteristics are superior to the other forms presented and that the map is naturally symmetric about $\Delta_{1/2} = (4\Delta x)$, demonstrating enhanced scale-discriminant damping away from this mid-wavenumber cut-off – as previously noted, a property desired for the maintenance of spectral sharpness towards low $k_{1/2}$. This shows great promise for its use in high-order LES and stabilization implementations, where one hopes to limit the impact of dissipative effects on the overall scheme accuracy.

Finally, Figure 2.12(f) plots the map for the tri-diagonal Lele scheme defined by Equations 2.51-2.52. Interestingly, the well-used method, has a very restricted range of applicability that results from the fact that stability is only achieved for the tuning parameter $\delta_\Delta \in \langle 0.5, 0.5 \rangle$. Therefore, as the filter order is increased, the minimum cut-off wavenumber shifts up and the range of attainable cut-offs is also reduced. The resulting map can furthermore be seen as a direct subset of the Long filter map shown in Figure 2.12(a). In order to mitigate such limitations of this tri-diagonal scheme, it may be beneficial to employ formulations with larger

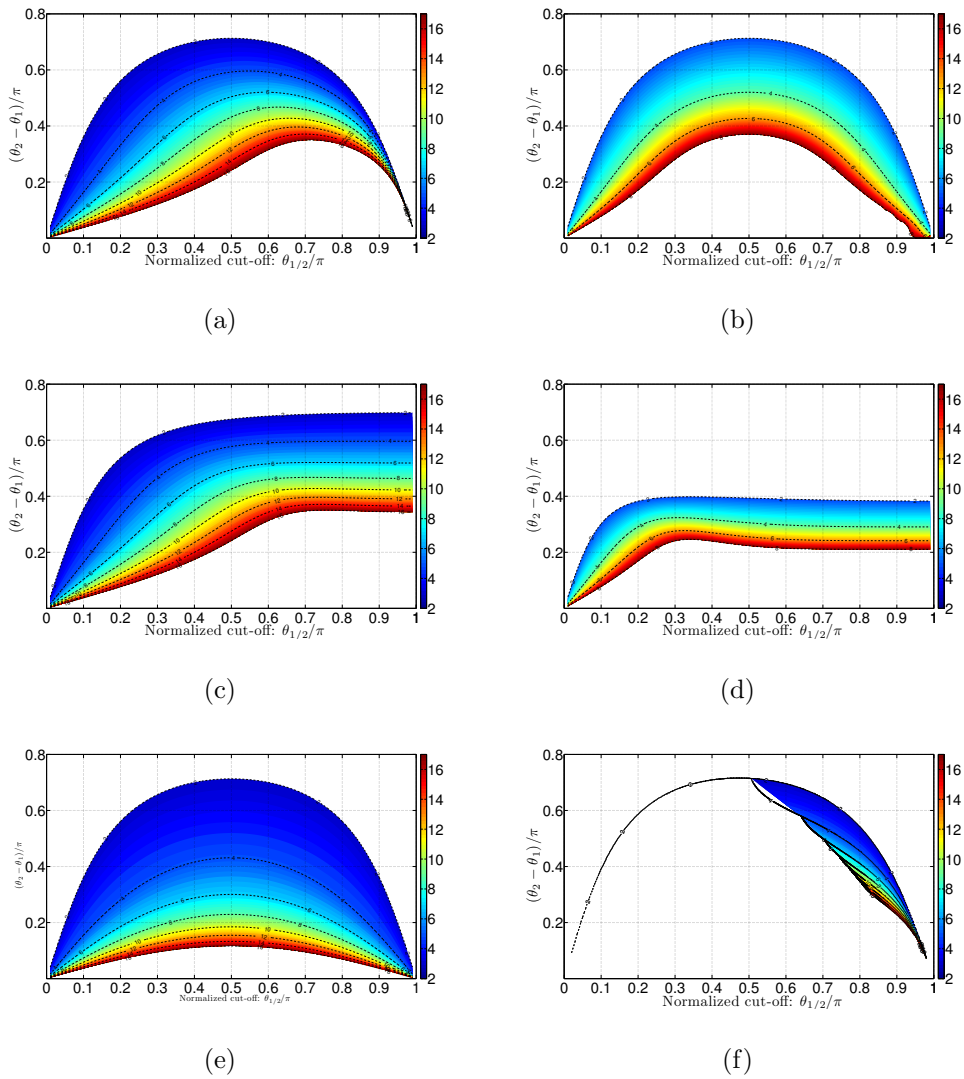

Figure 2.12: Map of filter transition width ($\Delta\theta/\pi$) (as defined by Equation 2.50) versus filter cut-off frequency ($k_{1/2}\Delta x/\pi$) for $2R$-th order filters (denoted by dashed contour lines), colored by stencil width ($2R+1$): a) Long filter, b) implicit Purser filter based on Equation 2.43 with $\mathcal{D}^{(C_R,C_S)}_{\text{Purser,exp}}$ tuned to $\Delta_{1/2} = (4\Delta x)$, c) implicit Sine filter based on Equation 2.45 with $\mathcal{D}^{(C_{R_2},C_{S_2})}_{\text{Purser,exp}} = 0$ and $[C_{R_1}, C_{S_1}] = [R-1, 0]$, d) generalized implicit Sine filter based on Equation 2.45 with $\mathcal{D}^{(C_{R_2},C_{S_2})}_{\text{Purser,exp}} = 0$ and $[C_{R_1}, C_{S_1}]$ tuned to $\Delta_{1/2} = (4\Delta x)$, e) Tangent filter, and (f) Lele tri-diagonal filter.

implicit operators (e.g., penta-diagonal formulations etc...); however, this detracts from the attractiveness of these Lele-type schemes: the ability to easily invert the stencils.

The featured performance maps give useful insight on the way in which one can employ the different classes of discrete filters. Again, it should be noted that the importance of scale-discriminant attenuation is non-linear in terms of the cut-off frequency. For example, a transition length of $\Delta\theta/\pi = 0.2$ entails very different things at high cut-off wavenumbers compared to low cut-off wavenumbers. For a low cut-off, this could correspond to making the distinction between a ten points-per-wave (e.g., $\theta/\pi = 0.2$) and a five points-per-wave (e.g., $\theta/\pi = 0.4$) feature. In this case, it would seem that most of the schemes are comparable in performance. Meanwhile, for a high cut-off, a transition length of $\Delta\theta/\pi = 0.2$ could correspond to making the distinction between a four points-per-wave (e.g., $\theta/\pi = 0.5$) and a three points-per-wave (e.g., $\theta/\pi = 0.7$) feature. The inability to make careful distinctions between $4\Delta x$ and $3\Delta x$ resolutions, for example, also has important consequences in terms of aliasing control. Based on the performance maps, it can then be surmized that the invertible sine filters are inadequate for deliminating between high-wavenumber features. Instead, it may be advisable to employ other stencils such as the Tangent or generalized implicit Purser filter idea featured in Figure 2.12(b). These can furthermore be rendered invertible by using a rescaling technique, as suggested in Equation 2.38 (also see Appendix A).

2.6 Further Practical Considerations

Several additional implementation specifics regarding discrete filtering could be addressed; however, the following section focuses on two: stencil invertibility and multi-dimensional use of the filter. With respect to the former, implicit schemes require one to consider complications that may arise from the need to invert a lin-

ear system; this may be assessed by considering the conditioning of the stencils, and doing so allows one to surmise useful guidelines for suggested use of the proposed discrete filter families. Secondly, practical application of filtering requires multi-dimensional formulations; therefore, it is important to understand different extensions of the one-dimensional schemes introduced thus far.

2.6.1 Stencil Invertibility

Inversion of a system introduces the potential for numerical complications associated with the stability of the procedure. In some cases, improperly formulated direct methods can allow small numerical error perturbations to drastically alter the final result; meanwhile, in the case of iterative procedures, one may encounter stalled convergence. The susceptibility to such inversion errors depends on the specific method but is also related to the system's conditioning. A condition number may be defined as,

$$\kappa = \max\{\lambda\}/\min\{\lambda\} \ , \quad \text{where} \quad \lambda = eig\{LHS\} \ . \tag{2.53}$$

Here, λ represents the eigenvalue spectrum of the linear system $[LHS]\cdot\overrightarrow{x} = RHS$. For the one-dimensional Padé filters on periodic boundary conditions, one may write the following:

$$\lambda(k) = 1 + \sum_{\ell=1}^{L} 2a_\ell \cdot \cos(\ell k \Delta x) \ . \tag{2.54}$$

Utilizing this information, it is possible to assess the conditioning associated with different implicit schemes as a function of filter order ($2R$) and filter width ($\Delta_{1/2}$). With respect to the stencils in consideration, it furthermore becomes evident that the conditioning of the system is strongly relatable to extreme values (either large or small) of the tuning parameter, δ_Δ.

Figure 2.13 shows color maps of $\log\{\kappa\}$ as a function of filter order and filter width for the different stencils. From these, it may be concluded that increased filter sharpness is concomitant with a larger condition number (i.e., an increas-

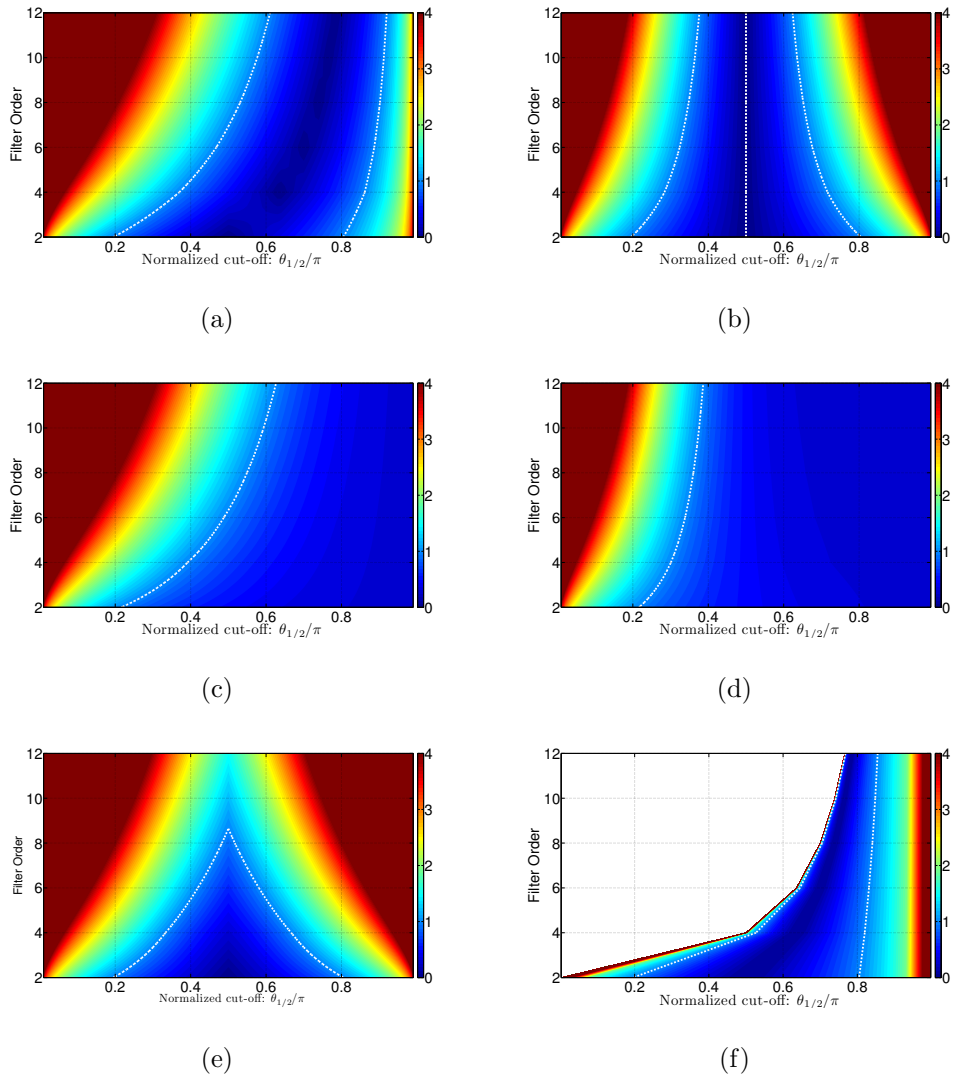

Figure 2.13: Map of filter order ($2R$) versus filter cut-off ($\theta_{1/2}/\pi$), colored by logarithm of condition number κ as defined by Equations 2.53-2.54 (dashed white line corresponding to $\kappa = 10$): (a)Long filter , (b) Tangent filter, (c) tri-diagonal Lele filter, and (d) implicit Purser filter with $[C_{R_\ell}, C_{S_\ell}] = [0, 0]$ and $C_R = C_S$.

ingly ill-conditioned system), thus communicating the difficulty of producing such scale-discriminant characteristics. Furthermore, the condition number maps replicate the previous skewness found in the scale-separation performance of the prior section. And so, while the spectral characteristics of the Tangent filter shown in Figure 2.12(e) are quite impressive, inspecting the condition number in Figure 2.13(e) reveals that practical use is limited to mid-wavenumbers and that the range of acceptable cut-offs is shrunk with increasing order of the filter. With respect to the Purser stencils, in Section 2.5.1, it is shown how the resulting sharpness performance can be manipulated by the underlying Purser stencil from which the final implicit filter is derived. Here, in Figures 2.13(a)-2.13(d), it is evident that the underlying explicit stencil also has an associated effect on the condition number of the associated implicit scheme. For example, the Long stencil has a minimal condition number for values of its tuning parameter ($\delta_\Delta = 1$) that are associated with its explicit scheme. By designing implicit schemes around explicit Purser filters with lower cut-offs, the area of optimal conditioning can thus be effectively shifted down. Of further note, it was previously observed that the performance map of Lele tri-diagonal stencils (Figure 2.12(f)) replicates that of the Long stencils (Figure 2.12(a)) at high-wavenumber cut-offs. But while this may be true, comparing their respective conditioning maps (Figures 2.13(f) and 2.13(a)) reveals that there are some differences in terms of conditioning, with the tri-diagonal Lele formulation demonstrating slightly poorer conditioning.

2.6.2 Implementation for Multi-dimensions

The familiarity gained with one-dimensional (1D) filter stencils may be extended to multi-dimensions. There are several ways in which this can be done, each offering a balance between performance and efficiency.

Perhaps most common, one can choose to employ the one-dimensional opera-

tors successively in each spatial direction:

$$\begin{aligned} \bar{u} &= \mathcal{G}_{\prod}\{u\} \ , \\ &= \mathcal{G}_1\{\mathcal{G}_2\{...\{u\}...\}\} \ . \end{aligned} \tag{2.55}$$

In spectral space, this amounts to a product of each directional transfer function and allows one to easily apply the desired attenuating properties to all directions. An example of this is plotted in Figure 2.14(a) which shows the magnitude of the resulting two-dimensional (2D) transfer function $\hat{\mathcal{G}}_{\prod}$ corresponding to a fourth-order 1D Shapiro filter, $\hat{\mathcal{G}}(k) = 1 - \sin^4(k\Delta x/2)$. The cost of performing multiple filtering operations is offset somewhat by the opportunity to leverage efficient banded matrix direct solvers. Nevertheless, such an implementation has a large stencil footprint, relating to the implied presence of cross-derivative terms.

As an alternative to the product approach, it is interesting to consider a one-step approach such as,

$$\begin{aligned} \bar{u} &= \mathcal{G}_{\sum}\{u\} \ , \\ \rightarrow \qquad & \left[\mathcal{I}_\ell + \sum_{n=1}^{\text{NDim}} \mathcal{D}_{\ell,n}\right]\{\bar{u}\} = \left[\mathcal{I}_r + \sum_{n=1}^{\text{NDim}} \mathcal{D}_{r,n}\right]\{u\} \ , \end{aligned} \tag{2.56}$$

where the one-dimensional dissipation operators are summed. This formulation avoids the footprint overhead by avoiding cross-derivative terms, which offers implementation efficiency. However, in terms of implicit filters, the resulting formulation is sparse and will likely require an iterative inversion technique (e.g., Krylov methods, Gauss-Siedel/Jacobi methods). An accompanying consequence of the $\mathcal{G}_{\sum}$ operator is the fact that the one-dimensional filtering properties are not extended to multi-dimensional features. For example, odd-even mode attenuation is not maintained for higher dimensional odd-even modes. This is demonstrated in Figure 2.14(b) which plots $|\hat{\mathcal{G}}_{\sum}|$. In addition, using the original one-dimensional coefficients can lead to under-damping (or even instability) of the multi-dimensional features and thus requires one to perform an appropriate

rescaling of the coefficients. The inconsistency arising between filters of differing dimensionality thus makes it harder to compare results derived in 1D versus 2D, etc... when employing this approach.

In order to recover the independence of the filter response relative to the dimensionality of the problem, it is possible to construct special mixed-derivative forms that preserve the 1D characteristics along mixed modes, $k = (k_1, k_2)$ with $k_1 = k_2$:

$$\begin{aligned} \bar{u} &= \mathcal{G}_{\sum^*}\{u\} \ , \\ \rightarrow \qquad & [\mathcal{I}_\ell + \mathcal{D}_{\ell,\mathrm{NDim}}]\{\bar{u}\} = [\mathcal{I}_r + \mathcal{D}_{r,\mathrm{NDim}}]\{u\} \ . \end{aligned} \tag{2.57}$$

Falissard (2013) works to extend the 1D attenuating properties to multi-dimensions and constructs an operator $\mathcal{D}_{\mathrm{NDim}}$ (for 2D) defined as,

$$\mathcal{D}_{2\mathrm{D}}^{\mathrm{Fal}} = \sum_{r=1}^{R} \epsilon_{EF,2r} \delta_{2\mathrm{D}}^{2r} \ , \tag{2.58}$$

$$\begin{aligned} \text{with} \quad \delta_{2D}^{2r} &= (\Delta x_1)^{2r} \delta_{x_1}^{2r} + (\Delta x_2)^{2r} \delta_{x_2}^{2r} - \\ & \sum_{n=1}^{r-1} \alpha_n (\Delta x_1)^{2n} (\Delta x_2)^{2r-2n} \ \delta_{x_1}^{2n} \ \delta_{x_2}^{2r-2n} \ , \end{aligned} \tag{2.59}$$

with additional constraints,

$$\alpha_n \geq 0, \quad \sum_{r-1}^{n=1} \alpha_n = 1, \quad \alpha_n = \alpha_{r-n} \ . \tag{2.60}$$

The reader is referred to the Falissard paper (Falissard, 2013) for the three-dimensional (3D) extension and further details. Relative to this difference form, the genuinely multi-dimensional formulation of Falissard has the advantage of utilizing the same 1D coefficients ϵ_{2r} as the 1D schemes it is derived from. Furthermore, the footprint of the resulting stencil can be manipulated based on the choice of α_n constraints (see Equation 2.60), which controls the influence of mixed derivative terms. Figure 2.14(c) shows the transfer function magnitude of $\hat{\mathcal{G}}_{\sum^{\mathrm{Fal}}}$ and confirms that the one-dimensional attenuation along $k = (k_1, 0)$ and $k = (0, k_2)$ are indeed valid for two dimensional features $k = (k_1, k_1)$.

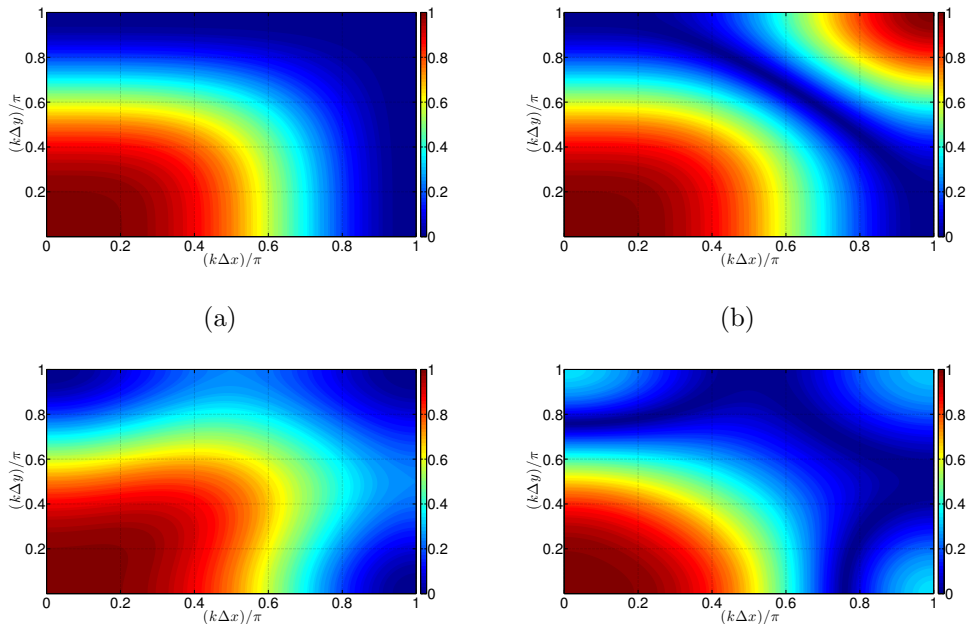

Figure 2.14: Color map of spectral responses of two-dimensional filtering implementations of a one-dimensional fourth-order Shapiro filter, $\hat{\mathcal{G}}(k) = 1 - \sin^4(k\Delta x/2)$: a) serial implementation $\mathcal{G}_{\prod}$, b) traditional single-step implementation $\mathcal{G}_{\sum}$, c) Falissard (Falissard, 2013) single-step implementation $\mathcal{G}_{\sum^{\text{Fal}}}$ (with $\alpha =$,), and d) anisotropic correction (AC) (Sescu et al., January 2009) single-step implementation $\mathcal{G}_{\sum^{AC}}$.

Although the Falissard implementation $\mathcal{G}_{\sum^{\text{Fal}}}$ successfully extends the 1D operators to multi-dimensions, the resulting function can still exhibit strong anisotropy in the mid-wavenumber range. Therefore the anisotropy-corrected (AC) schemes of Sescu et al. (January 2009) can be considered. These utilize information from the orthogonal and diagonal directions and seek to minimize the difference between the 1D and multi-dimensional spectral functions. The resulting scheme is

more simply expressed (for 2D) in terms in stencil form:

$$\bar{u}_{i,j} = \left(\frac{1}{1+\beta}\right)\sum_{r=-R}^{R} b_r \left\{\left[E_i^r + \frac{\beta}{2}D_x\right] + \left[E_j^r + \frac{\beta}{2}D_y\right]\right\}\{u_{i,j}\} \ ,$$

$$\text{where} \quad \begin{cases} D_x = (E_i^r E_j^r + E_i^{-r} E_j^r), \text{ and } D_y = (E_i^r E_j^r + E_i^r E_j^{-r}) \ , \\ E_i^r\{u_{i,j}\} = u_{i+r,j}, \text{ and } E_j^r\{u_{i,j}\} = u_{i,j+r} \end{cases} . \tag{2.61}$$

The coefficients b_r are identical to the 1D stencils and adhere to the following relations,

$$\sum_{n=0}^{R} (\epsilon_{EF,2n} \cdot c_{2n,r}) = b_r \ , \tag{2.62}$$

where $b_r = b_{-r}$ for symmetric schemes and $c_{r,2n}$ are coefficients describing the 1D second-order, dissipative, narrow-stencil representation:

$$\delta^{2n}\{u_i\} = c_{2n,0} + \sum_{j=1}^{n} c_{2n,j}\left(u_{i+j} + u_{i-j}\right) \ . \tag{2.63}$$

The definition of the multi-dimensional operator $\mathcal{D}_{2\mathrm{D}}^{\mathrm{AC}}$ in terms of difference operators may furthermore be determined by applying the relations of Equations 2.15, 2.17, and 2.62 to the transfer function representation for the 2D symmetric stencil of Equation 2.61:

$$\begin{aligned} \hat{\mathcal{G}}_{\sum^{\mathrm{AC}}}(k_1,k_2) &= 2b_0 + \left(\frac{2}{1+\beta}\right)\sum_{r=1}^{R} b_r \cdot [\cos(rk_1\Delta x) + \cos(rk_2\Delta x) \\ &\quad + \frac{\beta}{2}\left[\cos(r(k_1+k_2)\Delta x) + \cos(r(k_1-k_2)\Delta x)\right] \\ &\quad + \frac{\beta}{2}\left[\cos(r(k_1+k_2)\Delta x) + \cos(r(k_2-k_1)\Delta x)\right]] \ , \qquad (2.64) \\ &\quad \text{where} \quad \cos(\alpha \pm \beta) = \cos\alpha\cos\beta \mp \sin\alpha\sin\beta \ . \qquad (2.65) \end{aligned}$$

The anisotropic-correction parameter, β, is derived from an optimization procedure and is responsible for enhancing isotropic characteristics (see Sescu et al. (January 2009) for more details). Figure 2.14(d) demonstrates the improved

isotropy of the resulting method, although it is noted that the original 1D attenuation properties are altered – the alterations mostly being relegated towards the high wavenumbers.

The aforementioned multi-dimensional formulations are each shown to yield advantages and disadvantages. Henceforth, however, it can be assumed (unless otherwise noted) that $\mathcal{G}_{\prod}$ is employed for any multi-dimensional filtering operations.

2.7 Demonstration Test Case: Signal Post-Processing

The potential utility of the different discrete filter formulations previously studied is here demonstrated by the task of removing noise from a signal. By manipulating the spectral properties of the filter via an understanding of the stencil specification, it is possible to retrieve an acceptable approximation to the original signal.

An example in two-dimensions is chosen, where the Gaussian pulse is defined to be,

$$f_{\text{orig},2D}(x_1, x_2) = Ae^{-\left[\frac{(x_1-\mu_{x_1})}{2\sigma^2})+\frac{(x_2-\mu_{x_2})}{2\sigma^2})\right]} \ . \tag{2.66}$$

The parameters are chosen to be the following: $N = N_x = N_y = 100, L = 1, A = 1, \sigma = (0.05)L$. The perturbation used to derive the corrupt signal (see Equations 2.67-2.68) is then defined such as to induce both one-dimensional and multi-dimensional noise features. Also note that the parameters $C_{\kappa,x_n},\ \phi_{\mathsf{U},x_n}$ are varied only with respect to the corresponding x_n direction and that $[\kappa_{\min} = N/4, \kappa_{\max} = N/2]$ once again targets the perturbed modes.

$$f_{\text{corrupt},2D}(x_1, x_2) = f_{\text{orig},2D}(x_1, x_2) + \delta_{\text{pert},2D}(x_1, x_2) \tag{2.67}$$

$$\begin{aligned} \delta_{\text{pert},2D}(x_1, x_2) &= \sum_{\kappa_{\min}}^{\kappa_{\max}} C_{\kappa,x_2} \cdot \sin\left(\frac{2\pi\kappa x_1}{L} + \phi_{\mathsf{U},x_2}\right) \\ &\quad + \sum_{\kappa_{\min}}^{\kappa_{\max}} C_{\kappa,x_1} \cdot \sin\left(\frac{2\pi\kappa x_2}{L} + \phi_{\mathsf{U},x_1}\right) \ , \\ \text{with} \qquad C_{\kappa,x_n} &= \begin{cases} \left(\frac{1}{\kappa}\right) \cdot \mathsf{U}(\kappa), & (x_1 - \mu_{x_1})^2 + (x_2 - \mu_{x_2})^2 \leq (4\sigma)^2 \\ 0, & \text{otherwise} \end{cases} . \end{aligned} \tag{2.68}$$

Figure 2.15(a) and Figure 2.15(b) plot the original and perturbed signals respectively. Meanwhile Figure 2.16(a) and Figure 2.16(b) plot their corresponding normalized spectra $|\hat{f}(\kappa_1, \kappa_2)|^2/|\hat{f}(0,0)|^2$. It is confirmed that the perturbations are localized around the pulse in physical space as implied by C_{κ,x_n} and that notable noise is induced in wavenumber space starting at $(k_n\Delta x) = 0.25\pi$. Here the benefits of employing high-order and tuned filtering will be shown to become evident. Figures 2.15(c) shows results of the second-order Shapiro filter (equivalent to a second-order Tangent filter tuned to $(k_{1/2}\Delta x) = 0.5\pi$) and highlights the inability to remove enough of the noise due to improper tuning. Figure 2.15(e) plots the solution filtered by a second-order Tangent filter tuned to the spectral point at which the erroneous noise appears, $(k_{1/2}\Delta x) = 0.25\pi$; the result is shown to have removed much of the noise but to have depleted the original pulse height. Finally Figure 2.15(f) shows the solution filtered by a tenth-order Tangent filter also tuned to $(k_{1/2}\Delta x) = 0.25\pi$, which successfully removes erroneous modes while sufficiently maintaining content below the prescribed cut-off – in this case allowing the original pulse strength to be preserved. Analogous information is provided by the corresponding spectral plots shown in Figures 2.16(c), 2.16(e), and 2.16(f). It is clear that proper filtering can do well to retrieve the original signal, as shown

in Figure 2.16(f) for the tuned and scale-discriminant implementation.

Also interesting to witness is the impact of the multi-dimensional filter implementation. Figure 2.15(d) plots a filtered solution resulting from the use of a second-order Shapiro stencil implemented with the summation multi-dimensional form of Equation 2.56 (here used without re-scaling of the one-dimensional filter coefficients). Unlike the $\mathcal{G}_{\prod}$ implementation, the $\mathcal{G}_{\sum}$ operator is known to damp single- and multi-dimensional features differently. In this case, it is apparent that the resulting filtered signal is more noisy than the product $\mathcal{G}_{\prod}$ implementation shown in Figure 2.15(c). Inspection and comparison of the corresponding spectra in Figures 2.16(c) and 2.16(d) indeed confirms that the $\mathcal{G}_{\sum}$ operator, without re-scaling, preserves the one-dimensional filtering characteristics but does not address multi-dimensional components. The spectral plot of Figure 2.16(d) mirrors the previous theoretical perdition shown in Figure 2.14(b). With proper re-scaling, one can indeed apply the desired damping characteristics to multi-dimensional modes satisfying $k_1 = k_2$, however, the one-dimensional treatments become altered, usually diminished in their attenuating properties. Thus, this highlights the importance of employing isotropic multi-dimensional operators such as $\mathcal{G}_{\prod}, \mathcal{G}_{\sum^{\text{Fal}}}, \mathcal{G}_{\sum^{\text{AC}}}$ presented in Section 2.6.2.

Naturally, the performance of the different filter formulations is highly dependent on the nature of the erroneous content. The magnitude of the perturbation and the scales it affects can challenge the filtering methods in various ways. In addition, the practicality of certain implementations should be considered. For example, while the tenth-order Tangent filter used in the above examples performed well, it becomes rather ill-conditioned as one moves away from a $(k_{1/2}\Delta x) = 0.5\pi$ designation, a 4 PPW resolution (see Section 2.6.2). Nevertheless, the demonstrated success in being able to manipulate the spectral properties of the filtering methods is attractive. Leveraging of such schemes can become significant, not only in the context of LES but also in the more general arena of scheme stabilization,

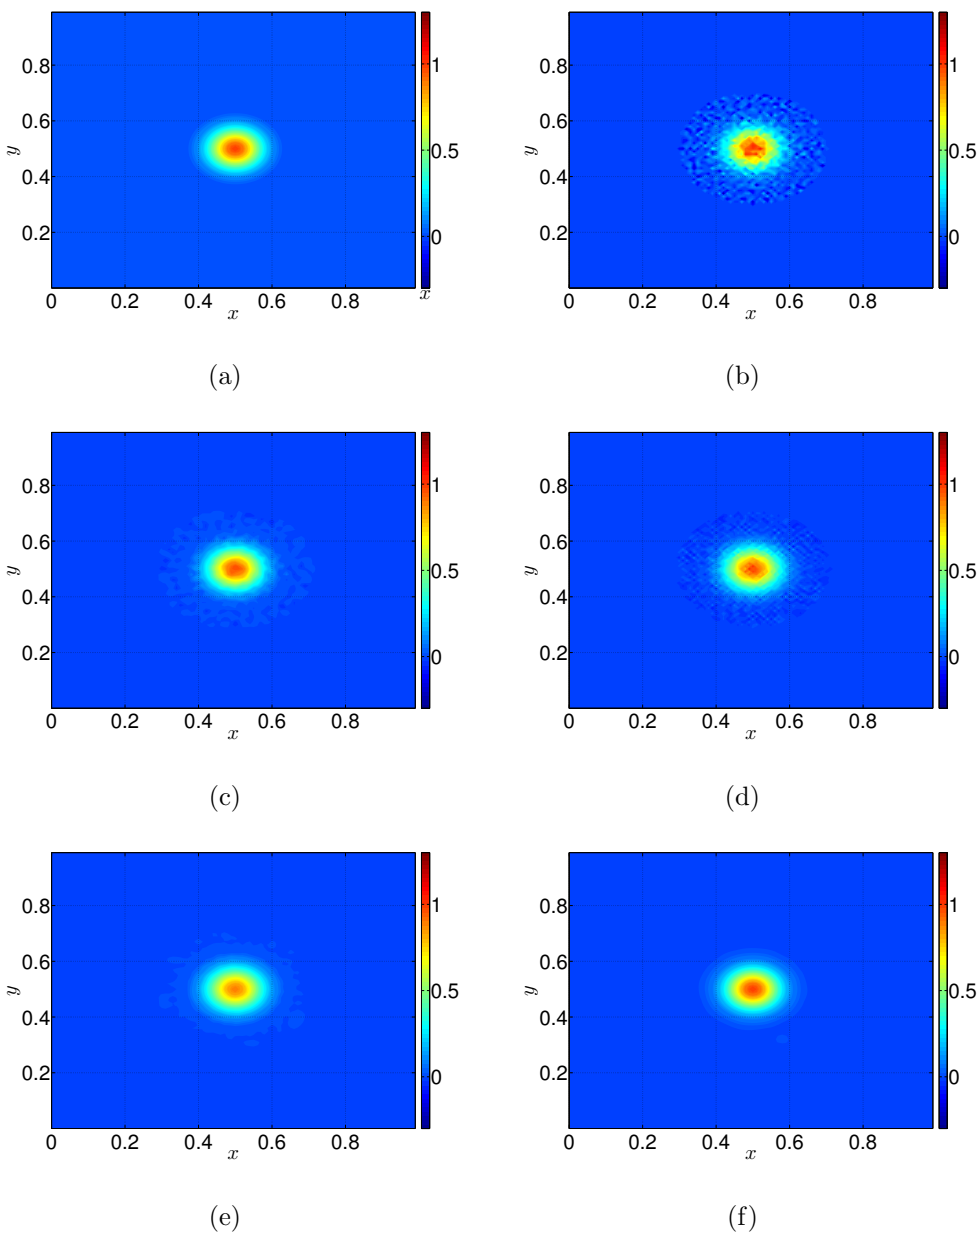

Figure 2.15: 2D Gaussian pulse signal: a) original signal, b) corrupt signal, c) $\mathcal{G}_{\prod}$ filtered result with second-order Shapiro stencil, d) $\mathcal{G}_{\sum}$ filtered result with second-order Shapiro stencil (without coefficient re-scaling), e) $\mathcal{G}_{\prod}$ filtered result with second-order Tangent stencil tuned to $(k_{1/2}\Delta x) = 0.25\pi$, and f) $\mathcal{G}_{\prod}$ filtered result with tenth-order Tangent stencil tuned to $(k_{1/2}\Delta x) = 0.25\pi$.

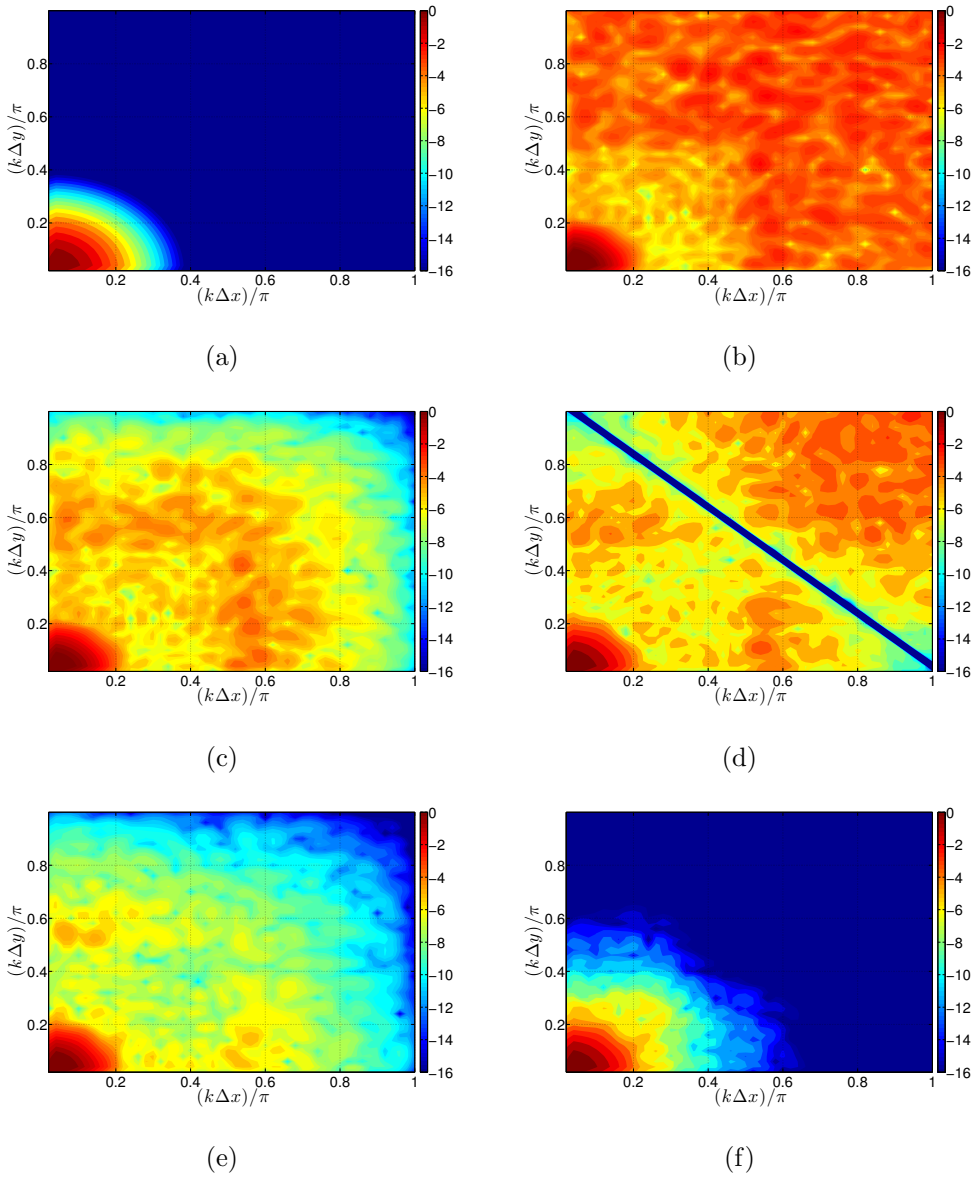

Figure 2.16: Normalized Power Spectral Density (PSD) of the 2D Gaussian pulse signals ($\log\{|\hat{f}(k_1,k_2)|^2/|\hat{f}(0,0)|^2\}$): a) original signal, b) corrupt signal, c) $\mathcal{G}_{\prod}$ filtered result with second-order Shapiro stencil, d) $\mathcal{G}_{\sum}$ filtered result with second-order Shapiro stencil (without coefficient re-scaling), e) $\mathcal{G}_{\prod}$ filtered result with second-order Tangent stencil tuned to $(k_{1/2}\Delta x) = 0.25\pi$, and f) $\mathcal{G}_{\prod}$ filtered result with tenth-order Tangent stencil tuned to $(k_{1/2}\Delta x) = 0.25\pi$.

as will be demonstrated in the following chapters.

CHAPTER 3

Filtering and Artificial Dissipation

This chapter is taken with some modifications from the AIAA conference proceedings articles "Highly-Accurate Filter-Based Artificial-Dissipation Schemes for Stiff Unsteady Fluid Systems" (Mundis et al., January 2016) and "Comparison of Artificial Dissipation and Filtering Schemes for Time-Accurate Simulations" (Edoh et al., 2015).

3.1 Background

The under-resolution associated with LES often requires one to consider incorporating suitable stabilization methods. These are required in order to maintain robustness of the calculation in the absence of physical diffusion mechanisms in the scale of typical LES meshes. In the absence of high-wavenumber dissipation, aliasing and discretization errors are able to accumulate and thus can threaten numerical stability (Phillips, 1959). In an attempt to maintain scheme robustness, one can consider formulations such as secondary conservation methods that include the class of non-dissipative kinetic energy or entropy preserving schemes (Morinishi et al., 1998; Tadmor and Zhong, 2006). These methods are known to enhance non-linear stability by implicitly satisfying discrete auxiliary conservation laws (e.g., kinetic energy, entropy) that both improve physical accuracy of the simulation and impose bounds on the admissible numerical solutions. Although less susceptible to numerical instability, such schemes are not devoid of accumulated high-wavenumber error and still require damping. Therefore, it is useful

to also consider numerical dissipation techniques for active removal of contaminated high-wavenumber content. Such methods include artificial dissipation (AD) and solution-filtering (SF). Being closely related, nuances in their similarities and differences are studied in the present chapter.

Perhaps more commonly used for stabilization, artificial dissipation (AD) methods – primarily developed with regards to shock capturing (von Neumann and Richtmyer, 1950) – introduce additional damping terms into the governing equations[1]. These artificial contributions emulate physical diffusion but can affect an arbitrarily wide range of the resolved scales, thus potentially hindering solution accuracy. Another consequence of altering the original equations is the way in which the new AD terms couple with other spatial contributions (e.g., convective discretizations) as well as the temporal integration method. In this regard, it is possible for the AD method to not only affect damping but to also alter the dispersive characteristics of the original method. Furthermore, the CFL stability limits can also be affected.

As an alternative to artificial dissipation, solution-filtering (SF) can be used for stabilization. These schemes originated from the need to control aliasing issues in meteorological computations, which are severely under-resolved (Shapiro, 1971). Unlike the artificial dissipation treatment, the damping of solution-filtering is administered as a type of corrector step to the incremental integration of the governing equations. In this way, the desired modes are attenuated while keeping the original spatial and temporal properties of the base scheme intact. A new challenge arises, however, with respect to determining how often the solution variables are filtered. For example, one can choose to filter after N time-steps or at each stage of a Runge-Kutta integration – the choice of which affects the overall dissipation accumulated over time. The potential for excessive modal attenuation

[1]Upwinding stencils are inherently dissipative and may be interpreted as the addition of a properly scaled artificial dissipation term to a central (non-dissipative) scheme.

furthermore requires judicious use of solution-filtering in the case of stiff problems, where there may be large disparities in the time-scales associated with different physical phenomena. In response to these ambiguities, algorithms based on kinetic energy criteria have been developed in order to automate the decision of when to filter relative to the flow evolution (Tantikul and Domaradzki, 2010; Flad et al., 2016).

Despite originating from different fields, there is a strong resemblance between artificial dissipation and solution-filtering methods. This is made evident by rewriting the filtering stencils as difference operators (see Equation 2.18) and noting the implied presence of dissipative operators. Naturally, these stabilization methods have been applied interchangeably. For instance, solution-filtering has been re-adapted for shock-capturing purposes (Visbal and Gaitonde; Bogey et al., 2009); meanwhile, Yee et al. (1999) re-cast dissipative portions of shock-capturing TVD (total variation diminishing) and WENO (weighted essentially non-oscillatory) stencils for use as characteristic-based filters. A better understanding of how these damping techniques relate to each other can therefore provide valuable insights on how to properly apply them to different situations.

For example, both artificial dissipation and solution-filtering can be seen to have direct application to LES modeling. The classic eddy-viscosity model concept corresponds to a second-order dissipative contribution δ^2 (scaled by a coefficient based on velocity gradient components Silvis et al. (2017)); thus its numerical impact is analogous to artificial dissipation methods. While the incorporation of the dynamic model procedure of Germano (Germano et al., 1991) and subsequent improvements (Lilly, 1992; Park and Mahesh, 2009) have been pivotal to the success of such eddy-viscosity models, basic linear analysis reveals inherent limitations relative to their suitability for high-order filter formulations (Pruett and Sochacki, 2001). In response, hyper-viscous eddy-viscosity models have also been studied and have shown improvements (Jeanmart and Winckelmans, 2007). The link to

such successes is likely tied to the high-order dissipative spectral characteristics of the models; thus, studying the design of artificial dissipation schemes can indirectly find relevance to the realm of functional LES modeling. With respect to solution-filtering, Mathew et al. (2003, 2006) have re-interpreted the Approximate Deconvolution method (ADM) of Stolz and Adams (Stolz and Adams, 1999) as a solution-filtering algorithm. The use of filtering as a model – termed relaxation filtering – has been explored (Berland et al., 2008; Aubard et al., 2013) and related to an effective hyper-viscous eddy-viscosity model that scales with the frequency of filtering applications (Fauconnier, 2013).

The inherent relation of solution-filtering and artificial dissipation thus provides different perspectives from which one can formulate improved scheme stabilization for LES methods. Most notable, perhaps, is the opportunity to leverage a well-established understanding of the spectral specification of filters and adapt these to a generalized artificial dissipation framework. This would further allow the incorporation of stiff-preconditioning techniques (Turkel, 1999; Venkataswaran and Merkle, 2000; Mundis et al., January 2016) as well as the derivation of generalized hyper-viscous LES models, to be used either independently or as proper stabilization for other LES models (i.e., mixed-modeling) (Vreman et al., 1996; Winckelmans et al., 2001) that would be amenable to high-order filter formulations. Therefore the following chapter focuses on comparing solution-filtering and artificial dissipation methods, highlighting their respective weaknesses and advantages.

3.2 Solution-Filtering

Solution-filtering (SF) is applied as a corrector step to the integration of a general governing equation, here taken to be the linear 1D advection equation integrated

using a forward Euler temporal scheme:

$$\begin{aligned} 1) \quad u^{*,n+1} &= u^n + (\Delta t) \cdot R_o(u^n) \ , \quad \text{where} \quad R_o\{u\} = -\lambda \delta_x\{u\} \ , \\ &\vdots \\ 2) \quad u^{n+N_f} &= \mathcal{G}\{u^{*,n+N_f}\} \ . \end{aligned} \tag{3.1}$$

In general, the filtering operation of step 2 can occur regularly after every N_f time-steps, after each stage of a Runge-Kutta method, or at arbitrary intervals ascribing to a particular indicator. The consequence of the filtering operations are made more evident by combining steps 1 and 2 into a single effective evolution equation. This is done most simply for the case of $N_f = 1$, wherein filtering occurs after each time-step[2]. Assuming a general implicit filter defined by the difference form of Equation 2.36, the resulting modified equation reads:

$$\begin{aligned} \left[\epsilon_{IF,0} + \sum_{\ell=1}^{L} \epsilon_{IF,2\ell} (\Delta x)^{2\ell} \delta_x^{2\ell} \right] \left\{ \frac{u^{n+1} - u^n}{\Delta t} \right\} &= \sum_{r=0}^{R} \epsilon_{EF,2r} (\Delta x)^{2r} \delta_x^{2r} \left\{ R_o\{u^n\} \right\} \\ &+ \left(\frac{1}{\Delta t} \right) \sum_{r=0}^{R} \epsilon_{EF,2r} (\Delta x)^{2r} \delta_x^{2r} \{u^n\} \\ &- \left(\frac{1}{\Delta t} \right) \sum_{\ell=0}^{L} \epsilon_{IF,2\ell} (\Delta x)^{2\ell} \delta_x^{2\ell} \{u^n\} \ , \end{aligned} \tag{3.2}$$

which uses the definition of $u^{n+1} = u^n + \Delta u$. Substituting $\Delta u \approx \Delta t \cdot R_o\{u^n\}$ (which implicitly assumes that $u^{n+1} \approx u^{*,n+1}$) and using $\epsilon_{IF,0}/\epsilon_{EF,0} = 1$ then

[2]See Fauconnier (2013) who addresses interval filtering with general N_f.

gives,

$$\begin{aligned}
\epsilon_{IF,0} \cdot \frac{u^{n+1} - u^n}{\Delta t} &= \epsilon_{EF,0} \cdot R_o\{u^n\} \\
&\quad + (\Delta x) \left[\sum_{r=1}^{R} \epsilon_{EF,2r} (\Delta x)^{2r-1} \delta_x^{2r} \right] \{R_o\{u^n\}\} \\
&\quad - (\Delta x) \left[\sum_{\ell=1}^{L} \epsilon_{IF,2\ell} (\Delta x)^{2\ell-1} \delta_x^{2\ell} \right] \{R_o\{u^n\}\} \\
&\quad + \left(\frac{\Delta x}{\Delta t} \right) \left[\sum_{r=1}^{R} \epsilon_{EF,2r} (\Delta x)^{2r-1} \delta_x^{2r} \right] \{u^n\} \\
&\quad - \left(\frac{\Delta x}{\Delta t} \right) \left[\sum_{\ell=1}^{L} \epsilon_{IF,2\ell} (\Delta x)^{2\ell-1} \delta_x^{2\ell} \right] \{u^n\} .
\end{aligned} \tag{3.3}$$

Relative to the original governing equation (step 1 of Equation 3.1), it is apparent that additional dissipative terms are induced in the equations. These, furthermore, are seen to scale according to a numerical wave speed, $|\lambda_{\text{num}}| = (\Delta x/\Delta t)$. Therefore, the relative influence of the artificial terms is unchanged in the case that $|\lambda_{\text{num}}|$ is constant. The consequences of this scaling relative to the physical nature of the equation is made more apparent by employing the definition $(\Delta x/\Delta t) = (|\lambda|/\text{CFL})$, which forms a stability condition for the advection equation in question. From this point of view, there is no reason that the CFL would be kept at a set ratio relative to the physical wave speed, λ. Instead, its prescription is typically determined based on accuracy considerations, where smaller time-steps (and thus smaller CFL) are concomitant with smaller temporal error. In the case of solution-filtering, however, it becomes apparent that smaller time-stepping results in increased damping effects associated with the filtering of u^n – the consequences of this administered dissipation depends on the relative magnitude of $|\lambda|$ and the CFL[3].

