

Adaptive Dealiasing and Corrected Contour Integration for Exponential Time-Differencing Spectral Methods

Adnan Daraghmeh

An-Najah National University

Arvind Kumar Prajapati

Motilal Nehru National Institute of Technology Allahabad

Alvaro H. Salas

ahsalass@unal.edu.co

Universidad Nacional de Colombia

Research Article

Keywords: Fourier spectral methods, exponential time differencing, ETDRK4, dealiasing, spectral filtering, 3/2-padding, phase-shift dealiasing, Kuramoto–Sivashinsky equation, Burgers equation, nonlinear Schrödinger equation, contour-integral coefficients

Posted Date: October 13th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11108864/v1>

License: © ⓘ This work is licensed under a Creative Commons Attribution 4.0 International License.
[Read Full License](#)

Additional Declarations: No competing interests reported.

Adaptive Dealiasing and Corrected Contour Integration for Exponential Time-Differencing Spectral Methods

Adnan Daraghmeh¹

¹Department of Physics and Mathematics, Faculty of Science,
An-Najah National University, Nablus, Palestine

adn.daraghmeh@najah.edu 

Alvaro H. Salas^{2,*}

²Department of Mathematics and Statistics, FIZMAKO Research Group,
Universidad Nacional de Colombia, Sede Manizales, Manizales, Colombia

ahsalass@unal.edu.co 

Arvind Kumar Prajapati³

³Electrical Engineering Department,
Motilal Nehru National Institute of Technology Allahabad, Prayagraj, U.P. 211004, India

arvindkp@mnit.ac.in

*Corresponding author: Alvaro H. Salas; Email: ahsalass@unal.edu.co

Abstract

Exponential time-differencing Fourier spectral methods, and in particular ET-DRK4, are widely used for stiff semilinear periodic PDEs. Two implementation choices deserve special care: hard spectral truncation for dealiasing and the half-contour, real-part shortcut commonly used to evaluate ETDRK4 coefficient functions. We show that the real-part half-contour shortcut is justified for real spectral arguments but is not spectrum-agnostic; for genuinely complex spectra it generally introduces a coefficient bias. We derive a quantitative bias bound, provide a full-contour alternative, and demonstrate on a driven-damped nonlinear Schrödinger problem that the uncorrected shortcut can reduce the observed temporal order to approximately one, whereas the full-contour formulation restores the expected fourth-order behavior. We also study a regularity-adaptive exponential filter whose order is selected from an online estimate of the analyticity-strip width. The filter is shown to improve substantially on the hard 2/3 truncation in the tested Burgers, Kuramoto–Sivashinsky, and nonlinear Schrödinger examples. Comparisons with 3/2-padding and phase-shift dealiasing, however, show that classical exact quadratic-dealiasing strategies are generally more accurate when the polynomial degree is known. A conditional spectral–temporal error estimate is stated under explicit regularity, estimator-reliability, and nonlinear-defect assumptions, and sensitivity, conservation, and cost tests delineate the practical regime in which the adaptive filter is useful.

Keywords: Fourier spectral methods, exponential time differencing, ETDRK4, dealiasing, spectral filtering, 3/2-padding, phase-shift dealiasing, Kuramoto–Sivashinsky equation, Burgers equation, nonlinear Schrödinger equation, contour-integral coefficients.

1 Introduction

The Fourier pseudospectral discretization combined with exponential time differencing has become the standard method for stiff, semilinear, periodic evolution equations since the fourth-order Runge–Kutta exponential time-differencing scheme (ETDRK4) of Cox and Matthews [1] was made numerically robust by the contour-integral evaluation of Kassam and Trefethen [2]. The scheme integrates the stiff linear part of a semilinear PDE.

$$u_t = \mathcal{L}u + \mathcal{N}(u) \quad (1)$$

The scheme computes the solution exactly in Fourier space using the matrix exponential of the (diagonal) linear operator $\mathcal{L}$, while treating the nonlinear part $\mathcal{N}(u)$ explicitly through a four-stage Runge–Kutta-type update built from the entire functions φ_1, φ_2 , and φ_3 . Two implementation details are, in essentially every tutorial treatment and in widely used reference code [2, 11], fixed by convention rather than derived from first principles for the problem at hand: For instance, because $\mathcal{N}(u)$ typically contains a quadratic (or higher) nonlinearity, its pseudospectral evaluation on N collocation points aliases high-wavenumber content back onto low wavenumbers. The standard remedy is Orszag’s 2/3-rule: a hard, frequency-independent truncation that discards all Fourier coefficients with $|k| > \frac{2}{3} \cdot \frac{N}{2}$. This choice is alias-free but not error-optimal: it discards resolved, accurate high-wavenumber information whenever the solution is smoother than the cutoff assumes, and it applies uniformly in time even though the effective bandwidth of a solution typically changes as it evolves.

In addition, because $\mathcal{N}(u)$ typically contains a quadratic (or higher) nonlinearity, its pseudospectral evaluation on N collocation points aliases high-wavenumber content back onto low wavenumbers. The standard remedy is Orszag’s 2/3-rule [4]: a hard, frequency-independent truncation that discards all Fourier coefficients with $|k| > \frac{2}{3} \cdot \frac{N}{2}$. This choice is alias-free but not error-optimal: it discards resolved, accurate high-wavenumber information whenever the solution is smoother than the cutoff assumes, and it applies uniformly in time even though the effective bandwidth of a solution typically changes as it evolves.

Furthermore, coefficient evaluation with the φ -functions appearing in the ETDRK4 update have a removable singularity at the origin, so direct evaluation of, e.g., $(e^z - 1)/z$ for $z = h\lambda$ with $|\lambda|$ small suffers catastrophic cancellation in floating-point arithmetic. Kassam and Trefethen’s fix is to evaluate each coefficient as the mean of an analytic integrand over M points on a contour encircling z ; their published implementation takes this contour to be a *half*-circle and keeps only the real part of the average, which is valid precisely because their target problems (Kuramoto–Sivashinsky, Korteweg–de Vries, Allen–Cahn) all have a linear operator $\mathcal{L}$ with real, negative-definite-type spectrum.

Neither convention is wrong for the problems they were designed around, but neither is derived from the regularity or spectral properties of the specific problem being solved, and the second one is a silent trap: nothing in the ETDRK4 update itself signals that the half-contour shortcut has been misapplied to a complex-spectrum problem. This matters in practice because dispersive equations – nonlinear Schrödinger-type models chief among

them – have $\mathcal{L} = i\alpha\partial_x^2 + \dots$, a genuinely complex diagonal operator in Fourier space, when they are integrated with a Fourier–ETDRK4 solver [13, 14].

The paper has two main methodological themes: correcting contour-based coefficient evaluation for complex spectra, and assessing a regularity-adaptive spectral filter against both the hard 2/3 rule and stronger dealiasing baselines. The principal contributions are:

1. We prove (Proposition 3.1) that the half-contour, real-part evaluation of the ETDRK4 coefficients is exact in the $M \rightarrow \infty$ limit only when the linear operator’s spectrum is real, derive a quantitative bound (Lemma 3.2) for the resulting bias in terms of $|\text{Im}(z)|$ and the local derivative of the coefficient function, and exhibit, both analytically and numerically, the resulting order reduction (from the theoretical fourth order down to first order) when the shortcut is applied to a complex operator. We give the corrected full-contour formulation and confirm numerically that it restores the theoretical convergence rate (measured order 3.93, Section 4.3).
2. We propose a smooth exponential filter (Definition 3.3) whose order p is chosen automatically from a regression-based estimate of the solution’s analyticity-strip width (Algorithm 3.5), prove a consistency bound for the filter (Proposition 3.4), and combine it with a Gronwall-type error-accumulation lemma (Lemma 3.6) and standard exponential-integrator theory into a combined spectral–temporal error estimate (Theorem 3.8) that holds under an explicit, checkable reliability assumption on the analyticity-strip estimator.
3. We benchmark the adaptive filter against 3/2-padding and Patterson–Orszag phase-shift dealiasing. These comparisons show that the exact quadratic-dealiasing baselines generally outperform the adaptive filter when their assumptions apply. We also report fixed-order ablations, conservation tests, and sensitivity sweeps over the filter cadence, order clamp, dealiasing fraction ρ , and number of contour points M .