[3]Relative to the concomitant filtering of the residual (R_o), one could also presume that solution-filtering also produces dispersive effects. This, however, would violate the notion that

3.2.1 Temporally-Consistent Solution-Filtering

The evolution equation relating to solution-filtering described by Equation 3.3 highlights the temporal inconsistency of the stabilization approach and conveys that the associated dissipation error of the scheme is inherently tied to the frequency of filtering (i.e., the time-step size, considering a constant N_f). Inspecting Equation 3.3, an apparent solution to this problem would be to re-scale the dissipative filter coefficients relative to the CFL, such that $\epsilon^*_{2n\geq 1} = \text{CFL} \cdot \epsilon_{2n\geq 1}$. In the one-dimensional example considered here, this modification is analogous to a scalar artificial dissipation scheme; for systems of equations, one could furthermore consider matrix dissipation (Swanson et al., 1998). The need to scale the filter coefficients by the time-step size has also been recognized in Asthana and Jameson (2014) who work to stabilize high-order Flux Reconstruction schemes by using convolution-based filtering, made to be consistent (to first order in time) with respect to artificial dissipation.

In terms of explicit filter stencils ($\epsilon_{IF,2\ell\geq 1} = 0 \; \forall \; \ell$), the suggested modifications are trivial. In the case of implicit filter stencils, however, there is a significant implementation overhead associated with the fact that the linear system to be inverted would vary in time according to $\text{CFL} = (|\lambda|\Delta t/\Delta x)$, where λ is typically a function of the solution variable u for general non-linear equations. This would then require real-time inversion of the system, which may constitute too large of a computational overhead. Ideally, the required system inversion would preferably be calculated as a pre-processing step (e.g. via a stored LU decomposition) to be applied during run-time. Possibly, one could choose to freeze the left hand matrix and to update it periodically, however, this introduces new sources of error (e.g.,

---

the filtering operator in step 2 of Equation 3.1 is purely dissipative. A modified equation analysis (see Appendix of Edoh et al. (2015)) reveals that the induced dispersive contributions in Equation 3.3 are necessary in order to offset the phase error generated through coupling of the temporal scheme and the new dissipative components. As a secondary effect, the new odd-derivative terms also couple with the integration schemes and contribute to the overall damping characteristics of the original solution-filtering scheme.

linearization) that could furthermore corrupt the time-accurate integrity of the scheme.

Another concern is the proposed scaling, intended to reduce damping across all wavenumbers while maintaining the general spectral response shape. The proposed re-scaling achieves this for explicit stencils, but fails to do so for implicit stencils. In the case of the Padé schemes, simple re-scaling of all coefficients can be shown to drastically alter the transfer response. This is demonstrated in Figures 3.1(a) and 3.1(b) which respectively plot the growth factor for a fourth-order explicit Shapiro scheme and a fourth-order implicit Long scheme (using $\delta_\Delta = 0.1$) with the proposed re-scaling. As mentioned, the explicit stencil maintains its general spectral characteristics; meanwhile, the implicit stencil is completely altered. The unintended alterations to the implicit filter's transfer function, furthermore, nullify previous efforts to understand and predictably prescribe its spectral characteristics.

A better understanding of why this happens can be garnered by inspecting the resulting transfer function modified by the proposed rescaling:

$$\hat{\mathcal{G}}^* = \frac{\hat{\mathcal{I}} + \mu\hat{\mathcal{D}}_r}{\hat{\mathcal{I}} + \mu\hat{\mathcal{D}}_\ell} = \hat{\mathcal{I}} + \frac{\mu\left(\hat{\mathcal{D}}_r - \hat{\mathcal{D}}_\ell\right)}{\hat{\mathcal{I}} + \mu\hat{\mathcal{D}}_\ell} \ . \tag{3.4}$$

This expression only recovers the intended re-scaling of,

$$\hat{\mathcal{G}}^* = \hat{\mathcal{I}} + \mu\left(\hat{\mathcal{G}} - \hat{\mathcal{I}}\right) \ , \tag{3.5}$$

in the case of an explicit stencil, where $\hat{\mathcal{D}}_\ell = 0$. Instead, spectral satisfaction of Equation 3.5 is attained when employing the operator scaling previously provided in Equation 2.38. Writing the resulting temporally-consistent form of Equation

3.3 gives,

$$\begin{aligned}
\epsilon_{IF,0} \cdot \frac{u^{n+1} - u^n}{\Delta t} &= \epsilon_{EF,0} \cdot R_o\{u^n\} \\
&\quad + (|\lambda|\Delta t) \left[\sum_{r=1}^{R} \epsilon_{EF,2r} (\Delta x)^{2r-1} \delta_x^{2r} \right] \{R_o\{u^n\}\} \\
&\quad - (|\lambda|\Delta t) \left[\sum_{\ell=1}^{L} \epsilon_{IF,2\ell} (\Delta x)^{2\ell-1} \delta_x^{2\ell} \right] \{R_o\{u^n\}\} \\
&\quad + |\lambda| \left[\sum_{r=1}^{R} \epsilon_{EF,2r} (\Delta x)^{2r-1} \delta_x^{2r} \right] \{u^n\} \\
&\quad - |\lambda| \left[\sum_{\ell=1}^{L} \epsilon_{IF,2\ell} (\Delta x)^{2\ell-1} \delta_x^{2\ell} \right] \{u^n\} \ ,
\end{aligned} \tag{3.6}$$

which still recovers the desired temporal consistency of the scheme, yet maintains the response qualities of the implicit scheme (see Figure 3.1(b)). Furthermore, the aforementioned computational overhead associated with inverting a time-varying system is avoided by relegating the CFL scaling to the righthand size of the system, which corresponds to differencing the original signal u. In this way, it should also be noted that the scaling should be done conservatively[4].

The implications of this CFL re-scaling for multi-dimensional use can also be considered. In the case where successive 1D operators are applied in each direction, the associated response can be written as,

$$\hat{\mathcal{G}}^*_{\prod,\mathrm{NDim}} = \prod_{n=1}^{\mathrm{NDim}} \left[\hat{\mathcal{I}} + \mu_n \left(\hat{\mathcal{G}}_n - \hat{\mathcal{I}} \right) \right] \ . \tag{3.7}$$

This naturally suggests that N-dimensional waves will be scaled as $\prod_n^N \{\mu_n\}$, which means that mixed-modes will have their coefficients further reduced than one-dimensional modes. This is consistent with the notion that mixed directions have

[4]Note, the scaling should be done conservatively (i.e., $|\lambda|\delta^2\{u\} \rightarrow \delta\{|\lambda|\delta\{u\}\}$) with special averaging (Kamakoti and Pantano, 2009) such that narrow-stencil representations are recovered, rather than wide stencil formulations that may be susceptible to odd-even decoupling.

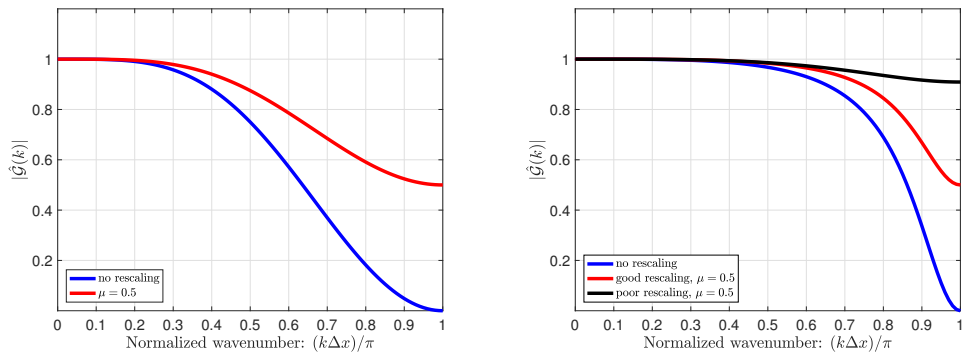

(a) (b)

Figure 3.1: Effects of proposed filter re-scaling methods: a) Growth factor $|\hat{\mathcal{G}}|$ of explicit fourth-order Shapiro stencil with re-scaling (using $\mu = 0.5$ using Equation 3.5) and without re-scaling; b) Growth factor of implicit fourth-order Long filter Shapiro stencil with re-scaling (with $\mu = 0.5$ using Equation 3.5 in red, versus Equation 3.4 in black) and without-rescaling.

a coarser effective grid length (e.g., the diagonal length of a rectangular grid), which entails a lower effective CFL.

3.3 Filter-based Artificial Dissipation

The ability to regain temporal consistency by re-scaling the solution-filtering coefficients by the CFL number can be very useful. Extension of these treatments to such things as pseudo-time iteration (Merkle and Athavale, 1987) (e.g., for stiff-physical systems), however, is not straight-forward and can lead to stalled convergence (Mundis et al., January 2016) and unintended dispersive errors. As a potential remedy to this, one may consider the use of artificial dissipation schemes rather than solution-filtering. These methods are well-established for use in pseudo-time and for stiff-regimes (Venkataswaran and Merkle, 2000). Perhaps lacking in the traditional use of artificial dissipation, however, is knowledge of the customized

design of the spectral damping characteristics. Though recent efforts have been made in developing optimized artificial dissipation stencils (Mansouri and Hixon, 2015), attaining a level of specification as found in the realm of filtering is still strongly desirable. In this way, it is attractive to leverage the stencils belonging to different filter families and to re-express them as artificial dissipation schemes. This is done here and termed filter-based artificial dissipation.

Equation 3.6 is the evolution equation relating to the temporally-consistent solution-filtering. It can be seen to take on an artificial dissipation form, and thus gives insight into deriving the new filter-based artificial dissipation schemes. It suggests that one would need to filter both the solution u^n as well as the residual R_o; the resulting scheme may be written in the following ordinary differential equation (ODE) form for a Forward (explicit) Euler integration method:

$$\frac{u^{n+1} - u^n}{\Delta t} = R_o\{u^n\} + (\Delta t) \cdot R_{AD,ex}\{R_o\{u\}\} + R_{AD,ex}\{u^n\} \ , \tag{3.8}$$

where,

$$R_{AD,ex}\{u\} = \left[\sum_{r=1}^{R} \epsilon_{EF,2r} (\Delta x)^{2r-1} \delta_x^{2r-1} - \sum_{\ell=1}^{L} \epsilon_{IF,2\ell} (\Delta x)^{2\ell-1} \delta_x^{2\ell-1} \right] \{|\lambda| \delta u\} \ . \tag{3.9}$$

Note, the above equation assumes that all coefficients $\epsilon_{2n\geq 1}$ are properly normalized by $\epsilon_{IF,0}$ and presumes that proper averaging of non-constant coefficients (e.g., $|\lambda|$) is performed in order to ensure telescoping properties for discrete conservation in practical use. According to the premise of solution-filtering from which it is derived, an artificial dissipation scheme as written in Equation 3.8 should only impart dissipation while avoiding any alterations to the base scheme's phase characteristics. This formulation, however, is non-traditional in that it includes both artificial dissipation and artificial dispersive terms. Adaptation to multi-stage schemes, where $R_o^s\{u\} = \sum_m^s a_{ms} R_o\{u^m\}$, can also become algorithmically complicated and may degrade temporal accuracy[5].

[5]the solution-filtering evolution equation shown in Equation 3.3 is based on the assumption that $\Delta u \approx \Delta t \cdot R_o$, which implies the possibility of splitting errors in the derivation of Equation 3.8.

In order to derive a more tractable and traditional implementation for the filter-based artificial dissipation, it is useful to consider exact integration and to assume all damping to be the result of adding a dissipative term to the original residual. Under this semi-discrete perspective, the new method takes the following form,

$$\frac{\partial u}{\partial t} = R_o\{u\} + R_{AD}\{u\} \ , \tag{3.10}$$

$$\text{where} \qquad R_{AD}\{u\} = u'' \ , \tag{3.11}$$

$$\text{such that} \quad \left[\mathcal{I} + \mathcal{D}_\ell\right]\{u''\} = R_{AD,ex}\{u\} \ . \tag{3.12}$$

In this way, extension to multi-stage schemes is straightforward. One simply needs to add R_{AD} to the original residual. The form of the system to solve in Equation 3.12 is derived by considering the spectral function $\hat{\mathcal{D}} = \hat{\mathcal{G}} - \hat{\mathcal{I}}$ in an operator form. In essence, the filter-based artificial dissipation extracts the damping component of the filtering scheme and adds it to the ODE residual. Multi-dimensional forms are derived by a similar exercise. For simplicity, however, the additive approach relating to $\mathcal{G}_{\sum}$ is employed henceforth, where $R_{AD} = \sum_n^{\text{NDim}} R_{AD,n}$.

A few consequences in the interpretation of Equation 3.10 is the fact that the artificial dissipation scheme is not precisely consistent with the solution-filtering. As a consequence, the phase characteristics of the base scheme can be altered and temporal coupling with the intended damping behavior is expected to be stronger. Additionally, the adaptation of implicit filter stencils to the artificial dissipation form now requires the inversion of the system in Equation 3.12; this is in contrast to Equation 3.8, where the implicit nature of the filter stencil is mathematically represented by proper filtering of the residual (R_o) and the previous time-step solution (u^n) – an allusion to the class of residual based compact schemes (RBC) (Lerat and Corre, 2001).

In order to better assess the subtle difference between the proposed schemes of Equation 3.8 versus Equation 3.10, it is helpful to write a multi-stage Runge-Kutta

integration scheme for the general ODE, $\partial_t u = R\{u\}$:

$$\begin{aligned} u(t^n + \Delta t) &= u^n + \int_{t^n}^{t^n+\Delta t} R\{u\} = u^{n+1}, && (3.13) \\ &\approx u^n + (\Delta t) \cdot \sum_{s=1}^{S} \tilde{b}_s R\{u^{n+\tilde{c}_s}\}, && (3.14) \\ &\quad \text{where} \quad u^{n+\tilde{c}_s} = u^n + (\Delta t) \cdot \sum_{j=1}^{S} \tilde{a}_{sj} R\{u^{n+\tilde{c}_s}\} \ . \end{aligned}$$

The more traditional interpretation offered by Equation 3.12 would produce the following stage-values:

$$u^{n+\tilde{c}_s} = u^n + (\Delta t) \cdot \sum_{j=1}^{S} \tilde{a}_{sj} \left[R_o\{u^{n+\tilde{c}_s}\} + R_{AD}\{u^{n+\tilde{c}_s}\} \right] \ . \tag{3.15}$$

Meanwhile, the scheme offered by Equation 3.8 – conjured to be more congruent with the temporally-consistent solution-filtering of Equation 3.3 – would have the stage calculations read as the following (assuming solution-filtering is done at each integration stage):

$$\begin{aligned} u^{n+\tilde{c}_s} &= u^n + (\Delta t) \cdot R_{AD,ex}\{u^n\} \\ &\quad + (\Delta t) \cdot \sum_{j=1}^{S} \tilde{a}_{sj} \left[R_o\{u^{n+c_s}\} + (\Delta t) \cdot R_{AD,ex} \left\{ R_o\{u^{n+\tilde{c}_s}\} \right\} \right] \ . \end{aligned} \tag{3.16}$$

This form does not naturally adhere to the multi-stage formulation due to the additional $R_{AD,ex}\{u^n\}$ term that is not scaled by the stage coefficients a_{sj}. Thus, this proposed formulation would need to be analyzed as a fully discrete scheme. The implementation of Equation 3.12 is chosen due to its more traditional form and straight forward extension to high-order temporal schemes.

3.4 Spectral Assessment of Stabilization Methods

The spectral performance of the proposed solution-filtering and filter-based artificial dissipation schemes is assessed in this section. Because the schemes funda-

mentally differ in their interaction with the integration method, focus is placed on consequences pertaining to temporal scheme coupling as well as relating to issues of temporal consistency.

In order to inspect the algorithms spectrally, von Neumann analysis (VNA) is employed, where we define the complex-valued amplification factor $\hat{\mathcal{G}}$ such that,

$$\hat{u}^s = \hat{\mathcal{G}}_s \cdot \hat{u}^n \ . \tag{3.17}$$

The amplification factor may be decomposed into growth and phase information as shown in Equation 2.4 and thus communicates the solution's evolution assuming linearity. Von Neumann analysis derives an expression for the amplification factor by considering the fully discrete equations. Substituting Fourier representations for the residual (i.e., discrete derivatives), yields the following system of equations to be solved in the case of multi-stage integration:

$$\hat{\mathcal{G}}_{n+1}(k) \ = \ \hat{\mathcal{I}} + (\Delta t) \cdot \sum_{s=1}^{S} \tilde{b}_s (\hat{\mathcal{G}}_s \cdot \hat{\mathcal{R}}) \ , \tag{3.18}$$

$$\text{such that} \quad \hat{\mathcal{G}}_s \ = \ \hat{\mathcal{I}} + (\Delta t) \cdot \sum_{j=1}^{S} \tilde{a}_{sj} (\hat{\mathcal{G}}_j \cdot \hat{\mathcal{R}}) \ . \tag{3.19}$$

In the current exercise, the one-dimensional linear advection equation is considered on a uniform and periodic domain:

$$\partial_t u - R_o\{u\} = 0 \text{ with } R_o\{u\} = -\lambda \partial_x u \ , \tag{3.20}$$

Substituting Fourier representations into the discrete representation of the respective derivatives yields,

$$\hat{\mathcal{R}}_o = -\lambda \cdot \imath k_{mod} = -\lambda \cdot \frac{\imath k'_{mod}}{\Delta x} \ , \tag{3.21}$$

with k_{mod} serving as the modified wavenumber of the differencing scheme used on the spatial derivatives (Kravchenko and Moin, 1997). In this manner, the solution-filtering scheme is analyzed to have an overall amplification factor of

$\hat{\mathcal{G}}_{n+1} = (\hat{\mathcal{G}}_{\text{fil}} \cdot \hat{\mathcal{G}}_{*,n+1})$ (assuming filtering at the end of each full time-step); meanwhile, the filter-based artificial dissipation scheme is analyzed with respect to $\hat{\mathcal{G}}_{n+1}$ based on subsequent spectral representations of the discretized residual $R\{u\} = R_o\{u\} + R_{AD}\{u\}$. The following considers a sixth-order standard discretization of $\partial_x u$, which yields $k_{mod}(k) = (1/30) \cdot [45\sin(k\Delta x - 9\sin(2k\Delta x) + \sin(3k\Delta x)]$. For integration, the third-order Runge-Kutta temporal scheme of Wray is considered (see the Butcher Tableaux of Table 4.2 in Chapter 4 for the integration coefficients, $[\tilde{c}_s, \tilde{a}_{sj}, \tilde{b}_s]$).

Figure 3.2 shows the growth factor $|\hat{\mathcal{G}}|$ of the different stabilization schemes as a function of the CFL number. Figure 3.2(a) shows the case of the base scheme without any stabilization. Here, it is evident that the integration method produces some mid-wavenumber damping, and that this behavior – which is unfounded for the purely hyperbolic system – is correctly diminished as the CFL is reduced and temporal error is decreased. Figure 3.2(b) plots the added effect of including solution-filtering to the base scheme. The overall growth factor is a product of the original damping and the filter operator. In this case, reducing the CFL has little effect on the imparted dissipation characteristics which are dominated by the solution-filtering contribution. While the extent to which this is true depends on the chosen cut-off wavenumber $k_{1/2}$ for the filter, one can anticipate that the relative decoupling between the time-step size (i.e., CFL number) and the effective dissipation could lead to temporal-consistency issues as $\Delta t \to 0$. Next, Figure 3.2(d) reveals how re-scaling the filter coefficients by the time-step size allows the damping to be reduced according to the CFL. The case of the filter-based artificial dissipation scheme in Figure 3.2(c) shows congruent damping characteristics, with discrepancies relative to the temporally-consistent filtering only arising under non-negligible temporal coupling effects (i.e., as the time-step size increases).

In order to understand the long-term implications of the stabilization schemes, Figure 3.3 plots the cumulative damping as $|\hat{\mathcal{G}}|^{1/CFL}$, thus highlighting any tem-

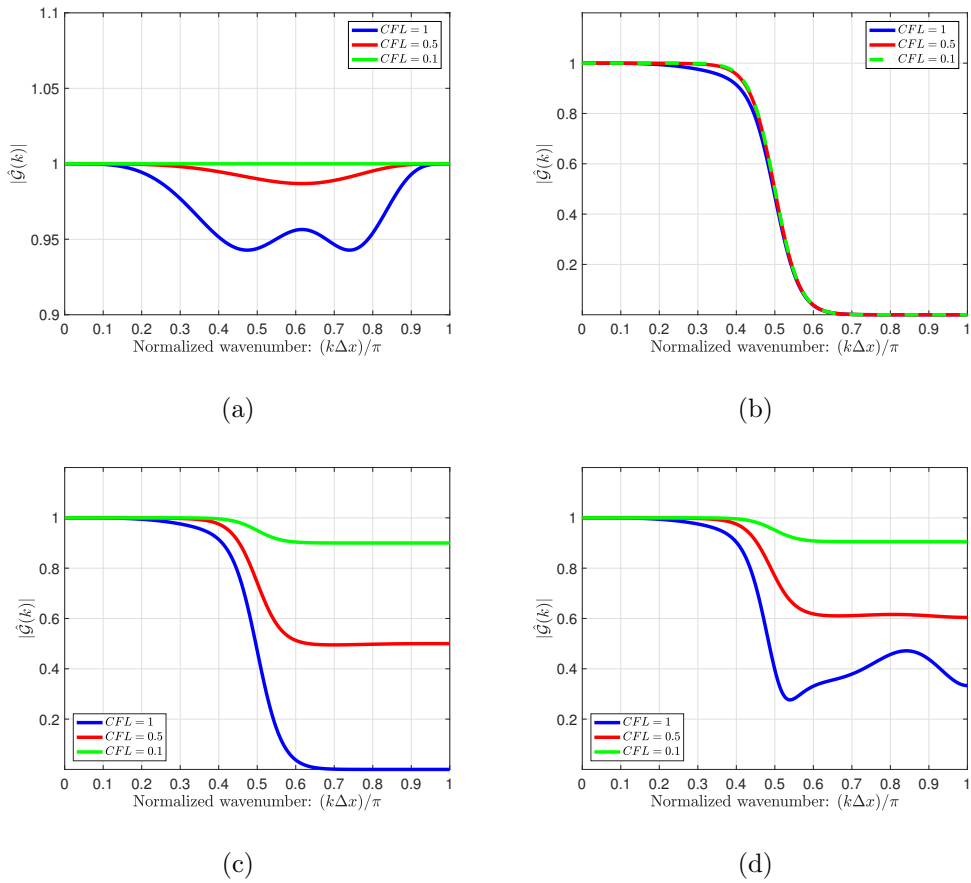

Figure 3.2: Growth factor $|\hat{\mathcal{G}}|$ as a function of CFL and corresponding to tenth-order Tangent stencil stabilization ($\theta_{1/2} = 0.5\pi$) used with the third-order Runge-Kutta integration scheme (see Table 4.2) and sixth-order central discretization of the 1D hyperbolic equation (see Equation 3.20): a) base scheme with no stabilization, b) solution-filtering, c) temporally-consistent solution-filtering as proposed in Equation 3.5, and d) filter-based artificial dissipation as proposed in Equation 3.15.

poral consistency issues with the different formulations. In the case of the base scheme shown in Figure 3.3(a), we note that the system is increasingly preserved as one decides to take smaller time-steps. Again, this is expected as the governing equations are hyperbolic and any damping is the result of errors in the integra-

tion method. Conversely, Figure 3.3(b) demonstrates how for solution-filtering, the effective cut-off is increasingly coarsened (in the cumulative sense) as the time-step is decreased. In this example, reducing the time-step size by two orders of magnitude has resulted in a coarsening of the computation, from $\Delta \approx 4\Delta x$ to $\Delta \approx 8\Delta x$. This underlines the potential issue with traditional solution-filtering implementations: decoupling the stabilization from the integration process makes the dissipation agnostic to the physics and thus calls for additional precautions with respect to how much damping is administered over time. The proposed re-scaling of the solution-filtering schemes by CFL in Figure 3.3(c), however, seems to maintain the desired cut-off under time-step refinement (note: the curves for CFL $= 10^{-2}$ and 10^{-4} are almost identical, thus suggesting asymptotic behavior). The fall-off in spectral response is largely preserved, although high-wavenumber damping is reduced. The reduction in cumulative high-wavenumber dissipation, however, is not believed to threaten stability since the damping effect is geometric; meanwhile, the accumulation of solution error (i.e., aliasing) is characteristically additive. Figure 3.3(d) reiterates that the filter-based artificial dissipation is able to maintain the spectral damping response in a temporally-consistent manner.

While the re-scaled solution-filtering and filter-based artificial dissipation formulations can be shown to be largely equivalent in 1D, their behavior deviates at moderate CFL, in which case the artificial dissipation method undergoes stronger coupling with the integration scheme. Comparing the damping responses in Figures 3.2(b) and 3.2(d), such coupling is seen to be mainly manifested at the high-wavenumbers. A secondary difference in the two formulations is the associated effect on phase characteristics. As previously discussed, solution-filtering is purely dissipative and thus would preserve the dispersive characteristics of the base scheme. Figure 3.4(a) shows the relative phase error of the base scheme for CFL $= 10^{0}$ and 10^{-2}. As expected, the dispersive error is dictated by the accuracy of the spatial scheme at smaller time-steps. But as the CFL is increased,

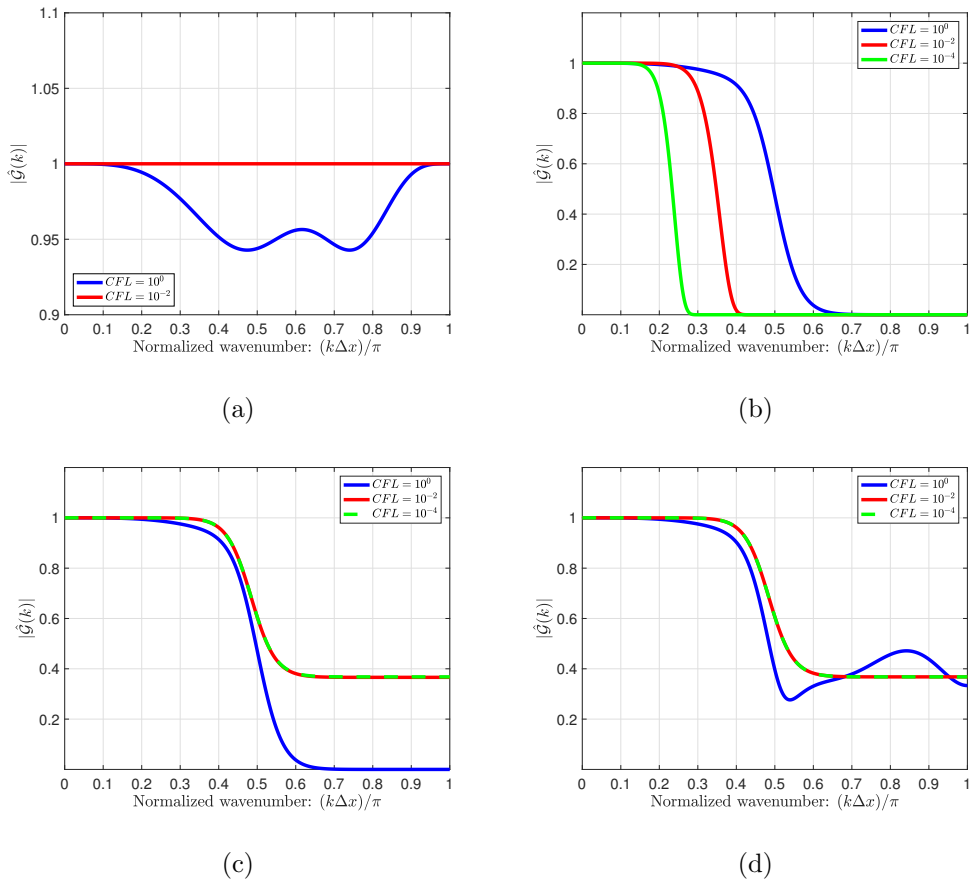

Figure 3.3: Cumulative Growth factor $|\hat{\mathcal{G}}|^{1/CFL}$ as a function of CFL and corresponding to tenth-order Tangent stencil stabilization ($\theta_{1/2} = 0.5\pi$) used with the third-order Runge-Kutta integration scheme (see Table 4.2) and sixth-order central discretization of the 1D hyperbolic equation (see Equation 3.20): a) base scheme with no stabilization, b) solution-filtering, c) filter-based artificial dissipation as proposed in Equation 3.5, d) temporally-consistent solution-filtering as proposed in Equation 3.15.

temporal error effects begin to take effect. As evidenced in Figure 3.4(b), the same is true in the case of the filter-based artificial dissipation scheme, except that the temporal coupling effects may be much more drastic (depending on the underlying filter stencil and the integration method). In this particular case,

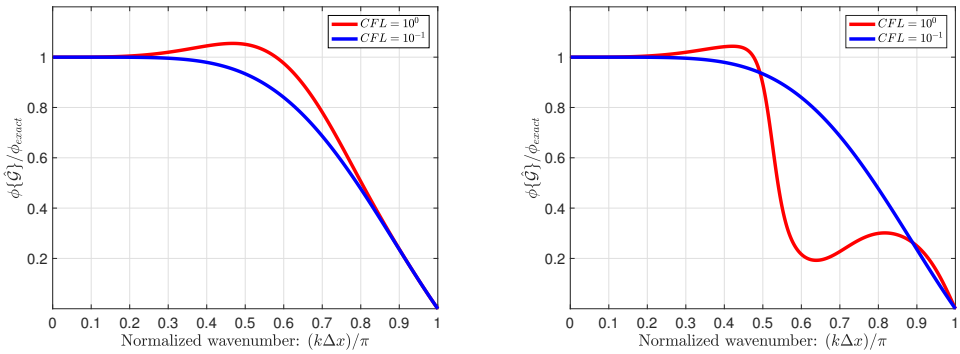

(a) (b)

Figure 3.4: Normalized phase ($\phi\{\hat{\mathcal{G}}\}/[CFL \times (k\Delta x)]$) as a function of CFL and corresponding to tenth-order Tangent stencil stabilization ($\theta_{1/2} = 0.5\pi$) used with the third-order Runge-Kutta integration scheme (see Table tbl:Butcher) and sixth-order central discretization of the 1D hyperbolic equation (see Equation 3.20): a) no stabilization, and b) filter-based artificial dissipation as proposed in Equation 3.15.

one notices a pronounced reduction in phase accuracy which somewhat mirrors the dissipative activation of the filter stencil. While this may be unwanted in terms of dispersive accuracy, such misgivings are tempered by the fact that the sharp increase in phase error is accompanied by an increase in the damping of the stabilization. There are no such phase complications, however, in the case of the temporally-consistent solution-filtering method; nevertheless, the filter-based artificial dissipation formulation will be used henceforth as it provides simpler extension to dual-time iterative procedures – necessary for implicit integration methods (Mundis et al., January 2016).

3.5 Demonstration Test Case: Feature Preservation

The inherent consequence of temporal consistency for stabilization is briefly demonstrated here. The spectral analysis of Section 3.4 has shown that the original solution-filtering algorithm imparts the same amount of damping regardless of the time-stepping. Therefore, it accumulates large amounts of dissipation, assuming a fixed frequency of filtering (with respect to number of time-steps). This leads to an issue of temporal consistency that is addressed by the CFL re-scaling or artificial dissipation formulations proposed.

Consider the following monochromatic signal,

$$u_(x) = \cos\left(\frac{2\pi\kappa_p x}{L}\right) \ , \tag{3.22}$$

to be convected by the advection equation of Equation 3.20 with $\lambda = 1$. The grid is uniform and periodic on $x \in [0, L = 1]$ such that $x(0) = x(L)$. Here, an integral wavenumber of $\kappa_p = N_x/4$ is chosen for the primary signal, corresponding to four points-per-wave (PPW). The signal is properly convected with minimal dissipation and dispersion error by optimized spatial and temporal methods: the fourth-order, eleven-point central discretization from Bogey and Bailly (2004), and the fourth-order, six-stage Low-Dissipation-Dispersion Runge-Kutta integration method of Hu et al. (1996).

In order to highlight the effectiveness of the stabilization schemes, erroneous high-wavenumber modes are added to the original signal according to,

$$u'(x) = u(x) + \sum_{\kappa=Nx/3}^{Nx/2} \left(\frac{1}{\kappa}\right) \cdot \sin\left(\frac{2\pi\kappa x}{L} + \phi_{\mathsf{U}}(\kappa)\right) \ . \tag{3.23}$$

The additional modes constitute waves between 3 PPW and 2 PPW, assigned with randomized phase $\phi_{\mathsf{U}}(\kappa)$. In applying the stabilization, the goal is to retrieve and maintain the original signal $u(x)$. A tenth-order Tangent stencil with $(k_{1/2}\Delta x) = (2/3)\pi$ is chosen due to its scale-discriminant properties which per-

mits preservation of the original signal with strong attenuation of the erroneous content.

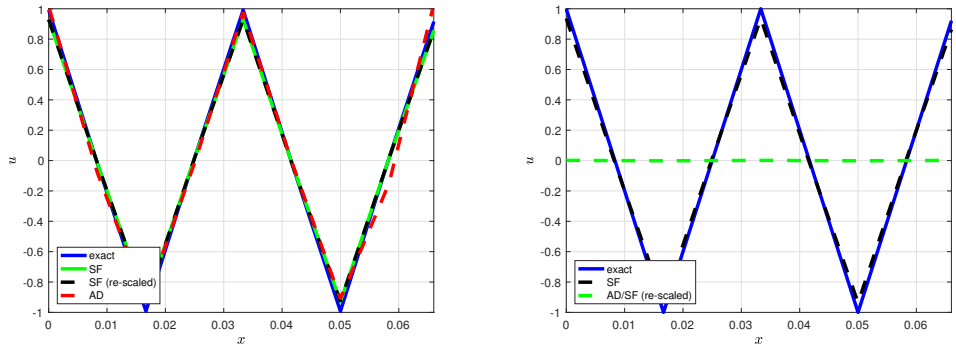

(a) (b)

Figure 3.5: Effect of the stabilization method signal after being convected for four wavelengths of the primary mode $u(x)$: a) CFL = 1.0 and b) CFL = 0.01.

Figure 3.5 plots $u(x)$ and shows the results from employing traditional solution-filtering and the filter-based artificial dissipation (see Equation 3.15) to $u'(x)$ after having convecting a distance of four wavelengths forward. In the case of CFL = 1, shown in Figure 3.5(a), the traditional solution-filtering and filter-based artificial dissipation schemes do equally well; they have removed the erroneous high-wavenumber content and have done a decently good job at maintaining the primary signal. On the other hand, Figure 3.5(b) shows the results for CFL = 0.01 and highlights the temporal inconsistency of traditional solution-filtering; while the erroneous modes have been eliminated, the main signal is almost completely removed. This has to do with the large effective coarsening that results from the cumulative dissipation, as previously highlighted in Figure 3.3(b). On the other hand, the filter-based artificial dissipation scheme is temporally consistent and is able to maintain the primary signal and eliminate the noise; it performs in similar fashion to the CFL = 1.0 scenario. Although not included, tests of the temporally-consistent filtering formulation based on the scaling of Equation 3.5

were shown to agree exactly with the filter-based artificial dissipation scheme – although some discrepancies (i.e., phase error) may occur in general situations, at least for moderate CFL.

The ability to utilize the spectrally-tuned filtering stencils within a stabilization context gives great flexibility for the customization of dissipation schemes based on accuracy considerations, as will be covered in Chapter 4. Further introducing temporally-consistent implementations such as the filter-based artificial methods also allows for proper extension to stiff-systems and pseudo-time iteration methods.

CHAPTER 4

Design of Scale- and Scheme-Discriminant Dissipation Methods

This chapter is taken with some modifications from the AIAA conference proceedings article "The Role of Dispersion and Dissipation on Stabilization Strategies for Time Accurate Simulations" (Edoh et al., January 2016).

4.1 Background

Accurate calculation of the fluid equations requires proper representation of the corresponding transport and diffusive mechanisms. The careful selection of the numerical scheme – both spatial and temporal – is thus crucial. In the case of under-resolved flow such as LES, the characteristics of the equations become primarily hyperbolic and thus significant emphasis must be placed on capturing convection properties while minimizing numerical damping. This is particularly important when considering time-accurate simulations, as compared to steady-state methods which focus on the removal of error through damping mechanisms.

With respect to spatial methods, the emphasis on transport and minimal numerical damping has called into question the suitability of dissipative upwind-type spatial discretizations of the convective terms (Sagaut, 2006). Instead, centralized stencils are preferred. The performance of these different spatial schemes may be assessed by a modified wavenumber spectral analysis (Kravchenko and Moin, 1997) and judged relative to Fourier spectral methods, often taken to be

an exact reference. Relatedly, high-order methods have traditionally been preferred since spectral accuracy can be seen to improve in accordance with improved asymptotic convergence. More recently, however, the potential superiority of spectrally-optimized discretizations has been acknowledged for sustaining feature propagation at low-resolutions. Such schemes, most commonly termed dispersion-relation-preserving (DRP) schemes, sacrifice asymptotic order convergence and look to extend scheme performance to higher wavenumbers for a given stencil size (Lele, 1992; Tam and Webb, 1993; Zingg et al., 1996; Bogey and Bailly, 2004; Linders and Nordstrom, 2015). While use of such stencils is predominantly found in aero-acoustics, the utility for computation of coarse features relative to the LES grid is natural. For instance, Fauconnier et al. (2009, 2011) formulate a dynamic finite difference method capable of switching between optimal-order (i.e., standard) and optimal-spectral (i.e., DRP) stencils; they mark improved results for LES calculations of transitional turbulence such as in the case of the Taylor-Green vortex problem. Keeping with the DRP concept, new compact methods that rely on auxiliary equations have also been developed. For example, the schemes from Liu et al. (2013a,b) use nodal and interface value information to calculate the derivative, which then requires both nodal and interface data to be advanced in time. Similarly, the "layered" scheme of Bai and Zhong (2017) solves auxiliary transport equations for the flux derivatives, in similar fashion to some Discontinuous-Galerkin (DG) concepts. While both of these new methods entail increased memory costs, they are able to attain spectral-like performance and also avoid the typical issue of stationary odd-even modes encountered with central discretizations on collocated grids. The gain in accuracy at high wavenumbers provides further utility for use in under-resolved simulations.

While attaining spectral-like performance via nodal methods is attractive, there is an increased susceptibility to aliasing error. This is a non-linear effect that can be significant in under-resolved situations (i.e., in the absence of phys-

ical diffusion mechanisms) It is furthermore known to be a de-stabilizing source of error (Phillips, 1959) that can be thought to artificially recycle energy into the large scales. Ghosal (1996) shows how aliasing error can overwhelm LES model contributions, and how the issue is made worse for high-order schemes. A similar conclusion can be drawn for DRP methods relating to their enhanced spectral fidelity. The presence of such non-linear errors might thus motivate one to also adopt skew-symmetric and split type representations (Blaisdell et al., 1996; Kennedy and Gruber, 2008), known to reduce aliasing effects, or to perform some type of de-aliasing – necessary even in the case of the Fourier spectral methods. For quadratic non-linearities, Orszag (1971) shows that the "2/3" de-aliasing rule is sufficient for removing all aliasing contributions; in this way, the upper third of the resolvable spectrum needs to be removed. Similar reasoning then generalizes to the fact that N-th order non-linearities produce a "$2/(N+1)$" de-aliasing rule.

Parallel to looking at spatial discretizations is the need to consider temporal integration methods. As previously alluded to, the choice of the time scheme can induce unwanted dispersion and dissipation errors for moderate to large CFL. Analogous to the topic of spatial schemes, high-order temporal implementations are also favored due to their improved accuracy. Efforts to extend time-accurate performance for larger time-stepping has thus also led to the development of optimized time schemes such as the Low-Dissipation-Dispersion Runge-Kutta (LD-DRK) methods (Hu et al., 1996; Bogey and Bailly, 2004). Unfortunately, the coupling of such schemes with the DRP discretizations does not necessarily enhance overall scheme accuracy. As a result, Ramboer et al. (2006) look to optimize the overall scheme, thus taking into account temporal-spatial coupling effects.

The mitigation of numerical error – whether spectral error or aliasing error – can largely be addressed by the removal of the problematic high wavenumbers. This is mainly relating to under performance of the modified wavenumber. The removal of the erroneous content can be done by judiciously introducing numer-

ical dissipation, which would furthermore aid in providing scheme stability. The decision on how best to supply this dissipation is imperative and should be made according to the characteristics of the base scheme, as doing so can minimize the impact of damping error. With respect to transport-dominated equations, it is thus useful to inspect the modified wavenumbers and consider limits implied by spectral and aliasing considerations. Incorporating temporal effects can provide additional insight. The following chapter looks at the performance characteristics of different spatial and temporal schemes, both individually and jointly. In using a scalar advection equation for subsequent analysis, emphasis is placed on characterizing numerical phase properties and selecting proper levels of numerical dissipation. In this way, the benefits of scale-discriminant dissipation – which maintains a clear distinction between the preserved and attenuated scales – are highlighted with respect to tuning damping characteristics relative to the base scheme.

4.2 Resolvability of Numerical Methods

In order to understand how dissipation may be used to properly supplement the base scheme, it is necessary to first characterize the nature of the numerical error. These errors, which represent differences between numerical approximation and the exact solution, are implicitly tied to notions of resolution, as consistency requires that the spatial and temporal discretizations converge to the continuous operators under refinement. Practical use, however, suggests the need to know how much resolution is required in order to support a given error tolerance. With this arises the notion of scheme resolvability. The following first addresses spatial and temporal schemes separately, then considers their coupled performance via von Neumann analysis in order to draw insight into how numerical dissipation may be incorporated into the overall scheme in order to enhance stability while

maintaining or improving accuracy.

4.2.1 Spatial Characterization

Spectral characterization of spatial schemes is possible through the inspection of the modified wavenumber, k_{mod}. First, one may establish an exact representation of the spatial operator by expressing the solution as a Fourier function, as previously shown in Equation 2.3. Taking the first derivative then gives,

$$\partial_x u(x_i) = \sum_k [\imath k] \cdot \hat{u}(k) e^{\imath k x_i} \ . \tag{4.1}$$

Here, it is evident that the representation of the derivative is simply a re-scaling of the original Fourier modes by their respective wavenumber, k. In seeking to derive the modified wavenumber of a discrete derivative stencil, we thus look to discern the nature of an analogous re-scaling of the original signal.

A general discrete implicit representation of $\partial_x u \approx \delta_x\{u\}$ may be written as,

$$u'_i + \sum_{\ell=-J_\ell, \ell\neq 0}^{K_\ell} a_\ell \ u'_{i+\ell} = \frac{1}{\Delta x} \sum_{r=-J_r}^{K_r} b_r u_{i+r} \ , \tag{4.2}$$

where $u'_i = \partial_x u(x_i)$. Substituting Equation 4.1 into Equation 4.2 and performing the necessary expansions about x_i can be shown to yield,

$$[ik_{mod}] = \frac{1}{\Delta x} \left\{ \frac{\sum_{r=-J_r}^{K_r} b_r \ [\cos(k\Delta x_r) + \imath \sin(k\Delta x_r)]}{1 + \sum_{\ell=-J_\ell, \ell\neq 0}^{K_\ell} a_\ell \ [\cos(k\Delta x_\ell) + \imath \sin(k\Delta x_\ell)]} \right\} \ . \tag{4.3}$$

As in the case of the discrete filters, the coefficients (a_ℓ, b_r) are chosen in order to establish consistency with the desired operator and may be optimized with respect to spectral or asymptotic convergence considerations. In the case of a symmetric stencil on a uniform grid, the above simplifies to the following:

$$k_{mod} = \frac{1}{\Delta x} \left\{ \frac{2 \sum_{r=1}^{R} b_r \sin(rk\Delta x)}{1 + 2 \sum_{\ell=1}^{L} a_\ell \cos(\ell k \Delta x)} \right\} \ . \tag{4.4}$$

In this case, the modified wavenumber k_{mod} is real in accordance with the exact Fourier representation of Equation 4.1. By comparing k_{mod} to the exact scaling

k, one is then able to assess how well the finite difference scheme can approximate exact differentiation. Furthermore, as such comparisons are a function of wavenumber, one may extract information about the spatial resolvability of the scheme.

scheme	CD02	CD04	CD06	CD04-7pt *	CD10	CD04-11pt *
b_1	1/2	8/12	45/60	0.79926643	1050/1260	0.87275699
b_2		$-1/12$	$-9/60$	-0.18941314	$-300/1260$	-0.28651117
b_3			1/60	0.02651995	75/1260	0.09032000
b_4					$-12.5/1260$	-0.02077940
b_5					1/1260	0.00248459

Table 4.1: Stencil coefficients for the first derivative, with $b_0 = 0$ and $b_r = -b_{-r}$ (*optimized coefficients, as written, are only listed to single-precision).

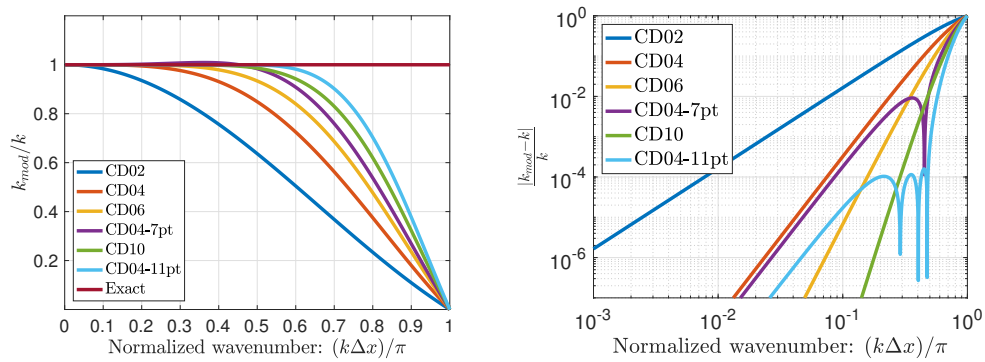

(a) (b)

Figure 4.1: Central-difference scheme comparison: a) normalized modified wavenumber (k_{mod}/k), and b) relative error of modified wavenumber ($|k_{mod}-k|/k$).

Table 4.1 includes coefficients for some explicit central difference stencil discretizations (i.e., $a_{\ell \geq 1} = 0$, written shortly as $CDXX$, with XX denoting the stencil order. Listed are standard schemes as well as spectrally-optimized (i.e., DRP) methods which include a specifier regarding stencil width (e.g, CD04-7pt

is a fourth-order, seven point stencil). The spectral performance of the respective schemes can easily be inspected by looking at the normalized modified wavenumber (k_{mod}/k). This quantity is plotted in Figure 4.1(a) as a function of $(k\Delta x)/\pi$. It is apparent that all discretizations provide accurate representation of the derivative at low wavenumbers (i.e., high resolutions) but become increasingly inaccurate at high wavenumbers (i.e., low resolutions). The decrease in accuracy at small-scales is shown to be less severe by employing higher-order stencils. The use of optimized schemes further improves the range of resolvability, wherein fewer points per wave (PPW) are needed to represent the derivative of a waveform. For example, it is apparent that the optimized fourth-order schemes (CD04-7pt and CD04-11pt) show superior spectral properties relative to the standard five-point, fourth-order scheme (CD04). This is expected, however, due to the fact that these optimized schemes utilize larger stencils and thus have more degrees of freedom (DOF) through which properties may be enhanced. Therefore, it is noteworthy to compare schemes of equivalent stencils size or DOF. The CD06 and CD04-7pt schemes each utilize a seven-point stencil. Again, it is evident that the CD04-7pt scheme accurately resolves a larger part of the spectrum. The same can be said when comparing CD10 and CD04-11pt, which each utilize an eleven-point stencil. Assessing the modified wavenumber in terms of relative error on a logarithmic scale provides an alternate perspective and additional insight. Figure 4.1(b) thus plots the relative error $|k_{mod} - k|/k$ logarithmically. Such a representation emphasizes the theoretical truncation error of the respective schemes, which becomes evident as $(k\Delta x) \to 0$ (i.e., at high resolutions). Here it is also evident that the featured optimized schemes (CD04-7pt and CD04-11pt) are indeed of fourth-order but have lower error constants relative to their standard counterpart. The perspective offered by this logarithmic error plot furthermore reveals the fact that for a given DOF, the standard schemes outperform their optimized counterparts at high-resolutions but are less accurate at higher wavenumbers. Looking at Padé

(i.e., implicit) stencil representations, though not done here, is also useful as these generally produce enhanced spectral characteristics for given DOF while maintaining compact stencil widths (Lele, 1992; Mahesh, 1998). By understanding the relative merits of the spatial discretization schemes, it then becomes possible to quantify the limits of spatial resolvability. This, in turn, can be used to develop a suitable numerical dissipation strategy that targets erroneously calculated modes in scale-discriminant fashion. The approach then results in a scheme-discriminant procedure that will serve as the base scheme for the eventual LES implementations to be proposed.

4.2.2 Temporal Characterization

Characterization of the temporal scheme also gives insight on sources of numerical error. To do this, one can assume a general linear ordinary differential equation (ODE),

$$\begin{aligned} \frac{du}{dt} &= f(u) \\ &\approx \beta \cdot u \ , \quad \text{where} \quad \beta = eig\{f(u)\} \ . \end{aligned} \tag{4.5}$$

Such a linear expression may be integrated exactly to yield,

$$u(t) = Ce^{\beta t} \ . \tag{4.6}$$

Based on the linearity assumption, solution advancement in time can furthermore be expressed as,

$$u(t + \Delta t) = e^{\beta \Delta t} u(t) \ . \tag{4.7}$$

Inspecting the form of Equation 4.7 and comparing it to Equation 3.17 reveals that $e^{\beta \Delta t}$ corresponds to an amplification factor. In the case of a multi-stage Runge-Kutta scheme (see Equations 3.14-3.15), the resulting amplification factor can be written as a "characteristic" polynomial $P(\beta \Delta t)$ based on the Butcher

Tableau coefficients (Sanderse, 2013):

$$P(\beta\Delta t) = \frac{\mathrm{Det}[I-(\beta\Delta t)\tilde{D}]}{\mathrm{Det}[I-(\beta\Delta t)\tilde{A}]}\ , \quad \text{with} \quad \tilde{D}=\tilde{A}-\tilde{e}\tilde{b}^T,\ \tilde{e}=[1\cdots 1]^T\ . \tag{4.8}$$

This characteristic polynomial can be understood as a Taylor-series approximation to the exact amplification factor,

$$e^{\beta\Delta t}=\sum_{n=0}^{\infty}\frac{(\beta\Delta t)^n}{n!}\approx P(\beta\Delta t)\ . \tag{4.9}$$

As noted in the case of spatial schemes, accuracy of the temporal method may be improved by considering high-order or optimized implementations. Implicit integrators can furthermore be understood to be more accurate, as they represent a rational function (i.e., Padé) approximation to the exponential function.