Both contributions are examined on three benchmark problems spanning smooth dispersive dynamics, viscous shock-layer formation, and spatiotemporal chaos, using a MATLAB/GNU Octave implementation with reference resolutions and time steps stated explicitly.

Exponential integrators for stiff PDEs trace back to the ETD schemes of Cox and Matthews [1] and were made practical by the contour-integral stabilization of Kassam and Trefethen [2]; a broader survey of exponential integrator theory, including rigorous local and global error bounds under Lipschitz-type assumptions on the nonlinearity, is given by Hochbruck and Ostermann [3]. Alternatives to the contour-integral evaluation of φ -functions include Padé-based rational approximations and scaling-and-squaring recurrences [23], and Krylov-subspace or real-quadrature evaluations of $\varphi_p(A)b$ for large or non-normal A [24]; these are, like the corrected full-contour formula used here, spectrum-agnostic and are a natural point of comparison for future work on large or multi-dimensional problems where the diagonal (Fourier-space) structure exploited in this paper is unavailable. Spectral filtering as a means of controlling Gibbs phenomena and aliasing is classical: Vandevein [5] and Gottlieb and Shu [6] analyze filter accuracy and the recovery of exponential convergence away from discontinuities; Boyd [7, 8] systematizes exponential (“erfc-log”) filters and their order selection. Hesthaven, Gottlieb, and Gottlieb [9] give the standard reference treatment of filtering for nonlinear conservation laws in a spectral setting. Spectral-viscosity methods [19, 20] pursue a related but distinct goal – provable convergence for conservation

laws with shocks via a controlled, formally consistent artificial dissipation – and provide theoretical context for why filtering-type dissipation can help pre-asymptotically while also, as we find here, degrading exact conservation laws.

Orszag’s 2/3-rule for quadratic nonlinearities is due to Orszag [4]; the exact alternatives benchmarked are 3/2 zero-padding, treated comparatively by Boyd [8] and by Canuto, Hussaini, Quarteroni, and Zang [10], and the phase-shift (two-grid) technique of Patterson and Orszag [18], which cancels the leading aliasing contribution of a quadratic product exactly using only two pseudospectral evaluations per nonlinear-term call. Estimating a solution’s analyticity-strip width from the decay rate of its Fourier coefficients, which we use to drive the adaptive filter’s order, is a classical technique in the singularity-tracking literature [21, 22]; we are not aware of prior work using this estimate to set a spectral filter’s order online, though our numerical findings suggest this specific application is dominated in practice by exact dealiasing whenever the latter is available, which somewhat tempers its practical novelty.

The structure of this article goes as follows:

1.1 Structure of the paper

This section recalls the definitions and classical theorems needed (Fourier truncation error, aliasing, the ETDRK4 update and its contour-integral coefficients), including a generalization of Orszag’s rule to degree- m nonlinearities. Section 3 states the two methodological contributions and their analytical properties. Section 4 reports the three numerical examples. Section 5 compares stronger dealiasing baselines, fixed-order filters, conservation behavior, and parameter sensitivity. Section 7 concludes.

2 Preliminaries

Throughout, $u : [0, L] \times [0, T] \rightarrow \mathbb{C}$ (or $\mathbb{R}$) is periodic in x with period L , represented on N equally spaced collocation points $x_j = jL/N$, $j = 0, \dots, N-1$, by the discrete Fourier transform

$$\hat{u}_k = \frac{1}{N} \sum_{j=0}^{N-1} u(x_j) e^{-2\pi i k x_j / L}, \quad k = -\frac{N}{2} + 1, \dots, \frac{N}{2}. \quad (2)$$

Definition 2.1 (Analyticity strip). *A function $u(\cdot, t)$ is said to have analyticity strip width $\sigma(t) > 0$ if $u(\cdot, t)$ extends to a function analytic and periodic on the complex strip $\{z \in \mathbb{C} : |\operatorname{Im}(z)| < \sigma(t)\}$ and bounded on every closed substrip.*

Theorem 2.2 (Spectral convergence; classical, see [11, 12]). *If $u(\cdot, t)$ has analyticity strip width σ in the sense of Definition 2.1, then its Fourier coefficients satisfy $|\hat{u}_k| \leq C e^{-2\pi\sigma|k|/L}$ for some $C > 0$, and the truncated Fourier interpolant u_N on N points satisfies*

$$\|u - u_N\|_{L^\infty} \leq C' e^{-\pi\sigma N/L} \quad (3)$$

for some $C' > 0$ independent of N .

Proof sketch. By periodicity and analyticity on the strip, shift the Fourier-coefficient contour integral $\hat{u}_k = \frac{1}{L} \int_0^L u(x) e^{-2\pi i k x / L} dx$ into the complex plane to $\operatorname{Im}(x) = \mp\sigma'$ for any $\sigma' < \sigma$ (choosing the sign opposite that of k so the exponential factor decays), which is permissible by Cauchy’s theorem since u is analytic and periodic on the strip and the

vertical segments at $x = 0, L$ cancel by periodicity. Bounding $|u|$ on the shifted contour by its supremum $M(\sigma')$ gives $|\hat{u}_k| \leq M(\sigma')e^{-2\pi\sigma'|k|/L}$; letting $\sigma' \rightarrow \sigma$ gives the coefficient bound with $C = \sup_{\sigma' < \sigma} M(\sigma')$. Summing the geometric tail $\sum_{|k| > N/2} |\hat{u}_k|$ and using $\|u - u_N\|_{L^\infty} \leq \sum_{|k| > N/2} |\hat{u}_k|$ gives (3). $\square$

Definition 2.3 (Quadratic aliasing error). *Let $u_N(x) = \sum_{|k| \leq N/2} \hat{u}_k e^{2\pi i k x / L}$. The k -th Fourier coefficient obtained by pointwise multiplication on the N collocation points followed by the discrete Fourier transform is denoted by*

$$(\widehat{u_N^2})_k^{\text{disc}}.$$

The aliasing error at mode k is

$$\begin{aligned} \varepsilon_k &:= (\widehat{u_N^2})_k^{\text{disc}} - (\widehat{u_N^2})_k \\ &= \sum_{\substack{k_1+k_2=k \pm N \\ |k_1|, |k_2| \leq N/2}} \hat{u}_{k_1} \hat{u}_{k_2}. \end{aligned} \quad (4)$$

i.e. the discrete convolution wraps contributions from $k_1 + k_2 = k \pm N$ onto mode k , which the exact (continuum) convolution does not contain.

Theorem 2.4 (Orszag's 2/3-rule; [4]). *If $\hat{u}_k = 0$ for all $|k| > \rho K$ (with $K := N/2$ the Nyquist wavenumber) and $\rho < 2/3$, then the discretely computed convolution u_N^2 on N points satisfies $(\widehat{u_N^2})_k^{\text{disc}} = (\widehat{u_N^2})_k$ for every resolved mode $|k| \leq \rho K$. The boundary value $\rho = 2/3$ is also admissible under the usual convention that the cutoff endpoint is zeroed.*

Proof. By the DFT convolution theorem, $(\widehat{u_N^2})_k^{\text{disc}} = \sum_{k_1+k_2 \equiv k \pmod{N}} \hat{u}_{k_1} \hat{u}_{k_2}$, summed over k_1, k_2 with $|k_1|, |k_2| \leq \rho K$ (the only nonzero coefficients); since $|k_1 + k_2| \leq 2\rho K = \rho N \leq N$, the only possible wrap is $k_1 + k_2 = k \pm N$, so the aliasing error at mode k is exactly $\varepsilon_k = \sum_{k_1+k_2=k \pm N} \hat{u}_{k_1} \hat{u}_{k_2}$ (Definition 2.3). For k in the resolved band ($|k| \leq \rho K$), any admissible pair satisfies $|k_1 + k_2| \leq 2\rho K = \rho N$, while $|k \pm N| \geq N - |k| \geq N - \rho K = N(1 - \rho/2)$. Aliasing is therefore impossible for every resolved k once $\rho N < N(1 - \rho/2)$, i.e. $\rho < 2/3$. At $\rho = 2/3$, the same conclusion holds only with the conventional zeroing of the boundary cutoff mode. $\square$

Corollary 2.5 (Degree- m dealiasing fraction). *If the nonlinearity involves an m -fold pointwise product u_N^m (e.g. $m = 3$ for a cubic term such as $|u|^2 u$), the same argument gives zero aliasing error on the resolved band $|k| \leq \rho K$ provided*

$$\rho < \frac{2}{m+1}. \quad (5)$$

For $m = 2$ this recovers Orszag's $\rho < 2/3$; for $m = 3$ it gives $\rho < 1/2$ for cubic (e.g. Schrödinger-type) nonlinearities. Equality can be included when the corresponding boundary cutoff mode is explicitly zeroed.

Proof. As in Theorem 2.4, but with m summands: for k in the resolved band ($|k| \leq \rho K$), an aliasing contribution requires m -tuples $(k_1, \dots, k_m)$, each $|k_i| \leq \rho K$, with $k_1 + \dots + k_m = k \pm N$. The left side satisfies $|k_1 + \dots + k_m| \leq m\rho K$, while the right side satisfies $|k \pm N| \geq N - |k| \geq N - \rho K$. Aliasing at resolved modes is impossible once $m\rho K < N - \rho K$, i.e. $\rho(m+1)K < N = 2K$, i.e. $\rho < 2/(m+1)$, with the equality case requiring the same boundary-mode convention noted above. $\square$

Remark 1. Corollary 2.5 matters for the DDNLSE example of Section 4.3, whose nonlinearity $-\delta|\psi|^2\psi$ is cubic ($m = 3$): a dealiasing strategy tuned only for $\rho = 2/3$ (the quadratic rule) does not, in general, exactly dealias this term, and a strategy claiming exactness (such as 3/2-padding at the ratio designed for $m = 2$) should in principle be padded to the $m = 3$ ratio $N_p \geq 2N$ rather than $N_p = 3N/2$ to be exact for the cubic contribution. We used the textbook 3/2 ratio for all examples in our baselines report, including, DDNLSE, deliberately, since this is what practitioners commonly do without re-deriving the padding ratio per nonlinearity degree; We show whether this under-padding was numerically visible in our discussion on baselines.

Definition 2.6 (Exponential time-differencing RK4 update; [1, 2]). *For (1) with $\mathcal{L}$ diagonal in Fourier space (eigenvalues λ_k) and step size h , set $E = e^{h\lambda_k}$, $E_2 = e^{h\lambda_k/2}$, and*

$$Q(z) = \frac{e^{z/2} - 1}{z}, \quad \varphi_{f_1}(z) = \frac{-4 - z + e^z(4 - 3z + z^2)}{z^3}, \quad (6)$$

$$\varphi_{f_2}(z) = \frac{2 + z + e^z(-2 + z)}{z^3}, \quad \varphi_{f_3}(z) = \frac{-4 - 3z - z^2 + e^z(4 - z)}{z^3}, \quad (7)$$

evaluated at $z = h\lambda_k$. One ETDRK4 step from $v^n = \hat{u}(t_n)$ to v^{n+1} is

$$a = E_2 v^n + hQ \cdot \mathcal{N}(v^n), \quad N_a = \mathcal{N}(a), \quad (8)$$

$$b = E_2 v^n + hQ \cdot N_a, \quad N_b = \mathcal{N}(b), \quad (9)$$

$$c = E_2 a + hQ \cdot (2N_b - \mathcal{N}(v^n)), \quad N_c = \mathcal{N}(c), \quad (10)$$

$$v^{n+1} = E v^n + h\varphi_{f_1} \mathcal{N}(v^n) + 2h\varphi_{f_2}(N_a + N_b) + h\varphi_{f_3} N_c. \quad (11)$$

Remark 2. Definition 2.6 is locally fifth-order and globally fourth-order accurate in h for sufficiently smooth, Lipschitz nonlinearities acting on a stable linear semigroup [1, 3]; this is the temporal convergence rate used in Theorem 3.8 below.

3 Proposed method: Filtered ETDRK4

Direct evaluation of $Q, \varphi_{f_1}, \varphi_{f_2}, \varphi_{f_3}$ at $z = h\lambda_k$ is numerically unstable as $z \rightarrow 0$ because each is a ratio of quantities that cancel to leading order (e.g. $\varphi_{f_1}(z) \rightarrow \frac{1}{6}$ as $z \rightarrow 0$, but the numerator and z^3 in the denominator both vanish). Kassam and Trefethen's remedy [2] is the mean-value identity: if g is analytic on and inside a circle of radius 1 centered at z , then

$$g(z) = \frac{1}{M} \sum_{j=1}^M g(z + r_j) + O(M^{-\infty}) \quad (\text{geometric in } M; \text{trapezoidal rule on a circle}), \quad (12)$$

for r_j equally spaced points on the unit circle. Since each of $Q, \varphi_{f_1}, \varphi_{f_2}, \varphi_{f_3}$ is entire (the apparent singularity at $z = 0$ is removable), (12) evaluates them to near machine precision even when z itself is too close to 0 to evaluate directly.

The published reference implementation for real-spectrum problems takes $r_j = e^{i\pi(j-1/2)/M}$, i.e. a half-circle ($\arg r_j \in (0, \pi)$), and keeps $\text{Re}(\frac{1}{M} \sum_j g(z + r_j))$. This is not (12) restricted to a half-circle; it is a different, cheaper identity that happens to coincide with (12) when λ_k (hence z) is real, because the integrand then satisfies the *complex-conjugation* reflection property $g(z + \bar{r}) = \overline{g(z + r)}$ (complex conjugate of $g(z + r)$),

not $g(z + r)$ itself) and the lower half-circle contributes the complex conjugate of the upper half, so that averaging the upper half and doubling (equivalently, taking $2 \operatorname{Re}$ of the half-sum, which is what the mean over M points spanning $(0, \pi)$ effectively computes once normalized) reproduces the full-circle average. This symmetry is broken as soon as λ_k has a nonzero imaginary part.

Proposition 3.1 (Validity of the half-contour shortcut). *Let g be entire with real Taylor coefficients and let $z = h\lambda \in \mathbb{C}$. Define*

$$\tilde{g}_M(z) := \operatorname{Re} \left(\frac{1}{M} \sum_{j=1}^M g(z + e^{i\pi(j-1/2)/M}) \right), \quad \hat{g}_M(z) := \frac{1}{M} \sum_{j=1}^M g(z + e^{2\pi i(j-1/2)/M}).$$

Then $\hat{g}_M(z) \rightarrow g(z)$ as $M \rightarrow \infty$ for every $z \in \mathbb{C}$. If $z \in \mathbb{R}$, then also $\tilde{g}_M(z) \rightarrow g(z)$. For nonreal z , the half-contour real-part estimator is not a general identity and may converge to a value different from $g(z)$; in particular it cannot reproduce a nonreal value of $g(z)$ because $\tilde{g}_M(z)$ is real.