$\tilde{c}$	$\underline{\tilde{A}}$
	$\tilde{b}$

General Butcher Tableau

(a)

0	0	0
1	1/2	1/2
	1/2	1/2

Implicit, Second-order Crank-Nicolson

(b)

0	0	0	0
8/15	8/15	0	0
2/3	1/4	5/12	0
	1/4	0	3/4

Explicit, Third-order RK (Wray)

(c)

Table 4.2: Runge-Kutta Butcher Tableaux of coefficients.

In order to illustrate the effect of using high-order Runge-Kutta methods, Figure 4.2 includes ODE magnitude thumbprints $|P(\beta\Delta t)|$ of some explicit and implicit integrator schemes. The magnitude thumbprints represent damping effects

and are compared to the real component of the exact linear integrator, $e^{\mathrm{Re}\{\beta\}\cdot\Delta t}$. In the case where $\mathrm{Re}\{\beta\} < 0$ (i.e., decaying modes), one expects dissipation of the signal in accordance with the magnitude of the real eigenvalue and independently of the imaginary component. This results in vertical lines of constant damping as shown in Figure 4.2(a) for the exact integration method. Inspecting the explicit methods in Figure 4.2(b) and Figure 4.2(c) reveals that such methods are conditionally stable in the negative real plane. This further results in the effective dissipation being notably dependent on both the real and imaginary components of the system eigenvalue. Although the third-order Runge-Kutta (RK3) method (see Table 4.2(c)) shown in Figure 4.2(b) is of higher asymptotic order than the second-order Low-Dissipation-Dispersion Runge-Kutta (LDDRK2-5) scheme featured in Figure 4.2(c) (see Hu et al. (1996); Stanescu and Habashi (1998) for corresponding Butcher Tableau coeffiecients), the latter is optimized for improved damping characteristics along the imaginary axis. The accompanying benefits to the optimization procedure include an enlarged stability region in addition to more accurate dissipation characteristics for decaying modes – at least for moderate radii of $|(\beta\Delta t)|$ from the origin. Therefore, the optimized LDDRK2-5 scheme is seen to be more accurate than the higher-order RK3 method for dissipation. On the other hand, the implicit schemes in Figure 4.2(d) and Figure 4.2(e) are seen to be unconditionally stable for decaying modes (i.e., A-stable) and thus are more accurate than the explicit methods for large time-steps. Relative to exact integration, however, it is evident that the second-order Crank-Nicolson scheme in Figure 4.2(d) does not replicate the fact that $|e^{\mathrm{Re}\{\beta\}\cdot\Delta t}| \to 0$ as $(\beta\Delta t) \to -\infty$ (i.e., L-stability). By contrast, the fourth-order optimized Explicit Singly-Diagonally Implicit Runge-Kutta (ESDIRK4) method (Bijl et al., 2002) satisfies this property and thus is better suited for large $(\beta\Delta t)$ values (either with respect to the time-step or system stiffness). This latter scheme is also seen to be superior in terms of emulating proper dissipative characteristics (i.e., nearly-vertical constant lines

of dampening), at least for moderate radii of $|(\beta\Delta t)|$ from the origin. However, with respect to the preservation of pure oscillatory modes (i.e., purely imaginary eigenvalues), the Crank-Nicolson scheme satisfies perfect preservation of the associated amplitudes regardless of the time-step size, while the ESDIRK4 scheme exhibits a limited region for accurate preservation.

As a complement to the ODE magnitude thumbprints, one may inspect the phase thumbprints ($\phi\{P(\beta\Delta t)\} = \tan^{-1}\{\mathrm{Im}\{P(\beta\Delta t)\}/\mathrm{Re}\{P(\beta\Delta t)\}\}$) in order to judge dispersion effects stemming from the integration method. Such plots are included in Figure 4.3 for explicit and implicit schemes. These are furthermore compared to the calculated phase from exact integration, $e^{\mathrm{Im}\{\beta\}\cdot\Delta t}$, which is shown to interpret phase independently of the real component of the eigenvalue. This would then suggest horizontal lines of constant phase. Such characteristics are generally not attainable by any of the temporal schemes in question; however, the high-order and optimized methods are able to extend their regions of accurate phase representation, at least for moderate radii of $|(\beta\Delta t)|$. Of further interesting note is the fact that $\phi\{P(\beta\Delta t)\}$ features branch cuts corresponding to evaluations of the inverse tangent function for determining phase angle; this results in instances where increasing the time-step size reverses the perceived phase direction, which can be interpreted as a temporal equivalent to aliasing.

Studying the magnitude and phase thumbprints of different integration methods thus provides useful information on the dissipation and dispersion proclivities of the schemes. These then help to anticipate potential benefits or shortcomings of the respective schemes, to be assessed with respect to the task at hand.

4.2.3 Spatio-Temporal Coupling Effects

While separately understanding the behavior of the spatial and temporal schemes provides insight into which methods one may wish to select, it is nevertheless im-

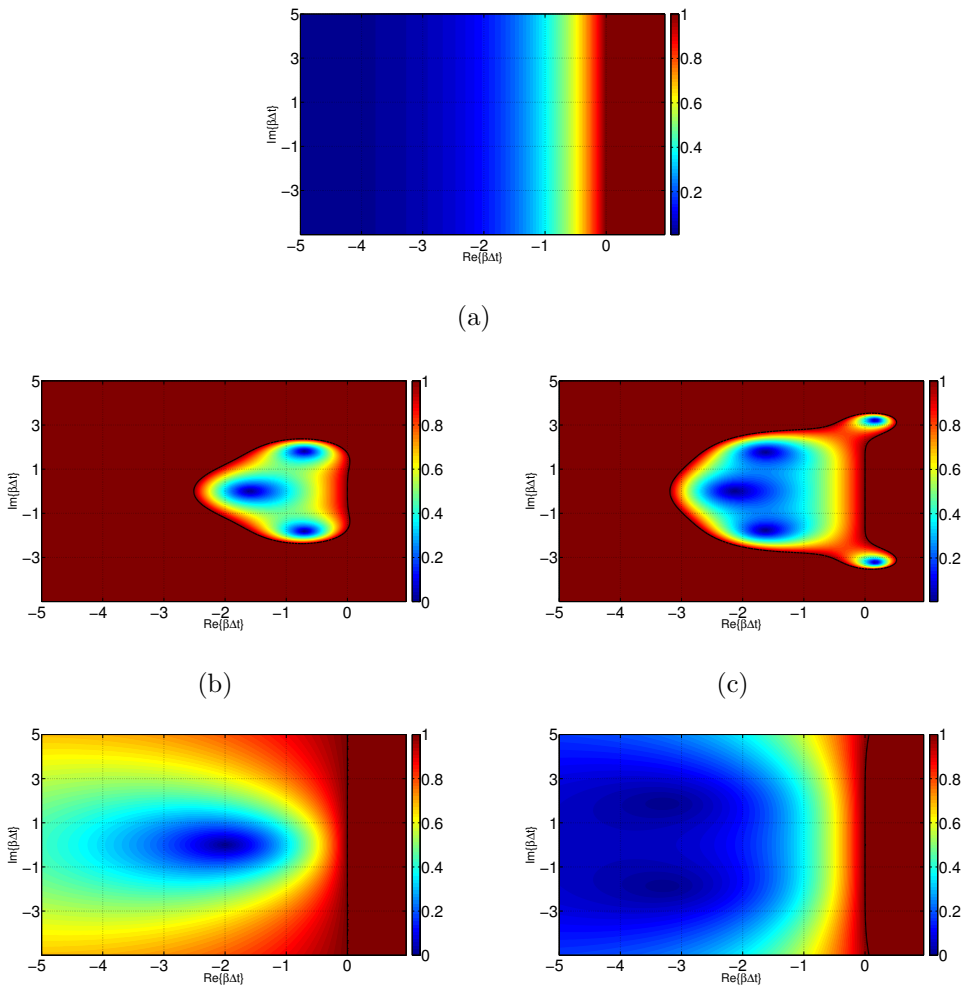

Figure 4.2: Magnitude ODE contours $|P(\beta\Delta t)|$: a) exact integration (Equation 4.7), b) third-order explicit Runge-Kutta (RK3) scheme (see Table 4.2), c) second-order, five-stage optimized explicit Low-Dissipation-Dispersion Runge-Kutta (LDDRK2-5) scheme (Hu et al., 1996; Stanescu and Habashi, 1998), d) second-order implicit Crank-Nicolson (CN) scheme (see Table 4.2), and e) fourth-order, six-stage optimized Explicit Singly-Diagonally Implicit Runge-Kutta (ESDIRK4) scheme (Bijl et al., 2002).

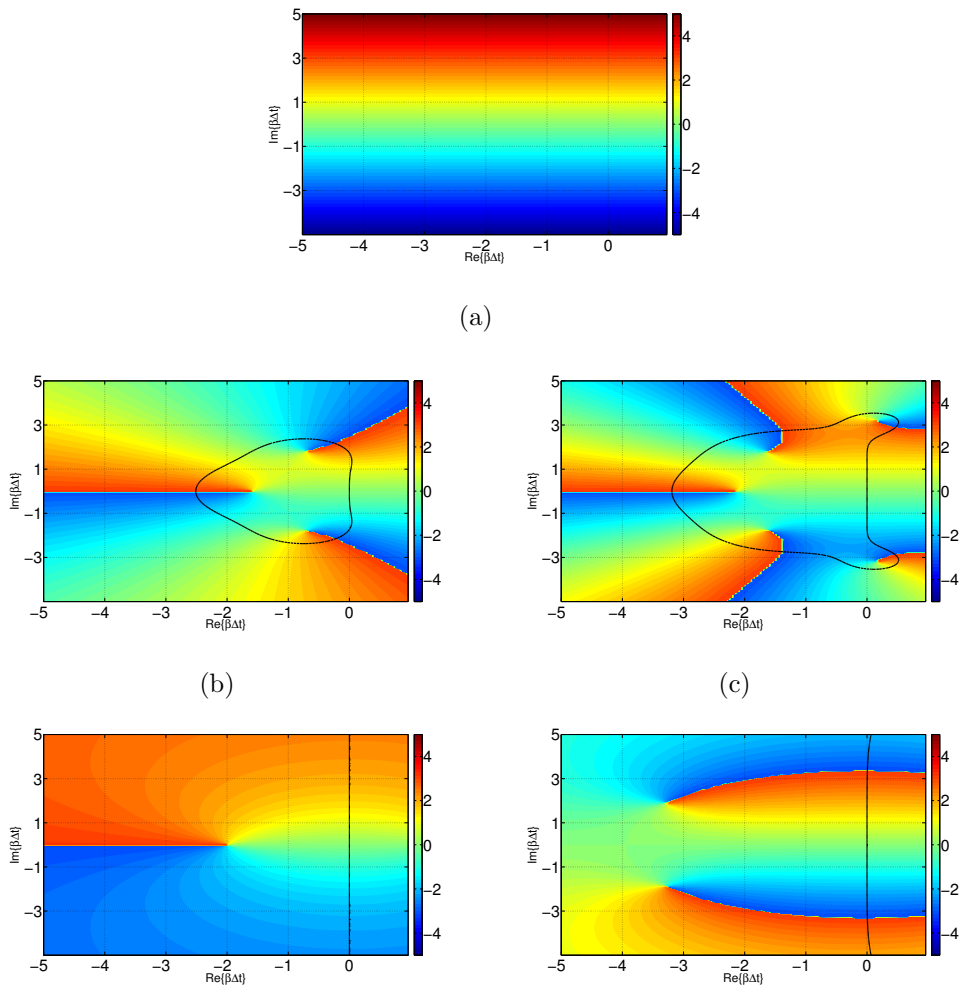

Figure 4.3: Phase ODE contours $\phi\{P(\beta\Delta t)\}$: a) exact integration (Equation 4.7), b) third-order explicit Runge-Kutta (RK3) scheme (see Table 4.2), c) second-order, five-stage optimized explicit Low-Dissipation-Dispersion Runge-Kutta (LDDRK2-5) scheme (Hu et al., 1996; Stanescu and Habashi, 1998), d) second-order implicit Crank-Nicolson (CN) scheme (see Table 4.2), and e) fourth-order, six-stage optimized Explicit Singly-Diagonally Implicit Runge-Kutta (ESDIRK4) scheme (Bijl et al., 2002).

portant to characterize the coupled behavior of the overall numerical treatment. In some instances, errors from each discretization component may get amplified when considered jointly; meanwhile, in other circumstances, these errors may balance each other out. In this way, there may be favorable combinations of spatial and temporal schemes, which has motivated research into the coupled optimization of the temporal and spatial methods (Ramboer et al., 2006). Here, we explore results relating to the simplified notion that the full potential of the base scheme is achieved when accurate spatial schemes are used with sufficiently accurate temporal schemes. While additional benefits may arise from a joint spatio-temporal optimization, synthesizing information from the previous individual spatial and temporal scheme analyses of Sections 4.2.1 and 4.2.2 provides useful guidance into how the components interact with each other.

The spatio-temporal coupling effects on dissipation and dispersion are captured in a spectral sense within the von Neumann analysis introduced in Equations 3.17-3.19 of Section 3.4. However, the framework as presented thus far makes it harder to intuitively understand trends associated with the selection of different spatial or temporal methods. Subsequent impacts of the CFL parameter (responsible for characterizing respective time and space discretization effects relative to the pertinent physics) on performance is also lost. When considering general classes of ODE integrators – rather than fully-discrete implementations – one gains the benefit of employing a methods-of-lines perspective, wherein the ODE eigenvalues are directly related to the spatial discretization. By doing this, the system eigenvalues then become an implicit function of the spatial scheme and thus can be parameterized spectrally by a wavenumber, $\beta(k)$. From here, one may simply inspect the path of the ODE eigenvalues against the ODE magnitude or phase thumbprints in order to ascertain the predicted dissipation and dispersion performance of different scheme combinations. This makes the process of selecting an appropriate integrator method more intuitively clear given the eigenvalue spectrum associated

with the spatial scheme and prescribed CFL.

Substituting the Fourier representation into the advection equation (see Equation 3.20) suggests that $\beta = -(\lambda \cdot \imath k'_{mod}/\Delta x)$. In the case of symmetric representations of the δ_x operator, the modified wavenumber is real therefore the system eigenvalue β is purely imaginary. This represents a strictly oscillatory signal – in other words, a transport equation. The characteristic polynomial, which captures dissipative and dispersive traits of the overall system, can then be described as a function of $(\beta\Delta t) = (\mathrm{CFL}_\lambda \cdot k'_{mod}(k))$. From this, von Neumann stability plots can then be made. The slight change in perspective offered by this alternate process, however, allows for some interesting observations to be made. For instance, a larger $\max\{k_{mod}\}$ is seen to result in a more restrictive CFL limit with respect to stability and accuracy considerations; in other words, high-order or optimized spatial discretizations will require more accurate integrators for a given time-step sizes, or may otherwise suffer the unwanted consequences of increased damping or phase errors. Such generalized statements are harder to make from performing a traditional von Neumann analysis procedure. In this new framework, the different spatial and temporal schemes can more easily be studied jointly.

In order to demonstrate the concomitant need for accurate spatial and temporal schemes, one can look at growth factor and phase error trends. In the case of pure advection, the exact amplification is unity ($|G|_{\text{exact}} = 1$), while the phase is taken to be $\phi_{\text{exact}} = \lambda k \Delta t$, which assumes exact exponential integration and a Fourier spatial scheme. Figure 4.4 demonstrates the way in which the choice of temporal scheme can at times hinder the presumed benefits of employing a more spectrally-accurate spatial scheme. Figure 4.4(a), for instance, explores the phase performance of the sixth-order standard central difference scheme (CD06) with different temporal integrators. Immediately, it is evident that the high-order and optimized time schemes impart less dispersion error. Next considering the fourth-order, eleven-point optimized stencil (CD04-11) (Bogey and Bailly, 2004)

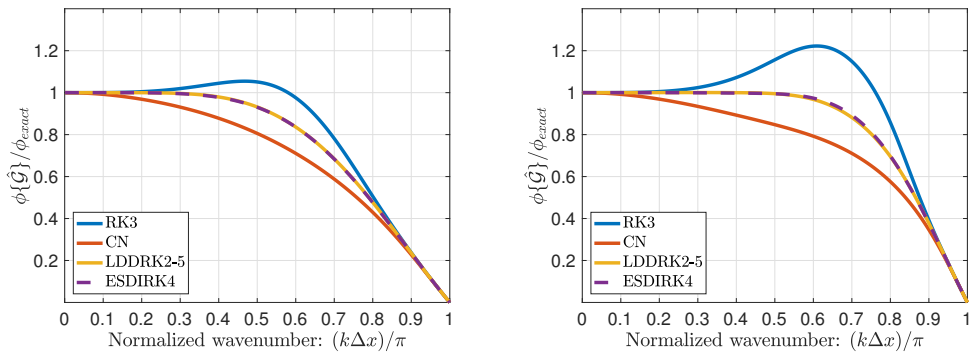

(a) (b)

Figure 4.4: Relative phase error $|\phi_{\text{num}} - \phi_{\text{exact}}|/\phi_{\text{exact}}$ of integration methods at CFL = 1.0 with respective spatial representations: a) CD06, and b) CD04-11.

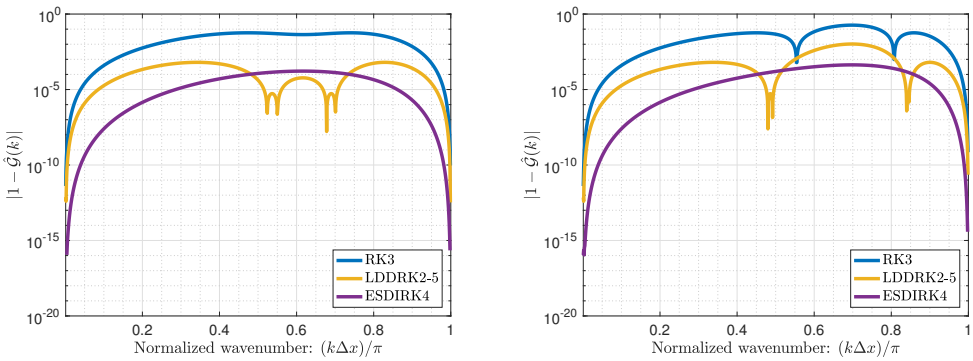

(a) (b)

Figure 4.5: Relative growth factor error $||G|_{\text{num}} - |G|_{\text{exact}}|/|G|_{\text{exact}}$ of integration methods at CFL = 1.0 with respective spatial representations: a) CD06, and b) CD04-11.

in Figure 4.4(b), it is apparent that the high-accuracy time schemes are again more accommodating. The less accurate time schemes are seen to either stifle the benefits of incorporating an improved spatial scheme or are seen to further degrade the overall scheme performance. Similar conclusions on performance are drawn when considering damping characteristics, as shown in Figure 4.5 for the growth factor error $|P(\beta\Delta t)| - |G_{\text{exact}}|$. In general, the high-accuracy time schemes are better able to preserve amplitudes of the different wavenumbers. And in the case of employing a more spectrally-accurate spatial scheme with a larger $\max\{k_{mod}\}$, the high-accuracy temporal schemes are less susceptible to imparting false temporal dissipation. In the case of the RK3 method, it is furthermore apparent that the overall scheme becomes unstable for the prescribed CFL = 1.0, which further highlights the temporal method's inability to adequately support spectrally-accurate spatial schemes[1]. A special exception to these trends are time-reversible or energy conserving schemes (Sanderse, 2013) which demonstrate perfect preservation for oscillatory systems, regardless of their formal order. Such an example is the second-order Crank-Nicolson scheme, excluded from Figure 4.5 for this reason.

In accompaniment to the above observations is the fact that the featured schemes preserve odd-even amplitudes. This can be seen as a direct consequence of the spatio-temporal coupling effects, where $|P(\beta\Delta t)| \to 1$ because $\beta(k) \to 0$ as k approaches the Nyquist frequency. The lack of high-wavenumber damping in the featured base schemes thus further motivates the need for suitable numerical dissipation.

[1]The LDDRK2-5 scheme is also seen to have a region of instability; however, this is believed to be a result of the finite precision of the reported Butcher Tableau coefficients (personal communication with Stanescu)

4.3 Pairing Stabilization with Base Discretizations

Having gained useful insight on the pairing of spatial and temporal methods that comprise a base scheme, the task of providing suitable dissipation for abating numerical error may be addressed. Solution of a hyperbolic-type system suggests selection of a base scheme that minimizes both dissipation and phase error. As a result, the incorporation of damping needs to be made with careful deliberation.

In most cases, issues relating to phase characteristics dominate and are generally driven by the spatial representation, which primarily imparts dispersion error. As previously noted, the benefits associated with the spectral-fidelity of such schemes is tempered by the instigation of aliasing effects in the case of non-linear problems. Therefore, in the spirit of enhancing the overall numerical accuracy of the base scheme as applied to transport-dominated situations, it is reasonable to judge the need for numerical dissipation based on the presence of phase error. To do this, a tolerance on an acceptable level of dispersion error can first be established.

Table 4.3 tabulates normalized cut-off wavenumbers ($k_c\Delta x/\pi$) corresponding to a 1% relative error in the phase of the base scheme. This is shown for different couplings of spatial and time integration methods evaluated at a CFL = 1.0. As expected, higher cut-offs are attainable for the high-order and optimized spatial schemes. In other words, it is possible to represent modes with fewer points-per-wave (PPW) by utilizing these more spectrally-accurate stencils. The influence of the temporal schemes on the effective resolution capacity is furthermore evident. Again, implementing a high-order or optimized temporal method limits further deleterious impacts of temporal error on the phase.

The established cut-offs based on phase characteristics should then be assessed relative to the system's non-linearities and any relevant aliasing limits. From these, the most conservative cut-off can be selected for tuning of the numerical

	exact	RK3	CN	LDDRK2-5	ESDIRK4
CD02	0.08	0.08	0.06	0.08	0.08
CD04	0.24	0.37	0.11	0.24	0.24
CD06	0.35	0.25	0.11	0.35	0.34
CD04-7pt	0.49	0.22	0.11	0.48	0.48
CD10	0.48	0.24*	0.11	0.47	0.46
CD04-11pt	0.58	0.24*	0.11	0.53	0.54

Table 4.3: Cut-off frequencies ($k_c\Delta x/\pi$) for an error tolerance $|(1-\phi_{\text{num}}/\phi_{\text{exact}})| = 0.01$ at CFL = 1.0 (* corresponding scheme is unstable at the prescribed CFL).

dissipation. Pulling from the conclusions of Chapter 3, a filter-based artificial dissipation formulation is chosen henceforth. Specific selection of the dissipation stencils is made according to scale-discriminant considerations. In this way, the dissipation scheme maintains scheme accuracy by preserving as much of the well-resolved modes as possible; meanwhile, it aggressively attenuates the higher modes that are prone to instigating non-linear numerical instability and spectral-error. As the phase-based cut-offs listed above ascribe to a 1% error rating, an analogous ($k_{0.99}\Delta x$) prescription of the filter cut-off is adopted, with the anticipation that the corresponding attenuation will match the fall-off in accuracy of the base-scheme when a filter scheme of equivalent merit (i.e., formal order or DOF) is utilized.

4.4 Demonstration Test Case: Isentropic Vortex Propagation (2D Euler Equations)

The two-dimensional isentropic vortex is considered here in order to demonstrate the potential consequences of pairing the dissipation scheme to the dispersion characteristics of the base method. This test case is often used by authors to show the stability and preservation properties of numerical schemes. The governing

system is given by the compressible Euler equations for an ideal gas:

$$\frac{\partial Q}{\partial t} + \frac{E_j}{\partial x_j} = 0 \ , \tag{4.10}$$

$$Q = \begin{bmatrix} \rho \\ \rho u_i \\ \rho e \end{bmatrix}, E_j = \begin{bmatrix} \rho u_j \\ \rho u_i u_j + P\delta_{ij} \\ (\rho e + P)u_j \end{bmatrix} \ , \tag{4.11}$$

$$e_o = \tfrac{p}{\rho(\gamma-1)} + \tfrac{1}{2}u_i u_i,$$
$$P = \rho R T,$$

with constants $\gamma = c_p/c_v$ and $R = c_p - c_v$. The inviscid nature of the equations thus requires numerical stabilization in order to control susceptibilities to non-linear errors. The flow is initialized as velocity and temperature perturbations made to a background uniform flow Q_∞ (Mundis et al., 2015):

$$\begin{aligned} \delta u &= -\sqrt{R_\infty T_\infty}\left(\tfrac{\alpha}{2\pi}\right)(y - y_o)e^{\phi(1-r^2)} \\ \delta v &= \sqrt{R_\infty T_\infty}\left(\tfrac{\alpha}{2\pi}\right)(x - x_o)e^{\phi(1-r^2)} \\ \delta T &= -T_\infty\left[\tfrac{\alpha^2(\gamma-1)}{16\phi\gamma\pi^2}\right]e^{2\phi(1-r^2)} \\ &\text{with} \quad r^2 = (x - x_o)^2 + (y - y_o)^2 \ . \end{aligned} \tag{4.12}$$

The solution definition is then made complete by employing an isentropic relation, $P = P_\infty(T/T_\infty)^{\gamma/(\gamma-1)}$.

As the ensuing perturbation is a solution to the governing inviscid equations, one would expect the feature to be perfectly preserved. The presence of numerical error stemming from dispersion or aliasing, however, often causes distortion of the vortex and can threaten robustness of the simulation. Therefore, in order to achieve long-time preservation of the vortex, one needs to employ accurate and stable discretizations. Here, the choice of employing numerical dissipation for long-time preservation of the vortex is studied relative to a base scheme – noting

the consequences to both robustness and high-accuracy, which may at times be in conflict.

Here, the vortex is initialized within a background flow propagating in the x-direction:

$$\begin{array}{lll} \rho u_\infty = 200.0 \ \left[\frac{kg}{m^2 \cdot s}\right] & \rho v_\infty = 0.0 \ \left[\frac{kg}{m^2 \cdot s}\right] & \rho w_\infty = 0.0 \ \left[\frac{kg}{m^2 \cdot s}\right] \\ \rho_\infty = 1.0 \ \left[\frac{kg}{m^3}\right] & \rho e_{0,\infty} = 305714.3 \ \left[\frac{kg}{m \cdot s^2}\right] & \end{array} . \quad (4.13)$$

The thermodynamic constants are $R_\infty = 287.11 \ \left[\frac{J}{kg \cdot \mathrm{K}}\right]$ and $\gamma = 1.4$ that yield $c_\infty = 400.00 \ [m/s]$, $T_\infty = 398.06 \ [K]$, $P_\infty = 114285.6 \ [Pa]$. The strength of the resulting vortex is largely controlled by the α parameter, while the gradient or width of the feature is strongly dependent on the ϕ parameter designation. The following investigation relates to a relatively weak vortex ($\alpha = 1.0, \phi = 1.0$), which produces a density perturbation of approximately 4% of the uniform field. The computational grid is uniform and periodic with length $L = 21.0 \ [m]$, which corresponds to approximately ten vortex widths as judged by a relaxation to uniform conditions ($\delta T / T_\infty = 0.001$). The resolution of the domain ($N_x = N_y = 60$) is chosen such that a majority of the spectrum is represented with at least a 4 points-per-wave (PPW) resolution. And this corresponds to a $6\Delta x$ resolution across the vortex. As the goal of the present exercise is preservation of the vortex feature, density is chosen as the observable variable. Figure 4.6 plots the density solution along with a normalized two-dimensional plot of the resulting spectrum; it furthermore highlights the fact that much of the signal is captured below $(k\Delta x) = 0.5\pi$, which also pertains to the aliasing limit for cubic non-linearities as pertaining to the compressible equations (relative to primitive variables).

Given the spectral qualities of the problem, one can then choose a base discretization scheme accordingly. Referencing the phase error-based resolvability metrics included in Table 4.3, it would seem that either tenth-order standard discretization (CD10) or the optimized fourth-order schemes (CD04-7pt, CD04-11pt)

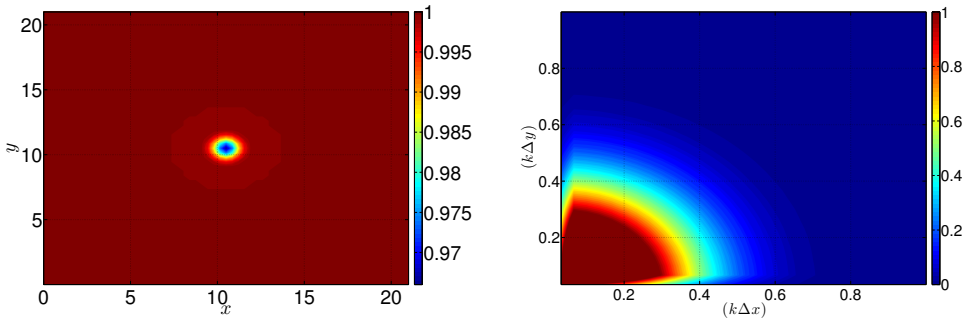

(a) (b)

Figure 4.6: Initial condition for isentropic vortex test case: a) density solution, and b) normalized two-dimensional power-spectral density of density solution ($|\hat{f}(k_1, k_2, t)|^2/|\hat{f}(0, 0, t = 0)|^2$).

would be appropriate for the 4 PPW performance target. The following investigation treats the governing system in conservative form with the inviscid terms expressed in divergence form. The RK3 integration scheme is considered for integration purposes. The chosen time-step size of $\Delta t = 1.75 \times 10^{-4}$ $[s]$ corresponds to a one-dimensional convective $\text{CFL}_{u,1D} = 0.1$, therefore temporal error effects are understood to be minimal. Taking into account an effective Mach number of 0.5 and multi-dimensional effects, the above is conservatively estimated to correspond to approximately one-sixth of the two-dimensional stability-based acoustic stability limit, $\text{CFL}_{(|u_{mag}|+c)}$.

	CD02	CD04	CD06	CD04-7pt	CD10	CD04-11pt
$\text{CFL}_{u,1D} = 0.1$	1538	890	996	940	1076	925
$\text{CFL}_{u,1D} = 0.01$	15262	8794	9798	9270	10615	9128

Table 4.4: Number of time-steps before which respective simulations of the 2D isentropic vortex integrated by the RK3 method failed.

Left to operate without any stabilization, the central difference methods, both high- and low-order, quickly go unstable; Table 4.4 lists the approximate num-

ber of time-steps as a function of $\mathrm{CFL}_{u,1D}$ before the respective simulations fail. Comparing the results of $\mathrm{CFL}_{u,1D} = 0.1$ and 0.01, it is apparent that the chosen time-step size of $\Delta t = 1.75 \times 10^{-4}$ $[s]$ is adequately small and justifies the subsequent neglect of temporal error effects in this analysis. The table, however, does not necessarily provide any insight on the role that scheme order, stencil size, or spectral-accuracy may play in the natural robustness of the methods. The implementations are seen to go unstable after approximately fifteen vortex widths. Comparing the density solutions, as shown in Figure 4.7, furthermore highlights that robustness does not necessarily imply accuracy. For example, the second-order centrally-differenced method (CD02) is able to run for the longest time but is seen to have significantly more error after having traveled only ten vortex-widths. The stability of the simulation may in part be tied to modified wavenumber considerations and the related abatement of aliasing; however, this cannot entirely be the explanation as the fourth-order CD04 method is seen to be the most unstable of the high-order methods, despite its small stencil width and subdued modified wavenumber relative to the higher-accuracy schemes. In general, the density solution plots of Figure 4.7 suggest that higher spectral fidelity provides better preservation of the vortex shape, at least at early times. This is corroborated by the power spectral density plots of Figure 4.8 that show a more coherent spectrum for the high-accuracy methods. As suggested, however, these high-order and optimized schemes are seen to have more erroneous energy in the high wavenumbers – likely resulting from the effects of an enhanced modified wavenumber that supports non-linear cascading of modes – which eventually threaten stability. The challenge of maintaining a robust and accurate solution can then be seen as being able to preserve the defining spectrum while simultaneously avoiding the generation of new modal content that can de-stabilize the solution. This is a similar concept to that encountered in explicitly-filtered LES, where one would expect the model contributions to implicitly enforce the proper

LES resolution.

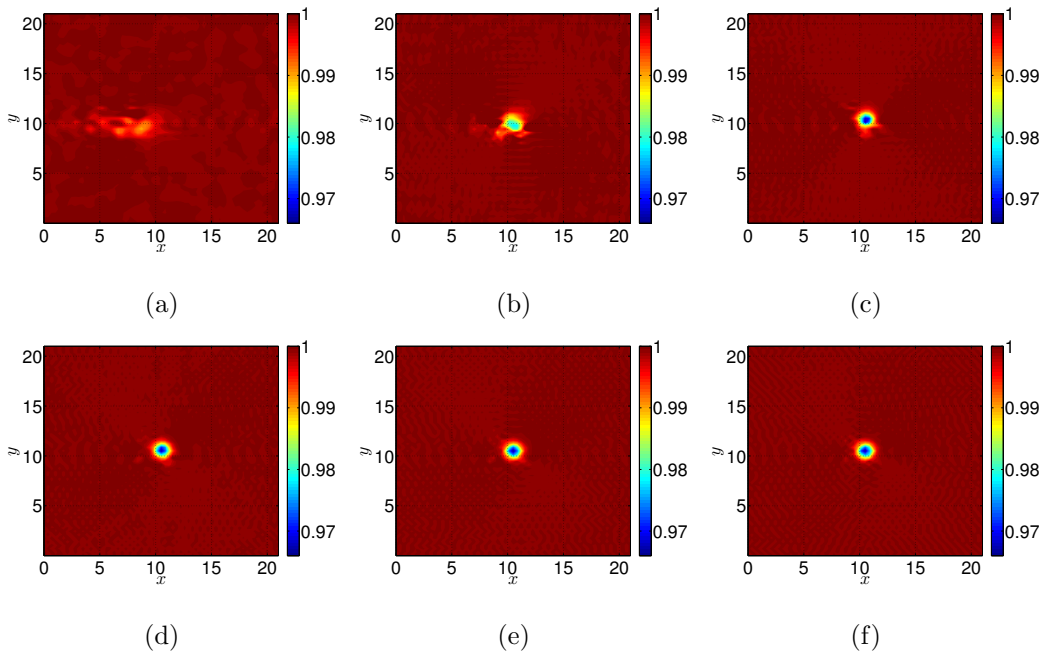

Figure 4.7: Isentropic vortex density solution after having traveled ten vortex widths with respect to different spatial discretization schemes: a) CD02, b) CD04, c) CD06, d) CD04-7pt, e) CD10, and f) CD04-11pt.

In order to achieve this goal, we seek to proceed under the guidance of the following proposed rationale: robustness is enhanced through the addition of numerical dissipation, but concomitant accuracy is achieved by tuning the dissipation relative to the solution spectrum and the resolvability of the base scheme. To demonstrate this, the performance of the CD04, CD06, and CD04-7pt schemes are compared when paired with different tuned filter-based artificial dissipation methods. These three schemes are chosen such as to help delimitate any effects associated with scheme order and spectral resolvability. Figure 4.9 plots the respective normalized modified wavenumbers k_{mod}/k on top of the responses of two sixth-order Tangent filter responses, respectively tuned to $(k_{99}\Delta x)/\pi = 0.35$ and 0.5. These cut-offs are chosen such as to correspond to the resolvability metric of the CD06 and CD04-7pt schemes, respectively, based on a 1% relative phase error

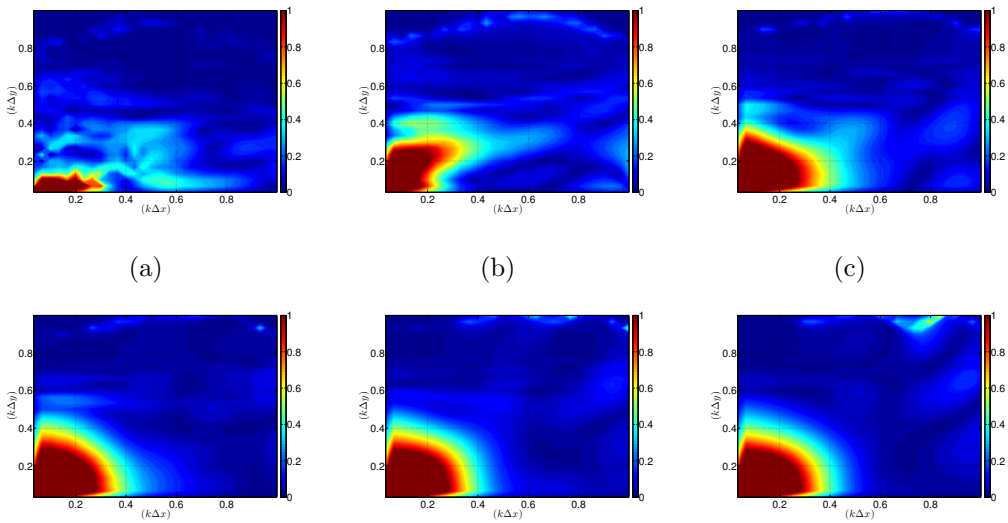

Figure 4.8: Power special density (normalized by energy of the constant mode) of isentropic vortex density signal after having traveled ten vortex widths with respect to different spatial discretization schemes: a) CD02, b) CD04, c) CD06, d) CD04-7pt, e) CD10, and f) CD04-11pt.

as marked in Table 4.3. In this way, one may seek to remove problematic content while preserving properly-calculated portions of the spectrum.

Figures 4.10, 4.12, and 4.14 show density contours of the respective density solutions after having traveled ten, thirty, and fifty vortex-widths, where dissipation has been tuned to the target 4 PPW resolution for the CD04-7pt, CD06, and CD04 discretization schemes, respectively. Clearly seen, is the fact that the optimized scheme is most capable of preserving the vortex. Additional insight is gained by analyzing the corresponding two-dimensional spectral contours in Figures 4.11, 4.13, and 4.15. Here it becomes apparent that the tendency of the vortex to stay coherent corresponds to the preservation of the spectral distribution, while maintenance of the vortex strength naturally has to do with preservation of the modal magnitudes. In this situation, the poorly-resolving schemes show a non-linear modulation of the modes that tends to push out the spectral content towards

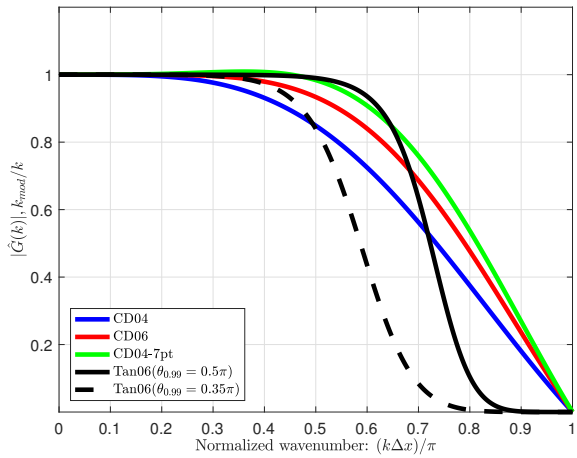

Figure 4.9: Normalized modified wavenumbers ($k_{\text{mod}}/k_{\text{exact}}$) for CD04, CD06, and CD04-7pts spatial schemes, plotted on top of growth factor $|\hat{\mathcal{G}}|$ for sixth-order Tangent scheme tuned to $(k_{0.99}\Delta x) = 0.35\pi$ and 0.5π.

higher modes, where it is met by the dissipation and promptly removed. The combined effect of these two occurrences results in the vortex breaking up into weaker sub-structures. Figure 4.16 and Figure 4.17 show similar information, this time with the dissipation tuned to a target ~ 6 PPW resolution and used with the optimized CD04-7pt scheme. Here the vortex strength is more quickly depleted, although coherence is decently maintained. Very similar results (not shown) were observed for the sixth-order CD06 scheme. The fact that different methods are seen to perform comparatively is understood to result from the dissipation being a dominant source of error. By employing alternate discretizations, such as kinetic energy preserving (Morinishi et al., 1998) or vorticity preserving methods (Lerat et al., 2007), one may be able to further reduce irregular modulation of the spectrum towards high-wavenumbers and decrease the possibility of a subsequent loss of vortex coherence that results from the sudden removal of content by the dissipation scheme.

In all instances, the base schemes are rendered stable by the addition of the

artificial dissipation terms; however, one is able to achieve superior accuracy by properly resolving the spectrum in question. Simultaneously satisfying the need for robustness and accuracy then requires proper selection of the base scheme and the accompanying dissipation, decisions of which should be made relative to the target spectral content one seeks to represent. Though the current example has been designed to highlight these notions[2], such lessons are extendable to the context of LES where one seeks to accurately and robustly represent a subset of the complete turbulent spectrum. In this setting, un-resolved effects are introduced via modeling; however, it has been shown that the impact and efficacy of these closures is related to the presence of numerical error (Sagaut, 2006). In this way, the basic concepts explored thus far – which relate to the design of filters and the spectral enforcement of solution content in light of numerical error considerations – is directly tied to LES applications. The consequence of incorporating these lessons in explicitly-filtered LES implementations is explored in the subsequent chapters.

[2]Stronger vortex perturbations may be analyzed by increasing the α and ϕ parameters. In such instances, a similar exercise may be carried out, assuming there is sufficient resolution to represent a majority of the spectral content. Anything less would be tantamount to a LES type problem, which introduces new challenges.

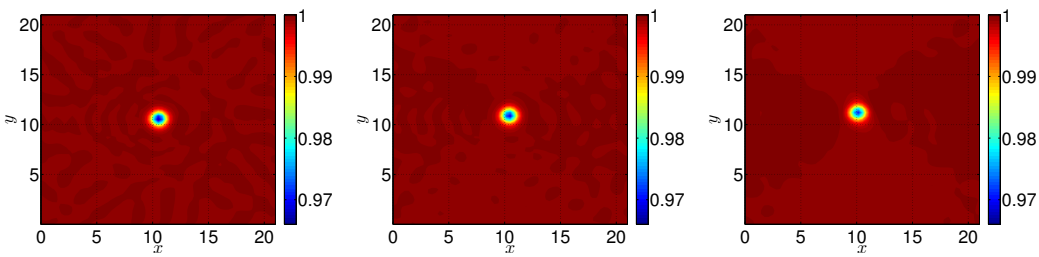

(a) (b) (c)

Figure 4.10: Effect of filter-based artificial dissipation (sixth-order Tangent stencil with $(k_{0.99}\Delta x) = 0.5\pi$,) on evolution of the density solution solved with CD04-7pt spatial scheme: a) ten vortex widths, b) thirty vortex widths, and c) fifty vortex widths.

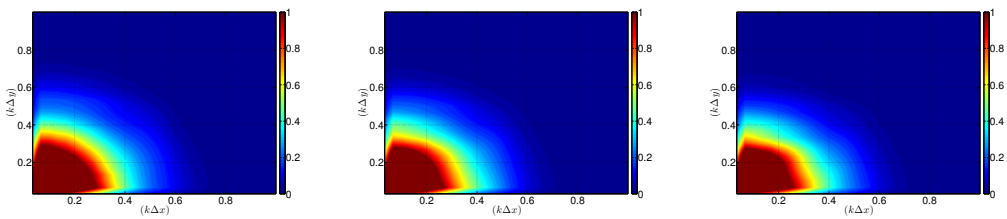

(a) (b) (c)

Figure 4.11: Effect of filter-based artificial dissipation (sixth-order Tangent stencil with $(k_{0.99}\Delta x) = 0.5\pi$,) on evolution of the density solution power spectral density ($|\hat{f}(k_1, k_2, t)|^2/|\hat{f}(0, 0, t = 0)|^2$) solved with CD04-7pt spatial scheme: a) ten vortex widths, b) thirty vortex widths, and c) fifty vortex widths.

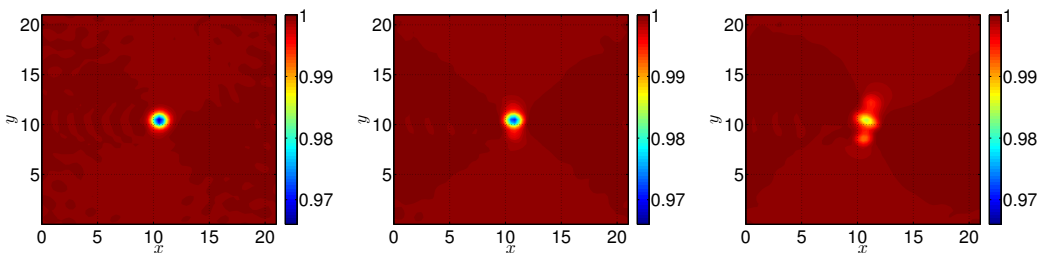

(a) (b) (c)

Figure 4.12: Effect of filter-based artificial dissipation (sixth-order Tangent stencil with $(k_{0.99}\Delta x) = 0.5\pi$,) on evolution of the density solution solved with CD06 spatial scheme: a) ten vortex widths, b) thirty vortex widths, and c) fifty vortex widths.

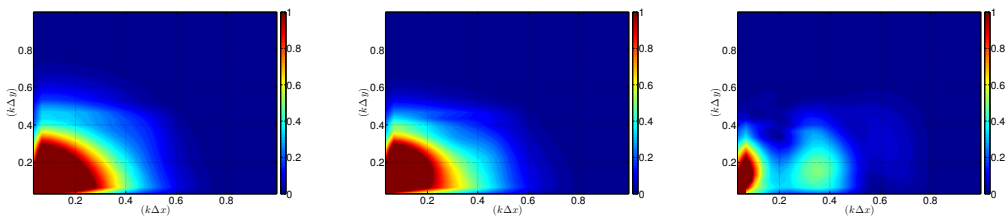

(a) (b) (c)

Figure 4.13: Effect of filter-based artificial dissipation (sixth-order Tangent stencil with $(k_{0.99}\Delta x) = 0.5\pi$,) on evolution of the density solution power spectral density ($|\hat{f}(k_1, k_2, t)|^2/|\hat{f}(0, 0, t = 0)|^2$) solved with CD06 spatial scheme: a) ten vortex widths, b) thirty vortex widths, and c) fifty vortex widths.

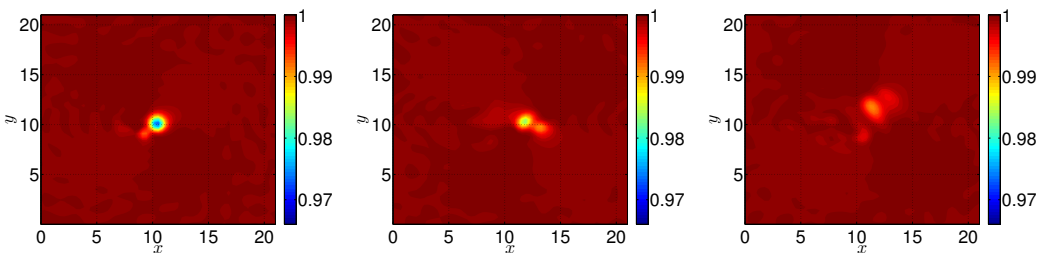

(a) (b) (c)

Figure 4.14: Effect of filter-based artificial dissipation (sixth-order Tangent stencil with $(k_{0.99}\Delta x) = 0.5\pi$,) on evolution of the density solution solved with CD04 spatial scheme: a) ten vortex widths, b) thirty vortex widths, and c) fifty vortex widths.

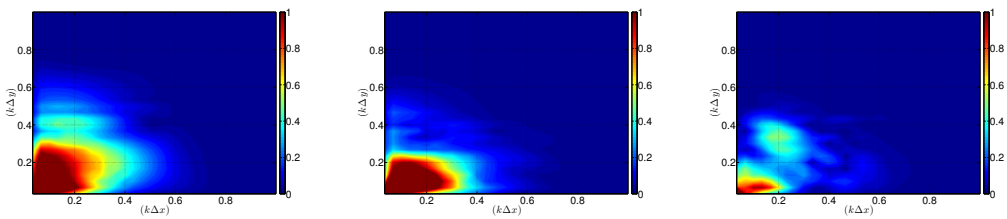

(a) (b) (c)

Figure 4.15: Effect of filter-based artificial dissipation (sixth-order Tangent stencil with $(k_{0.99}\Delta x) = 0.5\pi$,) on evolution of the density solution power spectral density ($|\hat{f}(k_1, k_2, t)|^2/|\hat{f}(0, 0, t = 0)|^2$) solved with CD04 spatial scheme: a) ten vortex widths, b) thirty vortex widths, and c) fifty vortex widths.

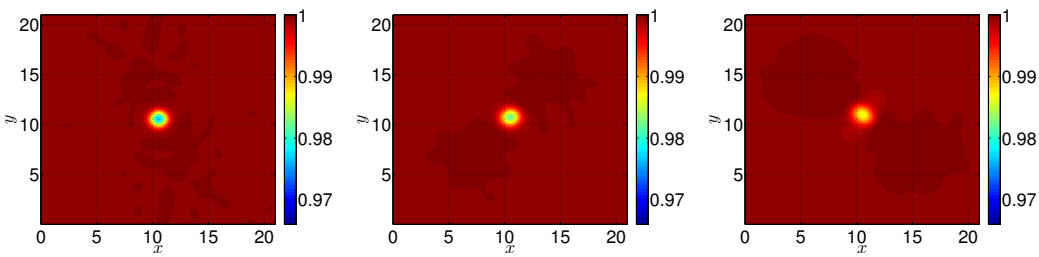

(a) (b) (c)

Figure 4.16: Effect of filter-based artificial dissipation (sixth-order Tangent stencil with $(k_{0.99}\Delta x) = 0.35\pi$,) on evolution of the density solution solved with CD04-7pt spatial scheme: a) ten vortex widths, b) thirty vortex widths, and c) fifty vortex widths.