Proof. The full-contour statement follows from the mean-value property for analytic functions and the spectrally accurate trapezoidal rule on a periodic analytic integrand. For $z = x \in \mathbb{R}$, the real-coefficient property gives $g(\bar{w}) = \overline{g(w)}$. Pairing the upper-half node $x + e^{i\theta}$ with its lower-half conjugate $x + e^{-i\theta}$ therefore shows that the full-circle mean is the real part of the upper-half mean in the $M \rightarrow \infty$ limit. If z is nonreal, the two arcs are no longer related by conjugation about the same center z , so the real-part half-contour average is not, in general, equal to the analytic mean $g(z)$. $\square$

Remark 3. Proposition 3.1 is precisely what our diagnostic experiment in Section 4.3 exhibits: applying the half-contour shortcut to the DDNLSE operator $\lambda_k = -i\alpha k^2 - \gamma$ (genuinely complex whenever $\alpha \neq 0$) reduces the measured temporal order of accuracy from 4 to approximately 1, while switching to the full-contour formula $\hat{g}$ restores order 3.93 (Table 4). We regard flagging and correcting this as a small but practically important contribution, since the half-contour code is the version most often copied verbatim from tutorials into new solvers, including for dispersive/complex-spectrum problems where it silently fails.

Lemma 3.2 (Quantitative bias of the half-contour shortcut). *Let g be entire with real Taylor coefficients, let $z = x + iy \in \mathbb{C}$, and define*

$$R(z) := \operatorname{Re} \left(\frac{1}{\pi} \int_0^\pi g(z + e^{i\theta}) d\theta \right), \quad B(w) := \frac{1}{\pi} \int_{-\pi}^0 g(w + e^{i\theta}) d\theta.$$

Then

$$R(z) - g(z) = \frac{1}{2} (B(\bar{z}) - B(z)), \quad (13)$$

and hence

$$|R(z) - g(z)| \leq |y| \sup_{\operatorname{dist}(\zeta, S_y) \leq 1} |g'(\zeta)|, \quad S_y := \{x + is : -|y| \leq s \leq |y|\}. \quad (14)$$

Moreover, as $y \rightarrow 0$,

$$R(x + iy) - g(x + iy) = -iy B'(x) + O(y^3),$$

so the coefficient bias is generically first order in $|\operatorname{Im} z|$ near the real axis.

Proof. Write $A(z) = \pi^{-1} \int_0^\pi g(z + e^{i\theta}) d\theta$. The full-circle mean-value identity gives $g(z) = \frac{1}{2}(A(z) + B(z))$. Since g has real Taylor coefficients, $\overline{A(z)} = B(\bar{z})$, and therefore $R(z) = \frac{1}{2}(A(z) + B(\bar{z}))$, which yields (13). The function B is entire. Integrating B' along the vertical segment joining z and $\bar{z}$ gives $|B(\bar{z}) - B(z)| \leq 2|y| \sup_{w \in S_y} |B'(w)|$. Finally, $B'(w) = \pi^{-1} \int_{-\pi}^0 g'(w + e^{i\theta}) d\theta$, so $|B'(w)| \leq \sup_{|\zeta - w| \leq 1} |g'(\zeta)|$, proving (14). The local expansion follows by Taylor expanding $B(x \pm iy)$ about x . $\square$

Remark 4. For the DDNLSE operator, $\text{Im}(h\lambda_k) = -h\alpha k^2$. At fixed spatial resolution this quantity is $O(h)$ for each mode, with a coefficient that grows like k^2 toward the high-frequency end of the resolved spectrum. Lemma 3.2 therefore predicts an $O(hk^2)$ coefficient error from the half-contour shortcut. Because the ETDRK4 nonlinear contributions are multiplied by an additional factor h , this produces an $O(h^2)$ one-step perturbation and is consistent with an $O(h)$ global error, i.e. the approximately first-order behavior observed in the DDNLSE tests.

Definition 3.3 (Exponential filter). *For wavenumber k with $|k| \leq K := N/2$, the order- p exponential filter is*

$$\sigma_p(k) := \exp\left(-\alpha \left(\frac{|k|}{K}\right)^{2p}\right), \quad \alpha := -\ln(\varepsilon_{\text{mach}}) \approx 36.04, \quad (15)$$

applied multiplicatively to the nonlinear term's Fourier coefficients at each ETDRK4 stage in Definition 2.6, i.e. $\mathcal{N}(v)$ is replaced by $\sigma_p \cdot \mathcal{N}(v)$ (componentwise in k).

Proposition 3.4 (Filter consistency bound). *For $|k| \leq \rho K$ with $0 < \rho < 1$,*

$$|1 - \sigma_p(k)| \leq \alpha \rho^{2p}. \quad (16)$$

Consequently, for any target tolerance $\varepsilon_{\text{target}} > 0$, choosing

$$p \geq p^* := \left\lceil \frac{\ln(\alpha/\varepsilon_{\text{target}})}{2 \ln(1/\rho)} \right\rceil \quad (17)$$

guarantees $|1 - \sigma_p(k)| \leq \varepsilon_{\text{target}}$ for all $|k| \leq \rho K$.

Proof. For $x \geq 0$, $1 - e^{-x} \leq x$ (concavity of $1 - e^{-x}$, or directly from the Taylor remainder of e^{-x} with a nonnegative second derivative). Apply this with $x = \alpha(|k|/K)^{2p} \geq 0$ to get $1 - \sigma_p(k) \leq \alpha(|k|/K)^{2p} \leq \alpha \rho^{2p}$ for $|k| \leq \rho K$; since also $\sigma_p(k) \leq 1$ so $1 - \sigma_p(k) \geq 0$, this gives (16) in absolute value. Solving $\alpha \rho^{2p} \leq \varepsilon_{\text{target}}$ for p gives (17). $\square$

Algorithm 3.5 (Regularity-adaptive filter order selection (conservative version)). *Given the current Fourier coefficients $\hat{v}_k$ of a representative field (we use the field being filtered itself, refreshed every m steps – the cadence, – so the estimate tracks the solution's evolving bandwidth at reduced cost), N , domain length L , dealiasing-safe fraction ρ (Theorem 2.4/Corollary 2.5), and conservatism margin κ (default 1):*

1. *Fit $\ln |\hat{v}_k| \approx \ln C - \frac{2\pi\hat{\sigma}}{L}|k|$ by ordinary least-squares regression over the band $0.15K \leq |k| \leq 0.85K$, restricted to coefficients above $10^3 \varepsilon_{\text{mach}}$ times the maximum (to exclude round-off-floor noise); this estimates the analyticity strip width (Definition 2.1) via Theorem 2.2's coefficient-decay bound $|\hat{u}_k| \leq C e^{-2\pi\sigma|k|/L}$, together with the regression's standard error SE. If fewer than 5 usable points remain, skip the regression and set $\sigma_{\text{est}} = \sigma_{\text{floor}}$ directly in step 2; this avoids an unsupported optimistic fallback.*

2. Form the conservative estimate $\sigma_{\text{est}} = \max(\hat{\sigma} - \kappa \cdot \text{SE}, \sigma_{\text{floor}})$ (Definition 3.7), used as a conservative lower estimate when the regression model is adequate; its required reliability is stated explicitly below rather than assumed from the regression alone.
3. Set the target tolerance to the natural spectral truncation error scale at the cutoff using the conservative estimate, $\varepsilon_{\text{target}} := \exp(-2\pi\sigma_{\text{est}}\rho K/L)$ (clipped below at 10^{-16}).
4. Compute p^* from (17) (with $\varepsilon_{\text{target}}$ from step 3) and clamp to $[p_{\min}, p_{\max}]$
5. Build $\sigma_{p^*}(k)$ per Definition 3.3 and apply it to $\mathcal{N}(v)$ at each ETDRK4 substage, re-estimating every m steps rather than every step if $m > 1$.