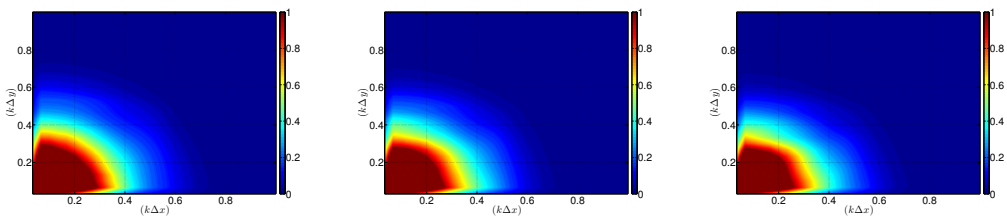

(a) (b) (c)

Figure 4.17: Effect of filter-based artificial dissipation (sixth-order Tangent stencil with $(k_{0.99}\Delta x) = 0.35\pi$,) on evolution of the density solution power spectral density ($|\hat{f}(k_1, k_2, t)|^2/|\hat{f}(0, 0, t = 0)|^2$) solved with CD04-7pt spatial scheme: a) ten vortex widths, b) thirty vortex widths, and c) fifty vortex widths.

CHAPTER 5

1D Synthetic Turbulence with the Viscous Burgers Equation

This chapter is written in part from modifications to a previous AIAA conference proceedings article, "Stabilized Scale-Similarity Modeling for Explicitly-Filtered LES" (Edoh and Karagozian, June 2017a).

5.1 Background

In order to observe performances of the numerical treatments (i.e., discretization and stabilization) in the context of Large-Eddy Simulation (LES), the one-dimensional (1D) Burgers equation is analyzed:

$$\partial_t u + \frac{1}{2}\partial_x uu = \nu \partial_x^2 u. \tag{5.1}$$

Notably simplified in comparison to the full, multi-dimensional Navier-Stokes system, the Burgers equation is often used as a surrogate because of its prototypical non-linear characteristics, which combine hyperbolic (i.e., convection) and parabolic (i.e., diffusion) components. This allows it to replicate a spectral redistribution of energy that is reminiscent of the cascade processes of turbulent fluid flow. Theoretical studies of the Burgers equation identify its ability to develop an inertial range and the termination of energy towards a viscous-dominated Kolmogorov-like small scale. The analogous characteristics shared between this equation and the fluid system has thus motivated use of the Burgers equation as a testbed for LES model assessment and development (Love, 1980; LaBryer et al.,

2015; Li and Wang, 2016).

While such LES investigations typically involve comparing *a posteriori* evaluations with filtered reference data (either derived from fully-resolved DNS or experiments), it can become difficult to identify the different sources of errors that impact a the LES solution. These erroneous contributions may stem from modeling or discretization choices, for instance. In order to address these difficulties, one may consider LES computations that employ consistent closures derived from filtered reference data. This type of approach – here termed quasi *a priori* – is introduced by DeStefano and Vasilyev, who supply "perfect" closures to a running LES computation in order to highlight the effects of filter sharpness on model robustness (Stefano and Vasilyev, 2002). The tractability of the 1D Burgers equation accommodates this quasi *a priori* approach. In this chapter, the analysis technique is employed in order to aid in understanding the roles that scheme discretization, filtering formulation, and filter cut-off have on the LES solution. Similar studies looking at the impact of such factors on the accuracy of the LES solution have been performed by others; however, these typically: 1) presume a specific closure model within an *a posteriori* (i.e., real-time calculations of the model) setting, which can introduce errors arising from the coupling of modeling and numerical error; or 2) are based on *a priori* analysis (i.e., study of the models based on the filtering of full-field, "exact" data) that overlook dynamical effects of the evolution of the LES field. As a result, employing the LES computation in the quasi *a priori* sense can supplement these more traditional approaches and give targeted insight into the impact of different algorithmic choices.

The following chapter thus seeks to employ concepts explored in Chapters 2-4 and to demonstrate the impact of discretization and stabilization within the context of LES. Additional care is then taken with respect to the presumed LES formulation and its impacts on scale-resolution and scale-separation, as they pertain to interactions with the numerics and overall solution accuracy. The quasi *a*

priori approach for analysis serves as an intermediate mode of interrogation that is meant to provide useful insight towards practical (i.e., *a posteriori*) LES computations. In this vein, the following investigation also looks to different classes of LES closures: the Approximate Deconvolution method (ADM) of Stolz and Adams (1999), along with a Bardina-type Scale-Similarity model (SSM). These, along with "perfect" models, are assessed in light of the base scheme, the stabilization approach, and the accompanying LES filter formulation.

5.1.1 Numerical Set-up

The subsequent LES investigations seek to calculate a turbulence-like solution of the 1D Burgers problem. A uniform and periodic domain is considered on $x \in [0, L = 2\pi]$ and the field is initialized by the following energy spectrum (San et al., 2016),

$$E(k) = Ak^4 exp[-(k/k_o)^2], \tag{5.2}$$

such that the signal modes have magnitudes,

$$|\hat{u}(k)| = \sqrt{2E(k)}. \tag{5.3}$$

The constant $A = \frac{2}{3\sqrt{\pi}} k_o^{-5}$ is chosen in order to prescribe an initial mean energy $E_{mean} = \int E(k)dk = 1/2$. The peak value of this initial spectrum is set to $k_o = 5$, and the signal is built as a Fourier series:

$$u(x) = 2 \sum_{n=0}^{Nx/2} \left[\mathrm{Re}\{\hat{u}(k_n)\} \cos(k_n x) + \mathrm{Im}\{\hat{u}(k_n)\} \sin(k_n x) \right], \tag{5.4}$$

$$\text{with} \quad \hat{u}(k) = \sqrt{2E(k)} \cdot exp\{\imath 2\pi \mathsf{U}_k\} \ . \tag{5.5}$$

Random phase angles are assigned by a uniform distribution $\mathsf{U}_k \in [0, 1]$, according to a conjugate wavenumber relationship $\mathsf{U}_k = -\mathsf{U}_{-k}$.

Due to the non-linearity of the Burgers equation, the initial solution generates higher wavenumbers and the energy cascades down to smaller scales until

terminating at a viscous-dominated scale, k_η. The extent of the scale separation between the forcing scale k_o and the viscous-dominated regime depends on the magnitudes of the initial perturbation relative to the physical viscosity. An estimate for this is provided by Kolmogorov-type approximations based on the integral Reynolds number, where $Re_{u_o} = \frac{u_o \ell_o}{\nu} \sim (k_o/k_\eta)^{-4/3}$, with $u_o \approx 2\sqrt{E_{mean}}$.

Determination of the grid required to constitute a fully-resolved computation is then controllable by assigning the viscosity, here chosen to be $\nu = 10^{-3}$. As a result, a DNS resolution of $N_x = 8192$, which corresponds to $\eta \sim 8\Delta x_{DNS}$, is chosen. Computations of the DNS solution is performed using fourth-order standard central finite difference schemes for both convection (in divergence form) and diffusion. Advancement in time is employed using the optimized fourth-order, six-stage low dissipation-dispersion Runge-Kutta (LDDRK) method of Hu *et al.* (Hu et al., 1996), with a time-step $\Delta t = \Delta x_{DNS}/10$.

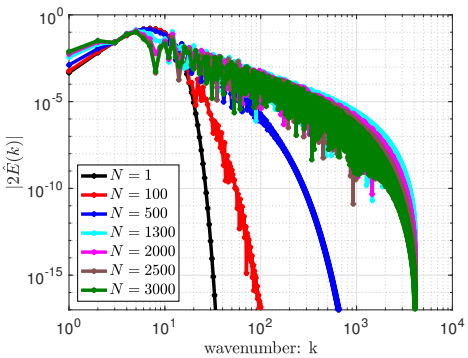

Figure 5.1: Kinetic energy spectrum, $2|\hat{E}(k)| = |\hat{u}(k)|^2$, for the DNS solution at various time-steps ($\Delta t = 7.66 \times 10^{-5}$). The Burgers inertial range scaling of k^{-2} is shown as a dashed line.

Figure 5.1 shows temporal evolution snapshots of the ensuing solution's kinetic energy spectrum, $2E(k) = |\hat{u}(k)|^2$: the energy cascades towards higher wavenumbers, an internal range with the Burgers k^{-2} scaling is established (Gurbatov et al., 1997), and the solution eventually begins to decay due to the lack of active turbulent forcing. The current investigations will thus focus on the transition and

development of a fully turbulent-like state with $t \in [0, 0.23]$. The large variations featured in the spectrum towards high-wavenumbers – particularly at later times – relate to intermittent behavior; these can be eliminated by considering ensemble averages of different turbulent realizations as done in San et al. (2016).

With respect to the LES computations, the following explicitly-filtered governing equations are considered,

$$\partial_t \bar{\tilde{u}} + \frac{1}{2}\partial_x \bar{\tilde{u}}\bar{\tilde{u}} = \nu \partial_x^2 \bar{\tilde{u}} - \frac{1}{2}\partial_x \tau^{SFS,inv}, \tag{5.6}$$

with the sub-filter-scale closure responsible for enforcing the LES field being decomposed into resolvable sub-filter scale (RSFS) and sub-grid scale (SGS) components:

$$\begin{aligned} \tau^{SFS,inv} &= \overline{\widetilde{uu}} - \bar{\tilde{u}}\bar{\tilde{u}} \\ &= \left(\overline{\bar{\tilde{u}}\bar{\tilde{u}}} - \bar{\tilde{u}}\bar{\tilde{u}}\right) + \left(\overline{\widetilde{uu}} - \overline{\bar{\tilde{u}}\bar{\tilde{u}}}\right) \\ &= \tau^{RSFS,inv} + \overline{\tau^{SGS,inv}} \ . \end{aligned} \tag{5.7}$$

Here, the overall filter operator $\mathcal{G}_{\bar{\tilde{\Delta}}}$ is applied to Equation 5.1 and represents the combined effect of projecting the continuous solution onto a coarser and discrete grid ($\mathcal{G}_{\tilde{\Delta}}$) and then explicitly defining an LES filter ($\mathcal{G}_{\bar{\Delta}}$). In the case of *a posteriori* evaluations, specific models that approximate $\tau^{model} \approx \tau^{SFS}$ are employed and the equation is integrated. In the special case of the quasi *a priori* analysis, the closure is pre-computed from the "exact" DNS calculation and projected onto the LES grid in a consistent manner with the filter formulation at hand. Therefore, relative to the LES equation, this "perfect" model contribution acts as an independent one-way source term that seeks to drive the simulation towards the exact filtered solutions. In this way, the closure model does not react to the current LES solution and one can designate any deviations relative to the reference field as non-model sources of error. This then allows one to characterize the effects of discretization error on accuracy[1]. In order to minimize temporal-inconsistencies,

[1]Note that while the convective schemes are varied in the current analysis, the viscous terms

the LES computations are carried out at the same time-step size as the DNS, and the exactly-built closure is available and supplied at each stage of the numerical integration scheme. Furthermore, the filtered-DNS reference solution is used as the initial condition to the LES computation.

The following steps are taken in order to derive filtered quantities $\bar{\bar{q}}$ from the DNS data (note that the exact projection to coarser grids and its dynamical consequences are not explicitly defined, but are approximated by steps 1 and 2 below):

1. The DNS solution is filtered to a wavenumber $k_{\Delta x_{LES}} = L/N_{x,LES}$ relating to the Nyquist frequency of the LES grid in question; this is done using a spectrally sharp cut-off filter, thus producing an intermediate filtered solution, u^*_{DNS} on the DNS grid.

2. The resulting field u^*_{DNS} is projected onto the LES grid using a physical "comb" filter such that $\tilde{u}(x_k)_{LES} = \tilde{u}^*(x_k)_{DNS}$ (this assumes all LES points to be collocated with the DNS grid).

3. The LES field $\tilde{u}(x_k)_{LES}$ is then filtered according to the explicit LES filter $G^{\bar{\Delta}}$, thus yielding the final representation on the LES grid, $\bar{\tilde{u}}_{ref}$ (also referred to as the filtered-DNS solution, fDNS).

As relating to the ambiguity of the projection procedure outlined above, it is important to understand that the choice of quasi *a priori* implementation may introduce some inconsistencies relating to $[(\bar{\tilde{u}}\bar{\tilde{u}})_{LES} \neq (\bar{\tilde{u}}\bar{\tilde{u}})_{fDNS}]$. While these terms are expected to cancel out in actual *a posteriori* computations, it is not necessarily the case with the projected reference data (see Equation 5.8). On one hand, in the case of a "perfect" model, one can employ the following implementation for

---

are kept to be fourth-order accurate.

the quasi *a priori* investigation:

$$\partial_t \bar{\tilde{u}} + \frac{1}{2}\delta_x \bar{\tilde{u}}\bar{\tilde{u}} = \nu\delta_x^2 \bar{\tilde{u}} - \frac{1}{2}\delta_x \left[\overline{\widetilde{uu}} - \bar{\tilde{u}}\bar{\tilde{u}}\right]_{fDNS}. \tag{5.8}$$

This, however, can lead to erroneous contributions in the DNS-assisted analysis relating to the aforementioned inconsistency. Alternatively, one could avoid these incongruities and choose to compute the following:

$$\partial_t \bar{\tilde{u}} + \frac{1}{2}\delta_x \left[\overline{\widetilde{uu}}\right]_{fDNS} = \nu\delta_x^2 \bar{\tilde{u}}. \tag{5.9}$$

To demonstrate the difference in these implementations, Figure 5.2 plots the energy spectrum for the filtered-DNS result along with solutions from the two proposed implementation options used in the reference-model-assisted setting. It is apparent that both methods impressively enforce the LES field quite well, but that the choice of Equation 5.8 introduces additional high-wavenumber error. Nevertheless, the implementation of Equation 5.8 is selected in the case of "perfect" modeling in order to incorporate more realism in the evolved quasi *a priori* solution that may stem from the dynamic feedback of the convective terms calculated at run-time.

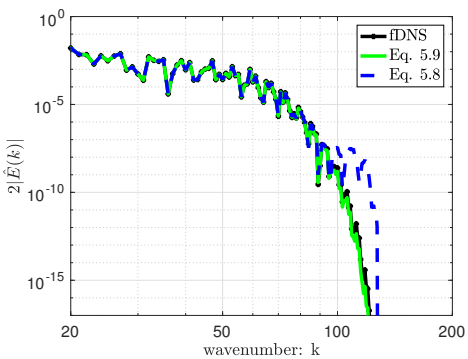

Figure 5.2: Comparison of the filtered-DNS (fDNS) reference (black line) kinetic energy spectrum, $2|\hat{E}(k)| = |\hat{u}(k)|^2$, and the quasi *a priori* evaluations using Equation 5.8 (dashed blue line) versus Equation 5.9 (solid green line) at time-step $N = 1300$ ($t \sim 0.1$).

In order to assess the accuracy of the LES computations, the evolution of the L_2 error of the solution is tracked relative to the appropriately-filtered reference solution. This corresponds to an overall assessment of the computational accuracy, but is supplemented by analysis of the energy spectra along with different metrics of the kinetic energy dissipation rate overaged over the domain Ω:

$$\frac{dE_\Omega}{dt} = -\epsilon_\Omega \ , \tag{5.10}$$

$$\text{where} \qquad E_\Omega = \text{avg}\left\{\frac{1}{2}u\right\}_\Omega^2 = \frac{1}{2}\int_{-k_m}^{k_m} |\hat{u}(k)|^2 dk \ , \tag{5.11}$$

$$\epsilon_\Omega = \text{avg}\left\{\nu(\partial_x u)^2\right\}_\Omega = \nu\int_{-k_m}^{k_m} k^2|\hat{u}(k)|^2 dk \ . \tag{5.12}$$

With respect to these metrics, one can choose to inspect the temporal derivative of energy, dE/dt, or the inner product of strain, $\epsilon(t)$, as defined above. In the case of an unfiltered and fully-resolved computation, these two metrics are identical and communicate the conservation of energy from large to small scales. For example, a DNS calculation of the unfiltered equations shows the two metrics to be almost identical, as shown in Figure 5.3. In the case of LES equations, however, these definitions are expected to differ, providing different perspectives. The strain-based definition, $\epsilon(t)$, estimates the dissipation rate based on small-scale activity (note the k^{-2} scaling of its definition in Equation 5.12); meanwhile, the dE/dt metric puts more focus on the large wavenumber component. As the filtering procedure attenuates high-wavenumber content, one should thus expect dE/dt and $\epsilon(t)$ to differ.

These contrasting behaviors are shown in Figure 5.4 which compares the energy spectra of the DNS solution and filtered-DNS (i.e., reference solutions) on two LES grids of $N_{x,LES} = [256, 512]$, each filtered with a sixth-order Tangent stencil with cutoff $\bar{\Delta}_{1/2} = 4\Delta x_{LES}$. As expected by the scale-discriminant behavior of the discrete filter in question, the low-wavenumber content is well-preserved while scales in the range of $k \in [k_{1/2}, k_{\Delta x_{LES}}]$ are strongly attenuated. Figure 5.3 includes the evolution of the respective energy dissipation metrics for the DNS

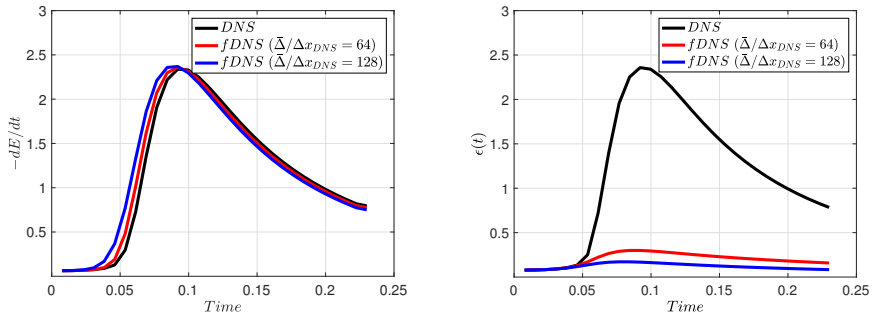

(a) (b)

Figure 5.3: Evolution of alternate kinetic energy dissipation rate metrics for DNS and filtered-DNS (fDNS) reference solution. The explicit LES filter is taken to be a sixth-order Tangent stencil, tuned according to $\bar{\Delta}_{1/2} = 4\Delta x_{LES}$: Energy dissipation rate $-dE/dt$ (see Equation 5.11), strain-rate $\epsilon(t)$ (see Equation 5.12).

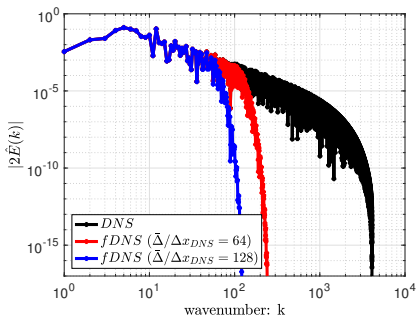

Figure 5.4: Kinetic energy spectrum, $2|\hat{E}(k)| = |\hat{u}(k)|^2$, for the DNS solution (with grid resolution $Nx_{DNS} = 8192$) and filtered-DNS (fDNS) reference solutions (with grid resolution $Nx_{LES} = [256, 512]$) at time-step $N = 1300$ ($t \sim 0.1$). The explicit LES filter is taken to be a sixth-order Tangent stencil, tuned according to $\bar{\Delta}_{1/2} = 4\Delta x_{LES}$.

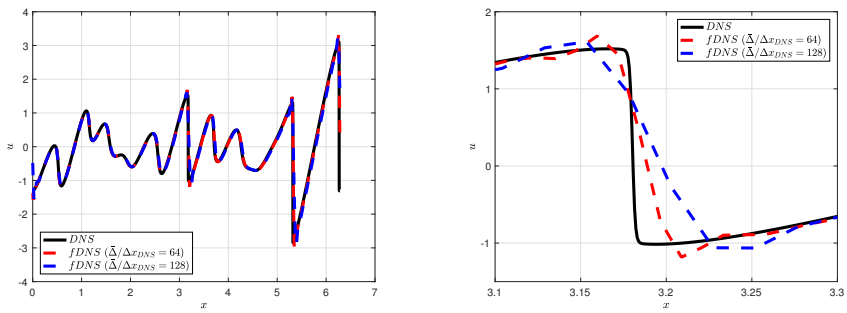

(a) (b)

Figure 5.5: Comparison of the DNS and the filtered-DNS (fDNS) solutions at time-step $N = 1300$ ($t \sim 0.1$): a) full domain, and b) magnified portion of the domain.

and filtered-DNS solutions. It is evident that the dE/dt definition of the filtered fields is very similar to that of the fully-resolved DNS, with only slight differences in the peak dissipation rate. This suggests that energy is removed from the large-scales in similar fashions. On the other hand, the strain-rate $\epsilon(t)$ definition gives starkly different characterizations of the filtered solutions relative to the DNS. As mentioned, this is due to the amount of high-wavenumber content, or lack thereof. Depending on the desired emphasis for scrutiny, one may find more insight from one metric over the other. For instance, in evaluating the presence of high-wavenumber error, the strain rate definition will be more sensitive to small-scale noise and can offer more critical assessments of the respective models at hand.

Also useful to note in Figure 5.3(b) is the time at which the curves of the filtered solutions deviate from that of the DNS ($t \sim 0.05$). This corresponds approximately to the time at which the true spectrum evolves past the LES filter width $\bar{\Delta}$. Going forward, this event will show itself to be consequential as it marks the point at which LES models need to begin extrapolating the dynamics of the small-scales onto the LES field. Therefore, this event is anticipated to coincide

with a rise in solution error, as will be shown in subsequent analyses.

In terms of observing the respective solutions, Figure 5.5 plots the DNS and filtered DNS results at $t \sim 0.1$, which corresponds to the time of peak energy dissipation rate for the fully-resolved flow. Here, one notices the shock-like structures in the DNS solution that are characteristic of the viscous Burgers equation. Filtering this solution, however, removes high-wavenumber content that would otherwise be responsible for representing these sharp features. Therefore, the filtered solutions are seen to contain oscillations near the shocklets. While one may typically interpret the LES solution as a projection of the continuous result onto a coarser grid (e.g., comb filtering the exact solution in physical space), the derived LES solution in Figure 5.5(b) suggest more nuanced consequences of filtering flows that contain sub-grid discontinuities[2].

5.2 Assessing Numerical Error Effects in LES with Exact Modeling

As previously indicated, the Burgers equation will be utilized in order to study the impact of numerical discretization and scheme stabilization on solution accuracy, while also looking into the effects of the LES filter formulation as dictated by the filter-to-grid ratio $FTGR = \bar{\Delta}/(2\Delta x_{LES})$ and the spectral characteristics of the presumed LES filter. Here, the consequences of these algorithmic choices is considered free of modeling errors by consistently deriving the "exact" LES model closure to be employed in the quasi *a priori* procedure according to Equation 5.6-5.7. While the conclusions drawn from using a "perfect" sub-filter scale term $\tau^{SFS} = \left[\widetilde{\overline{uu}} - \tilde{\bar{u}}\tilde{\bar{u}}\right]_{\widetilde{\overline{DNS}}}$ may become moot when one is required to employ practical

[2]The consequences of filtering flows containing discontinuities is theoretically investigated in Sagaut and Germano (2005), who note the appearance of parasitic contributions and the disappearance of jump conditions. The authors seek to resolve these issues by developing pseudo-jump relations for sub-grid models.

models in an *a posteriori* setting, this idealized mode of inspection provides a best-scenario characterization of numerical and LES formulation effects on solution accuracy.

5.2.1 Impacts of Discretization and Filter Width

The energy spectra of the dynamically-computed quasi *a priori* results are shown in Figure 5.6 for the case of a spectral sharp LES formulation. Figure 5.6(a) constitutes implicitly-filtering (i.e., filter-to-grid ratio of unity or $FTGR = 1$) with no modeling and shows an increased energy pile up towards the smallest resolved scales. This behavior is seen to be more pronouced for the high-order schemes which do less to attenuate high wavenumbers, a consequence of their accurate modified wavenumber. Figure 5.6(b) then incorporates the "exact" closure contribution, again for $FTGR = 1$, and shows a slight reduction in energy pile-up. There is still, however, noticeable error in the mid- to high-wavenumber range and little distinction between the discretization schemes (CD02, CD06, CD10, CD04-11pt). Subsequently increasing the filter-to-grid ratio is then seen to further reduce these numerical errors (see Figures 5.6(c) and 5.6(d)). Here, we note that the "perfect" closure is able to enforce a spectral-sharp attenuation of the sub-filter scales. Any deviation from the reference filtered DNS solution is then clearly attributable to the choice in discretization scheme. For example, the low-order CD02 method shows persistent error for $FTGR = 2$ and requires higher filter-to-grid ratios. Meanwhile, the high-accuracy discretizations are more suitable at modest FTGR and are thus preferable in terms of resolution efficiency (e.g., the required degrees of freedom use to calculate a desired LES resolution $\bar{\Delta}$). This confirms the fundamental analysis of Ghosal (1996) that advocates employing LES formulations with $FTGR > 1$. Deliberately taking into consideration the resolvability performance of the base numerical scheme can thus be advantageous in producing more accurate results. Furthermore, an important observation

is made through these exercises: a model that is consistent with the explicit LES filter will not only seek to represent proper comportment of the resolved scales, but will also work to enforce the presumed LES spectral resolution.This notion shall be employed later on in considering the incorporation of stabilization via filter-based artificial dissipation.

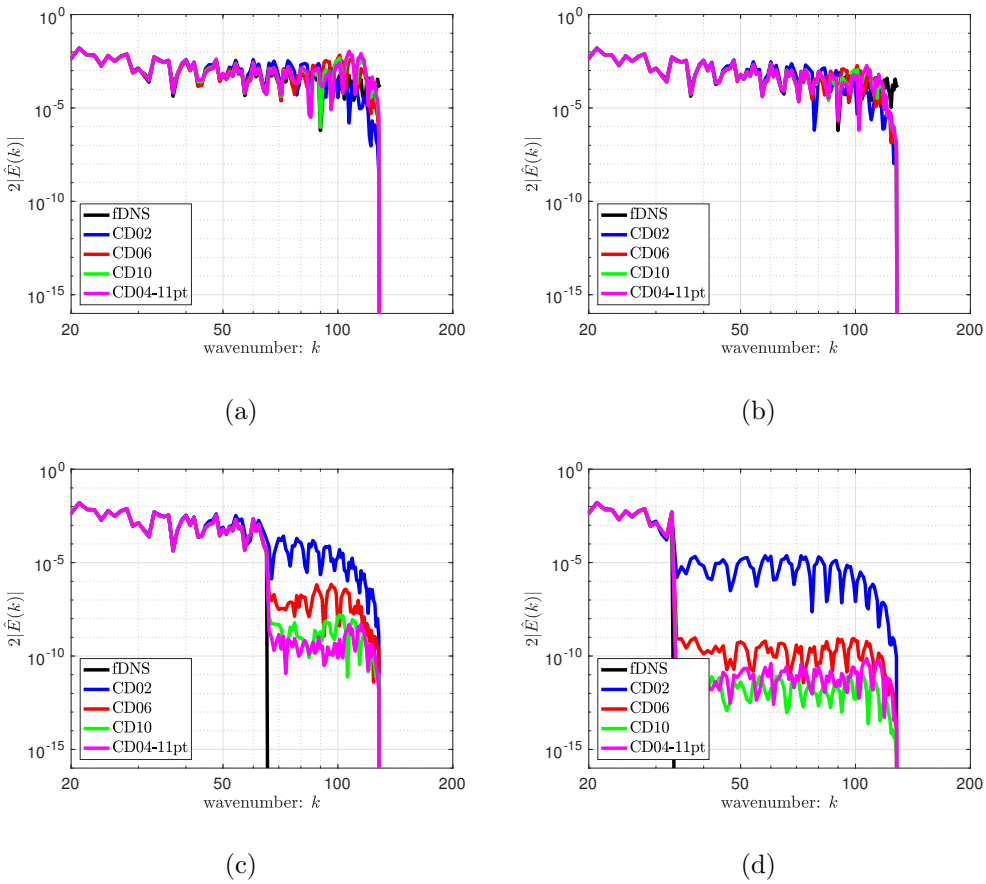

Figure 5.6: Effect of discretization scheme and filter-to-grid ratio $FTGR$ on the kinetic energy spectrum, $2|\hat{E}(k)| = |\hat{u}(k)|^2$, at time-step $N = 1300$ ($t \sim 0.1$). LES solutions assume a spectral-sharp filter formulation: $FTGR = 1$ (no model), $FTGR = 1$ (perfect model), $FTGR = 2$ (perfect model), $FTGR = 4$ (perfect model).

Figure 5.7 again demonstrates these concepts by plotting the quasi *a priori* solution error, normalized relative to the reference fields. As before, the influence

of the discretization scheme on accuracy is shown to be dependent on the filter-to-grid ratio. At $FTGR = 1$, exclusion of a model makes it such that scheme choice has little impact; incorporating the "perfect" model then improves this only marginally. Upon enforcing $FTGR > 1$, however, the model contribution becomes more prominent at the small-scales and one sees pronounced benefits to employing high-accuracy discretizations, with errors decreasing by orders of magnitude. Also of interesting note is the way in which the optimized fourth-order scheme (CD04-11pt) is seen to outperform the standard tenth-order discretization (CD10) for $FTGR = 2$. This trend is tied to the spectral fidelities of the respective methods at different resolutions (see Figure 4.1). At $FTGR = 4$, the standard scheme is more accurate than the optimized stencil, as expected by the resurgent importance of asymptotic order on scheme performance for high resolutions; in the context of LES, these trends are thus only recovered at sufficiently high filter-to-grid ratios.

As another point of inspection, it is interesting to consider the implication of the spectral characteristics of the explicit LES filter on sensitivities to discretization accuracy. In order to highlight these effects, one can consider alternate explicit LES formulations based on second, sixth, and tenth-order Tangent filter formulations. Figure 5.8 plots the respective spectral characteristics of these filters (each tuned according to $k_{1/2}$) in comparison with the spectral-sharp filter. Noting the scale-discriminant behavior of each response, it becomes evident that smoother formulations employ less scale-separation and would require the closure model to provide information regarding the large-scales. This, for example, is the case of the low-order filter stencil. As previously noted, this additional requirement may hinder the generality of closure models and increase the challenge of model development (Stefano and Vasilyev, 2002).

Figure 5.9 plots quasi *a priori* normalized solution error at $FTGR = 2$ and 4 for these Tangent filter LES formulations, measured relative to the appropriate reference field; the closure model is again built to be consistent with the presumed

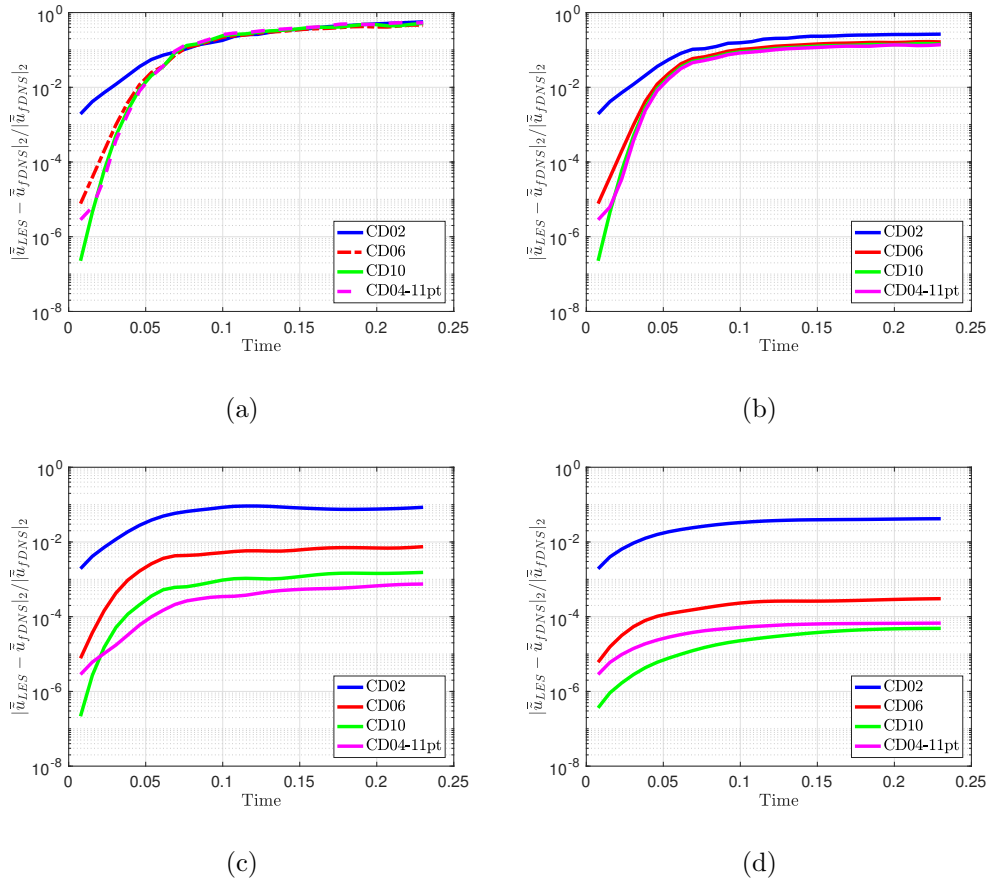

Figure 5.7: Effect of discretization scheme and filter-to-grid ratio $FTGR$ on the time evolution of the relative L_2 LES solution error computed by the quasi *a priori* evaluation, $|\bar{\tilde{u}}_{LES} - \bar{\tilde{u}}_{fDNS}|_2/|\bar{\tilde{u}}_{fDNS}|_2$. LES solutions assume a spectral-sharp filter formulation: $FTGR = 1$ (no model), $FTGR = 1$ (perfect model), $FTGR = 2$ (perfect model), $FTGR = 4$ (perfect model).

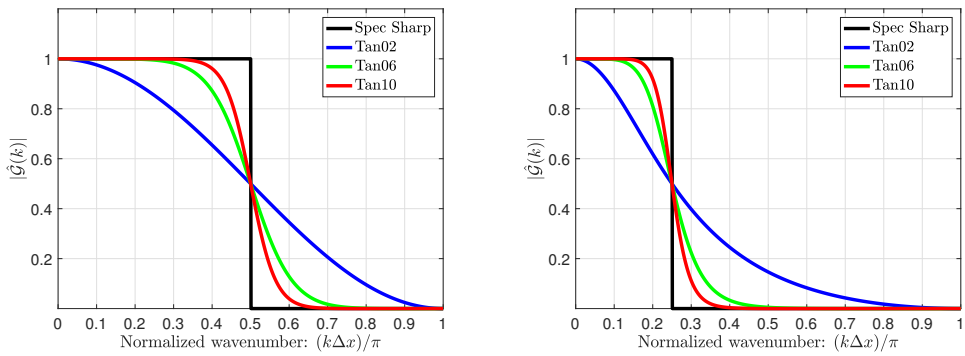

(a) (b)

Figure 5.8: Comparison of spectral-sharp and discrete Tangent growth factor $|\mathcal{G}(k)|$ for different filter-to-grid ratios (FTGR) as tuned according to $k_{1/2}$, as: $FTGR = 2$ ($\bar{\Delta}_{1/2} = 4\Delta x$), and $FTGR = 4$ ($\bar{\Delta}_{1/2} = 8\Delta x$).

explicit LES filter. For both filter-to-grid ratios considered, one witnesses a reduction of the solution error as the LES filter formulation is made more spectrally sharp. This, in particular, holds for the high-accuracy methods and enforces the intuitive notion that such schemes are necessary for supporting LES calculations that presume a spectral-like sharp cut-off. On the other hand, the second-order discretization scheme CD02 demonstrates less sensitivity to the underlying LES filter. In addition to these is the previous observation that increasing the filter-to-grid ratio reduces error in the solution.

While the current assessments evaluate the solution relative to the consistently-filtered DNS data, one may also be interested in replicating the spectrally-sharp-filtered reference solution, which represents the "idealized" case of scale-separation. In doing so, one would naturally conclude that the respective accuracies would directly correspond to how well the discrete filter matches the spectral-sharp filter. With respect to the Tangent stencils shown in Figure 5.8, it then becomes apparent that the scale-discriminant high-order schemes would yield more favorable results. Alternatively, one could seek to employ different cut-off designations

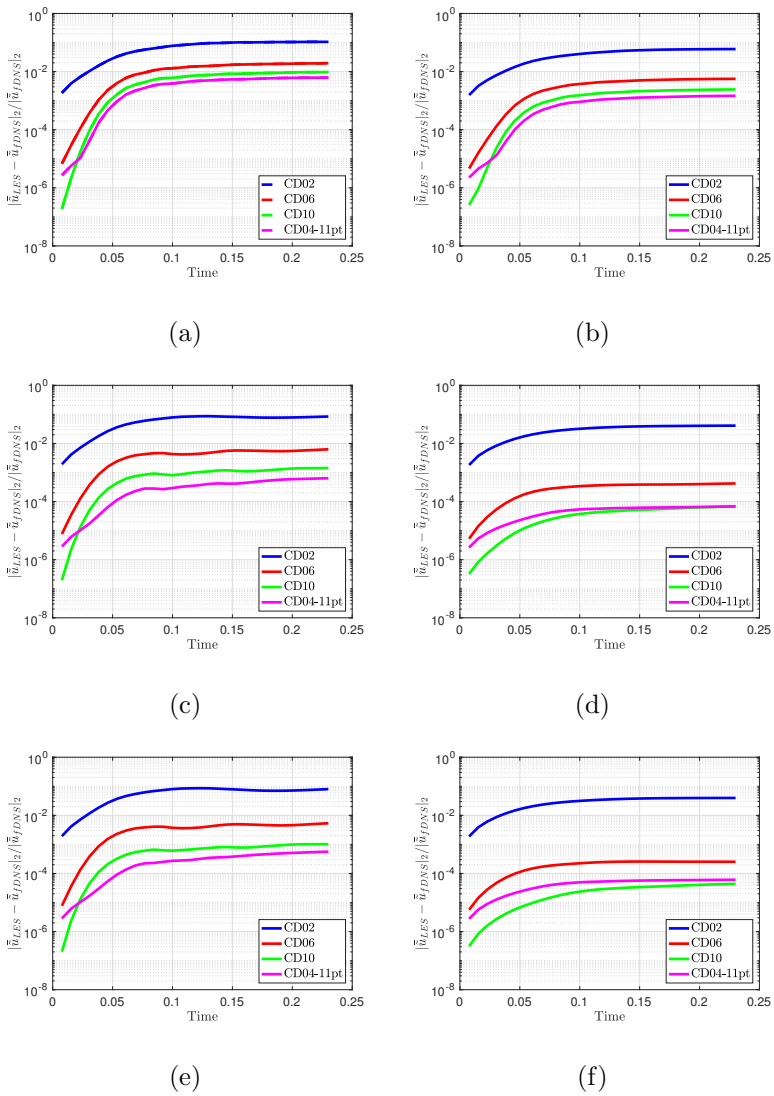

Figure 5.9: Effect of discretization scheme and filter-to-grid ratio (FTGR) on time evolution of the relative L_2 LES solution error computed by the quasi *a priori* evaluation, $|\bar{\tilde{u}}_{LES} - \bar{\tilde{u}}_{fDNS}|_2/|\bar{\tilde{u}}_{fDNS}|_2$. Assumes a discrete Tangent filter formulation, tuned according to $k_{1/2}$: a) second-order filter, $FTGR = 2$, b) second-order filter, $FTGR = 4$, c) sixth-order filter, $FTGR = 2$, d) sixth-order filter, $FTGR = 4$, e) tenth-order filter, $FTGR = 2$, f) tenth-order filter, $FTGR = 4$.

(e.g., $k_{0.99}$ versus $k_{1/2}$) for a given filter and note which minimizes the ensuing error. This, for instance, may constitute a basis by which the effective discrete filter cut-off can be characterized.

5.2.2 Impact of Stabilization

As in the calculation of an isentropic vortex (see Section 4.4), one can consider adding dissipation as a way of removing poorly-resolved components of the discretization. This also serves as a stabilizing agent in instances where such errors threaten the long-time robustness of the simulation. The deliberate addition of numerical dissipation is often avoided in the context of LES, as turbulent flow dynamics may be highly sensitive to artificial damping. The idea of including numerical dissipation, therefore, tends to be more accepted within a rationalization that the stabilization serves as a closure model. This is the idea behind implicit LES, including solution filtering (also referred to as relaxation filtering). But even in the realm of explicit LES, eddy-viscosity models are often incorporated with other closures under the pretense of a mixed-modeling approach. Nevertheless these eddy-viscosity terms are functional in representing the cascade and loss of energy to the small-scales; Yet, they effectively behave as a stabilizing agent of the primary model (e.g., a scale-similarity term). For now, such abstractions are avoided. Instead, this section seeks to demonstrate how the LES solution can be improved by judicious selection of an artificial dissipation contribution, even in the case of the "perfect" modeling employed as part of the quasi *a priori* analysis. Here, the tenth-order Tangent filter, being a close approximation to the spectral-sharp filter, is considered for the explicit LES formulation.

As with other algorithmic components, the choice of damping should take into account the underlying LES filter formulation – which in turn can be chosen to reflect the resolvability performance of the base scheme. In this way, the dissipation should work to provide appropriate assistance to the primary model for

enforcement of the LES resolution and will not be tasked with representing more intricate closure dynamics (e.g., back scatter). Here, stabilization is considered in the form of filter-based artificial dissipation (see Section 3.3). Focusing on a filter-to-grid ratio of $FTGR = 2$ with the optimized CD04-11pt discretization, the implicit Tangent stencil (see Section 2.4.2) is employed as the stabilization stencil.

Figure 5.10 plots the solution errors for different stabilization stencils (second-, sixth-, and tenth-order) tuned to a range of cut-offs based on a $k_{1/2}$ designation. Evidently, both the order of the dissipation scheme and its tuning affect the solution accuracy. In sum, the spectral characteristics of the dissipation are paramount relative to the LES filter formulation and the underlying base scheme. As $k_{1/2}$ is shifted towards the grid cut-off, the amount of damping applied to the LES resolution is reduced. Figure 5.11 shows this by plotting the growth factors of the dissipation schemes in question.

In the case of the second-order Tangent dissipation in Figure 5.10(a), the dissipative stabilization is seen to systematically add error to the baseline solution calculated without stabilization. This is in contrast to the high-order dissipation schemes shown in Figures 5.10(b) and 5.10(c) for the sixth- and tenth-order Tangent stencils, respectively. In these latter instances, there exists an intermediate tuning of the dissipation that yields a more accurate result than the non-stabilized scenario. For the sixth-order dissipation, this occurs for $k_{1/2}\Delta x = 0.8\pi$; meanwhile, for the tenth-order dissipation, this occurs for $k_{1/2}\Delta x = 0.7\pi$. Upon careful consideration, it can be shown that these tuning parameters both correspond to $k_{0.99}\Delta x \approx 0.6\pi$. Recalling that the modified wavenumber analysis of the optimized CD04-11pt scheme yielded a 1% error at $k\Delta x = 0.58$ (see Table 4.3), these observations gain some rationale. As the current model is "perfect", improvement of solution accuracy can be assumed to come with respect to mitigating the lingering influences of numerical error. A first step to address this is the selection

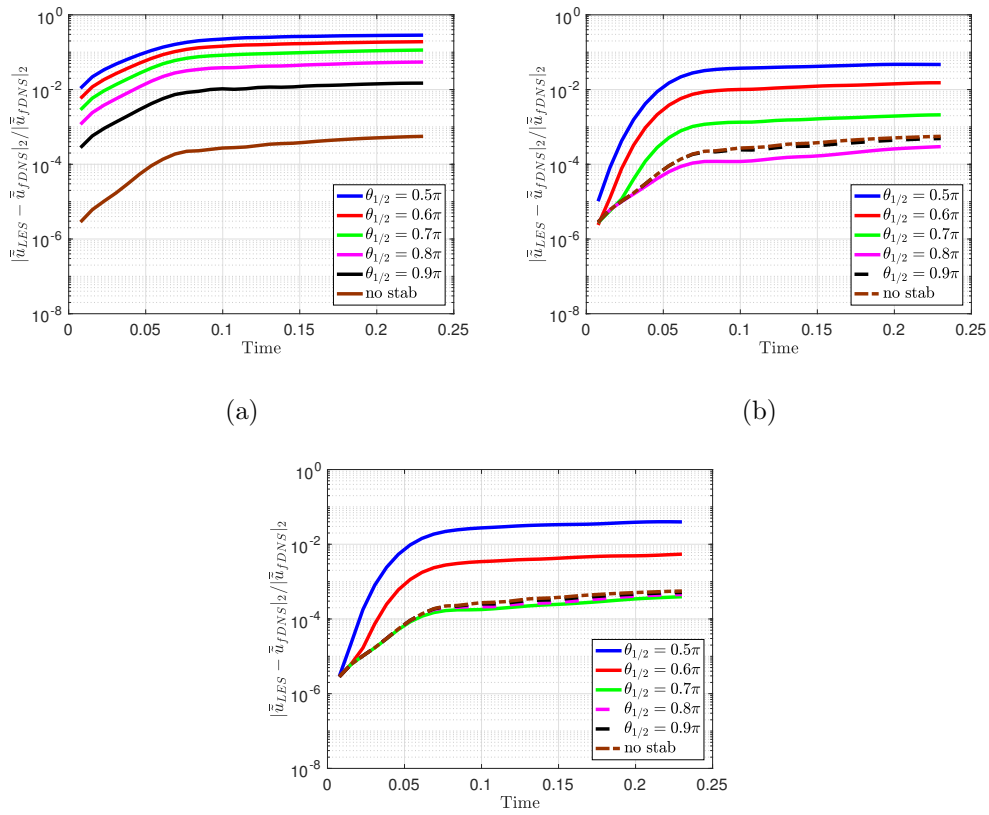

(a) (b)

(c)

Figure 5.10: Effect of filter-based artificial dissipation on the time evolution of the relative L_2 LES solution error computed by the quasi *a priori* evaluation, $|\bar{\tilde{u}}_{LES} - \bar{\tilde{u}}_{fDNS}|_2/|\bar{\tilde{u}}_{fDNS}|_2$. LES solutions assume a spectral-sharp filter formulation with a CD04-11pt scheme discretization: a) second-order Tangent stabilization stencil, b) sixth-order Tangent stabilization stencil, and c) tenth-order Tangent stabilization stencil.

of a proper filter-to-grid ratio, which allows the model contribution to dominate the discretization error. However, remnants of these errors are still present as evidenced by how the choice of scheme plays a role in Figure 5.9. Therefore, one can seek to directly target this source of error through the dissipation scheme.

Figure 5.12 shows the solution error when tuning the dissipation scheme to the discretization scheme according to the $k_{0.99}$ specification. Here, the sixth-order

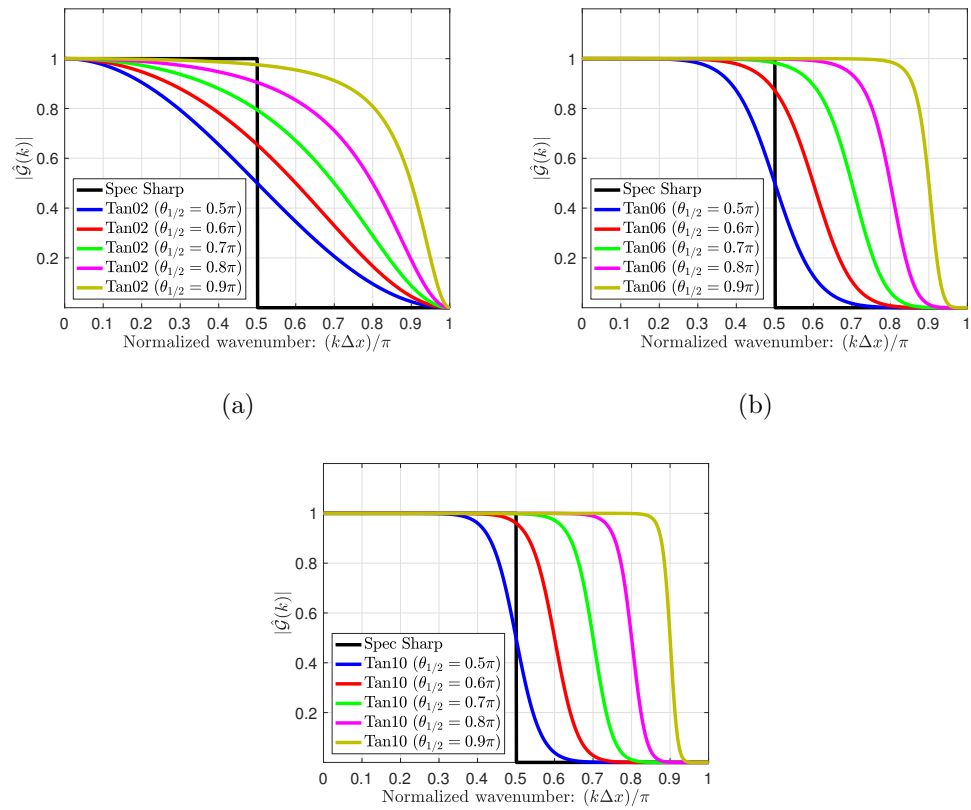

Figure 5.11: Comparison of spectral-sharp and discrete Tangent growth factor $|\mathcal{G}(k)|$ for different filter-to-grid ratios (FTGR) as tuned according to $k_{1/2}$, as: second-order stencil, b) sixth-order stencil, and c) tenth-order stencil.

stabilization performs best. The reason for this is then evident when comparing the dissipation growth factor relative to the normalized modified wavenumber of the CD04-11pt scheme, as demonstrated in Figure 5.13, and noticing that they are in closer agreement compared to the second- and tenth-order dissipation stencils. In the case of the second-order Tangent dissipation, too much erroneous content is preserved; meanwhile, the tenth-order Tangent dissipation may be removing too much of the high-wavenumber content, including needed model contributions. In other words, at least for this situation, the removal of error should properly balance the error generated by the discretization scheme. Employing this philosophy, one

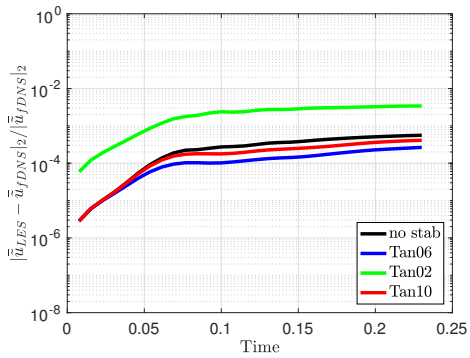

Figure 5.12: Effect of filter-based artificial dissipation (tuned to $k_{0.99}\Delta x = 0.6\pi$) on the time evolution of the relative L_2 LES solution error computed by the quasi *a priori* evaluation, $|\bar{\tilde{u}}_{LES} - \bar{\tilde{u}}_{fDNS}|_2/|\bar{\tilde{u}}_{fDNS}|_2$. LES solutions employ a CD04-11pt discretization and assume a tenth-order Tangent filter formulation tuned to $\bar{\Delta}_{1/2} = 4\Delta x$.

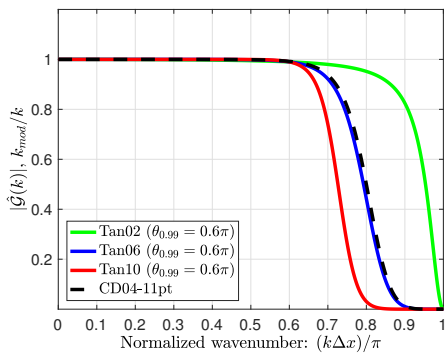

Figure 5.13: Growth factor $|\hat{\mathcal{G}}(k)|$ of Tangent dissipation stencil tuned to $k_{0.99}\Delta x = 0.6\pi$, plotted in comparison to the normalized modified wavenumber (k_{mod}/k) of the optimized CD04-11pt scheme (dashed black).

may then seek a stabilization scheme by consulting available performance maps (e.g., Figures 2.11 and 2.12) and matching the dissipation response characteristics of transition width and cut-off to that of the base scheme (e.g., the normalized wavenumber).

5.3 Assessing Numerical Error Effects in LES with Scale-Similarity Modeling

The quasi *a priori* analysis employed in the previous section takes advantage of a "perfect" DNS-derived model in order to single out effects relating to discretization error, as well as to observe impacts of LES parameters such as the filter-to-grid ratio (FTGR) on solution accuracy. Notable conclusions are drawn from this idealized investigation, such as the need to tune explicit LES formulations (i.e., filter response, including cut-off) relative to the spectral fidelity of the base numerical scheme. Neglecting to do so is shown to neutralize the potential benefits of applying high-accuracy discretization methods. In order to transition these insights to a more practical setting, the following analysis considers specific closure models of the scale-similarity type, which are known to exhibit high correlations when compared to the exact closure (Liu et al., 1995; Li and Wang, 2016; Carati et al., 2001). Here, the Approximate Deconvolution Method (ADM) of Stolz et al. (2001) in addition to the Scale-Similarity Model (SSM) of Bardina (Bardina et al., 1983) are chosen as candidate models.

To understand the rationale of the Approximate Deconvolution Method, one can first consider the sub-filter-scale component (SFS) and its subsequent decomposition into resolved sub-filter-scale (RSFS) and sub-grid-scale (SGS) components, as shown in Equation 5.7. This decomposition explicitly differentiates actual sub-gird contributions (i.e., unresolved dynamics) from scales that are resolved but fall below the cut-off of the explicitly-defined LES filter (i.e., the re-

solved sub-filter scales). For the common case of implicitly-filtered LES (i.e., grid-filtered, with $\bar{\Delta} = 2\Delta x$), the τ^{RSFS} component is negligible and exactly zero in the case of a spectral-sharp filtering of the equations. In turn, one recovers the more classical explicit LES notion, wherein the sub-filter and sub-grid scales are synonymous. In the case where filter-to-grid ratios are greater than unity or for sufficiently smooth filter responses, however, the RSFS component can become important. As a result, decomposition of the SFS term can reduce ambiguities in decisions regarding the modeling approach to be taken.

As the LES filter is explicitly defined with respect to $\bar{\Delta}$, the RSFS term $\tau^{RSFS} = [\overline{\tilde{\bar{u}}\tilde{u}} - \tilde{\bar{u}}\tilde{\bar{u}}]$ can be built by first recovering $\tilde{u} = \mathcal{G}_{\bar{\Delta}}^{-1}\{\tilde{\bar{u}}\}$ via a deconvolution procedure[3] (see Appendix A) or by solving a consistent auxiliary equation (Bull and Jameson, 2016). Having the RSFS component well-defined[4] then diminishes the opportunity for erroneous modeling of the total SFS term. For instances where $\bar{\Delta} > \tilde{\Delta}$ (where $\tilde{\Delta} = 2\Delta x_{LES}$), the resolved SFS component can be used as a surrogate to the entire closure. In this case, one then considers $\tau^{SFS} \approx \tau^{RSFS}$, with the assumption presumably becoming increasingly valid for larger filter-to-grid ratios. Substituting this approximation into the governing Burgers LES equation then reproduces the Approximate Deconvolution method:

$$\partial_t \tilde{\bar{u}} + \frac{1}{2}\partial_x \tilde{\bar{u}}\tilde{\bar{u}} = \nu\partial_x^2\tilde{\bar{u}} - \frac{1}{2}\partial_x \tau^{SFS \approx RSFS} \rightarrow \partial_t \tilde{\bar{u}} + \frac{1}{2}\partial_x \overline{\tilde{\bar{u}}\tilde{\bar{u}}} = \nu\partial_x^2\tilde{\bar{u}} \ . \tag{5.13}$$

The ADM approach ascribes to the scale-similarity philosophy in the sense that the RSFS uses the large-scales (i.e., resolved scales) of the SFS to estimate the complete closure. Therefore, because the ADM approach implicitly relies on the

[3]In the case of a zeroth-order deconvolution, where $\tilde{u} \approx \tilde{\bar{u}}$, then the ADM method is tantamount to residual filtering (Lund, 2003; Radhakrishnan and Bellan, 2012).

[4]As an alternative to deconvolution procedures, the RSFS is sometimes modeled by the gradient model (also referred to as the non-linear model, or the tensor diffusivity model), which is used for second-order LES filter formulations but may be generalized as a truncated Taylor-series approximation to the RSFS term (Pruett and Sochacki, 2001; Sagaut, 2006).

presence of a RSFS component, its use is largely limited to LES formulations that employ either smooth filtering or filter-to-grid ratios above unity.

As an alternative to the ADM procedure, one can consider the closure term corresponding to the filtered continuous equations and seek to emulate this by using the available LES variable; doing so with respect to the current working variable $\bar{\tilde{u}}$ gives $[\overline{uu} - \bar{u}\bar{u}] \approx [\overline{\bar{\tilde{u}}\bar{\tilde{u}}} - \bar{\tilde{u}}\bar{\tilde{u}}]$ and is here labeled as a Bardina-type Scale-Similarity model (SSM) for the explicitly-filtered LES equations. In essence, the SFS is approximated by the filter scales. Further generalizations of the modeling approach is provided by Liu et al. (1995) who write the proposed scale-similarity closure as $\tau^{SSM} = C_{ss}[\widehat{\bar{\tilde{u}}\bar{\tilde{u}}} - \hat{\bar{\tilde{u}}}\hat{\bar{\tilde{u}}}]$ and gives flexibility in the choice of a test filter $\mathcal{G}_{\hat{\Delta}}$ and scaling coefficient. Pruett *et al.* rationalize the selection of C_{ss} as a means of minimizing the error in the Taylor-series expansion of the true stress and show that this error is minimized for the Bardina model, where $C_{ss} = 1$ and $\mathcal{G}_{\hat{\Delta}=\bar{\Delta}}$. Substituting the Scale-Similarity model into the explicit LES formulation yields the following:

$$\partial_t \bar{\tilde{u}} + \frac{1}{2}\partial_x \bar{\tilde{u}}\bar{\tilde{u}} = \nu\partial_x^2\bar{\tilde{u}} - \frac{1}{2}\partial_x \tau^{SFS\approx SSM} \rightarrow \partial_t \bar{\tilde{u}} + \frac{1}{2}\partial_x\left[\overline{\bar{\tilde{u}}\bar{\tilde{u}}} + \left(\bar{\tilde{u}}\bar{\tilde{u}} - \bar{\bar{\tilde{u}}}\bar{\bar{\tilde{u}}}\right)\right] = \nu\partial_x^2\bar{\tilde{u}}\ . \tag{5.14}$$

In this representation, it also becomes evident that the proposed model includes a velocity-increment contribution $(\bar{\tilde{u}}\bar{\tilde{u}} - \bar{\bar{\tilde{u}}}\bar{\bar{\tilde{u}}})$ that only vanishes in the special case of a spectral filter, wherein $\bar{\tilde{u}} = \bar{\bar{\tilde{u}}}$. The presence of this non-linear velocity increment is expected to instigate high-wavenumber activity beyond the desired LES cut-off, as well as to increase aliasing effects.

Having a functional form of the models along with the exact DNS solution provides the opportunity to carry out *a priori* comparisons of the proposed closures and examine these relative to the exact LES closure contribution for examining such things as correlation and magnitude information. The following section considers these scale-similarity-type closures as candidates for practical

LES computations. As before, the effects of discretization and LES parameters (e.g., filter response and cut-off) are explored with respect to the overall solution error. The following considers explicit LES formulations based on the class of discrete Tangent filters, as opposed to a spectral-sharp filter, in order to further observe practical implementations.

5.3.1 Impacts of Discretization and Filter Width

Having candidate models at hand, it is interesting to compare both *a priori* and *a posteriori* assessments as a way of extracting possible trends that shed insight into limitations of the respective methods. This is done here with respect to the LES parameters such as filter-to-grid ratio, as well as with respect to numerical factors such as discretization.

5.3.1.1 The Approximate Deconvolution Method (ADM)

An *a priori* evaluation of the ADM model is performed by filtering the DNS data and manipulating it according to the τ^{RSFS} closure. Figures 5.14(a), 5.15(a), and 5.16(a) plot the correlation coefficient of the RSFS component relative to the total exact SFS and also analyze these qualities relative to the LES filter, where second-, sixth-, and tenth-order Tangent stencils are considered. As a general trend, we note that the model correlation improves as the filter-to-grid ratio increases. This is the case for all the LES filters considered. This makes sense, as the RSFS is the resolved sub-component of the SFS closure and its dynamical relevance would increase as the presence of these scales is increased. Even at a filter-to-grid ratio $FTGR = 2$, one observes correlation coefficients above 0.95 for the tenth-order Tangent formulation. Next, inspecting the effects of the LES filter, there is the interesting observation that low-order formulations induce higher correlations. Reverting back to the spectral response plot of Figure 5.8, the reason perhaps

becomes more clear: smooth filters increase the relevance of the RSFS towards the larger scales where greater energy-content resides. Therefore, in this way, the RSFS again becomes more accurate.

Somewhat similar trends are observed when comparing the magnitude of the RSFS closure to that of the "perfect" model. Figures 5.14(b), 5.15(b), and 5.16(b) plot this for each LES filter formulation. Again, one notices that increasing the filter-to-grid ratio improves accuracy of the model. However, unlike the correlation coefficient, the magnitude of the RSFS relative to the SFS closure is increased for the more scale-discriminant tenth-order filter.

Naturally, the extent of the model's success is related to the specific problem and such things as the LES resolution as compared to the resolution of the full dynamics. This is noted by the sudden decrease in performance near $t \sim 0.05$, which was previously highlighted as the time at which the fully-resolved DNS spectrum evolves past the prescribed LES filter width $k_{\bar{\Delta}}$. Thereafter, the models face a greater challenge as they seek to extrapolate the sub-filter scale dynamics.

Having somewhat characterized the ADM model, parametric *a posteriori* evaluations of different discretization schemes are used to provide additional insight into the impact of numerical errors. Here, the operator re-scaling as detailed in Appendix A is employed. Figure 5.17 plots the solution error of the LES results as compared to the filtered reference solution for filter-to-grid ratios $FTGR = 2$ and 4. Prior to the threshold time $t \sim 0.05$, there is a clear demarcation wherein the high-accuracy schemes exhibit less error. Eventually, however, all of the solutions are seen to rapidly increase in error likely owing to the surge in model inaccuracy. The respective dissipation rate metric $\epsilon(t)$ is plotted in Figure 5.18 and corroborates the inability of the model to curtail the accumulation of small-scale error. Figure 5.19 plots the accompanying spectra at $t \sim 0.1$ and reveals that the ADM method inherently tries to enforce the LES field by filtering the non-linear residual. Nevertheless, this does not serve to actively remove any high-wavenumber

content that may gather over time. Cascade of this energy is furthermore seen to be tied to the spectral resolvability of the numerical scheme, as previously established, with the high-accuracy discretizations exhibiting more content towards the grid cut-off. The presence of this content past the prescribed LES cut-off is also a result of the smooth characteristics of the LES filter; in the limit of a spectral-sharp filter, the presence of such sub-filter error would not occur since such content is unilaterally removed from the residual (Edoh et al., June 2016).

In general, the above figures confirm the previously-established notion that increasing the filter-to-grid ratio improves results. However, this is shown to improve the solution error only marginally. In addition, impacts of the numerical scheme are not evident – at least for the higher-order methods. This is in contrast to the "perfectly-modeled" quasi *a priori* analysis in Section 5.2 which reveals stark differences in accuracy, provided a suitable filter-to-grid ratio $FTGR > 1$. The current lack of differentiation in numerical scheme performance suggests the overwhelming presence of modeling error over discretization effects. Similar conclusions are drawn when considering quasi *a priori* analysis of the current models (not shown).

Referring to the *a priori* model evaluations in Figures 5.14-5.16, one recalls a predicted improvement in the ADM model correlation in accordance with the LES filter smoothness (e.g., lower-order discrete filter formulations). This idea is indeed observed in the *a posteriori* evaluations – although marginally – when comparing the second- (see Figures 5.17(a)- 5.17(b), 5.18(a)- 5.18(b) 5.19(a)- 5.19(b)) and tenth-order (see Figures 5.17(c)- 5.17(d), 5.18(c)- 5.18(d) 5.19(c)- 5.19(d)) Tangent formulations.

Particular to the ADM model is the role of the deconvolution procedure on error effects. Figures 5.17(c)- 5.17(d), 5.18(c)- 5.18(d), 5.19(c)- 5.19(d) consider a high-accuracy deconvolution ($k_{cond} = 10^3$ per the re-scaling approach of Appendix A), while Figures 5.17(e)- 5.17(f), 5.18(e)- 5.18(f), 5.19(e)- 5.19(f) show results of

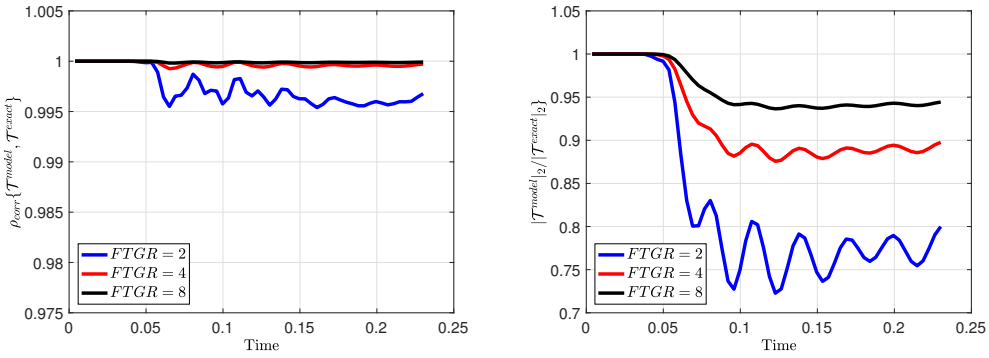

(a) (b)

Figure 5.14: A priori comparisons of the Approximate Deconvolution method (ADM)for a second-order Tangent LES filter formulation: a) correlation coefficient $\rho(\tau^{model}, \tau^{SFS})$, and b) magnitude $|\tau^{model}|_2/|\tau^{SFS}|_2$.

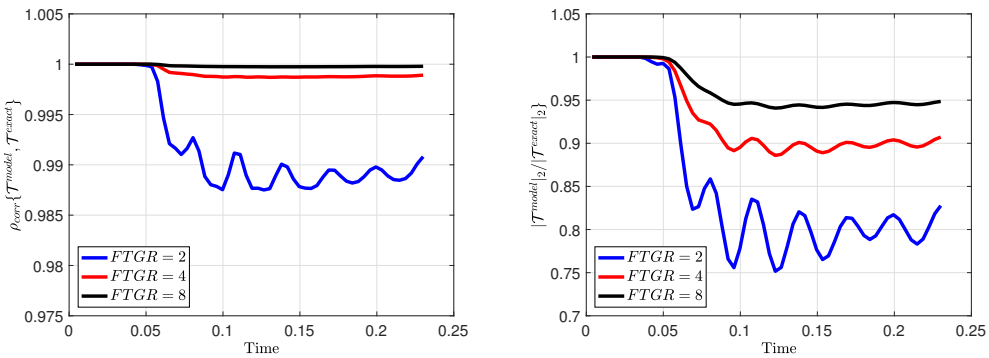

(a) (b)

Figure 5.15: A priori comparisons of the Approximate Deconvolution method (ADM) for a sixth-order Tangent LES filter formulation: a) correlation coefficient $\rho(\tau^{model}, \tau^{SFS})$, and b) magnitude $|\tau^{model}|_2/|\tau^{SFS}|_2$.

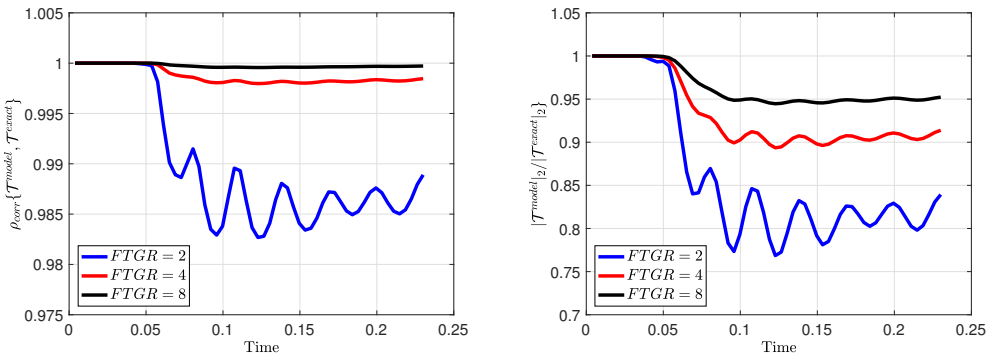

(a) (b)

Figure 5.16: A priori comparisons of the Approximate Deconvolution method (ADM) for a tenth-order Tangent LES filter formulation: a) correlation coefficient $\rho(\tau^{model}, \tau^{SFS})$, and b) magnitude $|\tau^{model}|_2/|\tau^{SFS}|_2$.

a zeroth order deconvolution, where $\tilde{u} \approx \bar{\bar{u}}$ ($k_{cond} = 1$). The latter is clearly seen to be less accurate to the extent that discretization effects for the high-order schemes on solution error are even less apparent during the starting phase $t < 0.05$ (see Figures 5.17(c)- 5.17(d) versus Figures 5.17(e)- 5.17(f)). Filter-to-grid ratio effects are also nullified. With respect to the energy spectra (see Figures 5.19(c)- 5.19(d) versus Figures 5.19(e)- 5.19(f)), it is seen that the "low-order" deconvolution tends to redistribute energy near the LES cut-off. This is then also manifested as much larger values of the normalized dissipation rate metric $\epsilon(t)$ (see Figures 5.18(c)- 5.18(d) versus Figures 5.18(e)- 5.18(f)). And so, it is apparent that the accuracy of the deconvolution can induce important modeling errors that further obfuscate the impact of discretization and LES parameter choices. Nevertheless, such low-accuracy deconvolutions have the advantage of potentially inhibiting aliasing effects through the non-linear term $\partial_x \overline{\tilde{u}\tilde{u}}$.

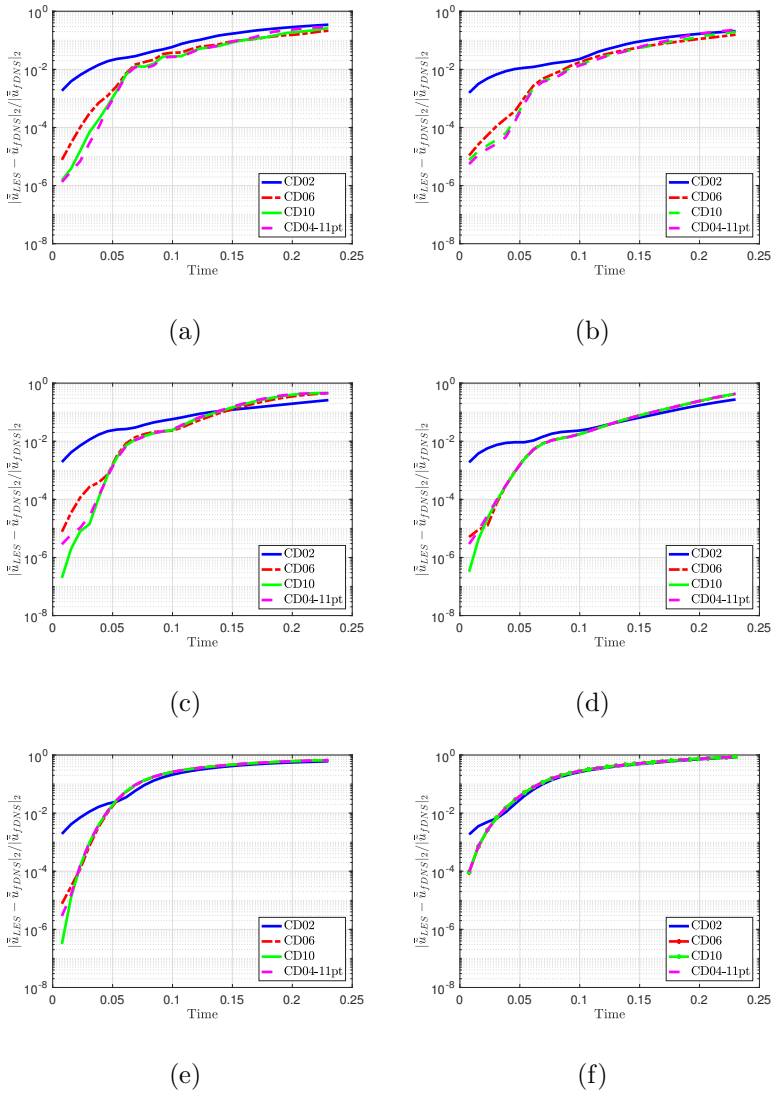

Figure 5.17: Effect of discretization scheme and filter-to-grid ratio (FTGR) on the time evolution of the relative L_2 LES solution error computed by *a posteriori* evaluation, $|\bar{\tilde{u}}_{LES} - \bar{\tilde{u}}_{fDNS}|_2 / |\bar{\tilde{u}}_{fDNS}|_2$. LES solutions assume a discrete Tangent filter formulation, tuned according to $k_{1/2}$. The Approximate Deconvolution method (ADM) is applied using re-scaled deconvolution (see Appendix A): a) second-order LES filter ($FTGR = 2$, $k_{cond} = 1.e3$), b) second-order LES filter ($FTGR = 4$, $k_{cond} = 1.e3$), c) tenth-order LES filter ($FTGR = 2$, $k_{cond} = 1.e3$), d) tenth-order LES filter ($FTGR = 4$, $k_{cond} = 1.e3$), e) tenth-order LES filter ($FTGR = 2$, $k_{cond} = 1.e0$), and f) tenth-order LES filter ($FTGR = 4$, $k_{cond} = 1.e0$).

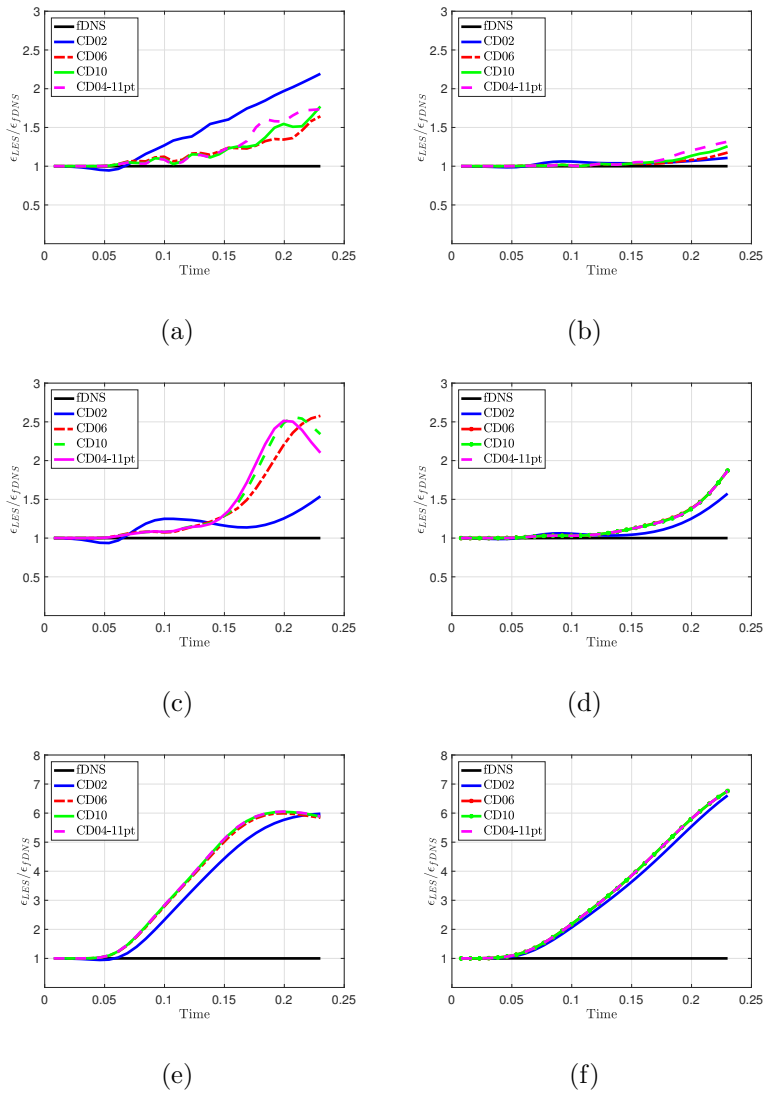

Figure 5.18: Effect of discretization scheme and filter-to-grid ratio (FTGR) on the time evolution of the normalized kinetic energy dissipation rate $\epsilon(t)$. LES solutions assume a discrete Tangent filter formulation, tuned according to $k_{1/2}$. The Approximate Deconvolution method (ADM) is applied using re-scaled deconvolution (see Appendix A): a) second-order LES filter ($FTGR = 2$, $k_{cond} = 1.e3$), b) second-order LES filter ($FTGR = 4$, $k_{cond} = 1.e3$), c) tenth-order LES filter ($FTGR = 2$, $k_{cond} = 1.e3$), d) tenth-order LES filter ($FTGR = 4$, $k_{cond} = 1.e3$), e) tenth-order LES filter ($FTGR = 2$, $k_{cond} = 1.e0$), and f) tenth-order LES filter ($FTGR = 4$, $k_{cond} = 1.e0$).

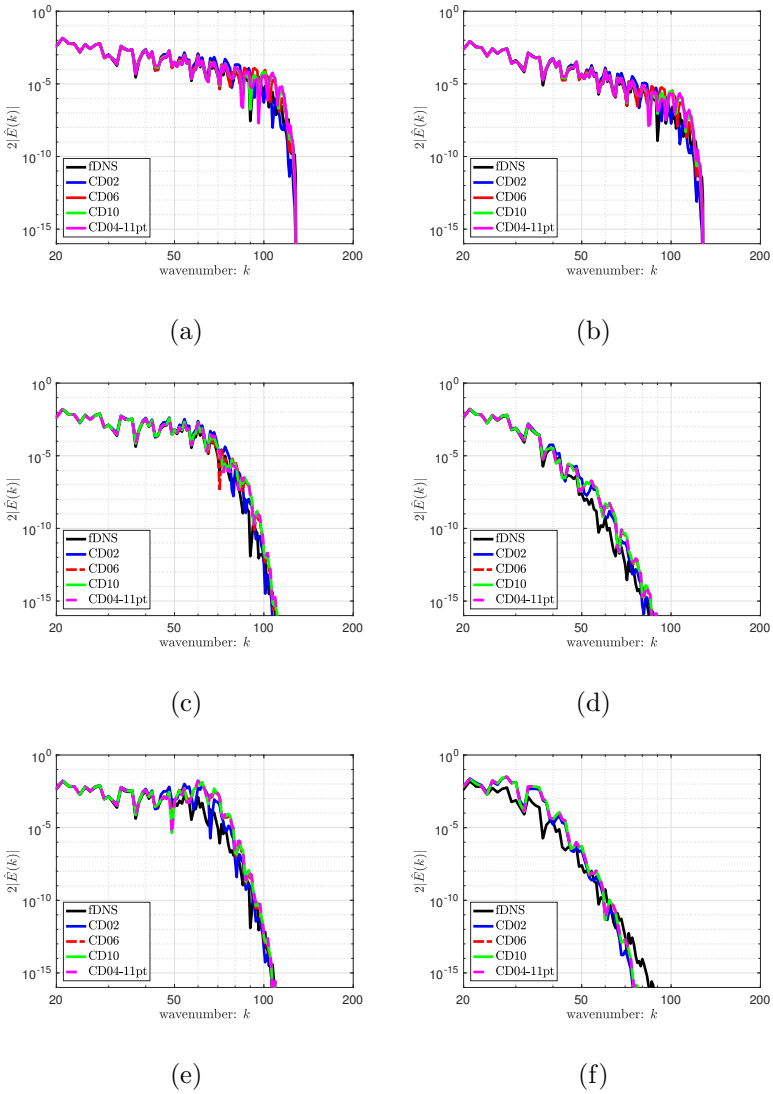

Figure 5.19: Effect of discretization scheme and filter-to-grid ratio (FTGR) on the kinetic energy spectrum, $2|\hat{E}(k)| = |\hat{u}(k)|^2$, at time-step $N = 1300$ ($t \sim 0.1$). LES solutions assume a discrete Tangent filter formulation, tuned according to $k_{1/2}$. The Approximate Deconvolution method (ADM) is applied using re-scaled deconvolution (see Appendix A): a) second-order LES filter ($FTGR = 2$, $k_{cond} = 1.e3$), b) second-order LES filter ($FTGR = 4$, $k_{cond} = 1.e3$), c) tenth-order LES filter ($FTGR = 2$, $k_{cond} = 1.e3$), d) tenth-order LES filter ($FTGR = 4$, $k_{cond} = 1.e3$), e) tenth-order LES filter ($FTGR = 2$, $k_{cond} = 1.e0$), and f) tenth-order LES filter ($FTGR = 4$, $k_{cond} = 1.e0$).

5.3.1.2 A Bardina-type Scale-Similarity Model (SSM)

Next the Bardina-type Scale-Similarity model (SSM) is considered. Again, the DNS data is filtered and manipulated in order to conduct an *a priori* evaluation of the τ^{SSM} closure. Figures 5.20(a), 5.21(a), and 5.22(a) plot the correlation coefficient of the SSM model as a function of time. Unlike the RSFS/ADM closure, the SSM model does not constitute a subset of the true SFS, but rather a filtered-scale (e.g., large-scale) approximation to the small-scale dynamics. As a consequence, the observed performance of SSM is much degraded in comparison to the previous ADM *a priori* analysis. Inspecting the correlation coefficient relative to the exact closure, the SSM model shows decent performance at low filter-to-grid ratios (FTGR), but degrades when the LES filter width becomes too large relative to the gird. This is in contrast to the ADM method which is seen to improve under such circumstances. The reason for this may lie in the way the SSM model is made to use increasingly larger scales in order to approximate the true closure – this eventually becomes ill-founded as one eventually approaches forcing or geometry-dependent scales. Interestingly, however, correlations are improved as the smoothness of the filter reformulation is increased. In this way, the second-order Tangent implementation shows a ~ 0.95 correlation for $FTGR = 2$, compared to the tenth-order filter which has a correlation of ~ 0.8.

With respect to assessing the model magnitude relative to the true closure, Figures 5.20(b), 5.21(b), and 5.22(b) show incremental degradation of performance as the filter-to-grid ratio is increased. Also, smoother filter formulations are revealed to have lower representations of the closure magnitude during the start-up phase $t < 0.05$, which again may be a consequence of building the model off of the large scales, even while the dynamics may be fully represented on the LES resolution. Contrastingly, the ADM procedure makes no assumptions while the dynamics reside within the LES cut-off range. As previously noted, further inspection of the effect of the SSM closure in Equation 5.14 also reveals that it approaches the

RSFS with zeroth-order deconvolution at the limit of a spectral-sharp filter. In this way, the model's shortcomings can be understood as an abatement of the nonlinear term $\widetilde{\overline{u}\overline{u}} \approx \overline{\tilde{\bar{u}}\tilde{\bar{u}}}$ in addition to a velocity increment relating to $\bar{\bar{\tilde{u}}} \neq \tilde{\bar{u}}$ that can generate unwanted content past the LES cut-off and threaten algorithmic stability.

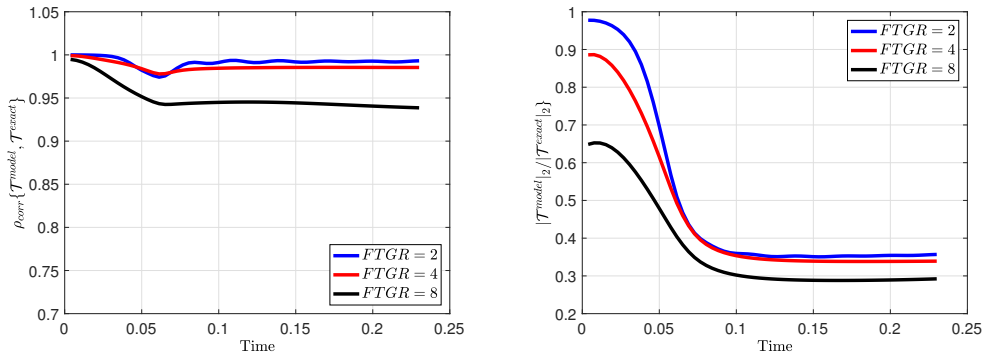

(a) (b)

Figure 5.20: A priori comparisons of the Bardina-type Scale-Similarity model (SSM) for a second-order Tangent LES filter formulation: a) correlation coefficient $\rho(\tau^{model}, \tau^{SFS})$, and b) magnitude $|\tau^{model}|_2/|\tau^{SFS}|_2$.

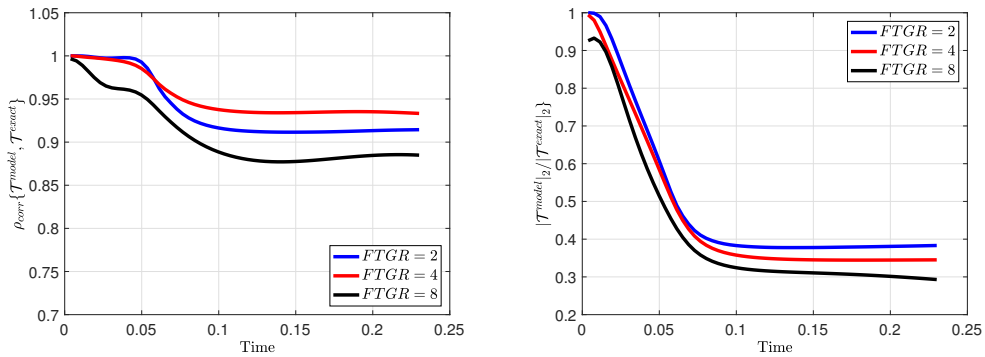

(a) (b)

Figure 5.21: A priori comparisons of the Bardina-type Scale-Similarity model (SSM) for a sixth-order Tangent LES filter formulation: a) correlation coefficient $\rho(\tau^{model}, \tau^{SFS})$, and b) magnitude $|\tau^{model}|_2/|\tau^{SFS}|_2$.

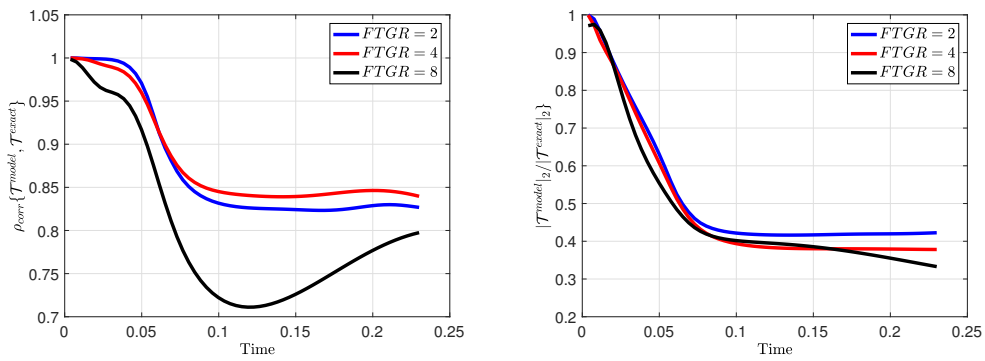

(a) (b)

Figure 5.22: A priori comparisons of the Bardina-type Scale-Similarity model (SSM) for a tenth-order Tangent LES filter formulation: a) correlation coefficient $\rho(\tau^{model}, \tau^{SFS})$, and b) magnitude $|\tau^{model}|_2/|\tau^{SFS}|_2$.

The anticipated short-comings of the SSM method are confirmed by *a posteriori* LES calculations. Figure 5.23 pots the solution error and indeed reveals the anticipated result of decreased performance at a higher filter-to-grid ratio. In addition, there is diminished differentiation between the numerical schemes (at least beyond the low-order CD02 stencil), especially in comparison to the ADM closure. This again suggests the pervasiveness of model error. Inspecting the evolution of the $\epsilon(t)$ metric in Figure 5.24 shows that the SSM treatment admits much more high-wavenumber activity than desired. And unlike the ADM method, this error is exacerbated by increasing the filter-to-grid ratio (at least in a relative sense). The spectra of Figure 5.25 provide additional insight, demonstrating the large amount of content past the target LES cut-off. Under certain conditions, such developments can lead to unmitigated pile-up and can eventually lead to numerical instability. As anticipated, the velocity increment identified in Equation 5.14 actively induces this high-wavenumber activity, which in turn is maintained based on the discretization scheme's modified wavenumber characteristics. For example the low-order CD02 scheme has less energy near the grid

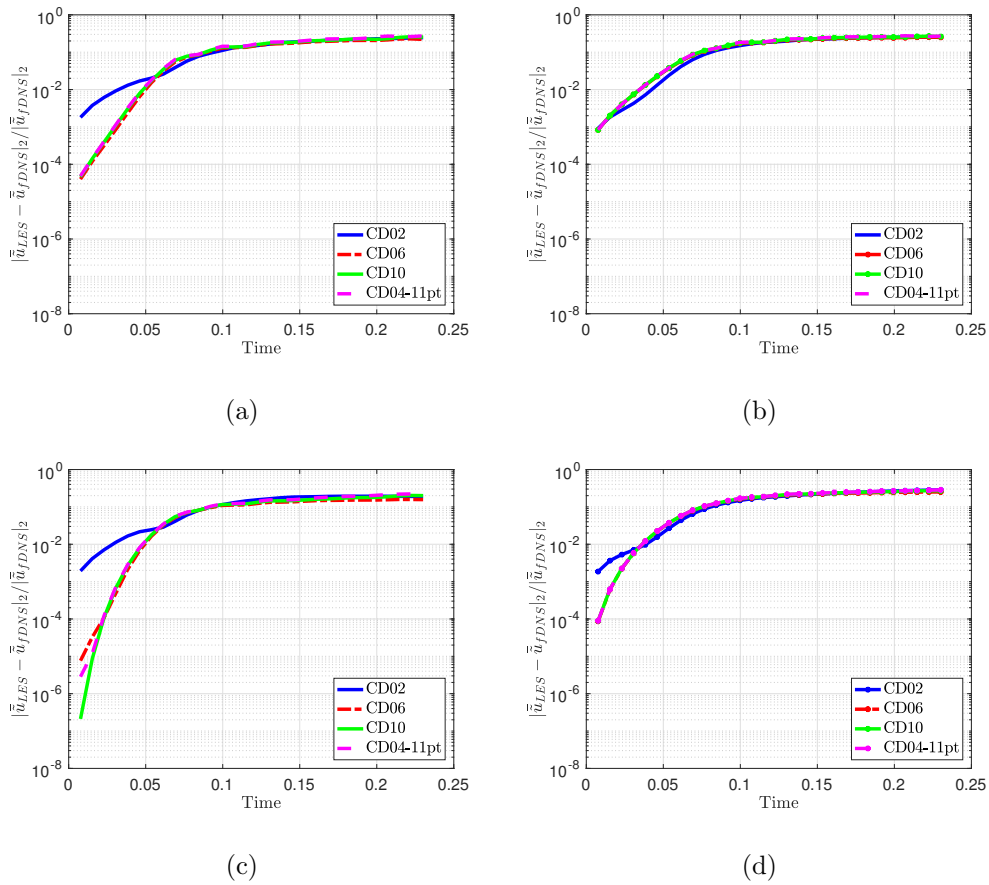

Figure 5.23: Effect of discretization scheme and filter-to-grid ratio (FTGR) on the time evolution of the relative L_2 LES solution error computed by *a posteriori* evaluation, $|\bar{\tilde{u}}_{LES} - \bar{\tilde{u}}_{fDNS}|_2/|\bar{\tilde{u}}_{fDNS}|_2$. LES solutions assume a discrete Tangent filter formulation, tuned according to $k_{1/2}$, employed with the Bardina-type Scale-Similarity model (SSM closure: a) second-order Tangent LES filter ($FTGR = 2$, $k_{cond} = 10^3$), b) second-order Tangent LES filter ($FTGR = 4$, $k_{cond} = 10^3$), c) tenth-order Tangent LES filter ($FTGR = 2$, $k_{cond} = 10^3$), and d) tenth-order Tangent LES filter ($FTGR = 4$, $k_{cond} = 10^3$).

cut-off in comparison to the higher-order schemes; the same observation is supported by inspecting the evolution of $\epsilon(t)$. Meanwhile, the ADM procedure is seen to inherently enforce the explicit LES filter and to be less susceptible to the proliferation of small-scale content based on the discretization scheme. Employing

increasingly scale-discriminant LES formulations for SSM, however, is shown to reduce such errors, likely owing to $(\bar{\tilde{u}}\bar{\tilde{u}} - \bar{\bar{\tilde{u}}}\bar{\bar{\tilde{u}}}) \rightarrow 0$.

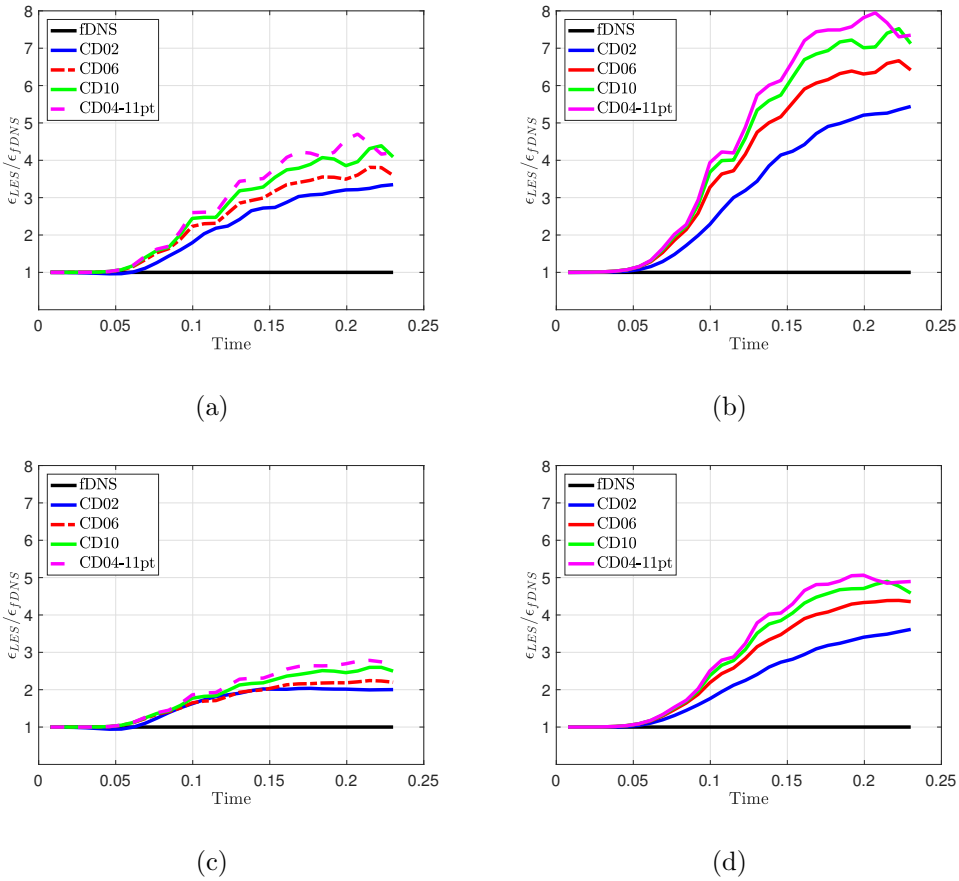

Figure 5.24: Effect of discretization scheme and filter-to-grid ratio (FTGR) on the time evolution of the normalized kinetic energy dissipation rate based on the strain-rate metric $\epsilon(t)$ (see Equation 5.12). LES solutions assume a discrete Tangent filter formulation, tuned according to $k_{1/2}$, employed with the Bardina-type Scale-Similarity model (SSM closure: a) second-order Tangent LES filter ($FTGR = 2$, $k_{cond} = 10^3$), b) second-order Tangent LES filter ($FTGR = 4$, $k_{cond} = 10^3$), c) tenth-order Tangent LES filter ($FTGR = 2$, $k_{cond} = 10^3$), and d) tenth-order Tangent LES filter ($FTGR = 4$, $k_{cond} = 10^3$).

Because of the presence of high-wavenumber content with the SSM model, it is often implemented with an additional eddy-viscosity model (Winckelmans et al.,

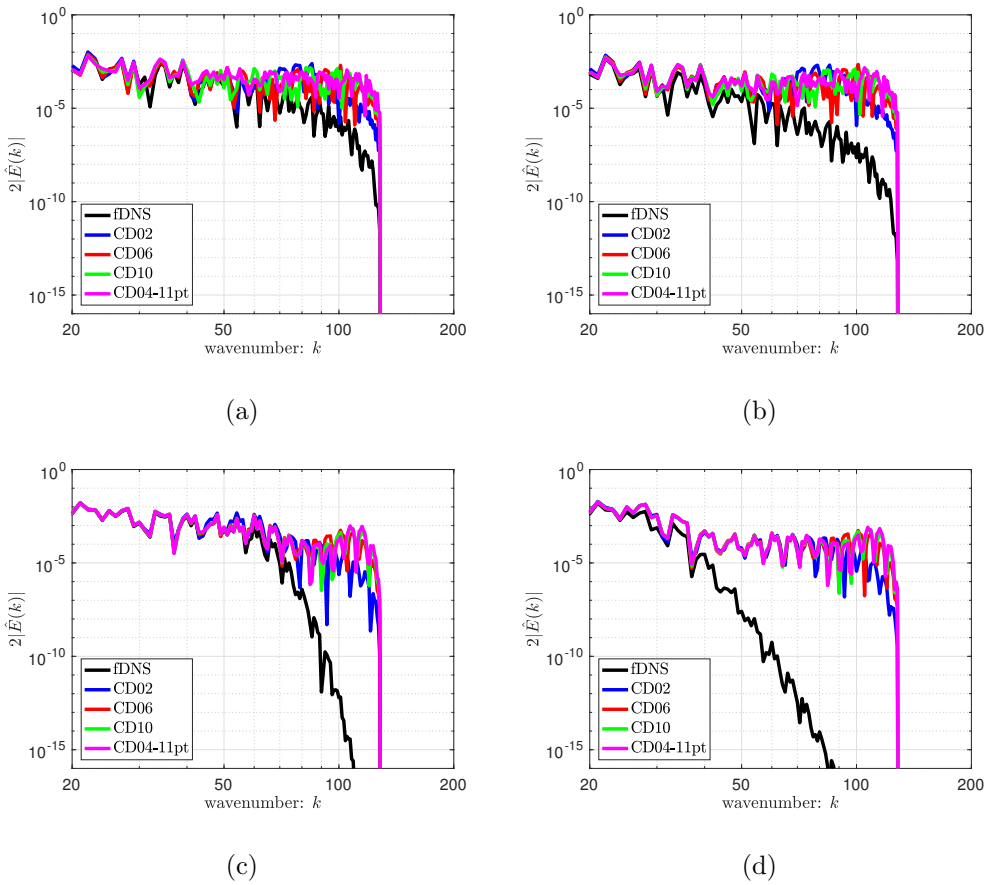

Figure 5.25: Effect of discretization scheme and filter-to-grid ratio (FTGR) on the kinetic energy spectrum, $2|\hat{E}(k)| = |\hat{u}(k)|^2$, at time-step $N = 1300$ ($t \sim 0.1$). LES solutions assume a discrete Tangent filter formulation, tuned according to $k_{1/2}$, employed with the Bardina-type Scale-Similarity model (SSM closure: a) second-order Tangent LES filter ($FTGR = 2$, $k_{cond} = 10^3$), b) second-order Tangent LES filter ($FTGR = 4$, $k_{cond} = 10^3$), c) tenth-order Tangent LES filter ($FTGR = 2$, $k_{cond} = 10^3$), and d) tenth-order Tangent LES filter ($FTGR = 4$, $k_{cond} = 10^3$).

2001). While the reasoning for this mixed-model approach can be rationalized by the sub-filter scale decomposition of Equation 5.7, wherein respective models address the RSFS and SGS components (Carati et al., 2001), a more practical explanation stems from the fundamental need to removed the velocity increment contributions. In this way, one can consider adding the filter-based artificial dis-

sipation of Section 3.3. The impact of such a modification is explored in the following section as a stabilization to both the ADM and SSM models.

5.3.2 Impact of Stabilization

Here, the influence of dissipation in its ability to mitigate computational errors is examined. In the quasi *a priori* evaluation of Section 5.2, a "perfect" model was assumed; as a result, the dissipation was tuned such as to target the lingering numerical error. In the context of practical LES modeling, however, one may need to quell more urgent sources of inaccuracy. The non-stabilized study of the ADM and SSM closures in Section 5.3 demonstrates that a significant portion of the error stems from the LES model, to the point that none of the high-accuracy schemes (e.g., CD06, CD04-11pt, and CD10) are much differentiable. Therefore, in the current context of stabilizing these scale-similarity-type models, tuning the dissipation relative to the numerical scheme is not effective. Instead, it is more advantageous to reduce model error with the expectation that numerical scheme choice will become identifiable under perfect-like modeling. Inspecting the dissipation rate metric $\epsilon(t)$ and the energy spectra of the *a posteriori* studies suggests that a significant source of the solution error stem from the inability to enforce the LES resolution – in other words, the models are inappropriately allowing high-wavenumber content past the designated cut-off. As a result, the influence of numerical dissipation (administered in the form of filter-based artificial dissipation) is chosen here to be identical to the explicit LES filter and satisfies the purpose of enforcing the prescribed LES field. A tenth-order Tangent LES filter formulation is considered at a filter-to-grid ratio $FTGR = 2$ ($\bar{\Delta}_{1/2} = 4\Delta x_{LES}$).