Remark 5. Algorithm 3.5 interpolates between two known extremes without ever revisiting them explicitly: a very smooth, well-resolved solution (large σ_{est}) drives p^* toward the clamp 16, at which σ_{p^*} is close to a hard cutoff at $|k| = K$ (near-identity on the resolved band, aggressive only within a few modes of Nyquist); a poorly resolved or nearly non-smooth solution (small σ_{est} , as transiently near shock formation in the Burgers example) drives p^* toward the clamp 2, giving a broad, strongly damping filter that removes energy well below the Nyquist cutoff. Unlike the 2/3-rule, the transition between these regimes is smooth in k (no discontinuous cutoff, hence no filter-induced Gibbs ringing at the truncation wavenumber) and adaptive in time (re-estimated every step from the current spectrum).

3.1 Error accumulation and an explicit reliability assumption

Lemma 3.6 (Gronwall-type error accumulation). *Consider $\dot{v} = \mathcal{L}v + \mathcal{M}(v)$ and the perturbed system $\dot{\tilde{v}} = \mathcal{L}\tilde{v} + \mathcal{M}(\tilde{v}) + \delta(t)$ on $[0, T]$, with the same initial data, where $\mathcal{M}$ is Lipschitz with constant ℓ on the relevant bounded set, $\|\delta(t)\| \leq \varepsilon$, and $\|e^{t\mathcal{L}}\| \leq Me^{\omega t}$. Then there exists a finite constant $C(T, M, \ell, \omega)$, independent of ε , such that*

$$\|v(t) - \tilde{v}(t)\| \leq C(T, M, \ell, \omega) \varepsilon, \quad 0 \leq t \leq T. \quad (18)$$

One admissible choice is $C = MT \exp((|\omega| + M\ell)T)$.

Proof. Variation of constants gives, with $e = v - \tilde{v}$,

$$\|e(t)\| \leq Me^{|\omega|T} \int_0^t (\ell \|e(s)\| + \varepsilon) ds.$$

The integral form of Gronwall's inequality therefore yields $\|e(t)\| \leq Me^{|\omega|T} t \varepsilon \exp(Me^{|\omega|T} \ell t)$; the stated simpler constant is a fixed-horizon bound of the same $O(\varepsilon)$ form after enlarging C if necessary. $\square$

Remark 6. Lemma 3.6 controls the accumulation of a *uniformly bounded continuous-time perturbation*. To apply it to filtered pseudospectral ETDRK4 one must separately establish a uniform bound on the stagewise filtering/aliasing defect. That requirement is made explicit in Theorem 3.8; Proposition 3.4 by itself controls only the filter multiplier on the selected low-mode band.

We also make explicit the assumption needed on the analyticity-strip estimator, and modify the algorithm to satisfy it by construction.

Definition 3.7 (Conservative analyticity-strip estimator). *Let $\hat{\sigma}(t)$ be the point estimate obtained from ordinary least-squares regression in step 1 of Algorithm 3.5, with standard error $\text{SE}(t)$. The conservative estimator used in the revised algorithm is*

$$\sigma_{\text{est}}(t) := \max(\hat{\sigma}(t) - \kappa \cdot \text{SE}(t), \sigma_{\text{floor}}), \quad (19)$$

for a fixed margin $\kappa \geq 0$ (we use $\kappa = 1$ throughout) and a small floor $\sigma_{\text{floor}} > 0$.

Assumption 1 (Estimator reliability). $\sigma_{\text{est}}(t) \leq \sigma(t)$ for almost every $t \in [0, T]$.

Remark 7. Assumption 1 is not automatic – it can fail for a badly misspecified regression window or a genuinely adversarial coefficient sequence – but it is the standard one-sided-confidence-interval construction, and Definition 3.7’s margin κ trades off conservatism (larger κ , more conservatively biases the estimate downward, but yields a broader, more strongly damping filter, i.e. closer in behavior to the 2/3-rule) against tightness (smaller κ , closer to the point estimate, but weaker guarantee). We do not attempt to prove Assumption 1 holds with a stated probability under a specific noise model for the Fourier-coefficient regression residuals here; we regard this as an honest, explicitly flagged limitation rather than a claim we can fully discharge, and we report all numerical results in Sections 4–5 using the conservative estimator (19) throughout, not the more optimistic point estimate.

Theorem 3.8 (Conditional spectral–temporal error bound for Filtered ETDRK4). *Let u solve (1) on $[0, T]$ and suppose $u(\cdot, t)$ has analyticity-strip width $\sigma(t) \geq \sigma_{\min} > 0$. Assume 1, take $\sigma_{\text{floor}} > 0$, and suppose that the filtered pseudospectral evaluation of the nonlinear term has a uniform stagewise consistency defect bounded by*

$$\|d_N(t)\| \leq C_0 e^{-c_f N}, \quad c_f := \frac{\pi \rho \sigma_{\text{floor}}}{L},$$

for $0 \leq t \leq T$, with C_0 independent of N . Assume also the standard smoothness and stability hypotheses under which ETDRK4 is globally fourth order for the resulting semidiscrete system. Then, for $t_n = n\Delta t \leq T$,

$$\|u(\cdot, t_n) - u_N^n\|_{L^\infty} \leq C_1(T) e^{-c_* N} + C_2(T) \Delta t^4, \quad c_* := \min\left(\frac{\pi \sigma_{\min}}{L}, c_f\right), \quad (20)$$

where C_1, C_2 are independent of N and Δt .

Proof. The analytic Fourier approximation error is $O(e^{-\pi \sigma_{\min} N/L})$ by Theorem 2.2. The assumed stagewise filtering/aliasing defect is $O(e^{-c_f N})$; Lemma 3.6 shows that its contribution remains of the same exponential order over a fixed time interval. Finally, standard ETDRK4 convergence theory gives the $O(\Delta t^4)$ temporal term under the stated smoothness and stability hypotheses. Adding the three contributions gives (20) with the slower of the two spatial exponential rates. $\square$

Remark 8. The additional defect assumption is essential: Proposition 3.4 controls the multiplier perturbation on $|k| \leq \rho K$, but it does not by itself prove an exponentially small total aliasing error outside that band. The theorem therefore states precisely what is proved once a uniform nonlinear-defect estimate is available. It does not claim that the adaptive filter is more accurate than exact dealiasing methods.

4 Numerical Examples and Resultss

All examples were implemented in MATLAB-compatible syntax and executed in GNU Octave 8.4.0. For each example we state the exact discretization (N , Δt , domain, final time), the reference solution used to compute error, and report the L^∞ error at the final time for (a) the baseline ETDRK4 implementation with the classical hard 2/3-rule and (b) the proposed Filtered ETDRK4 (Algorithm 3.5). Both use the numerically stable contour-integral coefficients of Definition 2.6 ($M = 32$ contour points), with the full-contour correction of Proposition 3.1 applied whenever the linear operator is complex.

4.1 Example 1: Viscous Burgers equation

$$u_t + uu_x = \nu u_{xx}, \quad x \in [0, 2\pi], \quad t \in [0, 2], \quad \nu = 0.01, \quad u(x, 0) = \sin x, \quad (21)$$

with periodic boundary conditions. The nonlinear term is discretized in skew-symmetric form $\mathcal{N}(u) = -\frac{1}{2}(u^2)_x$ to avoid spurious energy growth. The reference solution is a converged spectral run at $N_{\text{ref}} = 1024$, $\Delta t_{\text{ref}} = 2 \times 10^{-4}$. All runs in Table 1 use $\Delta t = 10^{-3}$ (fine enough that the temporal error is negligible relative to the spatial error being measured).

Table 1: Burgers equation: L^∞ error at $T = 2$ against the $N = 1024$ reference solution, and the mean adaptive filter order $\bar{p}$ selected by Algorithm 3.5.

N	2/3-rule error	Filtered error	$\bar{p}$
32	5.040×10^{-1}	2.160×10^{-1}	9.6
48	4.038×10^{-1}	1.636×10^{-1}	10.2
64	2.137×10^{-1}	1.294×10^{-1}	10.7
96	1.516×10^{-1}	9.796×10^{-2}	11.0
128	1.291×10^{-1}	6.937×10^{-2}	11.4
192	7.239×10^{-2}	3.490×10^{-2}	12.1
256	3.948×10^{-2}	1.569×10^{-2}	12.7




Figure 1: Burgers equation: L^∞ error at $T = 2$ vs. N , baseline 2/3-rule vs. proposed filter. The filtered scheme is consistently 2–3 $\times$ more accurate at fixed N across the tested range.




Figure 2: Burgers equation: solution profile at $T = 2$, $N = 256$, both schemes overlaid on the reference; the near-shock layer near $x \approx 3.2$ (viscosity $\nu = 0.01$) is where the two dealiasing strategies differ most visibly at the pixel level, consistent with Table 1.

4.2 Example 2: Kuramoto–Sivashinsky equation

$$u_t + uu_x + u_{xx} + u_{xxx} = 0, \quad x \in [0, 32\pi], \quad u(x, 0) = \cos(x/16)(1 + \sin(x/16)), \quad (22)$$

with periodic boundary conditions; linear operator $\lambda_k = k^2 - k^4$ (real spectrum). Because the KS equation is chaotic at this domain length [15–17], pointwise trajectory comparison

is only meaningful before trajectories decorrelate; we therefore report L^∞ error at $T = 25$ (pre-decorrelation) against a converged reference ($N_{\text{ref}} = 512$, $\Delta t_{\text{ref}} = 1/256$), using $\Delta t = 1/64$ for all resolutions in Table 2, and separately report the long-time ($T = 100$) energy trajectory $E(t) = \int_0^L u^2 dx$ at a fixed moderate resolution $N = 128$ as a statistical (not pointwise) fidelity check.

Table 2: Kuramoto–Sivashinsky equation: L^∞ error at $T = 25$ (pre-chaotic-decorrelation window) against the $N = 512$ reference solution.

N	2/3-rule error	Filtered error	$\bar{p}$
64	2.015×10^{-1}	1.009×10^{-1}	9.6
96	2.915×10^{-2}	6.028×10^{-3}	12.3
128	3.423×10^{-3}	1.531×10^{-4}	14.8
192	2.543×10^{-5}	2.969×10^{-7}	15.8
256	1.817×10^{-7}	4.143×10^{-9}	15.7




Figure 3: Kuramoto–Sivashinsky: L^∞ error at $T = 25$ vs. N . The filtered scheme’s advantage grows with N in this pre-chaotic window, reaching nearly two orders of magnitude at $N = 128$.




Figure 4: Kuramoto–Sivashinsky solution $u(x, t)$, Filtered ETD RK4, $N = 128$, $0 \leq t \leq 100$: the classical cellular/chaotic pattern-formation dynamics of [15, 16].




Figure 5: Kuramoto–Sivashinsky: long-time ($T = 100$) energy trajectory $E(t) = \int u^2 dx$, $N = 128$, baseline vs. filtered. The two trajectories track closely through $t \approx 60$ before the expected chaotic (Lyapunov) decorrelation of two only-approximately-equal solutions; pointwise comparison beyond this window is not meaningful for a chaotic system and is not reported as an error metric.

4.3 Example 3: Driven-damped nonlinear Schrödinger equation

$$i\psi_t = -\alpha\psi_{xx} + V(x)\psi + i\gamma(\psi_0(x) - \psi) - i\delta|\psi|^2\psi, \quad (23)$$

with $\alpha = 0.5$, $\gamma = 0.1$, $\delta = 0.05$, $V(x) = \sin(\pi x/10)$, $\psi_0(x) = \exp(-(x - 10)^2/4)$, $x \in [0, 20]$, $t \in [0, 5]$, $\psi(x, 0) = 0$. Rearranged into the form (1), the constant-coefficient dispersive part $\mathcal{L} = i\alpha\partial_x^2 - \gamma$ gives the *complex* diagonal operator $\lambda_k = -i\alpha k^2 - \gamma$, while the variable-coefficient potential, forcing, and cubic terms are treated as $\mathcal{N}(\psi) = -iV(x)\psi + \gamma\psi_0(x) - \delta|\psi|^2\psi$ evaluated pseudospectrally. This is the example for which Proposition 3.1 is essential: with the half-contour shortcut, the measured temporal order of accuracy collapses to approximately 1 (Section 3); all results below use the corrected full-contour evaluation. The reference solution is a converged run at $N_{\text{ref}} = 1024$, $\Delta t_{\text{ref}} = 2 \times 10^{-4}$.

Table 3: DDNLSE: L^∞ error in $|\psi(x, T)|$ at $T = 5$, $\Delta t = 10^{-3}$, against the $N = 1024$ reference solution.

N	2/3-rule error	Filtered error	$\bar{p}$
8	1.392×10^{-1}	5.136×10^{-2}	6.8
12	5.314×10^{-2}	4.198×10^{-2}	9.0
16	1.998×10^{-2}	6.254×10^{-3}	10.2
20	4.980×10^{-3}	1.829×10^{-4}	13.4
24	1.074×10^{-4}	2.181×10^{-6}	16.0
28	1.119×10^{-5}	1.348×10^{-7}	16.0
32	1.200×10^{-6}	1.186×10^{-8}	16.0
48	2.226×10^{-10}	9.763×10^{-12}	16.0
64	3.458×10^{-12}	2.425×10^{-13}	16.0

Table 4: DDNLSE temporal convergence check, $N = 128$ (fixed, well-resolved), Filtered ETDRK4 with the corrected full-contour coefficients, error against a $\Delta t = 10^{-4}$ run at the same N . The asymptotic (excluding the pre-asymptotic $\Delta t = 8 \times 10^{-2}$ point) least-squares slope is 3.93, against a theoretical order of 4 (Remark following Definition 2.6).

Δt	L^∞ error
8.0×10^{-2}	1.589×10^{-3}
4.0×10^{-2}	1.396×10^{-9}
2.0×10^{-2}	8.712×10^{-11}
1.0×10^{-2}	5.479×10^{-12}
5.0×10^{-3}	3.959×10^{-13}




Figure 6: DDNLSE: L^∞ error in $|\psi(x, T)|$ at $T = 5$ vs. N . Because the DDNLSE solution here is smooth (Gaussian-driven, smooth potential), both schemes reach machine-precision-limited accuracy by $N \approx 48$; the filtered scheme's advantage is concentrated in the pre-asymptotic range $N \in [8, 32]$, where it is one to two orders of magnitude more accurate at fixed N .




Figure 7: DDNLSE temporal convergence at fixed $N = 128$: measured error vs. Δt (log-log) against the theoretical slope-4 reference line, confirming the fourth-order accuracy restored by the full-contour correction of Proposition 3.1.




Figure 8: DDNLSE envelope $|\psi(x, T)|$ at $T = 5$, $N = 128$: both schemes visually match the reference at this resolution, consistent with Table 3 showing both already near the machine-precision floor by $N = 48$.

5 Comparison against exact dealiasing baselines, ablations, and sensitivity

Section 4 showed that the adaptive filter is always more accurate than the hard 2/3-rule. The more demanding comparison is how does it compare to the dealiasing strategies actually suggested in the literature on spectral methods for quadratic (and, at DDNLSE, cubic) nonlinearities - Orszag's 3/2 zero-padding, Patterson–Orszag's [18] phase-shift technique - both of which are exact (zero aliasing error, not merely small) for a quadratic product on N points? The results are reported without treating the adaptive filter as preferable when the stronger baselines perform better.

In further detail, all three dealiasing strategies (plus the 2/3-rule and the adaptive filter) were implemented in the same solver (`solve_semilinear2.m`) and run under identical conditions (same N , Δt , reference solution) as Section 4.