In order to further bring generality to the stabilization scheme, a dynamic procedure is implemented, where the dissipation is scaled by a coefficient (C/C_o)

such that the integrable equation reads,

$$\partial_t \bar{\tilde{u}} = R_o\{\bar{\tilde{u}}\} + (C/C_o) \cdot R_{AD}\{\bar{\tilde{u}}\} \ . \tag{5.15}$$

$R_o\{\bar{\tilde{u}}\}$ corresponds to the original LES residual (including any closure models), while $R_{AD}\{\bar{\tilde{u}}\}$ constitutes the filter-based artificial dissipation. In order to satisfy the LES resolution, the dissipation scheme is chosen to be identical to the explicit LES filter. However this enforcement should ideally be active as necessary, thus also limiting the possibility of un-necessary dissipative error. The following definitions of the scaling coefficients are applied:

$$C = \frac{|\bar{\tilde{u}} - \mathcal{G}_{\bar{\Delta}}\{\bar{\tilde{u}}\}|_2}{|\mathcal{G}_{\bar{\Delta}}\{\bar{\tilde{u}}\} - \mathcal{G}_{\hat{\Delta}}\{\bar{\tilde{u}}\}|_2} \quad \text{such that} \quad \hat{\Delta} > \bar{\Delta} \ , \tag{5.16}$$

$$C_o = \frac{\int_{(k_{\bar{\Delta}}\Delta x)/\pi}^{(k_c^*\Delta x)/\pi} [(1 - |\hat{\mathcal{G}}_{\bar{\Delta}}|) \cdot k^\beta] \ d(k\Delta x/\pi)}{\int_{(k_{\hat{\Delta}}\Delta x)/\pi}^{(k_{\bar{\Delta}}\Delta x)/\pi} [(|\hat{\mathcal{G}}_{\bar{\Delta}}| - |\hat{\mathcal{G}}_{\hat{\Delta}}|) \cdot k^\beta] \ d(k\Delta x/\pi)} \ . \tag{5.17}$$

Equation 5.16 is a measure of the extraneous energy contained past the primary LES cut-off $\bar{\Delta}$, normalized for sake of dimensionality by the energy contained within a band of adjacent scales $k \in [k_{\hat{\Delta}}, k_{\bar{\Delta}}]$. Next, this measure is adjusted according to the amount of energy expected in each of the bands $k \in [k_{\hat{\Delta}}, k_{\bar{\Delta}}]$ and $k \in [k_{\bar{\Delta}}, k_c^*]$ (where $k_c^* \in \langle k_{\bar{\Delta}}, k_{\Delta x_{LES}}]$), assuming a k^β inertial range scaling (see Equation 5.17). Based on this, the filter-based artificial dissipation is fully "on" with $(C/C_o) \geq 1$, when the $k \in [k_{\bar{\Delta}}, k_c^*]$ portion of the resolved sub-filter scales are completely saturated in proper proportion to the expected unfiltered field. In order to make the damping more aggressive, the upper bound of the integral in the numerator of Equation 5.17 can be made to approach $(k_{\bar{\Delta}}\Delta x)/\pi$, thus making the dissipation less tolerant of content past the LES cut-off, $\bar{\Delta}$. This, however, comes with the risk of adding too much dissipation. The current scaling, being based on an L_2 norm of the solution, represents a measure of kinetic energy. It is inherently global and does not require a least-squares optimization in multi-dimensions, unlike the dynamic procedure of Germano (Germano et al., 1991; Lilly, 1992; Park

and Mahesh, 2009). Instead, a level of consistency is achieved by embedding the dissipation operators within the sensors. In this way, anisotropic effects can also naturally be accounted for. In the current exercises $\beta = -2$ and $\hat{\mathcal{G}}_{\hat{\Delta}}$ is chosen to also be a tenth-order Tangent stencil, tuned to a filter-to-grid ratio $FTGR = 4$ ($\hat{\Delta}_{1/2} = 8\Delta x_{LES}$). This adaptive formulation is reminiscent of the automatic filtering criterion of Tantikul and Domaradzki (2011), which considers the kinetic energy ratios between different wavenumber bands; other related schemes include the adaptive filtering of Flad et al. (2016) which incorporates aliasing considerations, as well as the original relaxation scheme employed by Stolz et al. (2001) for stabilizing the ADM closure.

Figure 5.26 shows the impact of the stabilization in terms of the solution error. As before, there is little-to-no differentiation between the performance of the higher-order schemes which suggests that modeling error is still prevalent. Nevertheless, the accuracy of the computations is seen to be improve relative to the non-stabilized scenarios of ADM (see Figures 5.17(c) and 5.17(e)) and SSM (see Figure 5.23(c)). The adaptive procedure is shown to hold the solution error relatively constant and hold off the accumulation of model error past $t > 0.05$. These observations are once more confirmed by the strain-rate metric $\epsilon(t)$ in Figure 5.27 that suggests a reduction in the small-scale activity. For example, when considering the CD04-11pt discretization, the high-accuracy deconvolution exhibits $\max\{\epsilon_{LES}/\epsilon_{fDNS}\} \sim 1.15$ compared to the unstabilized case, where $\max\{\epsilon_{LES}/\epsilon_{fDNS}\} \sim 2.4$. For the zeroth-order deconvolution, incorporating the present adaptive dissipation yields $\max\{\epsilon_{LES}/\epsilon_{fDNS}\} \sim 2.5$ in comparison to $\max\{\epsilon_{LES}/\epsilon_{fDNS}\} \sim 6.0$ without. The same is true for improvements to the SSM closure; with stabilization, one observes $\max\{\epsilon_{LES}/\epsilon_{fDNS}\} \sim 1.5$ compared to $\max\{\epsilon_{LES}/\epsilon_{fDNS}\} \sim 2.8$ without dissipation. Another interesting trend relates to the fact that the high-accuracy schemes now show improvement – albeit somewhat marginal – compared to the low-order CD02 scheme, at least for the

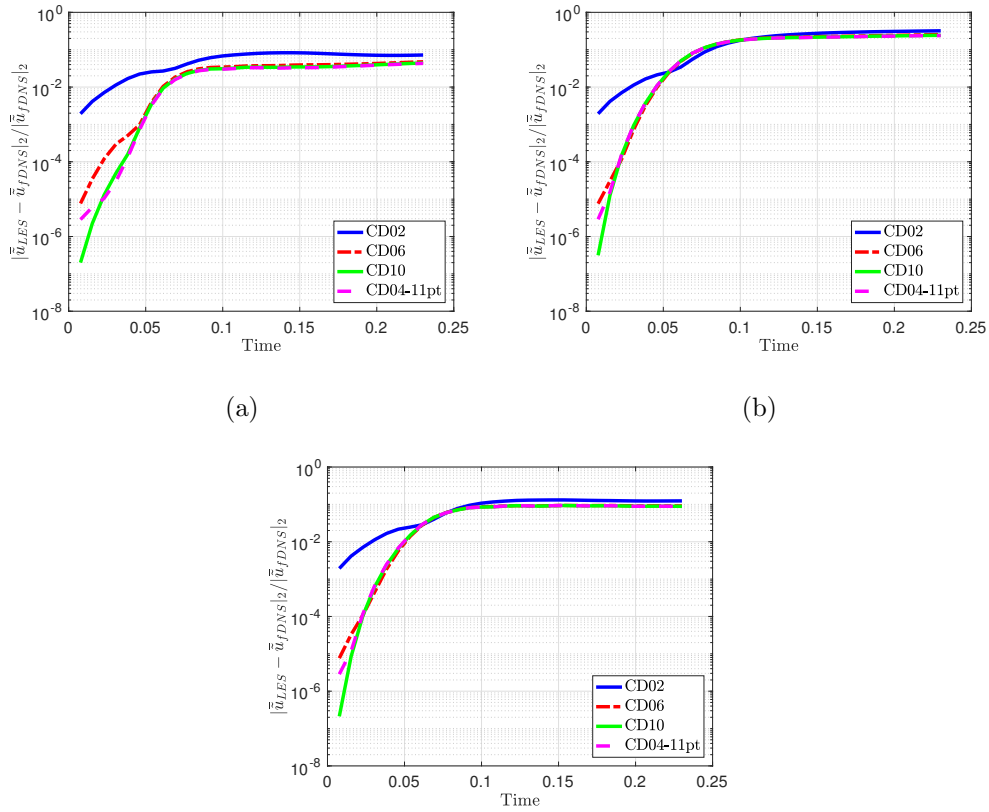

Figure 5.26: Effect of discretization scheme and filter-to-grid ratio (FTGR) on the time evolution of the relative L_2 LES solution error computed by *a posteriori* evaluation, $|\bar{\bar{u}}_{LES} - \bar{\bar{u}}_{fDNS}|_2/|\bar{\bar{u}}_{fDNS}|_2$. LES solutions assume a tenth-order discrete Tangent filter formulation, tuned according to $\bar{\Delta}_{1/2} = 4\Delta x_{LES}$. A filter-based artificial dissipation stencil consistent with explicit LES filter is employed with the adaptive procedure (see Equations 5.15-5.17 with $k_c^* \Delta x = \pi$): a) ADM model (deconvolution: $k_{cond} = 10^3$), b) ADM model (deconvolution: $k_{cond} = 1$), and c) Scale-Similarity model.

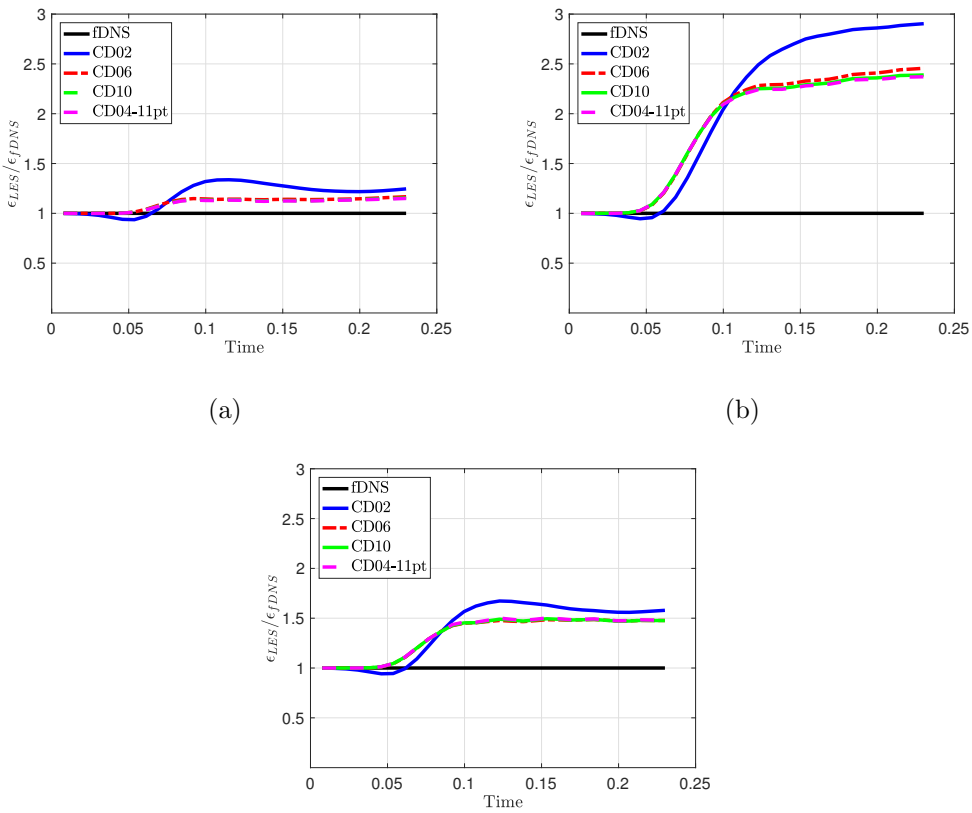

Figure 5.27: Effect of discretization scheme on the time evolution of the normalized kinetic energy dissipation rate based on the strain-rate metric $\epsilon(t)$ (see Equation 5.12). LES solutions assume a tenth-order discrete Tangent filter formulation, tuned according to $\bar{\Delta}_{1/2} = 4\Delta x_{LES}$. A filter-based artificial dissipation stencil consistent with explicit LES filter is employed with the adaptive procedure (see Equations 5.15-5.17 with $k_c^* \Delta x = \pi$): a) ADM model (deconvolution: $k_{cond} = 10^3$), b) ADM model (deconvolution: $k_{cond} = 1$), and c) Scale-Similarity model.

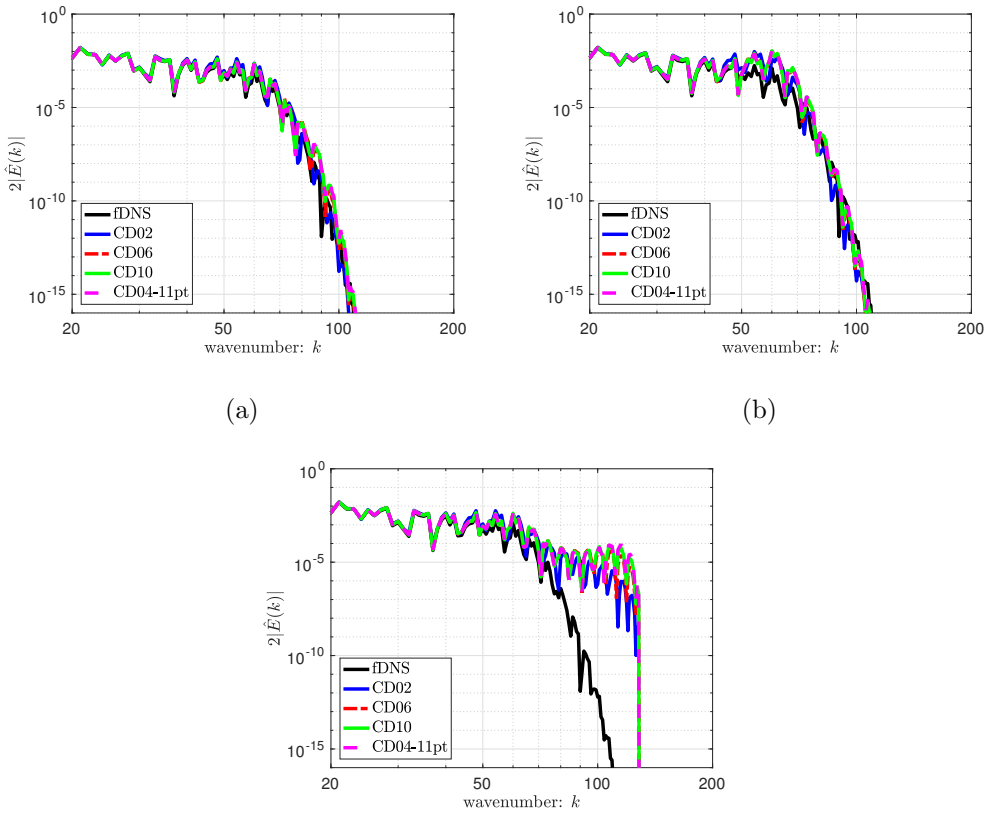

Figure 5.28: Effect of discretization scheme and filter-to-grid ratio (FTGR) on the kinetic energy spectrum, $2|\hat{E}(k)| = |\hat{u}(k)|^2$, at time-step $N = 1300$ ($t \sim 0.1$). LES solutions assume a tenth-order discrete Tangent filter formulation, tuned according to $\bar{\Delta}_{1/2} = 4\Delta x_{LES}$. A filter-based artificial dissipation stencil consistent with explicit LES filter is employed with the adaptive procedure (see Equations 5.15-5.17 with $k_c^* \Delta x = \pi$): a) ADM model (deconvolution: $k_{cond} = 10^3$), b) ADM model (deconvolution: $k_{cond} = 1$), and c) Scale-Similarity model.

$\epsilon(t)$ metric. This is the case for all the closures considered and encourages the potential benefits of employing high-accuracy schemes. Comparing the respective spectra in Figure 5.28 to the non-stabilized implementations of ADM (see Figures 5.19(c) and 5.19(e)) and SSM (see Figure 5.25(c)), one again notes that the dissipation indeed reduces some of the small-scale content and thus is helping to enforce the LES resolution.

5.4 Synopsis of Trends

The investigations of the current Chapter into the interacting effects of discretization, LES formulation, modeling, and stabilization provide useful insight which one may seek to employ and verify with respect to the more advanced fluid system. The list of broad observations can be summarized as follows for the special case of "perfect modeling":

1. Reduction in solution error is achieved by proper selection of the LES filter-to-grid ratio (FTGR) relative to the resolvability (i.e., spectral accuracy) of the base discretization scheme. At low filter-to-grid ratios (e.g, grid- or implictly-filtered LES, $FTGR = 1$), the ability of the model to enforce the LES field is significantly impeded by the overwhelming presence of discretization error near the grid cut-off. Larger filter-to-grid ratios approach grid-converged LES, and solution error is diminished as the influence of the exact model overshadows numerical artifacts.

2. Scale-discriminant LES filter formulations make it such that the choice of numerical scheme is more impactful, with high-accuracy schemes being preferable. At modest filter-to-grid ratios, optimized schemes are superior; however, standard schemes regain merit for sufficiently large filter-to-grid ratios.

3. Even in the scenario of "perfect modeling", one may further improve accu-

racy by removing lingering numerical error contributions. Therefore a numerical dissipation scheme can be tuned such as to target the scales poorly resolved the discretization. In this case, matching the dissipation scheme's growth-factor response to the normalized modified wavenumber of the convective scheme has shown to yield the best results. This is believed to stem from the fact that sufficient erroneous content is removed without also undermining the exact model contributions.

The lessons learned from the "perfect model" analysis are then balanced by considering practical LES models such as the ADM and SSM closures, which reveal the following trends:

1. Incorporation of real closures for *a posteriori* use introduces large amounts of modeling error that obfuscates discretization effects; a clear demarcation is made between the low-order CD02 scheme and the higher-order methods (CD06, CD04-11pt, CD10) but not much difference is seen between the high-accuracy methods (contrary to the case of "perfect modeling").

2. Improvement of solution accuracy largely resides in decreasing the model error. As most closures are inadequate at enforcing the LES spectral resolution, one may incorporate numerical dissipation and formulate it relative to the presumed explicit LES filter. In this way, the stabilization is made to help the primary model in satisfying the LES field, and thus reduces solution error. In doing so, intuitive trends regarding the benefits of employing high-oder methods are recovered (although still marginal) when comparing amongst the high-accuracy discretizations.

3. The Approximate Deconvolution method (ADM) is generally superior to the the Scale-Similarity model (SSM), provided sufficiently accurate deconvolution procedures are employed. The former model inherently tries to enforce the LES spectral resolution while the latter instigates content past

the desired LES cut-off due to the contribution of a velocity-increment that is inherent in its formulation.

The proceding chapter will thus consider the 3D compressible Navier-Stokes equations and will look at simulation of the Taylor-Green vortex as an analogous testbed to the current prototypical 1D viscous Burgers investigation, where analysis will seek to confirmation the preliminary conclusions garnered here.

CHAPTER 6

3D Taylor Green Vortex with the Compressible Navier-Stokes Equations

This chapter is written in part from modifications to a previous AIAA conference proceedings article "Inspecting Interactions of Discretization, Filter Formulation, and Stabilization in LES: Lessons from the Taylor-Green Vortex" (Edoh and Karagozian, June 2017b).

6.1 Background

While the explorations of the previous chapter utilize the 1D Burgers equation in order to extract important generalizations regarding the overall LES treatment, it remains to see how these lessons extend to the full fluid equations. To this effect, the current chapter considers the three-dimensional (3D) Navier-Stokes equations (NSE) (see Equation 6.1). Incorporating variable density is necessary for considering non-negligible Mach effects and offers greater generality in the study of fluid behavior; numerical investigation of incompressible-like conditions is then made possible by proper low-Mach preconditioning formulations that typically employ an iterative pseudo-time procedure (Merkle and Athavale, 1987; Turkel, 1999; Venkataswaran and Merkle, 2000). Specifically, the compressible form of the equations includes the ability to simulate high-pressure reactive flow – a long-term direction of the current research initiative – and thus deserves proper familiarity.

$$\frac{\partial Q}{\partial t} + \frac{\partial E_j}{\partial x_j} = \frac{\partial V_j}{\partial x_j},$$

$$Q = \begin{bmatrix} \rho \\ \rho u_i \\ \rho e_o \end{bmatrix}, \quad E_j = \begin{bmatrix} \rho u_j \\ \rho u_i u_j + P\delta_{ij} \\ (\rho e_o + P) u_j \end{bmatrix}, \quad V_j = \begin{bmatrix} 0 \\ \tau^d_{ij} \\ \tau^d_{jk} u_k + q_j \end{bmatrix},$$

$$\begin{aligned} P &= (\gamma - 1)\left[\rho e_o - \frac{1}{2}\rho u_i u_i\right], \\ \tau^d_{ij} &= \mu\left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i}\right) - \frac{2}{3}\mu\frac{\partial u_k}{\partial x_k}\delta_{ij}, \\ q_i &= \kappa\frac{\partial T}{\partial x_i}. \end{aligned} \tag{6.1}$$

As apparent, simultaneous solution of the different conservation laws constituting the system (i.e., mass, momentum, energy) introduces several non-linear equations, with dynamics being further complicated by thermodynamic coupling. As a result, the NSE system admits far more symmetries and represents a significant increase in complexity relative to the prototypical Burgers equation previously studied. In this way, the threats to numerical accuracy and stability are further heightened. For example, accumulation of error in the evolution of momentum can directly feed into the mass and energy equations through the convective terms. And in the instance where diffusive mechanisms are under-represented (as they are in the LES of high-Reynolds number flow) or are entirely absent (as in the continuity equation equation), such potential for error-inheritance increases computational susceptibilities to instability and inaccuracy.

Despite the new physical and algorithmic intricacies introduced by the fluid system, a tractable assessment of numerical scheme appropriateness can be made by considering canonical flows. One such example, explored herein, is the Taylor-Green vortex (TGV) first introduced for the incompressible equations (Taylor

and Green, 1937; Brachet et al., 1983). Subsequent investigations of the test case have included low Mach number initializations for the compressible system (DeBonis, 2013). Most intriguing about the Taylor-Green vortex is the natural transition from a laminar to a turbulent state, followed by a subsequent decay process. The initialization (see Equation 6.15) essentially constitutes a single large vortex, which then eventually breaks down into smaller features. In this sense, the test case offers the opportunity to witness a scheme's ability to predict the proliferation of an initial large-scale perturbation towards turbulence through the generation of small-scales (see Figure 6.1). In addition, calculations at the inviscid limit ($Re \rightarrow \infty$) have been used to highlight non-linear scheme stability properties such as secondary-conservation (Gassner et al., 2016).

While the triply-periodic problem domain is amenable to spectral computations which are often cited as reference solutions (Brachet et al., 1983), a popular objective in the community has been the task of recovering the proper dynamics with physical-space methods on an under-resolved mesh. Quite often, implicit LES (ILES) implementations are studied, wherein the inherent scheme dissipation – either native to the algorithm or employed with the sole intention of maintaining solution stability – is considered (Drikakis et al., 2007). For example, DeBonis (2013) studies effects of the convective discretization and the grid resolution on the evolution of the Taylor-Green vortex; yet, this is done while administering solution filtering with a scaling coefficient that is chosen to sustain the computation's stability with minimal damping. As such, in the context of relaxation filtering, many have looked at the effect of the filter order and filter strength on the ability to properly predict the TGV (Aubard et al., 2013; Fauconnier, 2013). Alternatively, particularly in the realm of Discontinuous Galerkin (DG) and Flux Reconstruction (FR) methods, authors have also considered the utility of energy-stable schemes (Bull and Jameson, 2014) and polynomial de-aliasing (e.g., over-integration) (Gassner and Beck, 2013).

In most instances, the performance of these algorithms is judged relative to the kinetic energy dissipation rate as calculated by differentiating energy with respect to time,

$$\epsilon(E_k) \;=\; -\frac{dE_k}{dt}, \text{ with } E_k = \frac{1}{\rho_o \Omega}\int_\Omega \rho \frac{u_i u_i}{2} d\Omega \; . \tag{6.2}$$

Such a metric – as previously explained in Section 5.1.1 – is biased towards large-scale depletion of kinetic energy and is less sensitive to enforcement of the LES field at high-wavenumbers. Although most of the TGV studies assume grid-filtering with $FTGR = 1$ (i.e., $\bar{\Delta} = 2\Delta x$), the choice of this particular metric makes it particularly limited in assessing LES formulations based on filter-to-grid ratios above unity. In addition, $\epsilon(E_k)$ as calculated by Equation 6.2 can be said to represent the net dissipation of the algorithm, which may be suitable for implicit-modeling but is less telling of explicit modeling efforts. As a direct analogue to the 1D Burgers investigation of Chapter 5, one may judge the kinetic dissipation rate by inspecting source terms of the corresponding conservation equation. In the case of compressible flow, this amounts to calculating,

$$\frac{1}{2}\frac{\partial \rho u_i u_i}{\partial t} \;=\; u_i \frac{\partial \rho u_i}{\partial t} - \frac{1}{2}(u_i u_i)\frac{\partial \rho}{\partial t} \tag{6.3}$$

$$\begin{aligned} &= -\frac{1}{2}\frac{\partial \rho u_i u_i u_j}{\partial x_j} - \frac{\partial u_j P}{\partial x_j} + P\frac{\partial u_j}{\partial x_j} \\ &\quad + \frac{\partial u_i \tau_{ij}^d}{\partial x_j} - \tau_{ij}^d \frac{\partial u_i}{\partial x_j} \end{aligned} \tag{6.4}$$

Assuming zero net flux on the boundaries, the non-conservative terms thus correspond to body source/sink contributions to the change in kinetic energy. These contributions stem from the deviatoric stress tensor, as well as from pressure ef-

fects, their average global quantities being defined as,

$$\begin{aligned}
\epsilon_1 &= 2\frac{\mu}{\rho_o}\frac{1}{\Omega}\int_\Omega \tau_{ij}^d \tau_{ij}^d d\Omega \ , \\
\epsilon_2 &= \frac{\mu_v}{\rho}\frac{1}{\Omega}\int_\Omega (\partial_{x_i} u_i)^2 d\Omega \ , \\
\epsilon_3 &= -\frac{1}{\rho_o \Omega}\int_\Omega P \cdot (\partial_{x_i} u_i)\ d\Omega \ .
\end{aligned} \tag{6.5}$$

Assuming low compressibility and zero bulk viscosity (i.e., $\mu_\nu = 0$), the dissipation to the kinetic energy equation is thus contained within ϵ_1 in Equation 6.5 above. In LES studies of the TGV, however, this quantity is seldom analyzed. And if it is, it is judged relative to the strain-rate-based dissipation of the fully-resolved DNS field, rather than a consistently filtered reference solution. This, however, presents a biased perspective as the metric measures small-scale activity, which cannot be equivalent to that of the full field for the case of under-resolved flows.

Keeping with the task of evaluating explicit LES formulations with filter-to-grid ratios above unity, the following analysis will thus utilize both metrics ($\epsilon(E_k)$ and ϵ_1) in assessment of the overall algorithm. Looking to the previous Burgers example as an antecedent to the current TGV test case with the Navier-Stokes equations, the effects of numerical discretization, modeling, and the filter formulation on the overall LES accuracy are studied.

6.1.1 Compressible Explicit LES Formulation

As the following investigation considers the explicitly-filtered LES formulation (e.g., $\bar{\tilde{\phi}}$), it is important to define the specific form of the filtered compressible

equations, along with the implied closure contributions:

$$\begin{aligned}
\partial_t \bar{\tilde{\rho}} + \partial_{x_j} \bar{\tilde{\rho}} \left[\widetilde{\overline{u_j^\star}} \right] &= 0 \ , \\
\partial_t \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] + \partial_{x_j} \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_j^\star}} \right] &= -\partial_{x_i} \left[\widetilde{\overline{P'}} \right] + \partial_{x_j} \left[\widetilde{\overline{\tau_{ij}^{d,\prime}}} \right] - \partial_{x_j} \left(\mathcal{T}_{ij}^{\rho uu} - \mathcal{T}_{ij}^{\tau^d} \right) , \\
\partial_t \bar{\tilde{\rho}} \left[\widetilde{\overline{e_o^\star}} \right] + \partial_{x_j} \bar{\tilde{\rho}} \left[\widetilde{\overline{e_o^\star}} \right] \left[\widetilde{\overline{u_j^\star}} \right] &= -\partial_{x_j} \left[\widetilde{\overline{u_j^\star}} \right] \left[\widetilde{\overline{P'}} \right] + \partial_{x_j} \left[\widetilde{\overline{u_k^\star}} \right] \left[\widetilde{\overline{\tau_{kj}^{d,\prime}}} \right] \\
&\quad + \partial_{x_j} \left[\widetilde{\overline{q'}} \right] \\
&\quad - \partial_{x_j} \left(\mathcal{T}_j^{\rho e_o u} + \mathcal{T}_j^{uP} - \mathcal{T}_j^{u\tau^d} - \mathcal{T}_j^{q} \right) \ .
\end{aligned} \tag{6.6}$$

As summarized by Garnier et al. (2009), the choice of the filtered compressible equations can vary based on how one chooses to handle the thermodynamic relations, notably the definitions of the filtered pressure and temperature. By redefining these to be,

$$\left[\widetilde{\overline{P'}} \right] = (\gamma - 1) \left\{ \bar{\tilde{\rho}} \left[\widetilde{\overline{e_o^\star}} \right] - \frac{1}{2} \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_i^\star}} \right] - \frac{\mathcal{T}_{ii}^{\rho uu}}{2} \right\} \ , \tag{6.7}$$

$$\left[\widetilde{\overline{T'}} \right] = \frac{(\gamma - 1)}{R} \left\{ \left[\widetilde{\overline{e_o^\star}} \right] - \frac{1}{2} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_i^\star}} \right] - \frac{\mathcal{T}_{ii}^{\rho uu}}{2\bar{\tilde{\rho}}} \right\} \ . \tag{6.8}$$

The above then imply that Equation 6.6 solves for the filtered total energy quantity (rather than the resolved total energy),

$$\bar{\tilde{\rho}} \left[\widetilde{\overline{e_o^\star}} \right] = \frac{\left[\widetilde{\overline{P'}} \right]}{\gamma - 1} + \frac{1}{2} \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_i^\star}} \right] + \frac{\mathcal{T}_{ii}^{\rho uu}}{2} \ . \tag{6.9}$$

This choice, however, does not affect the equation of state $\left[\widetilde{\overline{P'}} \right] = \bar{\tilde{\rho}} R \left[\widetilde{\overline{T'}} \right]$. As the following LES investigations of the Taylor-Green vortex will mainly consider the kinetic energy dissipation rate, one can form the associated resolved-scale

equation in an analogous manner to the unfiltered case of Equation 6.4:

$$\begin{aligned}\frac{\partial \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_i^\star}} \right]}{\partial t} &= -\frac{1}{2}\frac{\partial \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_j^\star}} \right]}{\partial x_j} \\ &\quad -\frac{\partial \left[\widetilde{\overline{u_j^\star}} \right] \left[\widetilde{\overline{P'}} \right]}{\partial x_j} + \left[\widetilde{\overline{P'}} \right] \frac{\partial \left[\widetilde{\overline{u_j^\star}} \right]}{\partial x_j} \\ &\quad +\frac{\partial \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{\tau_{kj}^{d,\prime}}} \right]}{\partial x_j} - \left[\widetilde{\overline{\tau_{kj}^{d,\prime}}} \right] \frac{\partial \left[\widetilde{\overline{u_i^\star}} \right]}{\partial x_j} \\ &\quad +\frac{\partial \left[\widetilde{\overline{u_i^\star}} \right] \mathcal{T}_{ij}^{mom}}{\partial x_j} - \mathcal{T}_{ij}^{mom} \frac{\partial \left[\widetilde{\overline{u_i^\star}} \right]}{\partial x_j} \, .\end{aligned} \tag{6.10}$$

New source/sink terms thus appear due to the closure terms, with $\mathcal{T}_{ij}^{mom} = \mathcal{T}_{ij}^{\rho uu} + \mathcal{T}_{ij}^{\tau^d}$; these contributions correspond to the transfer of energy relative to the resolved scales. The equation further emphasizes the deviation between $\epsilon(E_k)$ and ϵ_1 as the effects of filtering – and thus the closure contributions – become increasingly important.

The above equations are largely identical to traditional compressible LES equations; however, additional care is made in expressing the LES and grid-filter operators, $\mathcal{G}_{\bar{\Delta}}$ and $\mathcal{G}_{\tilde{\Delta}}$ respectively. The choice of density-weighted quantities (Favre averaging) is invoked (e.g., $\left[\widetilde{\overline{u_i^\star}} \right] = \widetilde{\overline{\rho u}_i}/\bar{\tilde{\rho}}$, while $\widetilde{\left[\widetilde{\overline{u_i^\star}} \right]} = \mathcal{G}_{\tilde{\Delta}}\{\widetilde{\rho u}_i/\tilde{\rho}\}$) in order to avoid model contributions in the density equations and render closure treatment more akin to an incompressible system (i.e., quadratic non-linearities with respect to the convective terms). Furthermore, quantities denoted by a prime (e.g., P') are assumed to be calculated with the corresponding Favre-averaged quantities.

The related closures are thus defined as,

$$
\begin{aligned}
\mathcal{T}_{ij}^{\rho uu} &= \widetilde{\overline{\rho u_i u_j}} - \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_j^\star}} \right] = \bar{\tilde{\rho}} \left(\widetilde{\overline{u_i u_j}} - \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_j^\star}} \right] \right) \\
\mathcal{T}_{i}^{\rho e_o u} &= \bar{\tilde{\rho}} \left(\widetilde{\overline{e_o u_i}} - \left[\widetilde{\overline{e_o^\star}} \right] \left[\widetilde{\overline{u_i^\star}} \right] \right) \\
\mathcal{T}_{ij}^{\tau^d} &= \widetilde{\overline{\tau_{ij}^d}} - \left[\widetilde{\overline{\tau_{kj}^{d,\prime}}} \right] \\
\mathcal{T}_{i}^{uP} &= \widetilde{\overline{u_i P}} - \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{P^\prime}} \right] \\
\mathcal{T}_{i}^{u\tau^d} &= \widetilde{\overline{u_k \tau_{kj}^d}} - \widetilde{[\overline{u_k^\star}]} \left[\widetilde{\overline{\tau_{kj}^{d,\prime}}} \right] \\
\mathcal{T}_{i}^{q} &= \bar{\tilde{q}} - \left[\widetilde{\overline{q^\prime}} \right]
\end{aligned}
\tag{6.11}
$$

The model contributions are thus seen to account for both density-weighting as well as the non-linearities, which may place additional burdens on the model treatment[1]. Velocity-filtered LES formulations that explicitly take density closure effects into considerations have recently been proposed (GS et al., June 2016) as a possible alternative for representing compressibility effects in LES.

$$
\mathcal{T}_{ij}^{\rho uu,RSFS} = \widetilde{\overline{\tilde{\rho} \left[\widetilde{u_i^\star} \right] \left[\widetilde{u_j^\star} \right]}} - \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_j^\star}} \right] \tag{6.12}
$$

$$
\begin{aligned}
\mathcal{T}_{ij}^{\rho uu,SSM} &= C_{ss} \left\{ \widehat{\alpha} - \hat{\bar{\tilde{\rho}}} \left[\widehat{\widetilde{\overline{u_i^\star}}} \right] \left[\widehat{\widetilde{\overline{u_j^\star}}} \right] \right\}, \\
&\quad \text{where } \alpha = \bar{\tilde{\rho}} \left[\widetilde{\overline{u_i^\star}} \right] \left[\widetilde{\overline{u_j^\star}} \right] .
\end{aligned}
\tag{6.13}
$$

The diverse set of physical mechanism that need to be represented can pose a great challenge in terms of model selection. However, employing a scale-similarity perspective provides great generality in the representation of the respective closures. The Approximate Deconvolution Method (ADM) and the Bardina-type Scale-Similarity Model (SSM) (where $C_{ss} = 1$ and $\mathcal{G}_{\hat{\Delta}=\bar{\Delta}}$ in Equation 6.13) are thus considered. Examples of their respective implementation as relating to $\mathcal{T}_{ij}^{\rho uu}$

[1]Note that the pressure work closure $\mathcal{T}_i^{uP}$ is composed of $\mathcal{T}_i^{\rho e_o u}$ and will also require modeling of a quadruple correlation term $\widetilde{\overline{\rho u_i u_i u_j}}$. These complexities may be avoided by considering the Vreman System 1 that is based on the resolved total energy (Vreman et al., 1995; Garnier et al., 2009), although this then requires modeling of additional non-conservative terms that may cause issues in the case of flow discontinuities.

is given in Equations 6.12-6.13 and provide a general guideline for extension to the other model terms[2]. As previously elaborated in Chapter 5, the ADM procedure is believed to provide non-linear stabilization through filtering of the non-linear terms (akin to residual filtering) (Layton and Neda, 2007a), while the SSM implementation is suspected to suffer from high-wavenumber instability due to an implicit velocity increment contribution (see Equation 5.14).

In the case of ADM, two levels of deconvolution fidelity are considered based on inverting a re-scaled or preconditioned version of the original discrete filter stencil (see Appendix A). Written in one-dimension, this corresponds to the following:

$$\begin{aligned}\left[\epsilon_{IF,0}+\sum_{\ell=1}^{L}\epsilon_{IF,2\ell}(\Delta x)^{2\ell}\delta_x^{2\ell}\right]\{\bar{\tilde{\phi}}\} &= \epsilon_{EF,0}\tilde{\phi}+\mu\sum_{r=1}^{R}\epsilon_{EF,2r}(\Delta x)^{2r}\delta_x^{2r}\{\tilde{\phi}\} \\ &\quad +(1-\mu)\sum_{\ell=1}^{L}\epsilon_{IF,2\ell}(\Delta x)^{2\ell}\delta_x^{2\ell}\{\tilde{\phi}\}\ , \\ &\quad \text{where } \mu=(1-1/\kappa_{cond})\in[0,1]\ .\end{aligned} \tag{6.14}$$

Multi-dimensional deconvolution is then achieved by successive application of the procedure in each direction. The accuracy of the deconvolution is controlled by the condition number κ_{cond} with higher values yielding more faithful retrieval of the original signal. On the other hand, the case of $\kappa_{cond}=1$ corresponds to a zeroth-order estimate, where $\tilde{\phi}\approx\bar{\tilde{\phi}}$.

6.1.2 Numerical Set-up

The initialization of the compressible Taylor-Green vortex consists of a low-wavenumber perturbation to an otherwise uniform flow on a triply periodic domain $-\pi L\leq x,y,z\leq\pi L$:

[2]Recall that estimating the closure term with its resolve sub-filter stress (RSFS) component is algorithmically equivalent to the ADM method.

$$
\begin{aligned}
u &= V_o \sin\left(\tfrac{x}{L}\right) \cos\left(\tfrac{y}{L}\right) \cos\left(\tfrac{z}{L}\right) \ , \\
v &= -V_o \cos\left(\tfrac{x}{L}\right) \sin\left(\tfrac{y}{L}\right) \cos\left(\tfrac{z}{L}\right) \ , \\
w &= 0 \ , \\
P &= P_o + \tfrac{\rho_o V_o^2}{16} \left[\cos\left(\tfrac{2x}{L}\right) + \cos\left(\tfrac{2y}{L}\right)\right] \left[\cos\left(\tfrac{2z}{L}\right) + 2\right]
\end{aligned}
\tag{6.15}
$$

The problem is then characterized by the following non-dimensional parameters: $\mathrm{M}_o = (V_o/c_o) = 0.1$, $\mathrm{Re} = (\rho_o V_o L)/\mu = 1600$, $\mathrm{Pr} = (\mu c_p)/\kappa = 0.71$, and the specific heat ratio, $\gamma = c_p/c_v = 1.4$. The speed of sound c_o corresponds to the temperature $T_o = P_o/R\rho_o$, the initial temperature taken to be uniform $T = T_o$ and the density taken as $\rho = P/RT_o$. The span of the computation is measured relative to the characteristic convective time $t_c = L/V_o$ and is considered to end at $t_{final} = 20t_c$. As previously mentioned, assessment of the flow evolution is made relative to the kinetic energy dissipation parameters (Equations 6.2 and 6.5). However, here the compressibility effects are minimal and the bulk viscosity μ_ν is taken to be zero; thus ones has $\epsilon_2 = 0$ and $\epsilon_3 \approx 0$. Instead, we turn primarily to $\epsilon(E_k)$ and ϵ_1 in order to characterize the flow. As was the case for the Burgers test case, these two metrics are equivalent in the case of fully-resolved computations, in correspondence with the Equation 6.4.

While results are most commonly compared to spectral computations based on the incompressible equations (Brachet et al., 1983), recent evidence suggests that the compressible equations converge to a a slightly different solution (Atkins, June 2016). As a result, a full-resolution DNS compressible computation is carried out in order to get a baseline for the flow characteristics. In each case, a fourth-order standard central finite difference discretization is employed for the inviscid terms, while a second-order narrow-stencil is utilized for the diffusive terms. A third-order Runge-Kutta temporal scheme serves as the integration method.

Typically, three phases of the problem are identified as shown in Figure 6.1: vortical roll-up ($t \sim 5t_c$), coherent breakdown of the vortices ($t \sim 9t_c$), and tur-

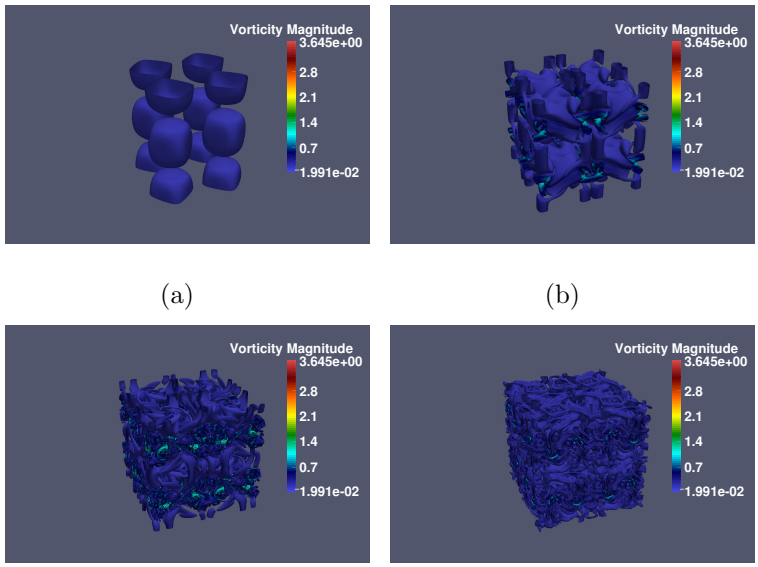

Figure 6.1: Demarcation of vortical structure relative to Q criterion (level = 8×10^{-4} $[1/s^2]$) colored by vorticity magnitude $[1/s]$ through different stages of the Taylor-Green vortex evolution cycle: (a) initial condition ($t_c = 0$), (b) vortical roll-up ($t_c = 5$), (c) coherent breakdown ($t_c = 9$) and (d) turbulent decay ($t_c = 15$).

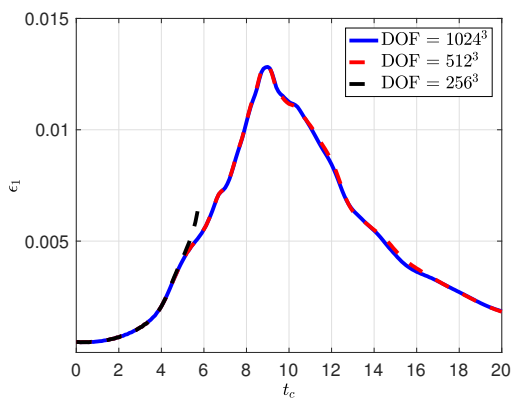

Figure 6.2: Kinetic energy dissipation rate for DNS (i.e., non-model) simulations as a function of grid resolutions.

bulent decay ($t > 9t_c$). Figure 6.2 plots the $-dEk/dt$ for different resolutions of the DNS (i.e., no modeling). The dissipation rate is seen to increase and peak at $t \sim 9t_c$, after which the rate decreases, signaling the turbulent decay-like process. It is apparent that the $DOF = 512^3$ grid is in good agreement with the $DOF = 1024^3$ grid, thus suggesting that the former is sufficiently converged. There is some minor discrepancy during $t > 10t_c$; however, this may be linked to the second-order diffusive terms under-performing during the decay process. At the $DOF = 256^3$ resolution, the viscous effects are not sufficiently resolved and numerical errors cause the simulation to go unstable[3]. At sufficiently low resolutions, one can assume that proper computation of the solution requires the inclusion of proper LES closures. Other studies have considered the computation of under-resolved Taylor-Green vortex simulations from the standpoint of employing additional scheme stabilization (e.g., artificial dissipation); these essentially correspond to implicit LES (ILES) implementation. Instead, the current study considers explicit modeling, relating to the scale-similarity type closures introduced in Section 5.3 (ADM and SSM). In addition, the filter-based artificial dissipation is also considered; and while its use might be categorized as implicit modeling, the conscious choice of having the dissipation be consistent with the LES filter $\mathcal{G}_{\bar{\Delta}}$ makes the approach more similar to a hyper-viscosity type approach (i.e., higher diffusive-type derivatives). The computations henceforth consider an LES grid of $DOF = 128^3$ with $\bar{\Delta} = [4\Delta x_{LES}, 8\Delta x_{LES}]$, which correspond to $DOF = 64^3$ and 32^3, respectively[4]. A Tangent discrete filter introduced in Section 2.4.2 is considered for the explicit LES filter $\mathcal{G}_{\bar{\Delta}}$.

[3]The current computations employ a divergence form of the convective terms. Using a skew-symmetric or other splittings (Kennedy and Gruber, 2008; Morinishi, 2010; Pirozzoli, 2010) can extend the stability of the simulation for lower resolutions such as $DOF = 256^3$, although the dissipation rates become erroneous.

[4]Note, 2 DOF are required per Δx for finite difference methods. Here, Δ is taken to be a full wavelength, rather than a half wavelength.

6.2 A Priori Assessments

As the LES filter is explicitly defined, it is possible to tractably derive reference solutions by which subsequent evaluations of the methods can be judged. This is presented here as part of an *a priori* analysis of both the LES field and the anticipated model contributions. Consequences of the filter operator (e.g., the filter order, the filter-to-grid ratio, etc ...) are inspected.

Reference solutions for the LES computations are derived by filtering the DNS solution in a three-step process:

1. The DNS field ($DOF_{DNS} = 512^3$) is filtered to the LES resolution ($DOF_{LES} = 128^3$), $\tilde{\phi}^*_{DNS} = \mathcal{G}_{\bar{\Delta}}\{\phi_{DNS}\}$, by using a tenth-order Tangent filter with $|\hat{\mathcal{G}}(k_c)| = 0.01$, such that $(k_c\Delta x) = 0.25\pi$.

2. The intermediate variable $\tilde{\phi}^*$ is then projected onto the LES grid via comb filtering in physical space (which presumes the LES grids to be collocated), wherein $\tilde{\phi}_{LES}(x_i) = \tilde{\phi}^*_{DNS}(x_i)$.

3. The final LES solution is then derived by employing the explicit filter on the LES grid variable, $\bar{\tilde{\phi}} = \mathcal{G}_{\bar{\Delta}}\{\tilde{\phi}_{LES}\}$, based on a $|\hat{\mathcal{G}}(k_{1/2})| = 0.5$ characterization.

As previously suggested, higher-order LES filter formulations are able to preserve more of the large-scale features above the prescribed cut-off. Such attenuation not only affects the range of scales present in the LES solution, but is also expected to affect the perceived kinetic energy dissipation rate curves. Figures 6.3 and 6.4 include both $\epsilon(E_k) = -d(Ek)/dt$ and ϵ_1 metrics for the different discrete filters tuned to the filter-to-grid ratios (FTGR) of 2 and 4 ($\bar{\Delta} = 4\Delta x_{LES}$ and $8\Delta x_{LES}$, respectively).

In general, one notices that the $-dEk/dt$ metric of the LES is quite similar to that of the DNS curve. This measurement, typically employed in assessments

of the Taylor-Green vortex, tends to focus on large-scale energy depletion[5]. Similarities with the DNS curve are supportive of the fact that the LES solution can replicate global characteristics of the fully-resolved field. Employing filter-to-grid ratios above unity is seen to shift the peak dissipation rate to an earlier time. The peak rate coincides with the time at which the energy spectrum is fullest; therefore, lower wavenumber cut-offs reach this point earlier on. Also, it is seen that sharper (i.e., higher-order) filter formulations have an increased peak owing to a more sudden transition of the dynamics past the LES cut-off; by contrast, the smoother filter (i.e., low-order) has a more gradual transfer of flow dynamics into the sub-filter range of scales, which causes the dissipation rate to appear more spread out and subdued.

Unlike the $\epsilon(E_k)$ assessment, the strain-rate-based dissipation component ϵ_1 for the LES solutions in Figure 6.4 is shown to be quite different from the corresponding DNS curve. Because this metric essentially focuses on the amount of small-scale activity, it is understandable that higher levels of attenuation – either with respect to smoother filter responses or with respect to higher filter-to-grid ratios – would yield lower ϵ_1 values. In the case of fully-resolved flow, however, $\epsilon(E_k) = \epsilon_1$ (assuming compressibility effects are negligible) which is evidenced by comparison of the DNS curves. Because ϵ_1 highlights high wavenumber activity, it is useful to look at in order to assess whether the LES computation is faithfully enforcing the expected LES field or whether an implementation may be susceptible to small-scale errors. As can be shown, it is possible for LES computations to yield decent $\epsilon(E_k)$ agreement while being grossly inaccurate with respect to ϵ_1. Consequently both metrics will be considered henceforth in the characterization of the LES performance.