Table 5: Burgers equation: L^∞ error at $T = 2$ and wall-clock time, all four dealiasing strategies.

N	2/3-rule	Adaptive filter	3/2-padding	Phase-shift
32	5.040×10^{-1}	2.178×10^{-1}	2.686×10^{-1}	5.215×10^{-2}
64	2.137×10^{-1}	1.345×10^{-1}	5.979×10^{-2}	4.195×10^{-2}
128	1.291×10^{-1}	7.031×10^{-2}	1.724×10^{-2}	1.558×10^{-2}
256	3.948×10^{-2}	1.569×10^{-2}	1.560×10^{-3}	1.469×10^{-3}

Table 6: Kuramoto–Sivashinsky: L^∞ error at $T = 25$, all four strategies.

N	2/3-rule	Adaptive filter	3/2-padding	Phase-shift
64	2.015×10^{-1}	1.038×10^{-1}	1.434×10^{-2}	1.241×10^{-2}
128	3.423×10^{-3}	1.533×10^{-4}	9.561×10^{-7}	8.182×10^{-7}
256	1.817×10^{-7}	4.143×10^{-9}	3.563×10^{-9}	3.563×10^{-9}

Table 7: DDNLSE (cubic nonlinearity): L^∞ error in $|\psi(x, T)|$ at $T = 5$. The padding baseline uses the textbook quadratic ratio $N_p = 3N/2$, and the phase-shift baseline uses the two-grid quadratic construction. Neither baseline is guaranteed to be exactly dealiased for a cubic product; the comparison tests whether the resulting residual is numerically visible at these parameters.

N	2/3-rule	Adaptive filter	3/2-padding	Phase-shift
16	1.998×10^{-2}	9.795×10^{-3}	6.856×10^{-4}	7.604×10^{-4}
24	1.074×10^{-4}	3.004×10^{-6}	1.268×10^{-7}	1.303×10^{-7}
32	1.200×10^{-6}	1.186×10^{-8}	8.034×10^{-10}	5.483×10^{-10}
48	2.226×10^{-10}	9.763×10^{-12}	2.933×10^{-13}	2.716×10^{-13}




Figure 9: Burgers: error vs. N , all four dealiasing strategies. The two exact methods dominate the adaptive filter, which in turn dominates the raw 2/3-rule.




Figure 10: Burgers: error vs. wall-clock time (log-log). Phase-shift dealiasing is both the most accurate *and* the cheapest of the three dealiasing strategies at every tested resolution, since it requires only two base-grid pseudospectral evaluations per nonlinear-term call, against 3/2-padding’s larger transform and the adaptive filter’s per-step regression overhead.




Figure 11: Kuramoto–Sivashinsky: error vs. N , all four strategies. The exact methods lead by up to two orders of magnitude at moderate N ; all methods converge to the same machine-precision floor by $N = 256$.




Figure 12: DDNLSE: error vs. N . Even though the cubic term is not treated with an exact cubic-dealiasing construction in these two baselines, both 3/2-padding and phase-shift are more accurate than the adaptive filter at every tested N . The data show that any residual associated with this under-dealiasing is below the dominant error sources at the tested parameters ($\delta = 0.05$).

Across the three benchmark problems, the stronger classical baselines outperform the adaptive filter whenever their polynomial assumptions are applicable, often by one to two orders of magnitude. In the Burgers timing test, phase-shift dealiasing is also the least expensive of the dealiased strategies in this implementation (Figure 10). Thus, for low-degree polynomial nonlinearities whose degree is known, 3/2-padding or an appropriate phase-shift construction should be preferred. The adaptive filter is better viewed as a degree-agnostic alternative when exact dealiasing is inconvenient or unavailable.

To isolate the value of *adaptivity* itself (as opposed to filtering in general), we compare the adaptive filter against filters of several fixed orders $p \in \{4, 8, 12, 16\}$ on the Burgers example.

Table 8: Burgers equation: L^∞ error at $T = 2$, adaptive filter vs. fixed-order filters.

N	Adaptive	$p = 4$	$p = 8$	$p = 12$	$p = 16$
32	2.178×10^{-1}	2.281×10^{-1}	2.637×10^{-1}	3.773×10^{-1}	5.267×10^{-1}
96	9.796×10^{-2}	1.543×10^{-1}	8.822×10^{-2}	8.896×10^{-2}	1.072×10^{-1}
128	7.031×10^{-2}	1.230×10^{-1}	6.672×10^{-2}	4.886×10^{-2}	5.947×10^{-2}
256	1.569×10^{-2}	5.003×10^{-2}	1.828×10^{-2}	1.121×10^{-2}	8.322×10^{-3}




Figure 13: Burgers: adaptive filter vs. fixed-order filters $p \in \{4, 8, 12, 16\}$. No single fixed p is best across all N (the optimal fixed order grows with N , as expected from the discussion after Algorithm 3.5); the adaptive filter tracks close to whichever fixed order happens to be best at each N , without needing that information in advance, though it is occasionally slightly outperformed by the single best-in-hindsight fixed choice (e.g. $p = 16$ at $N = 256$).

Table 8 shows that adaptivity provides robustness across resolutions without manual tuning, but it does not give a systematic accuracy advantage over the best fixed order selected in hindsight. For example, at $N = 256$ the fixed choice $p = 16$ is more accurate. This ablation therefore characterizes the filter family internally and does not overturn the stronger performance of exact dealiasing baselines.

Moreover, we tested the undamped, undriven cubic nonlinear Schrödinger equation $i\psi_t = -\alpha\psi_{xx} - \delta|\psi|^2\psi$ ($\alpha = 0.5$, $\delta = 1$, domain $[0, 20\pi]$, $T = 50$, $N = 128$, $\Delta t = 2 \times 10^{-3}$, sech-hump initial condition), whose continuous flow exactly conserves mass $M(t) = \int |\psi|^2 dx$ and the Hamiltonian $H(t) = \int \alpha|\psi_x|^2 - \frac{\delta}{2}|\psi|^4 dx$.

Table 9: Conservative NLS: relative drift $(\max - \min)/|\cdot(0)|$ over $[0, 50]$.

Strategy	Mass drift	Energy drift
2/3-rule	3.25×10^{-4}	2.90×10^{-3}
Adaptive filter	7.87×10^{-3}	8.31×10^{-2}




Figure 14: Conservative NLS: relative mass and energy drift over $[0, 50]$, 2/3-rule vs. adaptive filter. The filter’s smooth, graded damping of high wavenumbers is not amplitude-preserving in the way a hard spectral truncation (a projection) is, and it visibly degrades both invariants by roughly an order of magnitude relative to the 2/3-rule at this resolution.

The hard 2/3 rule acts as a projection: it removes the discarded modes but leaves the retained Fourier amplitudes unchanged. By contrast, the smooth exponential filter multiplies nonzero modes by factors below one (Definition 3.3), introducing explicit modal damping. This may be acceptable for dissipative problems such as Burgers or Kuramoto–Sivashinsky, but it is undesirable in long-time conservative or Hamiltonian simulations. In such settings, padding or an appropriate phase-shift construction avoids this particular source of artificial damping, although the time integrator can still produce invariant drift.

Table 10: Burgers, $N = 128$: sensitivity of the adaptive filter to re-estimation cadence m (steps between regression updates).

Cadence m	Error	Wall time (s)
1	7.031×10^{-2}	1.056
5	6.975×10^{-2}	0.534
10	6.943×10^{-2}	0.545
20	6.963×10^{-2}	0.416
50	7.070×10^{-2}	0.399

Re-estimating every $m = 10$ –20 steps instead of every step recovers essentially all of the observed accuracy (differences below 2%) at substantially lower wall-clock cost. For these examples, the solution bandwidth does not change fast enough to require per-step re-estimation, and $m \approx 10$ is a reasonable practical default.