As alluded to by the resolved kinetic energy equation in 6.10, the discrepancy

[5]Here, $\epsilon(E_k)$ is calculated by differentiating E_k with a standard central second-order discretization. Higher-order renditions are possible, but this second-order approximation uses a resolution of $\Delta T = 0.1t_c$ and is believed to be sufficient.

between $\epsilon(E_k)$ and ϵ_1 is encapsulated by the closure-induced dissipative effects which may be estimated by looking at $\epsilon_{\mathcal{T}} = \epsilon(E_k) - \epsilon_1$. The curve plotted in Figure 6.5 represents this contribution as assumed from an exact closure. As expected, effects on the dissipation rate due to the closures is negligible for the DNS calculation. As more of the scales are removed from the continuous field, however, these effects increase. In this way, the case of the larger filter-to-grid ratio results in high values of $\epsilon_{\mathcal{T}}$; furthermore, non-sharp filter formulations such as the second-order Tangent stencil also suggest increased model activity, likely to be tied to the implied notion of increased model input towards the large scales. Also interesting to note is the increased activity of the models towards the peak dissipation rate, wherein the energy spectrum is expected to be at its fullest. Comparing the $\epsilon_{\mathcal{T}}$ and ϵ_1 curves, it furthermore becomes evident that most of the variability in $\epsilon(E_k)$ results from the model contribution. This may be due to the way in which the *a priori* reference solution is calculated: the DNS field is filtered before computing ϵ_1, which may remove oscillations. Alternatively, this observed result may point to real intermittencies of the flow that get manifested through the model terms. The ability to identify interesting characteristics of the flow, such as these intermittencies, supports the need to analyze different metrics and the attempt to reconstruct the underlying mechanisms associated with the LES methodology.

Figure 6.6 shows the evolution of the actual kinetic energy $E_k(t)$. It is apparent that for the high-order filters tuned to $FTGR = 2$, a majority of the energy content is still resolved on the LES support. However, as more of the spectrum is removed – by employing less sharp filter formulations or considering higher filter-to-grid ratios – less of the energetic content is accounted for and one's reliance on proper modeling is expected to increase.

Before proceeding to incorporate the respective effects of the numerical discretization on the LES computation, a preliminary assessment of the ADM and

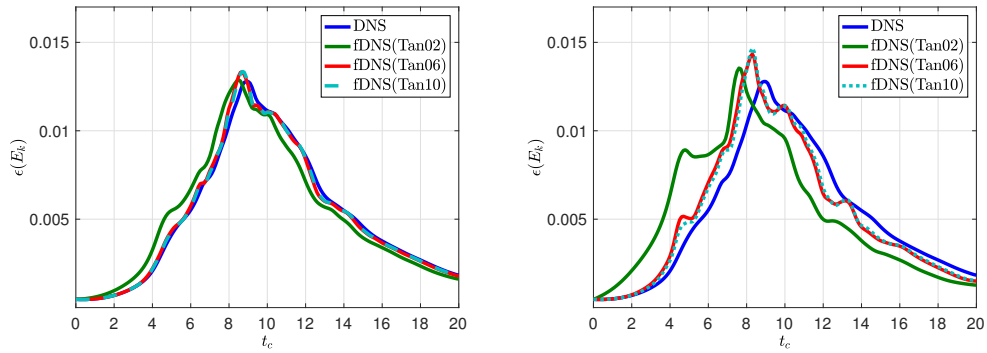

(a) (b)

Figure 6.3: Kinetic energy dissipation rate as measured by $\epsilon(E_k)$ for DNS and reference LES solutions (i.e., filtered DNS) as a function of the explicit LES filter (discrete Tangent stencils): a) $FTGR = 2$ and b) $FTGR = 4$.

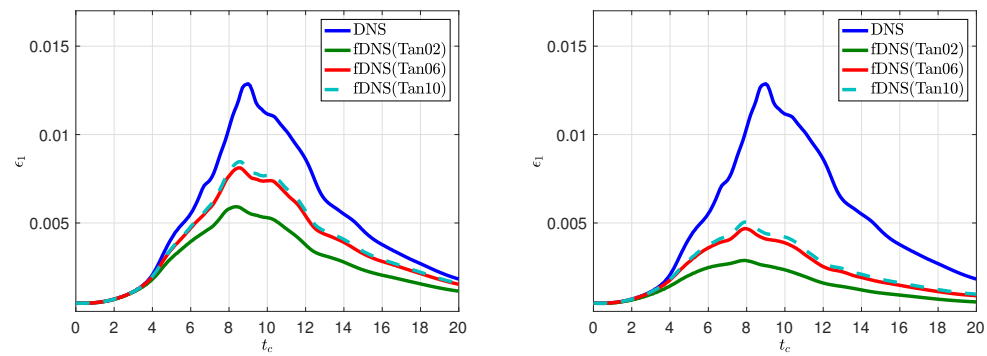

(a) (b)

Figure 6.4: Kinetic energy dissipation rate as measured by ϵ_1 for DNS and reference LES solutions (i.e., filtered DNS) as a function of the explicit LES filter (discrete Tangent stencils): a) $FTGR = 2$ and b) $FTGR = 4$.

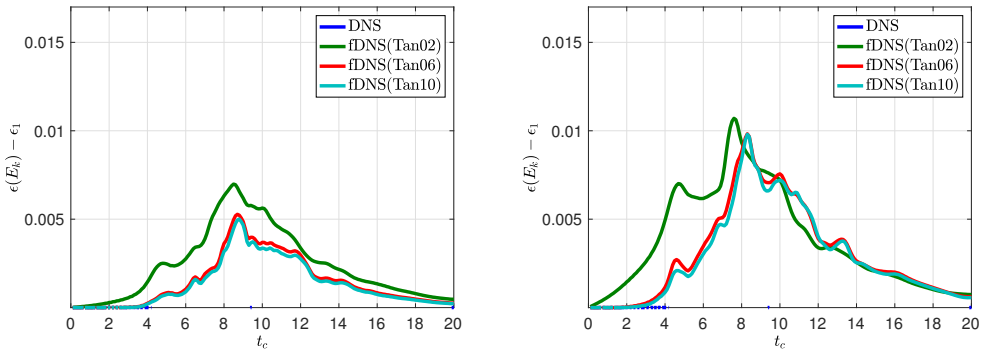

(a) (b)

Figure 6.5: Kinetic energy dissipation rate as accounted for by the closure model $\epsilon_{\mathcal{T}}$ for DNS and reference LES solutions (i.e., filtered DNS) as a function of the explicit LES filter (discrete Tangent stencils): a) $FTGR = 2$ and b) $FTGR = 4$.

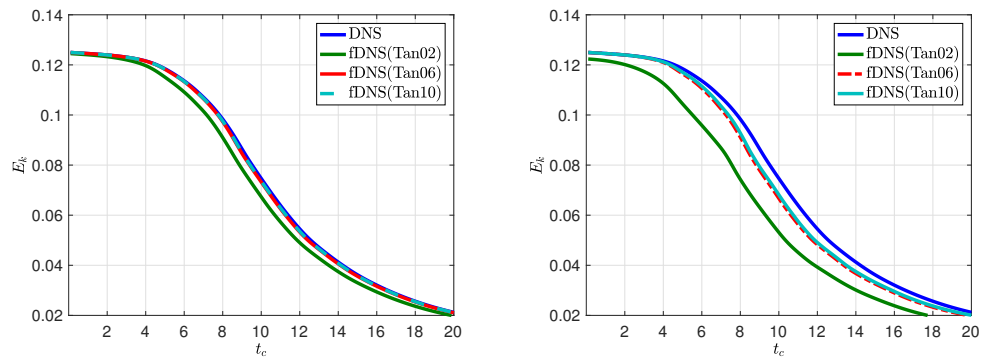

(a) (b)

Figure 6.6: Kinetic energy E_k in domain for DNS and reference LES solutions (i.e., filtered DNS) as a function of the explicit LES filter (discrete Tangent stencils): a) $FTGR = 2$ and b) $FTGR = 4$.

SSM models can be carried out by filtering the DNS data and forming the different closure terms (see Equations 6.11). Access to the "exact" DNS solution then also allows the formation of "perfect" closure terms (Stefano and Vasilyev, 2002; DeStefano and Vasilyev), which can be used as a baseline for evaluating the accuracy of the different closures. This corresponds to an *a priori* evaluation of the model terms as applied to the different non-linear physical processes.

Table 6.1 estimates the model closure contributions to the resolved kinetic energy dissipation rate at $t = 9tc$, comparing the ADM and SSM models to the exact model ($\epsilon_{\mathcal{T}^{model}}/\epsilon_{\mathcal{T}^{exact}}$). Referring to Equation 6.10, the ADM and SSM dissipation is constructed as

$$\epsilon_{\mathcal{T}^{model}} \sim \mathcal{T}_{ij}^{\rho uu,model} \partial_{x_j} \left[\widetilde{\overline{u_i^\star}} \right] \ , \tag{6.16}$$

owing to the fact that dilation and viscous closure effects are known to be minimal for the current test case. As before, the exact model is calculated from the filtered reference solution and is formed by $\epsilon_{\mathcal{T}^{exact}} = (\epsilon(E_k) - \epsilon_1)$. The table reveals tendencies of the models as corresponding to the different filter-to-grid ratios as well as to the explicit LES filter. This information is supplemented by correlation information $\rho_{corr}\{\epsilon_{\mathcal{T}^{model}}, \epsilon_{\mathcal{T}^{exact}}\}$ in Table 6.3. In combining these two perspectives, it may be possible to judge whether the models adhere to the general physical trends and whether they are sufficiently active or not. In viewing this information, one notices that for a given filter-to-grid ratio, increasing the order of the filter (i.e., its sharpness), results in an increase of both correlation and magnitude of the model's dissipation $\epsilon_{\mathcal{T}^{model}}$ relative to the exact closure. Also, it is apparent that the ADM consistently outperforms the current SSM interpretation. The same trends hold for the higher filter-to-grid ratio, although overall performance is seen to diminish somewhat. This is surprising as it conflicts with the *a priori* analysis of Section 5.2 which predicts higher correlations with increased filter-to-grid ratios, at least with respect to the ADM procedure (see Figures 5.14(a)-5.16(a)).

		$FTGR = 2$			$FTGR = 4$		
		Tan02	Tan06	Tan10	Tan02	Tan06	Tan10
$\epsilon'_{\mathcal{T}^{\rho uu}}$	ADM	0.34	0.65	0.72	0.30	0.62	0.70
	SSM	0.25	0.56	0.65	0.20	0.53	0.64

Table 6.1: Normalized kinetic energy closure contributions, $\epsilon'_{\mathcal{T}} = (\epsilon_{\mathcal{T}^{model}}/\epsilon_{\mathcal{T}^{exact}})$, at $t = 9t_c$ as calculated via the filtered DNS solution.

		$FTGR = 2$			$FTGR = 4$		
		Tan02	Tan06	Tan10	Tan02	Tan06	Tan10
$[\partial_{x_j}\mathcal{T}_{1j}^{\rho uu}]'$	ADM	0.31	0.63	0.69	0.23	0.63	0.70
	SSM	0.20	0.53	0.62	0.13	0.52	0.62

Table 6.2: Normalized coefficient of model contribution, $[\partial_{x_j}\mathcal{T}_j^{model}]' = (|\partial_{x_j}\mathcal{T}_j^{model}|_2/|\partial_{x_j}\mathcal{T}_j^{exact}|_2)$, for the ADM and SSM closure models at $t = 9t_c$ as calculated via the filtered DNS solution.

And while the current efforts are focused on properly representing the dissipation rate characteristics of the TGV, it is also perhaps illuminating to briefly consider the individual closure performances for the different physical non-linearities.

Focusing on the dissipation rate characteristics of the TGV, performance of the convective momentum closure force $\partial_{x_j}\mathcal{T}_1^{\rho uu}$ is considered in Tables 6.2 and 6.4, which show correlation coefficient and magnitude information at $t = 9t_c$. The momentum closure force is seen to abide by the same trends as the dissipation rate contribution: 1) the ADM treatment outperforms the SSM implementation, 2) increasing the filter order improves results, and 3) the same trends hold for higher filter-to-grid ratios but with a slight overall decrease in accuracy[6].

[6]Inspecting the actual closure ($\mathcal{T}$) rather than its force or divergence ($\partial_x\mathcal{T}$), would provide a more structural rather than functional assessment of the model performance. Such correlations are provided in Table 1 of Edoh and Karagozian (June 2017b) for the scalar $\mathcal{T}_{ij}^{\rho uu}\mathcal{T}_{ij}^{\rho uu}$, revealing slightly higher correlations and improved performance with a larger filter-to-grid ratio.

		$FTGR = 2$			$FTGR = 4$		
		Tan02	Tan06	Tan10	Tan02	Tan06	Tan10
$\rho_{corr}\{\epsilon_{\mathcal{T}^{\rho uu}}\}$	ADM	0.62	0.79	0.81	0.65	0.70	0.77
	SSM	0.57	0.74	0.78	0.56	0.64	0.73

Table 6.3: Correlation coefficient, $\rho_{corr}\{\epsilon_{\mathcal{T}^{model}}, \epsilon_{\mathcal{T}^{exact}}\}$, of kinetic energy dissipation rate due to modeling at $t = 9t_c$ as calculated via the filtered DNS solution.

		$FTGR = 2$			$FTGR = 4$		
		Tan02	Tan06	Tan10	Tan02	Tan06	Tan10
$\rho_{corr}\{\partial_{x_j}\mathcal{T}_{1j}^{\rho uu}\}$	ADM	0.45	0.77	0.80	0.46	0.72	0.79
	SSM	0.37	0.72	0.76	0.37	0.68	0.76

Table 6.4: Correlation coefficient, $\rho_{corr}\{\mathcal{T}^{model}, \mathcal{T}^{exact}\}$, for the ADM and SSM closure models at $t = 9t_c$ as calculated via the filtered DNS solution.

6.3 A Posteriori Assessments: ADM

Having established a baseline for the behaviors associated with the choice of model and the LES formulation (e.g., filter-to-grid ratio, filter stencil), it is now desirable to take into account the discretization effects. Here, focus is kept on the tenth-order Tangent stencil as the explicit filter because of its scale-discriminant characteristics. Employing it as the explicit LES filter thus allows one to resolve more of the spectral content that is representable by the grid; this makes the LES computation more efficient in terms of the prescribed filter-to-grid ratio and is also expected to make the benefits of high-order discretizations more evident, at least based on the perfectly-modeled Burgers equation exercises of Section 5.2. In addition the ADM procedure is chosen henceforth since the emphasis of the current study is on understanding the overall LES formulation – rather than specific model performance. Nevertheless, of interest to note is the fact that all SSM com-

putations, regardless of filter order, cut-off of discretization order, were observed to go unstable near $t \sim 5t_c$ – likely a result of the velocity increment contributions stated in Equation 5.14. Similarly, all ADM procedures utilizing the high-accuracy deconvolution ($\kappa_{cond} = 1.e3$) were seen to go unstable for the second-order Tangent filter formulation – again near $t \sim 5t_c$, which corresponds to vortical roll-up and a rise in non-linear effects; this in turn suggests that the residual filtering implied by ADM must be sufficient to quell the significant re-animation of small-scale content when using high-order deconvolution procedures. In consequence to these parameters, the following analysis for the tenth-order Tangent filter is relegated to the ADM procedures with zeroth-order ($k_{cond} = 1.e0$) and a high-order deconvolution($k_{cond} = 1.e3$). In addition, the impact of filter-based artificial dissipation is assessed, where the stabilization coefficients are chosen to be consistent with the explicit LES formulation.

Figures 6.7(a) and 6.7(b) plot $\epsilon(E_k)$ and compare performances of the different discretizations for both deconvolution procedures at filter-to-grid ratios of 2 and 4, respectively. Meanwhile Figures 6.8(a) and 6.8(b) show ϵ_1, while Figures 6.9(a) and 6.9(b) show the effective model dissipation $\epsilon(E_k) - \epsilon_1$. Inspecting the strain-rate-based metric (ϵ_1), one notices a stark difference between the low- and high-order deconvolution. The former, specifically, is seen to be overestimating, which suggests the presence of more small-scale activity than anticipated. This is also observed in the previous Burgers exercise. While inspecting the overall dissipation $\epsilon(E_k)$ does not reveal any particular trends, looking at the model dissipation $\epsilon(E_k) - \epsilon_1$ reveals that the low-order deconvolution can sporadically feed energy into the system (negative values of the metric). This again confirms the notion that high-wavenumber content is likely present. Interestingly, the occurrence is only associated with the high-order schemes. In general, the high-order deconvolution is seen to perform best in terms of ϵ_1, which speaks to the ADM model's ability to properly enforce explicit LES resolution. Mainly as a conse-

quence to this, the ADM also does a better job at representing the effective model dissipation ($\epsilon(E_k)-\epsilon_1$). In terms of discretization effects, its results further benefit from the use of high-order schemes (at least for $FTGR = 2$).

	ADM ($\kappa_{cond} = 1.e0$)				ADM ($\kappa_{cond} = 1.e3$)			
	CD02	CD06	CD10	CD04-11pt	CD02	CD06	CD10	CD04-11pt
$\rho_{corr}\{\rho\}$	0.79	0.87	0.87	0.87	0.75	0.93	0.92	0.92
$\rho_{corr}\{\rho u\}$	0.86	0.90	0.90	0.90	0.85	0.95	0.95	0.95
$\rho_{corr}\{\rho v\}$	0.86	0.90	0.90	0.90	0.85	0.95	0.95	0.95
$\rho_{corr}\{\rho w\}$	0.72	0.86	0.86	0.86	0.74	0.92	0.91	0.92
$\rho_{corr}\{\rho e_o\}$	0.79	0.85	0.86	0.85	0.77	0.93	0.92	0.92

Table 6.5: Correlation coefficients $\rho_{corr}\{\bar{\tilde{\phi}}_{LES}, \bar{\tilde{\phi}}_{fDNS}\}$ of the working variables at $t = 9t_c$ as calculated by ADM model closures and compared relative to the reference LES solution (i.e., filtered DNS). The LES filter is a tenth-order discrete Tangent stencil tuned according to $\bar{\Delta}_{1/2}$ for $FTGR = 2$.

	ADM ($\kappa_{cond} = 1.e0$)				ADM ($\kappa_{cond} = 1.e3$)			
	CD02	CD06	CD10	CD04-11pt	CD02	CD06	CD10	CD04-11pt
$\rho_{corr}\{\rho\}$	0.60	0.60	0.60	0.60	0.85	0.93	0.92	0.92
$\rho_{corr}\{\rho u\}$	0.78	0.75	0.74	0.74	0.91	0.95	0.95	0.95
$\rho_{corr}\{\rho v\}$	0.78	0.75	0.74	0.74	0.91	0.95	0.95	0.95
$\rho_{corr}\{\rho w\}$	0.68	0.63	0.63	0.63	0.83	0.92	0.91	0.91
$\rho_{corr}\{\rho e_o\}$	0.67	0.63	0.63	0.63	0.87	0.93	0.93	0.92

Table 6.6: Correlation coefficients $\rho_{corr}\{\bar{\tilde{\phi}}_{LES}, \bar{\tilde{\phi}}_{fDNS}\}$ of the working variables at $t = 9t_c$ as calculated by ADM model closures and compared relative to the reference LES solution (i.e., filtered DNS). The LES filter is a tenth-order discrete Tangent stencil tuned according to $\bar{\Delta}_{1/2}$ for $FTGR = 4$.

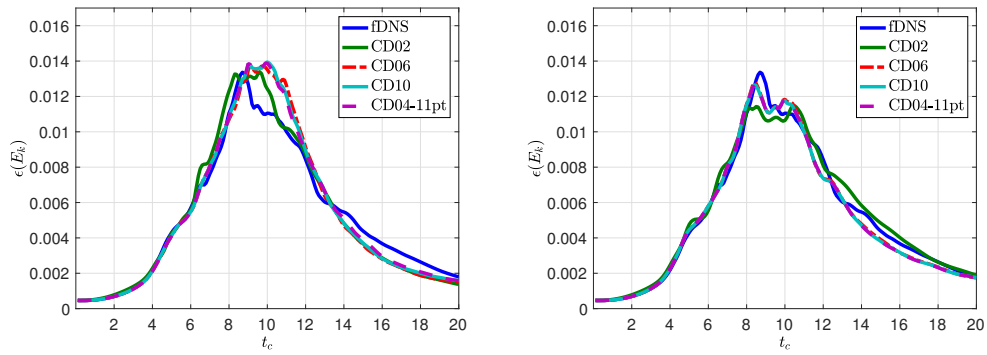

(a) (b)

Figure 6.7: Kinetic energy dissipation rate as measured by $\epsilon(E_k)$ for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$ to a filter-to-grid ratio $FTGR = 2$: a) ADM method with zeroth-order deconvolution ($\kappa_{cond} = 1.e0$) and b) ADM method with high-order deconvolution ($\kappa_{cond} = 1.e3$).

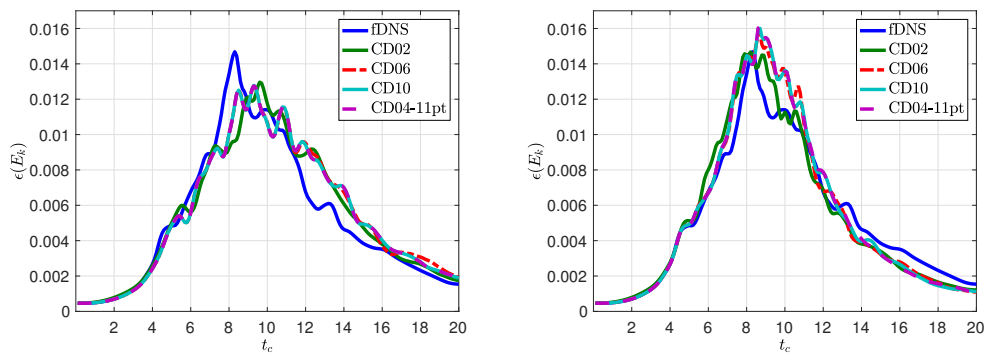

(a) (b)

Figure 6.8: Kinetic energy dissipation rate as measured by $\epsilon(E_k)$ for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$ to a filter-to-grid ratio $FTGR = 4$: a) ADM method with zeroth-order deconvolution ($\kappa_{cond} = 1.e0$) and b) ADM method with high-order deconvolution ($\kappa_{cond} = 1.e3$).

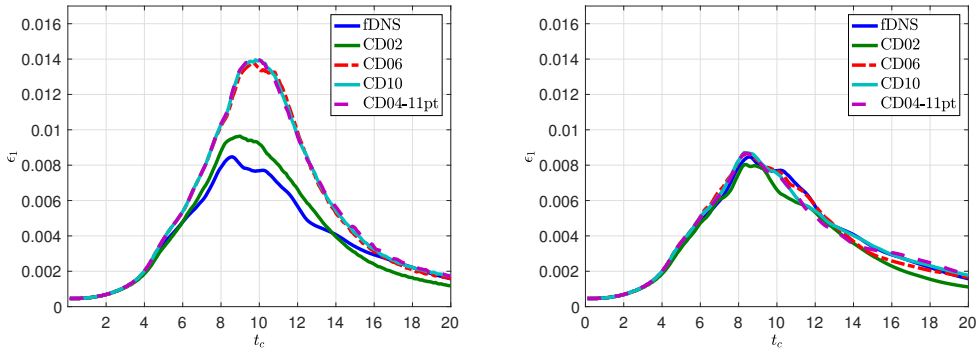

(a) (b)

Figure 6.9: Kinetic energy dissipation rate as measured by ϵ_1 for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$ to a filter-to-grid ratio $FTGR = 2$: a) ADM method with zeroth-order deconvolution ($\kappa_{cond} = 1.e0$) and b) ADM method with high-order deconvolution ($\kappa_{cond} = 1.e3$).

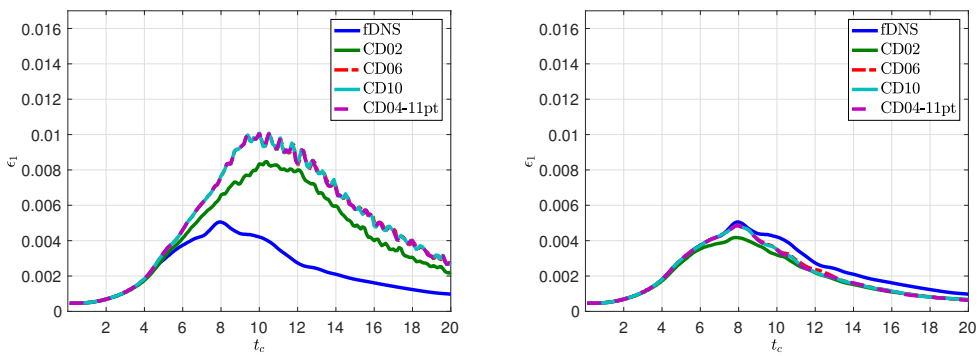

(a) (b)

Figure 6.10: Kinetic energy dissipation rate as measured by ϵ_1 for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$ to a filter-to-grid ratio $FTGR = 4$: a) ADM method with zeroth-order deconvolution ($\kappa_{cond} = 1.e0$) and b) ADM method with high-order deconvolution ($\kappa_{cond} = 1.e3$).

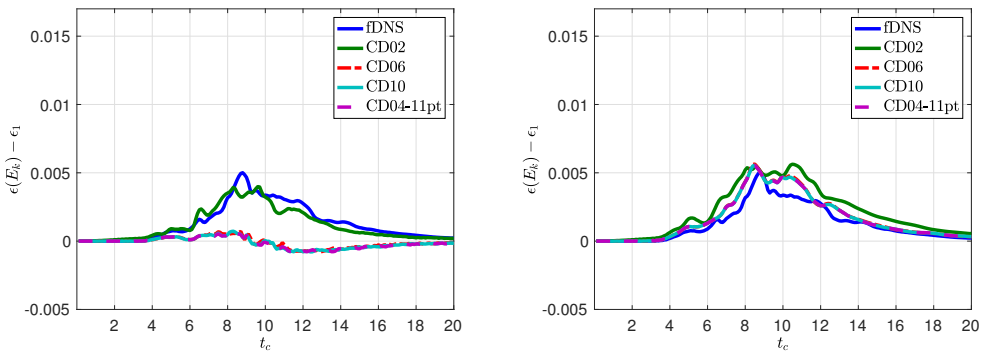

(a) (b)

Figure 6.11: Kinetic energy dissipation rate as measured by the closure model $\epsilon_{\mathcal{T}}$ for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$ to a filter-to-grid ratio $FTGR = 2$: a) ADM method with zeroth-order deconvolution ($\kappa_{cond} = 1.e0$) and b) ADM method with high-order deconvolution ($\kappa_{cond} = 1.e3$).

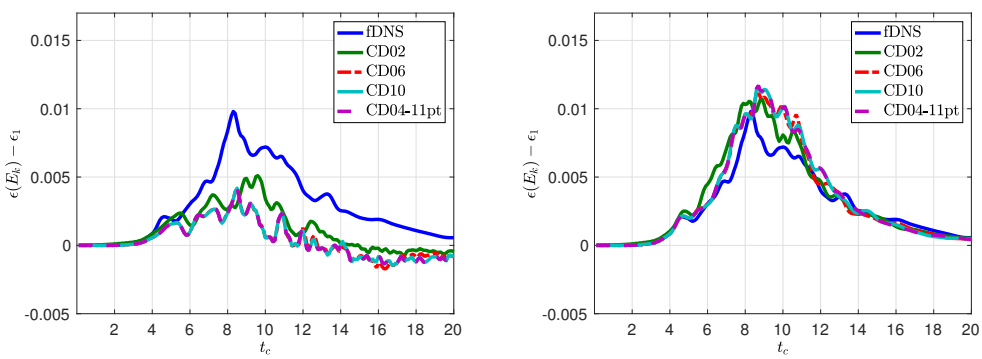

(a) (b)

Figure 6.12: Kinetic energy dissipation rate as measured by the closure model $\epsilon_{\mathcal{T}}$ for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$ to a filter-to-grid ratio $FTGR = 4$: a) ADM method with zeroth-order deconvolution ($\kappa_{cond} = 1.e0$) and b) ADM method with high-order deconvolution ($\kappa_{cond} = 1.e3$).

While the dissipation rate information focuses on energy within the flow, it is also useful to consider dispersion. This may be gathered from looking at correlation coefficients of the solution variables relative to the filtered DNS results as shown in Tables 6.5 and 6.6. In general, one would expect that increasing the discretization scheme would improve correlations by decreasing phase error, which is generally the case for $FTGR = 2$. In the case of $FTGR = 4$, however, only the high-order deconvolution is seen to maintain the anticipated trend. Once again, at the higher filter-to-grid ratio, employing high-oder discretization is seen to yield worse results with the low-order deconvolution. Employing a large filter-to-grid ratio is expected to reduce the effects of numerical error, which is only noticeable in the instance of the high-order deconvolution. These observations suggest that special precautions in the implementation of ADM are needed in order to take advantage of high-accuracy numerical schemes.

6.3.1 Impact of Stabilization

As the issue with ADM using low-order deconvolution can mainly be traced to excess high-wavenumber energy relative to the desired explicit LES resolution, it is then interesting to consider the addition of a stabilization term. Such an approach may be viewed as a type of mixed modeling. In this rendition, however, the stabilization is not employed in order to recover proper flow dynamics as is intended by the use of ADM, a structure-type model. Instead, the stabilization – which utilizes the coefficients of the explicit LES stencil in a filter-based artificial dissipation form – plays more of a phenomenological role (i.e., functional modeling) as typically done with eddy-viscosity type techniques (Sagaut, 2006). Yet, rather than necessarily seeking to represent the dissipative-like effect of the unresolved scales on the field (which, in the current explicitly-filtered case is in part handled by the ADM model) the artificial dissipation satisfies a more basic role of enforcing the intended LES spectra. Simultaneous stabilization of the

overall algorithm is presumed based on the notion that the LES filter is designed to mitigate the impact of numerical error (e.g., by use of a sufficient filter-to-grid ratio and spectral sharpness relative to the chosen discretization scheme etc).

The resulting system utilizing ADM and the filter-based stabilization is then written as the following:

$$\begin{aligned}
\partial_t \bar{\tilde{\rho}} + \partial_{x_j} \overline{\tilde{\rho}\left[\widetilde{u_j^\star}\right]} &= (C_\rho/C_o)\cdot \mathcal{D}_{\bar{\Delta}}^{AD}\left\{\bar{\tilde{\rho}}\right\} \ , \\
\partial_t \bar{\tilde{\rho}}\left[\ \widetilde{u_i^\star}\ \right] + \partial_{x_j} \overline{\tilde{\rho}\left[\widetilde{u_i^\star}\right]\left[\widetilde{u_j^\star}\right]} &= -\partial_{x_i}\overline{\left[\widetilde{P'}\right]} + \partial_{x_j}\overline{\left[\widetilde{\tau_{ij}^{d,\prime}}\right]} \\
&\quad +(C_{\rho u}/C_o)\cdot \mathcal{D}_{\bar{\Delta}}^{AD}\left\{\bar{\tilde{\rho}}\left[\widetilde{u^\star}\right]\right\} \\
\partial_t \bar{\tilde{\rho}}\left[\widetilde{e_o^\star}\right] + \partial_{x_j} \overline{\tilde{\rho}\left[\widetilde{e_o^\star}\right]\left[\widetilde{u_j^\star}\right]} &= -\partial_{x_j}\overline{\left[\widetilde{u_j^\star}\right]\left[\widetilde{P'}\right]} + \partial_{x_j}\overline{\left[\widetilde{u_k^\star}\right]\left[\widetilde{\tau_{kj}^{d,\prime}}\right]} \\
&\quad +\partial_{x_j}\overline{\left[\widetilde{q'}\right]} \\
&\quad +(C_{\rho e_o}/C_o)\cdot \mathcal{D}_{\bar{\Delta}}^{AD}\left\{\bar{\tilde{\rho}}\left[\widetilde{e_o^\star}\right]\right\} \ .
\end{aligned} \tag{6.17}$$

In this way, we see that the non-linear terms are filtered (i.e., residual filtering) due to ADM, which mitigates the accumulation of high-wavenumber content past the desired cut-off; meanwhile, the added stabilizations $\mathcal{D}_{\bar{\Delta}}^{AD}$ (the dissipative portion to the explicit LES filter, $\mathcal{G}_{\bar{\Delta}}$) are all assumed to be scaled by the spectral-radius of the inviscid flux Jacobian (i.e., scalar dissipation) and are responsible for damping content that has transgressed past the $\bar{\Delta}$ resolution. As in Section 5.3.2, this damping needs to be scaled such that it is active only when necessary. In the following, the normalizing coefficient C_o is determined similarly to Equation 5.17, assuming a spectrum that scales as $\beta = -5/3$:

$$C_o = \frac{\int_{(k_{\bar{\Delta}}\Delta x)/\pi}^{(k_{\Delta_x}\Delta x)/\pi} k^{-5/3}\ d(k\Delta x/\pi)}{\int_{(k_{\hat{\Delta}}\Delta x)/\pi}^{(k_{\bar{\Delta}}\Delta x)/\pi} k^{-5/3}\ d(k\Delta x/\pi)}. \tag{6.18}$$

Note, this also assumes that the LES formulation is based on a spectral sharp filter, as opposed to Equation 5.17 which is more general. Meanwhile the scaling

coefficient C_ϕ is determined by,

$$C_\phi = \frac{|\bar{\bar{\phi}} - \mathcal{G}_{\bar{\Delta}}\{\bar{\bar{\phi}}\}|_2}{|\mathcal{G}_{\bar{\Delta}}\{\bar{\bar{\phi}}\} - \mathcal{G}_{\hat{\Delta}}\{\bar{\bar{\phi}}\}|_2} \quad \text{such that} \quad \hat{\Delta} > \bar{\Delta}. \tag{6.19}$$

In the case of momentum, all components are based on evaluating kinetic energy ($\phi = \widetilde{\overline{[u_i^\star]}}\ \widetilde{\overline{[u_i^\star]}}$) in the testing procedure, as proposed in Stolz et al. (2001). This procedure, which is also similar to the dynamic procedure employed by Liu et al. (1995) for scaling of their Scale-Similarity model, essentially compares the amount of content in the resolved sub-filter scales (RSFS) to the content located in the band of neighboring scales defined by mid-pass range $\mathcal{G}_{\bar{\Delta}}\{\bar{\bar{\phi}}\} - \mathcal{G}_{\hat{\Delta}}\{\bar{\bar{\phi}}\}$. The normalization C_o is then used to tune the point at which the dissipation should be considered fully "on" (e.g, $C_\phi/C_o \sim 1$). While these adaptations may be labeled as dynamic, they are static in the sense that the C_o coefficient is held constant. As an example, even for a static eddy-viscosity model (e.g., $\mathcal{T}^{\rho u,\nu_e} = \nu_e \tau_{ij}^{d,\prime}$), the dissipative influence on the solution changes in time according to an eddy-viscosity $\nu_e = (C_{\nu_e}\bar{\Delta})^2(2|\overline{\tau_{ij}^{d,\prime}}^2|)^{1/2}$ that is scaled by the inner-product of strain; this is essentially responsible for gauging high-wavenumber content in similar fashion to ϵ_1, and is thus adjusted according to the flow. When the dynamic procedure of Germano is employed, wherein C_{ν_e} is determined on the fly (Germano et al., 1991), one is essentially re-adjusting the normalization coefficient C_o in order to recalibrate how strong the damping should be based on physical grounds of scale-similarity.

The Figures 6.13-6.15 demonstrate the effect of the stabilization on the ADM modeling based on low-order deconvolution. As little difference is seen amongst the high-order schemes, only the CD02 and CD06 discretizations are provided for comparison here. For both filter-to-grid ratios, the dissipation scheme is understood to have reduced high-wavenumber frequency relative to the non-stabilized case. One sees this by interpreting the ϵ_1 metric (see Figure 6.14(a) as compared to Figure 6.10(a)). In the case of $FTGR = 2$, almost perfect evolution is achieved for $t < 9t_c$, although the proceeding decay process is inaccurate. The adapting

procedure, however, does not seem to perform adequately for $FTGR = 4$; the stabilization becomes active prematurely, likely owing to issues with the choice of tuning parameters (C_ϕ/C_o). For instance, the normalization coefficient C_o of Equation 6.18 implicitly assumes a spectral sharp LES filter rather than the Tangent stencil, which would decrease the value of C_o and would thus make the dissipation overly aggressive. Also, as in all other instances, the second-order scheme CD02 is seen to reduce much of the activity – again mainly evident through inspection of ϵ_1. This is obviously an instance of the numerical scheme adding to the effective LES filter (Geurts and van der Bos, 2005a); the sub-optimal modified wavenumber characteristics of the discretization acts as an additional attenuation of the solution. As a result, the filter-like influence of the numerics effectively enhances the global attenuation properties, which then results in a larger perceived closure dissipation $\epsilon(E_k) - \epsilon_1$ relative to the higher-order schemes, as depicted in Figure 6.15(a).

Looking at the dissipation rate plots gives an idea of the energy and solution magnitude in the domain, however a second element of accuracy is the dispersion. To get an idea of these performances, one again can look at correlations as provided in Table 6.7. Here, there is a noticeable improvement due to the employment of the higher-order scheme. This is explained based on the modified wavenumber. However, once again unexpected is a general reduction in correlations when increasing the filter-to-grid ratio. This is contrary to the theory which predicts an overall reduction in numerical error through increased decoupling of the LES resolution and the computational grid.

Finally, it is helpful to observe the solution in order to query the field distribution relative to the reference LES, making discrepancies in magnitude and phase further evident. Figure 6.16 plots the kinetic energy evolution $E_k(t)$ for the different ADM implementations at $FTGR = 2$. In this particularly case, the stabilized ADM model performs the best. The high-order deconvolution is also

	$FTGR = 2$		$FTGR = 4$	
	CD02	CD06	CD02	CD06
$\rho_{corr}\{\rho\}$	0.83	0.90	0.73	0.81
$\rho_{corr}\{\rho u\}$	0.89	0.93	0.87	0.90
$\rho_{corr}\{\rho v\}$	0.89	0.93	0.87	0.90
$\rho_{corr}\{\rho w\}$	0.82	0.91	0.76	0.82
$\rho_{corr}\{\rho e_o\}$	0.83	0.89	0.74	0.83

Table 6.7: Correlation coefficients $\rho_{corr}(\phi_{LES}, \phi_{fDNS})$ of the working variables at $t = 9t_c$ as calculated by ADM model closures with adaptive filter-based artificial dissipation and compared relative to the reference LES solution (i.e., filtered DNS). The LES filter is a tenth-order discrete Tangent stencil tuned according to $\bar{\Delta}_{1/2}$.

shown to do decently well, in terms of properly locating high energy clusters, when employing high-order discretizations. Although, interestingly, when paired with dissipation, this scheme was seen to go unstable[7]. Finally, the non-stabilized ADM with zero-order deconvolution is shown to deplete kinetic energy the fastest.

Figures 6.17-6.20 offer volumetric plots of the kinetic energy at $t = 9t_c$ and help bring further context to the previous line plots. In general: 1) the use of a high-order scheme allows the features to be properly located (thus the large correlation between the LES and fDNS variables in Table 6.7); 2) the non-stabilized cases have much more noise, although this is reduced somewhat by the high-order deconvolution; and 3) the filter-based artificial dissipate removes the noise and

[7]Similar observations are reported in previous work (Edoh and Karagozian, June 2017a), noting an increase in mid-wavenumber content when combining the ADM procedure and filter-based artificial dissipation. This is a non-linear effect likely owing to an induced bottleneck and subsequent pile-up of energy when the two methods are combined. Reducing the order of the deconvolution procedure ameliorates this issues –as does reducing the artificial dissipation, whose effects on the spectrum and model performance should be further studied (Layton and Neda, 2007a,b).

thus can improve the overall performance[8].

[8]It should perhaps also be noted that very similar performance was achieved in the case of $FTGR = 2$ when coupling the adaptive dissipation method with the SSM model or even with a no-model implementation. This indifference to the model likely highlights that the current problem would be in fact well-handled by a properly-formulated ILES approach. While not always the case, the fact that the kinetic energy curve of the filtered solution in Figure 6.6 (in the case of a tenth-order filter and $FTGR = 2$) matches that of the DNS means that the corresponding LES resolution ($DOF_{LES} = 64^3$) is sufficiently resolved relative to the physics and that the issue of instability actually lies in the numerics rather than a particular need for a modeling. Considering higher Reynolds numbers at the current resolution, however, may indeed put more emphasis on the closure model despite the addition of stabilization and show more impact with respect to a particular closure choice.

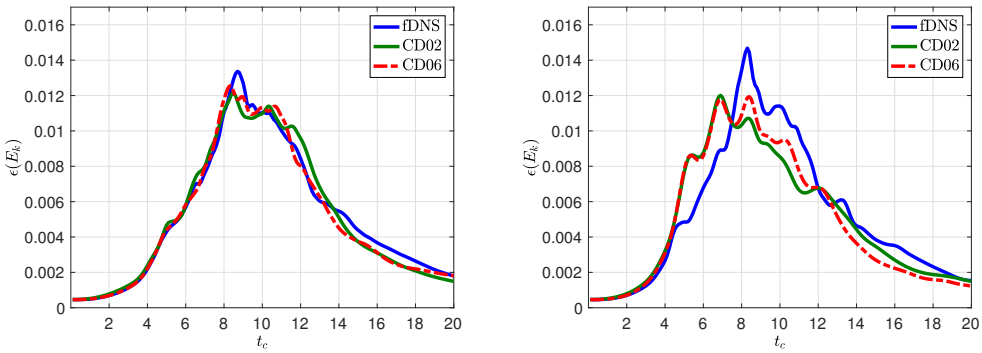

(a) (b)

Figure 6.13: Kinetic energy dissipation rate as measured by $\epsilon(E_k)$ for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$. LES closure is ADM ($k_{cond} = 1.e0$) with consistent filter-based artificial dissipation: a) $FTGR = 2$ and b) $FTGR = 4$.

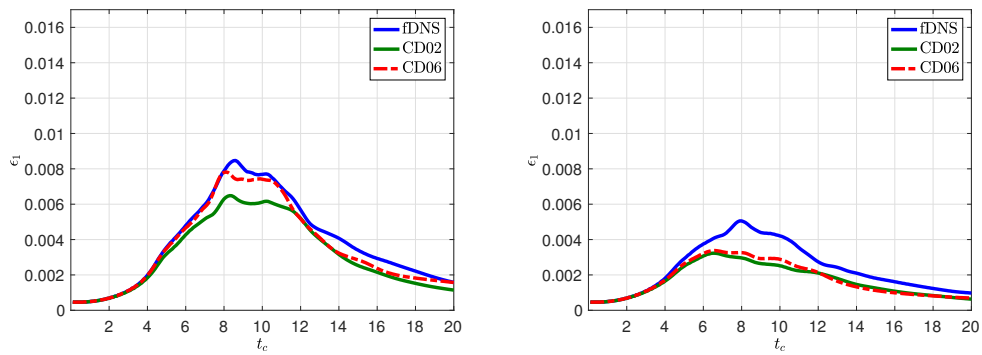

(a) (b)

Figure 6.14: Kinetic energy dissipation rate as measured by ϵ_1 for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$. LES closure is ADM ($k_{cond} = 1.e0$) with consistent filter-based artificial dissipation: a) $FTGR = 2$ and b) $FTGR = 4$.

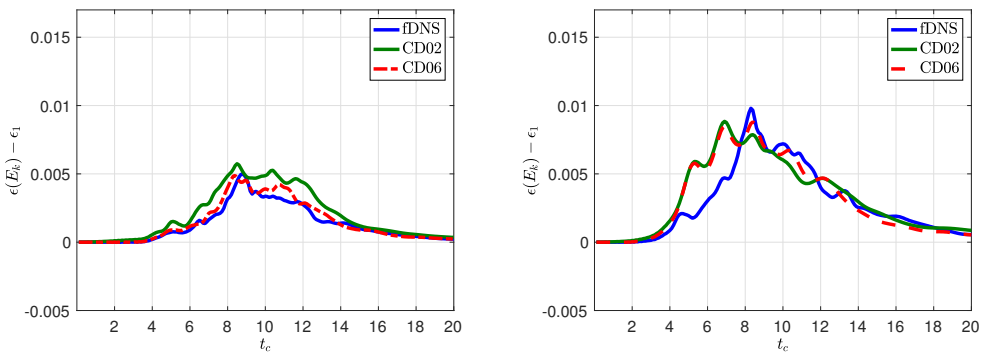

(a) (b)

Figure 6.15: Kinetic energy dissipation rate as measured by the closure model $\epsilon_{\mathcal{T}}$ for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $\bar{\Delta}_{1/2}$. LES closure is ADM ($k_{cond} = 1.e0$) with consistent filter-based artificial dissipation: a) $FTGR = 2$ and b) $FTGR = 4$.

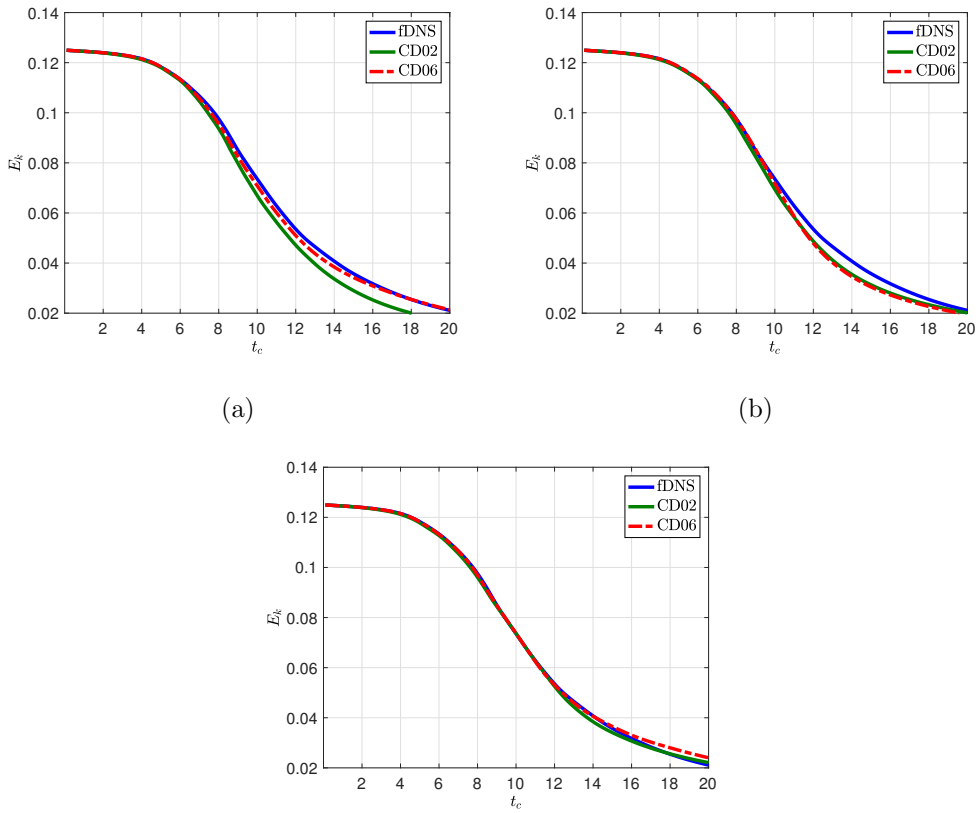

Figure 6.16: Kinetic energy E_k in domain for *a posteriori* LES solutions as a function of the inviscid discretization scheme. The explicit LES filter is assumed to be a tenth-order Tangent stencil tuned according to $FTGR = 2$: a) ADM ($k_{cond} = 1.e3$), b) ADM ($k_{cond} = 1.e0$), and c) ADM ($k_{cond} = 1.e0$) with LES-consistent adaptive filter-based artificial dissipation.

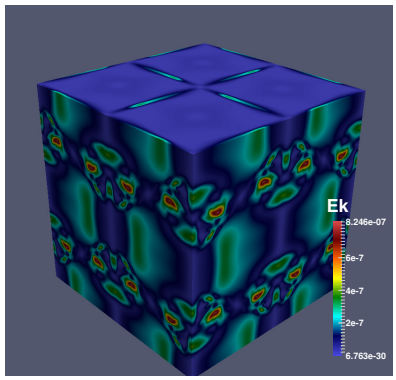

Figure 6.17: Kinetic energy E_k of reference LES solution (i.e., filtered DNS, fDNS) at $t = 9t_c$ for $FTGR = 2$, with the LES filter being a tenth-order discrete Tangent stencil tuned according to $\bar{\Delta}_{1/2}$.

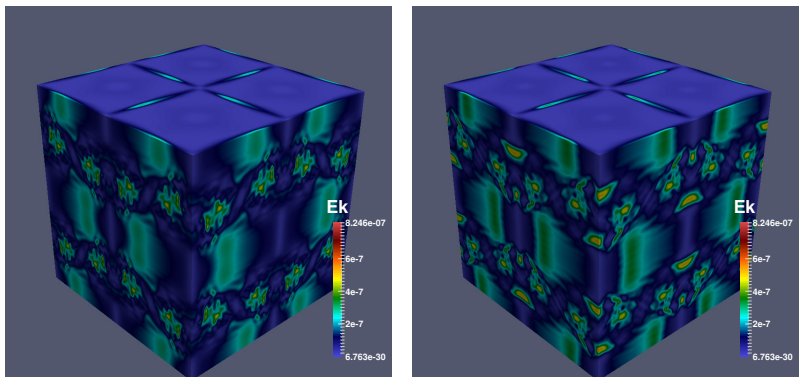

(a) (b)

Figure 6.18: Kinetic energy E_k at $t = 9t_c$ for $FTGR = 2$, with the LES filter being a tenth-order discrete Tangent stencil tuned according to $\bar{\Delta}_{1/2}$. Closure model is ADM with high-order deconvolution ($\kappa_{cond} = 1.e3$): a) CD02 and b) CD06.

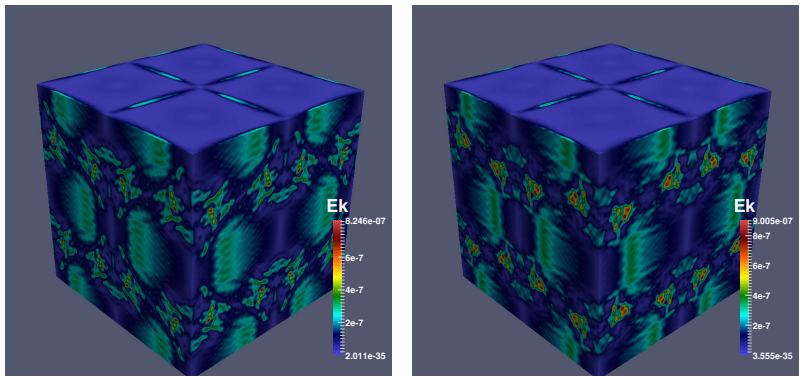

(a) (b)

Figure 6.19: Kinetic energy E_k at $t = 9t_c$ for $FTGR = 2$, with the LES filter being a tenth-order discrete Tangent stencil tuned according to $\bar{\Delta}_{1/2}$. Closure model is ADM with zeroth-order deconvolution ($\kappa_{cond} = 1.e0$): a) CD02 and b) CD06.

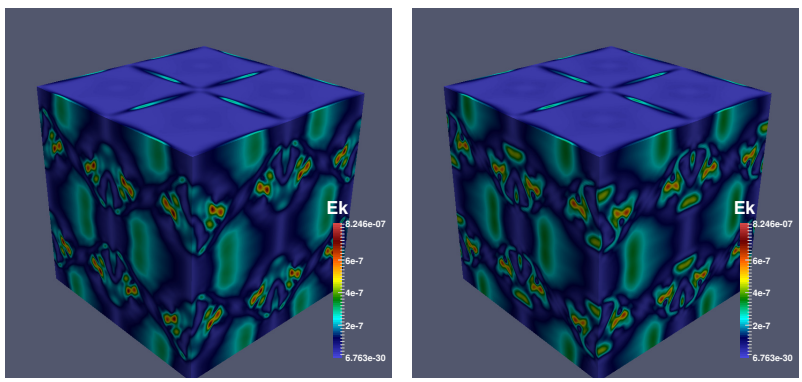

(a) (b)

Figure 6.20: Kinetic energy E_k at $t = 9t_c$ for $FTGR = 2$, with the LES filter being a tenth-order discrete Tangent stencil tuned according to $\bar{\Delta}_{1/2}$. Closure model is ADM with zeroth-order deconvolution ($\kappa_{cond} = 1.e0$) combined with LES-consistent and adaptive filter-based artificial dissipation: a) CD02 and b) CD06.

CHAPTER 7

Conclusions and Future Directions

7.1 Summary of Results

The call for Large-Eddy Simulations for the design of engineering systems derives from a need to capture time-accurate physical phenomena within a computationally feasible framework. The ability to achieve such a goal, however, can be seen as a balance between two parts. On one end are theoretical implications that include such aspects as the filter formulation (e.g., the type of spectral attenuation presumed) and its relation to the intended dynamics of interest. In this vein, one is thus required to consider proper modeling procedures responsible for expressing the underlying assumptions of the LES formulation at hand. The complement to this are the practical considerations regarding how the resulting equations should be solved, and matters such as the numerical treatment (e.g., discretization schemes) thus need to be considered. Achieving successful predictability via LES would thus deserve a holistic understanding of these different aspects of the methodology.

The present work has chosen to investigate these factors by tuning the LES filter formulation relative to the resolvability of the numerics. In such an effort, the spectral characterization of different spatial finite difference schemes is presented, and their respective impacts on time-accurate solutions have been studied. This is done in a basic non-LES context first. To this end, the benefits of incorporating tunable dissipation has also been explored with respect to the proposed filter-

based artificial dissipation method. Having established how one may construct a highly-accurate and stable base scheme, application to the LES problem is reviewed for the under-resolved calculation of test cases exhibiting turbulence. In the instance where exact and consistent modeling is provided using the so-called quasi *a priori* analysis, one observes theorized trends anticipating a reduction in the solution error by means of increasing the filter-to-grid ratio. In addition, employing discrete filtering stencils with enhanced spectral sharpness has been shown to highlight the benefits of high-order discretizations and thus provide added efficiency to the explicit-LES approach by reducing the extent to which the LES scale needs to be resolved (e.g., the ability to support lower filter-to-grid ratios). These clear trends, however, are obfuscated when employing practical models in an *a posteriori* evaluation. The particular choice of models considered – here, the Approximate Deconvolution Method (ADM) and the Bardina-type Scale-Similarity Model (SSM) – is shown to couple with the numerics and affect overall accuracy of the computation. And while the merit of employing lower-order (e.g., second-order) versus higher-order methods is evident, there is less discrimination amongst the latter schemes (e.g., sixth-, tenth-, and optimized fourth-order stencils). This is largely explained by considering their resolvability with respect to the spectral attenuation of the LES filters considered; further differentiation is expected when using sharper or lower filter-to-grid ratios. Nevertheless, benefits of the high-accuracy numerical methods become more evident once model error is reduced – demonstrated herein in terms of high-order deconvolution or the addition of a stabilization term that is consistent with the explicit-LES filter formulation. To this end, the notion that decoupling numerical error (e.g., by means of high-order discretizations and suitable filter-to-grid ratios) will lead to enhanced results is dependent on the assumption that modeling error is sufficiently low. But while this may support claims for implicitly-filtered LES implementations – or at least low filter-to-grid ratios, as all work herein has considered $FTGR \geq 2$ – the previ-

ous conclusion is balanced by the fact that the performance of respective models may not be properly scrutinized without the decoupling of numerical error effects. In this vein, one might be advised to select a particular implementation based on the intended purpose – for example, model development versus more practical computations.

The contributions of the current work may therefore be summarized as follows:

1. Establishment of spectral specification for discrete filtering methods with a focus on scale-discriminant (i.e., sharpness) and tunable properties for scale separation.

2. The formulation of a filter-based artificial dissipation schemes that naturally maintain consistent accuracy for small time-step sizes, unlike traditional solution-filtering procedures which are temporally inconsistent.

3. Combination of the above ideas towards tailoring numerical stabilization with respect to the discretization fidelity of different methods, where ambiguities in the design of dissipation characteristics are mitigated by considering the discretization method's resolvability, thus yielding robustness while preserving scheme accuracy.

4. Employment of such a stabilization approaches to LES and recognizing relation to hyper-viscosity closures that can be made consistent with the explicit LES formulation, thus aiding in closure efforts as well as reducing numerical error effects.

5. Use of the discrete filtering methods to explore different structural models (e.g., scale-similarity type closures) and provide enhanced closure procedures (e.g., high-accuracy deconvolution via stencil re-scaling). In this way, placing additional care in the definition of the filter yields a deliberate assembly of the explicitly-filtered LES formulations that can be used to inherently

consider both numerical and modeling contributions for improved solution accuracy.

6. The adoption of a three-level investigative procedure that includes: the characterization of reference LES solutions and closure behaviors by means of filtering exact data (*a priori* analysis); the assessment of discretization and filter choice on LES performance in the absence of model coupling effects (quasi *a priori* analysis); the quantification of overall accuracy of the LES algorithm with respect to modeling, filtering formulation, and discretization (*a posteriori* analysis).

7. Observations highlighting the need to employ numerical schemes with sufficient resolvability relative to the explicit LES filter, as well as the benefits of providing filter-consistent stabilization for reduction of modeling and/or numerical error.

7.2 Future Work

The current efforts in understanding the interplay of numerical methods and LES formulations are intended for the eventual study of turbulent reactive flow. Such physics are still not well understood, yet are required for predicting important phenomena such as ignition and blow-out as well as shock- or acoustic-related instabilities found in many propulsion applications. RANS, unfortunately, does not provide a viable way of capturing such time-sensitive phenomena. The inclusion of combustion dynamics introduces a new set of time and length scales that need to be considered relative to the LES resolution. For example, insufficiently resolving a flame front would require the LES formulation to effectively consider discontinuities, which would need to be reflected in the modeling methodology (Sagaut and Germano, 2005). As another aspect, the anticipation of backscatter effects (Towery et al., 2016; O'Brien et al., 2016) would also require suitable closures.

But while much knowledge still needs to be developed in terms of modeling, there is a concomitant call for understanding the role of the numerical treatment. The study of Cocks et al. (2015), for example, has demonstrated increased sensitivities to the characteristics of the base scheme under reacting conditions. Therefore, employing a quasi *a priori* type analysis as done in Chapter 5 may be enlightening. While this may be useful for understanding the relative effects of the filter-to-grid ratio, the filter shape, and the discretization, it can also be leveraged in order to assess model performance. In the case of the Navier-Stokes system, one can look to isolate the effects of the different closures – including the important combustion source term (Vreman et al., 2009) – through the sequential substitution of *a posteriori* and *a priori*-derived models. Considerations of geometry complexity can furthermore be postponed by studying simplified in-flow/out-flow premixed flame configurations (Poludnenko and Oran, 2010).

As reactive flow presents a more complex challenge both numerically and dynamically, it is also important for the schemes to be sufficiently robust while maintaining physical relevance. As a consequence, the incorporation of secondary-conservations schemes such as Kinetic Energy Preservation (Morinishi, 2010; Rozema et al., 2014) or Entropy-Conserving methods (Honein and Moin, 2004; Jameson, 2008) may be important and reduce the need to explicit add numerical dissipation for stabilization. The construction of these schemes corresponds to the proper reformulation of flux terms in order to satisfy auxiliary physical equations (Gassner et al., 2016). In doing this, one may force the system to behave in a more physically-relevant manner – this not only suggests improved accuracy, but in special cases, also implies robustness in the non-linear sense. While the current investigation has focused on spatial schemes, proper adherence to such physical symmetries will also require the use of proper temporal integration methods with [pseudo]-symplectic characteristics (Sanderse, 2013; Brouwer et al., 2014; Capuano et al., 2017), as well as the incorporation of provably-stable boundary conditions

such as summation-by-parts operators (Fernandez et al., 2014). Furthermore, adherence to secondary conservation anticipates implications with regards to the class of permissible filters, such as an adjoint property responsible for preserving global energy conservation (Vreman, 2004).

APPENDIX A

Deconvolution Procedures

A.1 Background

The ability to filter and attenuate modal content is often balanced by the need to undo alterations to the original signal. By retrieving the unaltered solution, one has the opportunity to apply alternative manipulations and potentially extract new information. As Section 2.2 highlights the association between filtering and convolution, the act of de-filtering may therefore be referred to as deconvolution. In the context of LES, inversion of an explicitly-defined filter gives one the opportunity to extrapolate the small-scales beyond the principal LES field (e.g, $\tilde{u} = \mathcal{G}_{\bar{\Delta}}^{-1}\{\bar{\tilde{u}}\}$). Although the available numerical resolution puts limits on the ability to retrieve the full unfiltered-field, the resolvable un-filtered variable $\tilde{u}(x)$ may serve as a surrogate to the true continuous field $u(x)$ and may be used to build new closure model approximations. This, for example, is the case for the Approximate Deconvolution method (ADM) introduced in Section 5.3. As a result, the ability to accurately and efficiently perform deconvolution is central to explicitly-filtered LES.

In the case of a defined filter function, one may seek to solve an inverse problem in order to recover the original signal. With respect to the difference operator forms of Equation 2.18 and 2.36, this involves solving the stencils "backwards". In this way, an explicit stencil would require one to solve an implicit system in order to recover the signal. In certain cases, however, the inversion procedure is not well-

defined and the true un-filtered solution cannot be fully retrieved. Such is the case for filtering operators that completely remove certain modes (i.e.: $|\hat{\mathcal{G}}(k)| = 0$) since typically $\mathcal{G}^{-1}\{0\} = 0$ without additional information. Because of such limitations, iterative procedures meant to estimate the true inversion procedure are adopted. The following sections revisit the well-used van Cittert iteration procedure (van Cittert, 1931) and provide useful generalizations. Additionally, a new re-scaling approach based on the re-conditioning of the original filter operator is presented as an alternative for the task of direct or iterative deconvolution.

A.2 Van Cittert Iteration

The van Cittert procedure is a useful iterative procedure often used for approximating the deconvolution of a filtered solution. Its derivation is motivated by considering the decomposition of discrete filter operator into preserving and differencing components: $\mathcal{G} = \mathcal{I} + \mathcal{D}$. The inverse operator $\mathcal{G}^{-1}$ may then be estimated by a Neumann series,

$$\mathcal{G}_N^{-1} \approx \sum_{\nu=0}^{N} (1 - \mathcal{G})^\nu = \sum_{\nu=0}^{N} (-\mathcal{D})^\nu \ , \tag{A.1}$$

where N relates the level or accuracy of the deconvolution. This relation can then be reformulated as a fixed-point iterative method,

$$\begin{aligned} u_0^* &= \bar{u} \ , \\ u_N^* &= \bar{u} - \mathcal{D}\{u_{N-1}^*\} \ , \\ \rightarrow u_N^* &= u_{N-1}^* + (\bar{u} - \mathcal{G}\{u_{N-1}^*\}) = \mathcal{G}_N^{-1}\{\bar{u}\} \ . \end{aligned} \tag{A.2}$$

By recognizing this iterative procedure to be in a residual form, one can further generalize by employing acceleration techniques such as over-relaxation (Maulik

and San, 2016), which seeks to enhance the intermediate approximations:

$$u_N^* = u_{N-1}^* + \omega_N \cdot (\bar{u} - \mathcal{G}\{u_{N-1}^*\}) \ , \tag{A.3}$$

$$\rightarrow \mathcal{G}_N^{-1} \approx \prod_{\nu=1}^{N}(\mathcal{I} - \omega_N \cdot \mathcal{G}) + \sum_{\nu=1}^{N-1}\left[\omega_N \cdot \prod_{\mu=\nu+1}^{N}(\mathcal{I} - \omega_\mu \cdot \mathcal{G})\right] + \omega_N \cdot \mathcal{I} \ . \tag{A.4}$$

As written above the relaxation parameter, ω_N, may be varied according to the iteration stage N as explored in Layton and Stanculescu (2009) for optimal iterations. Furthermore, the fact that Equation A.2 takes on a residual form allows one to abstract the procedure as a temporal smoothing scheme, with the relaxation parameter relating to a pseudo-stage time-step size. As a result, optimal smoothing schemes such as van Leer's predictor-corrector methods (Leer et al.) can be applied in order to enhance convergence. Furthermore, such methods can be combined with residual-smoothing and multi-grid techniques for additional acceleration (Edoh et al., June 2016; Haelterman et al., 2011).

Depending on the filter response function $\mathcal{G}$, higher deconvolution levels may be necessary in order to rebuild the original signal. This is particularly the case for sharp filters that strongly attenuate high-wavenumber content. Therefore, the efficiency of the deconvolution is essential for practicality. In the case of general implicit filter stencils that include evaluations of both the filtered and unfiltered variable, the procedure above suggests multiplied evaluations of $[\mathcal{I} - \mathcal{G}]\{u_{N-1}^*\}$ which implies repeated costly stencil inversion at each iteration. One remedy, proposed in a previous work (Edoh et al., June 2016), is to re-write $\bar{u} = \mathcal{G}\{u\}$ as $\mathcal{B}\{\bar{u}\} = \mathcal{T}\{u\}$ and to modify the procedure of Equation A.2 to now read,

$$\begin{aligned} u_0^* &= \bar{u} \ , \\ u_N^* &= u_{N-1}^* + \omega_N \cdot (\mathcal{B}\{\bar{u}\} - \mathcal{T}\{u_{N-1}^*\}) \ . \end{aligned} \tag{A.5}$$

then gives the benefit of solely evaluating explicit stencils relative to the iterated variable u_{N-1}^*.

In the instances where the original filter is non-invertible, the iteration process is not expected to converge, thus relating to the inability to fully recover the removed mode $|\hat{\mathcal{G}}|(k') = 0$. As a consequence, one is required to terminate the approximation after N iterations. Therefore, one would like the best answer in a given number of steps. Over-relaxation is often utilized in order to accelerate or enhance the quality of the estimates. However, as will be shown, special care should be employed.

Figure A.1 plots the estimate to the spectral response $|\hat{\mathcal{G}}| \cdot |\hat{\mathcal{G}}_N^{-1}| \sim \mathcal{I}$ for a sixth-order Tangent filter tuned to $\Delta_{1/2} = (4\Delta x)$ using the over-relaxation form of the van Cittert deconvolution procedure shown in Equation A.4 with uniform relaxation parameters $\omega_N = \omega$. It is evident that the estimate improves relative to the number of iterations. When the relaxation parameter is less than or equal to unity, the estimates approach identity from below. However, when the relaxation parameter is greater than unity, insufficient iterations produce an estimate that shows growths or depletions in certain lower modes. This can be concerning with respect to realizability constraints for simulations or could otherwise unintentionally affect the non-linear behavior by depleting or surging the energy in the low modes. Consequently, the observation suggests that over-relaxation techniques should be sufficiently converged in order to avoid such artificial issues. Understanding that the current scheme may be abstracted as a forward Euler temporal smoothing scheme for the iteration (with time-step, $\Delta\tau$, corresponding to the relaxation-parameter ω), it can be shown that stable iterative convergence occurs for $\omega \leq 2$ (assuming a stable filter with $|\hat{\mathcal{G}}| \leq 1$). The occurrence of over- and under-shoots in the intermediate estimates can be understood as uneven removal of error amongst the different modes in the iterative deconvolution estimate. In order to avoid such overshoots, the growth factor of the iterative scheme should match the monotonicity traits – or lack thereof – of the original filter response in question. In this exercise, the challenge of balancing efficiency and accuracy of

the procedure is introduced, where enhanced recovery of high-wavenumbers for a given a set iteration count is possible with over-relaxation but may result in less accuracy, particularly in the lower modes when the procedure is not sufficiently converged.

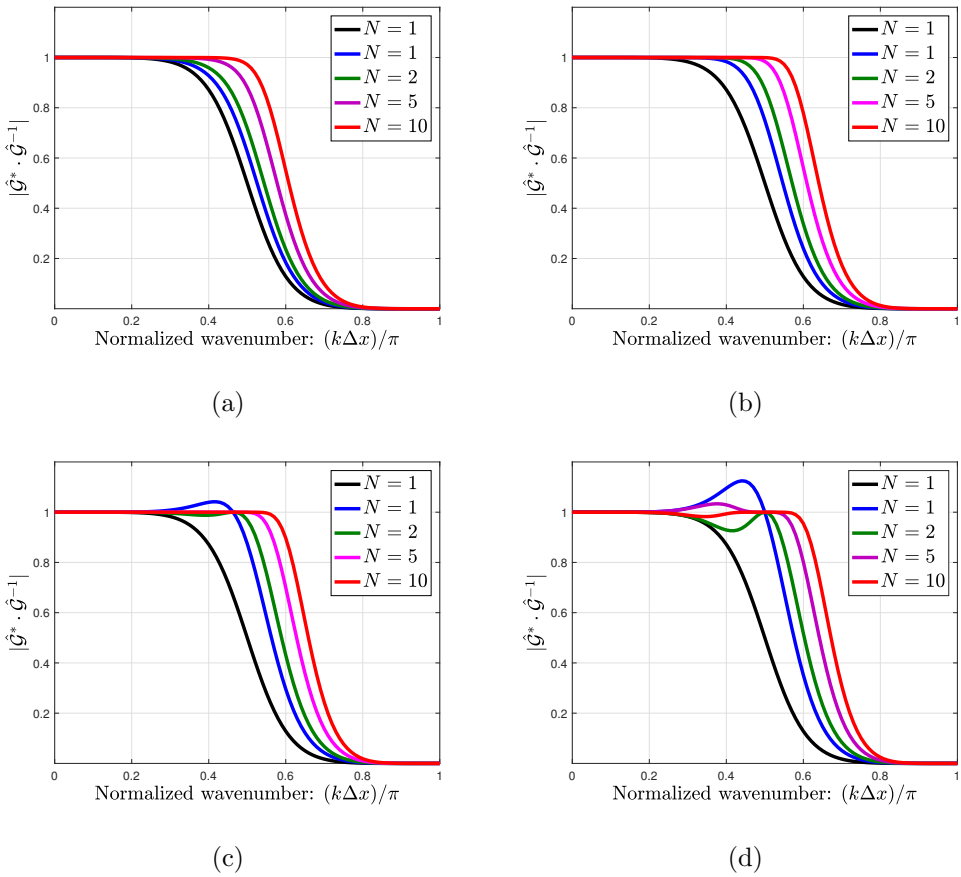

Figure A.1: Effect of relaxation-parameter ω and iteration count N on the growth factor $\{|\hat{\mathcal{G}}| \cdot |\hat{\mathcal{G}}_N^{-1}|\} \sim \mathcal{I}$ of an approximate deconvolution by the van Cittert procedure of Equation A.2 for the sixth-order discrete Tangent filter, tuned to a filter width $\Delta_{1/2} = (4\Delta x)$: a) $\omega = 0.5$, b) $\omega = 1.0$, c) $\omega = 1.5$, and d) $\omega = 2.0$.

A.3 An Operator Re-scaling Approach

In the case where the iteration procedure does not converge due to the lack of the $\mathcal{G}$ operator's invertibility, it becomes interesting to formulate a new operator that is solvable (i.e.: will have a convergent iterative procedure) and that approximates the original inverse problem, $\mathcal{G}' \approx \mathcal{G}$. As most filtering stencils inherit their lack of invertibility from the fact that they eliminate odd-even modes (i.e.: $|\hat{\mathcal{G}}(k_{\Delta x})| = 0$), an intuitive form of $\mathcal{G}'$ appears from re-scaling the stencil such that the Nyquist frequency is not completely removed. The new approximate operator can be defined implicitly as $\mathcal{B}\{\bar{u}\} = \mathcal{T}'\{u\}$ with,

$$\begin{aligned}\left[\epsilon_{IF,0} + \sum_{\ell=1}^{L} \epsilon_{IF,2\ell}(\Delta x)^{2\ell}\delta^{2\ell}\right]\{\bar{u}\} &= \left[\epsilon_{EF,0} + \mu\sum_{r=1}^{R} \epsilon_{EF,2r}(\Delta x)^{2r}\delta^{2r}\right]\{u\} \\ &\quad + \left[(1-\mu)\sum_{\ell=1}^{L} \epsilon_{IF,2\ell}(\Delta x)^{2\ell}\delta^{2\ell}\right]\{u\} \ ,\end{aligned} \tag{A.6}$$

where the scaling factor $\mu = (1 - 1/\kappa_{cond}) \in [0, 1]$ is related to the conditioning and accuracy of the new approximate operator. This new operator can then be used to solve the inverse problem with a direct method or as part of an iterative procedure that will now be expected to converge.

Figure A.2 considers a direct solve of the inverse problem and plots $|\hat{\mathcal{G}}|\cdot|\hat{\mathcal{G}}_N'^{-1}| \sim \mathcal{I}$ for a sixth-order Tangent filter tuned to $\Delta_{1/2} = (4\Delta x)$ as a function of different re-conditionings. Manipulating the condition number of the original non-invertible stencil is seen to tremendously improve estimates of the deconvolution, with higher conditioning numbers yielding more accurate results. As a point of comparison, the modest $\kappa_{cond} = 10$ re-scaling performs as well as $N = [10, 5, 5]$ van Cittert iterations with relaxation parameters $\omega = [0.5, 1, 1.5]$, respectively. In considering the original van Cittert procedure of Equation A.2, which suggests the need for matrix inversions of implicit stencils for each iteration, solution of the approximate problem via the current re-scaling approach is seen to be strongly preferable in

terms of both accuracy and efficiency. Previous issues relating to under- and overshoots in low-modal content of the reconstructed signal can still occur, however, when the approximate operator is solved iteratively and not permitted to converge sufficiently.

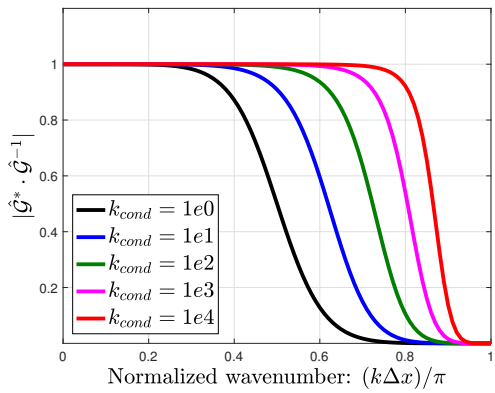

(a)

Figure A.2: Effect of conditioning number κ_{cond} on the growth factor $\{|\hat{\mathcal{G}}| \cdot |\hat{\mathcal{G}}_N^{-1}|\} \sim \mathcal{I}$ of an approximate deconvolution by the direct procedure of Equation A.6 for the sixth-order discrete Tangent filter, tuned to a filter width $\Delta_{1/2} = (4\Delta x)$.

Bibliography

Alpert. Implicit filtering in conjunction with explicit filtering. *Journal of Computational Physics*, 44:212–219, 1981.

K. Asthana and A. Jameson. High-order flux reconstruction schemes with minimal dispersion and dissipation. *Journal of Scientific Computing*, 2014.

H. L. Atkins. Simulations of compressible taylor-green flow by a discontinuous galerkin method. AIAA 2016-3330, June 2016.

G. Aubard, P. S. Volpiani, and X. Gloerfelt. Comparison of subgrid-scale viscosity models and selective filtering strategy for large-eddy simulations. *Flow Turbulence Combustion*, 91:497–518, 2013.

Z. Bai and X. Zhong. New very high-order upwind multilayer compact schemes with spectral-like resolution for flow simulations. AIAA 2017-0518, 2017.

J. Bardina, J. H. Ferziger, and W. C. Reynolds. Improved turbulence models based on large eddy simulation of homogeneous, incompressible, turbulent flows. Technical Report No. TF-19, 1983.

J. Berland, C. Bogey, and C. Bailly. Investigation using statistical closure theory of the influence of the filter shape on scale separation in large-eddy simulation. *Journal of Turbulence*, 9(21):1–22, 2008.

H. Bijl, M. H. Carpenter, V. N. Vatsa, and C. A. Kennedy. Implicit time integration schemes for the unsteady compressible navier-stokes equations: Laminar flow. *Journal of Computational Physics*, 179:313–329, 2002.

G. A. Blaisdell, E. T. Spyropoulos, and J. H. Qin. The effect of the formulation of nonlinear terms on aliasing errors in spectral methods. *Applied Numerical Mathematics*, 21:207–219, 1996.

C. Bogey and C. Bailly. A family of low dispersive and low dissipative explicit schemes for flow and noise computations. *Journal of Computational Physics*, 194:194–214, 2004.

C. Bogey, N. de Cacqueray, and C. Bailly. A shock-capturing methodology based on adaptive spatial filtering for high-order non-linear computations. *Journal of Computational Physics*, 228:1447–1465, 2009.

S. Bose, P. Moin, and D. You. Grid-independent large-eddy simulation using explicit filtering. *Physics of Fluids*, 22:1–11, 2010.

M. Brachet, D. Meiron, S. Orszag, B. Nickel, R. Morf, and U. Frisch. Small-scale structure of the taylor-green vortex. *Journal of Fluid Mechanics*, 130(41): 411–452, 1983.

J. Brouwer, J. Reiss, and J. Sesterhenn. Conservative time integrators of arbitrary order for skew-symmetric finite-difference discretizations of compressible flow. *Computers and Fluids*, 100:1–12, 2014.

J. R. Bull and A. Jameson. Simulation of the compressible taylor green vortex using high-order flux reconstruction schemes. AIAA 2014-3210, 2014.

J. R. Bull and A. Jameson. Explicit filtering and exact reconstruction of the sub-filter stresses in large-eddy simulation. *Journal of Computational Physics*, 306: 117–136, 2016.

S. Butterworth. On the theory of filter amplifiers. *Wireless Engineer*, 6:536–541, 1930.

F. Capuano, G. Coppola, L. Randez, and L. de Luca. Explicit runge-kutta schemes for incompressible flow with improved energy-conservation properties. *Journal of Computational Physics*, 328:86–94, 2017.

D. Carati, G. Winckelmans, and H. Jeanmart. On the modeling of the subgrid-scale and filtered-scale stess tensors in large-eddy simulation. *Journal of Fluid Mechanics*, 441:119–138, 2001.

P. A. T. Cocks, M. C. Soteriou, and V. Sankaran. Impact of numerics on the predictive capabilities of reactive flow les. *Combustion and Flame*, 162:3394–3411, 2015.

D. Corson, R. Jaiman, and F. Shakib. Industrial application of rans modeling: Capabilities and needs. *International Journal of Computational Fluid Dynamics*, 23(4):337–347, 2009.

J. R. DeBonis. Solutions of the taylor-green vortex problem using high-resolution explicit finite difference methods. Technical Report No. NASA/TM 217850, 2013.

G. DeStefano and O. V. Vasilyev. "perfect" modeling framework for dynamics sgs model testing in large-eddy simulation. *Theoretical Computational Fluid Dynamics, year = 2004, volume = 18, pages = 27-41.*

D. Drikakis, C. Fureby, F. F. Grinstein, and D. Youngs. Simulation of transition and turbulence decay in the taylor-green vortex. *Journal of Turbulence*, 8(20): 1–11, 2007.

A. Edoh, C. Merkle, A. Karagozian, and V. Sankaran. Comparison of artificial dissipation and filtering schemes for time accurate simulations. AIAA Paper 2015-0284, 2015.

A. K. Edoh and A. R. Karagozian. Stabilized scale-similarity modeling for explicitly-filtered les. AIAA 2017-1227, June 2017a.

A. K. Edoh and A. R. Karagozian. Inspecting interactions of discretization, fil-

ter formulation, and stabilization in les: Lessons from the taylor-green vortex. AIAA 2017-, June 2017b.

A. K. Edoh, N. L. Mundis, A. R. Karagozian, and V. Sankaran. The role of dispersion and dissipation on stabilization strategies for time-accurate simulations. AIAA Paper 2016-0071, January 2016.

A. K. Edoh, N. L. Mundis, A. R. Karagozian, and V. Sankaran. Discrete filtering formulations for large-eddy simulation. AIAA 2016-3794, June 2016.

F. Falissard. Genuinely multi-dimensional explicit and implicit generalized shapiro filters for weather forecasting, computational fluid dynamics and aeroacoustics. *Journal of Computational Physics*, 253:344–367, 2013.

F. Falissard. Uneven-order decentered shapiro filters for boundary filtering. *Journal of Computational Physics*, 292:168–175, 2015.

D. Fauconnier. On the performance of relaxation filtering for large-eddy simulation. *Journal of Turbulence*, 253(1):22–49, 2013.

D. Fauconnier, C. De Langhe, and E. Dick. A family of dynamic finite difference schemes for large-eddy simulation. *Journal of Computational Physics*, 228: 1830–1861, 2009.

D. Fauconnier, C. De Langhe, and E. Dick. Quality assessment of dynamic finite difference schemes on the taylor-green vortex. *Quality and Reliability of Large-Eddy Simulations II*, 16:287–296, 2011.

D. C. Del Rey Fernandez, J. E. Hicken, and D. W. Zingg. Review of summation-by-parts operators with simultaneous approximation terms for the numerical solution of partial differential equations. *Computers and Fluids*, 95:171–196, 2014.

D. Flad, A. Beck, and C. Munz. Simulation of underresolved turbulent flows by adaptive filtering using the high order discontinuous galerkin spectral element method. *Journal of Computational Physics*, 313:1–12, 2016.

D. V. Gaitonde and M. R. Visbal. Pade-type higher-order boundary filters for the navier-stokes equations. *AIAA Journal*, 38:2103–2112, 2000.

E. Garnier, N. Adams, and P. Sagaut. *Large Eddy Simulation for Compressible Flows*. Springer, 2009.

G. J. Gassner and A. D. Beck. On the accuracy of high-order discretizations for underresolved turbulence simulations. *Theoretical Computational Fluid Dynamics*, pages 222–237, 2013.

G. J. Gassner, A. R. Winters, and D. A. Kopriva. Split form nodal discontinuous galerkin schemes with summation-by-parts property for the compressible euler equations. *Journal of Computational Physics*, 327:39–66, 2016.

M. Germano. Differential filters for large-eddy numerical simulation of turbulent flow. *Physics of Fluids*, 29:1755–1756, 1986.

M. Germano, U. Piomelli, P. Moin, and W. H. Cabot. A dynamic subgridscale eddy viscosity model. *Physics of Fluids*, 3:1760–1765, 1991.

B. J. Geurts and F. van der Bos. Numerically induced high-pass dynamics in large-eddy simulation. *Physics of Fluids*, 17:1–12, 2005a.

B. J. Geurts and F. van der Bos. Commutator errors in the filtering approach to large-eddy simulation. *Physics of Fluids*, 17:1–20, 2005b.

S. Ghosal. An analysis of numerical errors in large-eddy simulations of turbulence. *Journal of Computational Physics*, 125(88):188–206, 1996.

Sidharth GS, A. Kartha, and G. V. Candler. Filtered velocity-based les of mixing in high speed recirculating shear flow. AIAA 2016-3184, June 2016.

S. N. Gurbatov, S. I. Simdyankin, E. Aurell, U. Frisch, and G. Toth. On the decay of burgers turbulence. *Journal of Fluid Mechanics*, 344, 1997.

R. Haelterman, J. Vierendeels, D. van Heule, S de Ridder, and H. Bruyninckx. Multi-stage solvers optimized for damping propagation. *Journal of Computational and Applied Mathematics*, 235:3775–3782, 2011.

M. E. Harvazinski, C. Huang, V. Sankaran, T. W. Feldman, W. E. Anderson, C. L. Merkle, and D. G. Talley. Coupling between hydrodynamics, acoustics, and heat release in a self-excited unstable combustor. *Physics of Fluids*, 27: 1–27, 2015.

A. Haselbacher and O. V. Vasilyev. Commutative discrete filtering on unstructured grids based on least-squares techniques. *Journal of Computational Physics*, 187:197–211, 2003.

C. Hirsch and B. Tartinville. Reynolds-averaged navier-stokes modeling for industrial applications and some challenging issues. *International Journal of Computational Fluid Dynamics*, 23(4):295–303, 2009.

A. E. Honein and P. Moin. Higher entropy conservation and numerical stability of compressible turbulence simulations. *Journal of Computational Physics*, 201: 531–545, 2004.

F. Q. Hu, M. Y. Hussaini, and J. L. Manthey. Low-dissipation and low-dispersion runge-kutta schemes for computational acoustics. *Journal of Computational Physics*, 124(52):177–191, 1996.

B. Jahne. *Digital Image Processing: Concepts, Algorithms, and Scientific Applications*. Springer-Verlag, Berlin, Heildelberg.

A. Jameson. Formulation of kinetic energy preserving conservative schemes for gas dynamics and direct numerical simulation of one-dimensional viscous compress-

ible flow in a shock tube using entropy and kinetic energy preserving schemes. *Journal of Scientific Computing*, 32:188–208, 2008.

H. Jeanmart and G. Winckelmans. Investigation of eddy-viscosity models modified using discrete filters: A simplified "regularized variational multiscale model" and an "enhanced field model". *Physics of Fluids*, 19:1–16, 2007.

N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous Univariate Distributions*, volume 2. Wiley, 1994.

S. A. Jordan. The spatial resolution properties of composite compact finite differencing. *Journal of Computational Physics*, 221:558–576, 2007.

R. Kamakoti and C. Pantano. High-order narrow stencil finite-difference approximations of second-order derivatives involving variable coefficients. *SIAM Journal of Scientific Computing*, 31(6):4222–4243, 2009.

C. A. Kennedy and A. Gruber. Reduced aliasing formulations of the convective terms within the navier-stokes equations for a compressible fluid. *Journal of Computational Physics*, 227:1676–1700, 2008.

J. Kerman. A closed-form approximation of the median of the beta distribution. *arXiv:1111.0433v1*, 2011.

J. Kim, P. Moin, and R. Moster. Turbulence statistics in fully developed channel flow at low reynolds number. *Journal of Fluid Mechanics*, 177:133 –166, 1987.

J. W. Kim. High-order compact filters with variable cut-off wavenumber and stable boundary treatment. *Computers and Fluids*, 39:1168–1182, 2010.

A. G. Kravchenko and P. Moin. On the effect of numerical errors in large eddy simulations of turbulent flows. *Journal of Computational Physics*, 131:310–322, 1997.

A. LaBryer, P. J. Attar, and P. Vedula. A framework for large eddy simulation of burgers turbulence based on spatial and temporal statistical information. *Physics of Fluids*, 27:1–22, 2015.

M. Lagha, J. Kim, J. D. Eldredge, and X. Zhong. A numerical study of compressible turbulent boundary layers. *Physics of Fluids*, 23:1–12, 2011.

J. Larsson and B. Gustafsson. Stability criteria for hybrid difference methods. *Journal of Computational Physics*, 227:2886–2898, 2008.

W. Layton and M. Neda. A similarity theory of approximate deconvolution models of turbulence. *Journal of Mathematical Analysis and Applications*, pages 416–429, 2007a.

W. Layton and M. Neda. Truncation of scales by time relaxation. *Journal of Mathematical Analysis and Applications*, pages 416–429, 2007b.

W. Layton and I. Stanculescu. Chebychev optimized approximate deconvolution models for turbulence. *Applied Mathematics and Computation*, 208:106–118, 2009.

B. Van Leer, C. H. Tai, and K. H. Powell. Design of optimally smoothing multi-stage scheme for the euler equations. AIAA Paper, 1989-1923, 1989.

S. K. Lele. Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics*, 103:16–42, 1992.

A. Lerat and C. Corre. A residual-based compact scheme for the compressible navier-stokes equations. *Journal of Computational Physics*, 170:642–675, 2001.

A. Lerat, F. Falissard, and J. Sides. Vorticity-preserving schemes for the compressible euler equations. *Journal of Computational Physics*, 225:635–651, 2007.

Y. Li and Z. J. Wang. A priori and a posteriori evaluations of sub-grid scale models for the burgers' equation. *Computers and Fluids*, 2016.

D. K. Lilly. A proposed modification of the germano subgrid-scale closure method. *Physics of Fluids A: Fluid Dynamics*, 4:633–634, 1992.

V. Linders and J. Nordstrom. Uniformly best wavenumber approximations by spatial central difference operators. *Journal of Computational Physics*, 300: 695–709, 2015.

S. Liu, C. Meneveau, and J. Katz. On the properties of similarity subgrid-scale models as deduced from measurements in a turbulent jet. *Journal of Fluid Mechanics*, 275:83–119, 1994.

S. Liu, C. Meneveau, and J. Katz. Experimental study of similarity subgrid-scale models of turbulence in the far-field of a jet. *Applied Scientific Research*, 54: 117–190, 1995.

X. Liu, S. Zhang, H. Zhang, and C. H. Zhu. A new class of central compact schemes with spectral-like resolution i: Linear schemes. *Journal of Computational Physics*, 248:235–256, 2013a.

X. Liu, S. Zhang, H. Zhang, and C. H. Zhu. A new class of central compact schemes with spectral-like resolution ii: Hybrid weighted nonlinear schmes. *Journal of Computational Physics*, 284:133–154, 2013b.

G. Lodato, P. Castonguay, and Antony Jameson. Discrete filter operators for large-eddy simulation using high-order spectral difference methods. *International Journal for Numerical Methods in Fluids*, 72:231–258, 2013.

P. E. Long, W. A. Shaffer, J. E. Kemper, and F. J. Hicks. The state of the techniques development laboratory's boundary layer model. Technical Report No. TDL66, 1978.

M. D. Love. Subgrid modeling studies with burgers equation. *Journal of Fluid Mechanics*, 87, 1980.

T. S. Lund. *On the Use of Discrete Filters for Large Eddy Simulations.* Annual Research Briefs. Center of Turbulence Research, NASA Ames-Stanford University, 1997.

T. S. Lund. The use of explicit filters in large eddy simulation. *Computers and Mathematics with Applications*, 46:603–616, 2003.

K. Mahesh. A family of high order finite difference schemes with good spectral resolution. *Journal of Computational Physics*, 145:332–358, 1998.

N. Mansouri and D. Hixon. Optimized explicit dissipation operators for computational aeroacoustics. AIAA Paper 2015-2539, 2015.

A. L. Marsden, O. V. Vasilyev, and P. Moin. Construction of commutative filters for les on unstructured meshes. *Journal of Computational Physics*, 175:584–603, 2002.

J. Mathew, R. Lechner, H. Foysi, J. Sesterhenn, and R. Friedrich. An explicit filtering method for large eddy simulation of compressible flow. *Physics of Fluids*, 15(8):2279–2289, 2003.

J. Mathew, H. Foysi, and R. Friedrich. A new approach to les based on explicit filtering. *International Journal of Heat and Fluid Flow*, 27:594–602, 2006.

R. Maulik and O. San. Evaluation of explicit and implicit les closures for burgers turbulence. *arXiv:1604.08649v1*, 2016.

C. L. Merkle and M. Athavale. A time accurate unsteady incompressible algorithm based on artificial compressibility. 1987.

H. Mirzaee, J. K. Ryan, and R. M. Kirby. Efficient implementations of smoothness-increasing accuracy-conserving (siac) filters for discontinuous galerkin solutions. *Journal of Scientific Computing*, 52:85–112, 2012.

Y. Morinishi. Skew-symmetric form of the convective terms and fully conservative finite difference scehemes for variable low-mach number flows. *Journal of Computational Physics*, 229:276–300, 2010.

Y. Morinishi, T. S. Lund, O. V. Vasilyev, and P. Moin. Fully conservative higher order finite difference schemes for incompressible flow. *Journal of Computational Physics*, 143:90–124, 1998.

J. S. Mullen and P. F. Fischer. Filtering techniques for complex geometry fluid flows. *Communications in Numerical Methods in Engineering*, 15:9–18, 1999.

N. L. Mundis, A. K. Edoh, and V. Sankaran. Optimal Runge-Kutta schemes for high-order spatial and temporal discretizations. AIAA Paper 2015-2752, 2015.

N. L. Mundis, A. K. Edoh, and V. Sankaran. Highly-accurate filter-based artificial-dissipation schemes for stiff unsteady fluid systems. AIAA Paper 2016-0072, January 2016.

J. O'Brien, C. A. Z. Towery, P.E. Hamlington, M. Ihme, A. Y. Poludnenko, and J. Urzay. The cross-scale physical-space transfer of kinetic energy in turbulent premixed flames. *Proceedings of the Combustion Institute*, pages 1–9, 2016.

S. A. Orszag. On the elimination of aliasing in finite-difference schemes by filtering high-wavenumber components. *Journal of Atmospheric Sciences*, 28:1074, 1971.

N. Park and K. Mahesh. Reduction of the germano-identity error in the dynamic smagorinsky model. *Physics of Fluids*, 21:1–16, 2009.

N. A. Phillips. *An Example of Non-Linear Computational Instability.* The Atmosphere and Sea in Motion. Oxford University Press, 1959.

S. Pirozzoli. Generalized conservatie approximations of split convective derivative operators. *Journal of Computational Physics*, 229:7180–7190, 2010.

A. Y. Poludnenko and E. S. Oran. The interaction of high-speed turbulence with flames: Global properties and internal flame structure. *Combustion and Flame*, 157:995–1011, 2010.

C. D. Pruett. Eulerian time-domain filtering for spatial large-eddy simulation. *AIAA Journal*, 38(9):1634–1642, 2000.

C. D. Pruett and J. S. Sochacki. On taylor-series expansions of residual stress. *Physics of Fluids*, 13(9):2578–2589, 2001.

C. D. Pruett, T. B. Gatski, C. E. Grosch, and W. D. Thacker. The temporally filtered navier-stokes equations: Properties of the residual stress. *Physics of Fluids*, 15(8):2127–2140, 2003.

R. J. Purser. The filtering of meteorological fields. *Journal of Climate and Applied Meteorology*, 26:1764–1769, 1987.

S. Radhakrishnan and J. Bellan. Explicit filtering to obtain grid-spacing-independent large-eddy simulation of compressible single-phase flow. *Journal of Fluid Mechanics*, 697:399–435, 2012.

J. Ramboer, T. Broeckhoven, S. Smirnov, and C. Lacor. Optimization of time integration schemes coupled to spatial discretization for use in caa applications. *Journal of Computational Physics*, 213:777–802, 2006.

W. H. Raymond. High-order low-pass implicit tangent filters for use in finite area calculations. *Monthly Weather Review*, 116:2132–2141, 1988.

W. H. Raymond and A. Garder. A review of recursive and implicit filters. *Monthly Weather Review*, 119:477–495, 1991.

W. Rozema, J. C. Kok, R. W. C. P. Verstappen, and A. E. P Veldman. A symmetry-preserving discretization and regularization model for compressible

flow with application to turbulent channel flow. *Journal of Turbulence*, 15: 386–410, 2014.

P. Sagaut. *Large Eddy Simulation for Incompressible Flows*. Springer, 2006.

P. Sagaut and M. Germano. On the filtering paradigm for les of flows with discontinuities. *Journal of Turbulence*, 6(23), 2005.

O. San, A. E. Staples, and T. Iliescu. Analysis of low-pass filters for approximate deconvolution in one-dimensional decaying burgers turbulence. *International Journal of Computational Fluid Dynamics*, 30(1):20–37, 2016.

B. Sanderse. Energy-conserving runge-kutta methods for the incompressible navier-stokes equations. *Journal of Computational Physics*, 233:100–131, 2013.

A. Sescu, R. Hixon, C. Sescu, and A. A. Afjeh. Stability investigation of multidimensional optimized spatial stencils. January 2009.

R. Shapiro. The use of linear filtering as a parametrization of atmospheric diffusion. *Journal of Atmospheric Sciences*, pages 523–531, 1971.

R. Shapiro. Linear filtering. *Mathematics of Computation*, 29(132):1094–1097, 1975.

M. H. Silvis, R. A. Remmerswaal, and R. Verstappen. Physical consistency of subgrid-scale models for large-eddy simulation of incompressible turbulent flows. *Physics of Fluids*, 29:1–16, 2017.

C. G. Speziale. Turbulence modeling for time-dependent rans and vles: A review. *AIAA Journal*, 36(2):173–184, 1998.

Stanescu and Habashi. 2n-storage low dissipation and dispersion runge-kutta schemes for computational acoustics. *Journal of Computational Physics*, 143: 674–681, 1998.

G. De Stefano and O. V. Vasilyev. Sharp cutoff versus smooth filtering in large eddy simulation. *Physics of Fluids*, 14:362–369, 2002.

S. Stolz and N. A. Adams. An approximate deconvolution procedure for large-eddy simulation. *Physics of Fluids*, 11(7):1699–1701, 1999.

S. Stolz, N. A. Adams, and L. Kleiser. The approximate deconvolution model for large-eddy simulation for compressible flows and its application to shock-turbulent-boundary-layer interaction. *Physics of Fluids*, 13(10):2985–3001, 2001.

R. C. Swanson, R. Radespiel, and E. Turkel. On some numerical dissipation schemes. *Journal of Computational Physics*, 147:518–544, 1998.

P. R. Tadikamalla. A look at the burr and related distributions. *International Statistical Review*, 48(3):337–344, 1980.

E. Tadmor and W. Zhong. Entropy stable approximations of navier-stokes equations with no artificial numerical viscosity. *Journal of Hyperbolic Differential Equations*, 3(3):529–559, 2006.

C. K. W. Tam and J. C. Webb. Dispersion-relation-preserving finite difference schemes for computational acoustics. *Journal of Computational Physics*, 107: 262–281, 1993.

T. Tantikul and J. A. Domaradzki. Large eddy simulations using truncated navier-stokes with the automatic filtering criterion. *Journal of Turbulence*, 11(21):1–24, 2010.

T. Tantikul and J. A. Domaradzki. Truncated navier-stokes equations with the automatic filtering criterion: Reynolds stress and energy budgets. *Journal of Turbulence*, 12(34):1–25, 2011.

G. Taylor and A. Green. Mechanisms of the production of small eddies from large ones. *Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 158(895):499–521, 1937.

C. A. Z. Towery, A. Y. Poludnenko, J. Urzay, J. O'Brien, M. Ihme, and P.E. Hamlington. Spectral kinetic energy transfer in turbulent premixed reacting flows. *Physical Review E*, 93:1–11, 2016.

E. Turkel. Preconditioning techniques in computational fluid dynamics. *Annual Review of Fluid Mechanics*, 31:385–416, 1999.

P. H. van Cittert. Zum einflub der spaltbreite auf die intensitatsverteilung in spektralinien. ii. *Z. Physics*, 69:298, 1931.

O. V. Vasilyev, T. S. Lund, and P. Moin. A general class of commutative filters for les in complex geometries. *Journal of Computational Physics*, 146:82–104, 1998.

S. Venkataswaran and C. L. Merkle. Efficiency and accuracy issues in contemporary cfd algorithms. AIAA Paper 2000-2251, 2000.

M. R. Visbal and D. V. Gaitonde. Shock capturing using compact-differencing-based methods. AIAA Paper 2005-1265.

J. von Neumann and R. D. Richtmyer. A method for the numerical calculation of hydrodynamic shocks. *Journal of Applied Physics*, 21, 1950.

H. von Storch. Construction of optimal numerical filters fitted for noise damping in numerical simulation models. *Contributions to Atmospheric Physics*, 51(2): 189–197, 1978.

A. W. Vreman. The adjoint filter operator in large-eddy simulation of turbulent flow. *Physics of Fluids*, 16(6):2012–2022, 2004.

A. W. Vreman, R. J. M. Bastiaans, and B. J. Geurts. A similarity subgrid model for premixed turbulent combustion. *Flow Turbulence Combustion*, 82:233–248, 2009.

B. Vreman, B. Geurts, and H. Keurten. A priori tests of large-eddy simulation of the compressible plane mixing layer. *Journal of Engineering Mathematics*, 29: 299–327, 1995.

B. Vreman, B. Geurts, and H. Keurten. Large-eddy simulation of the temporal mixing layer using the clark model. *Theoretical and Computational Fluid Dynamics*, 8:309–324, 1996.

G. S. Winckelmans, A. A. Wray, O. V. Vasilyev, and H. Jeanmart. Explicit-filtering large-eddy simulation using the tensor-diffusivity model supplemented by a dynamic smagorinksy term. *Physics of Fluids*, 13(5):1385–1403, 2001.

H. C. Yee, N. D. Sandham, and M. J. Djomehri. Low-dissipative high-order shock-capturing methods using characteristic-based filters. *Journal of Computational Physics*, 150:199–238, 1999.

D. W. Zingg, H. Lomax, and H. Jurgens. High-accuracy finite-difference schemes for linear wave propagation. *SIAM Journal of Scientific Computing*, 17(2): 328–346, 1996.