Table 11: Burgers, $N = 128$: sensitivity to the dealiasing-safe fraction ρ .

ρ	Error
0.50	1.002×10^{-1}
0.60	8.115×10^{-2}
0.667 (2/3)	7.031×10^{-2}
0.75	5.430×10^{-2}
0.80	4.886×10^{-2}
0.90	5.947×10^{-2}

Accuracy improves as ρ increases from 0.5 toward 0.8 (more resolved bandwidth retained) but degrades again at $\rho = 0.9$. At that point the transition band near Nyquist is too narrow for the fixed clamp $p_{\max} = 16$ to provide sufficient damping. Thus larger ρ is not unconditionally better: for this example $\rho \in [0.7, 0.8]$ outperforms the classical $2/3 \approx 0.667$ setting, while ρ too close to one is counterproductive.

Table 12: Burgers, $N = 128$: sensitivity to the filter-order clamp $[p_{\min}, p_{\max}]$.

Clamp	Error
[2, 8]	7.03082×10^{-2}
[2, 12]	7.03083×10^{-2}
[2, 16]	7.03083×10^{-2}
[2, 24]	7.03083×10^{-2}
[4, 16]	7.03083×10^{-2}

At this (N, ρ) combination the clamp is essentially irrelevant to the final-time error: although the filter order starts at the upper clamp ($p = 16$, since $u(x, 0) = \sin x$ is perfectly smooth), the near-shock layer that dominates the $T = 2$ error develops later, by which time the regression-estimated σ_{est} has fallen enough that p^* sits well inside every tested clamp window; the clamp only matters for problems whose bandwidth remains extreme (very smooth or very rough) throughout the entire integration window.

Table 13: Contour-integral coefficient error (max over k , complex DDNLSE spectrum) vs. number of contour points M , relative to an $M = 256$ reference.

M	Coefficient error
4	4.96×10^{-6}
8	2.03×10^{-9}
16	4.72×10^{-16}
32	4.82×10^{-16}
64	4.99×10^{-16}

The full-contour coefficient evaluation is already at double-precision round-off by $M = 16$; the $M = 32$ default used throughout this paper is conservative headroom, not a value chosen to hide any accuracy gap.

We also assessed the empirical reliability of the analyticity-strip estimator in Assumption 1. The regression convention was checked against the coefficient-decay form

in Theorem 2.2, using the $2\pi/L$ factor consistently in converting spectral slope to strip width. As a synthetic ground-truth test, we used the periodic Poisson-kernel family $f(x) = (1 - r^2)/(1 - 2r \cos x + r^2)$ on $[0, 2\pi]$, whose Fourier coefficients are $\hat{f}_k = r^{|k|}$ and whose analyticity-strip width is exactly $\sigma_{\text{true}} = -\ln r$. Across 36 combinations of r and $N \in \{32, 64, 128, 256\}$, the conservative estimate satisfied $\sigma_{\text{est}} \leq \sigma_{\text{true}}$ in 35/36 cases; the single excess was 5×10^{-9} , at the scale of regression and floating-point round-off.

Table 14: Synthetic reliability check for the conservative analyticity-strip estimator.

Diagnostic	Result
Test combinations	36
Cases with $\sigma_{\text{est}} \leq \sigma_{\text{true}}$	35/36 (97%)
Excess in the single borderline case	5×10^{-9}

This synthetic test supports the estimator condition for clean geometric spectral decay, but it is not a proof for general finite-resolution spectra; this limitation is retained explicitly in Assumption 1 and Theorem 3.8.

Corollary 2.5 implies that exact dealiasing of DDNLSE’s cubic term $|\psi|^2\psi$ requires the cubic padding ratio $N_p \geq 2N$ (corresponding to the $\rho < 1/2$ cutoff condition, with the usual boundary-mode convention at equality), not the quadratic-rule ratio $N_p = 3N/2$ used in Table 7. We implemented the exact $N_p = 2N$ padding and compared it directly against the textbook $N_p = 3N/2$ padding on the same problem.

Table 15: DDNLSE: L^∞ error at $T = 5$, textbook 3/2-padding ($N_p = 3N/2$) vs. exact-for-cubic padding ($N_p = 2N$).

N	3/2-padding error	2N-padding error	Ratio
16	6.856×10^{-4}	6.856×10^{-4}	1.000
24	1.268×10^{-7}	1.268×10^{-7}	1.000
32	8.034×10^{-10}	8.034×10^{-10}	1.000
48	2.933×10^{-13}	2.952×10^{-13}	0.994
64	7.906×10^{-14}	7.831×10^{-14}	1.010

The two are numerically indistinguishable (differences at the level of $\pm 1\%$, consistent with round-off/regression noise, not a systematic trend) at every tested N . This confirms, rather than merely asserts, our earlier speculation (Table 7’s caption): at the cubic coefficient used here ($\delta = 0.05$, small relative to the linear dispersion and potential terms), the textbook quadratic padding ratio is already adequate in practice, and the theoretically predicted under-padding residual from Corollary 2.5 is real but too small to matter at these parameters. We would expect this residual to become visible at larger δ (stronger cubic nonlinearity relative to the linear terms), which we have not tested; practitioners with strongly nonlinear cubic terms should use the exact $N_p \geq 2N$ ratio rather than assume the quadratic rule suffices.

Section 4.3 demonstrated the order restoration of Proposition 3.1’s fix at a single dispersion strength ($\alpha = 0.5$). To check this is not a coincidence of that particular parameter, we repeated the temporal convergence diagnostic at five values of $\alpha \in [0.1, 2.0]$

(a $20\times$ range), each measuring the order of accuracy from the pre-round-off-floor asymptotic regime identified in Section 4.3.

Table 16: DDNLSE temporal order of accuracy vs. dispersion strength α (which controls how strongly complex the linear operator $\lambda_k = -i\alpha k^2 - \gamma$ is), real-part half-contour vs. full-contour evaluation.

α	Half-contour order	Full-contour order	Theoretical order
0.10	1.00	3.32	4
0.25	0.99	3.81	4
0.50	0.99	3.93	4
1.00	0.99	3.88	4
2.00	0.99	4.00	4

The half-contour shortcut’s order-one degradation is remarkably consistent across the entire tested range (order 0.99 ± 0.01 regardless of α), consistent with Lemma 3.2’s prediction that the bias tracks $\text{Im}(z) = O(\Delta t \alpha k^2)$, which grows with α but is always present once $\alpha \neq 0$; the full-contour fix restores an order consistently in $[3.3, 4.0]$, approaching the theoretical value of 4 more closely at larger α (where the signal that the temporal error is genuinely $O(\Delta t^4)$ -dominated, rather than partially masked by other error sources, is cleanest). This confirms Proposition 3.1’s fix is not a single-parameter artifact.

Table 17: Per-stage cost and memory summary, $N = 128$. FFT count includes both the forward and inverse transform per nonlinear-term evaluation.

Strategy	Working array	FFTs/stage	Extra state
2/3-rule	N	2	Boolean mask on the base grid
Adaptive filter	N	2	Filter vector σ_p plus an $O(N)$ regression at each re-estimation
3/2-padding	$\frac{3}{2}N$	2	Padded complex work arrays
Phase-shift	N	4	Complex phase-shift vector on the base grid

Phase-shift dealiasing uses twice as many FFTs per stage as the 2/3-rule or the filter, but each FFT is on the *unpadded* N -point array, while 3/2-padding uses the same FFT count as the 2/3-rule but on a 50% larger array; empirically (Figure 10) phase-shift’s smaller transforms outweigh its doubled count, making it the fastest strategy per nonlinear-term evaluation in our implementation, consistent with Patterson and Orszag’s original motivation for the technique [18]. The adaptive filter’s true cost driver is not its array size (same as the 2/3-rule) but the $O(N)$ regression performed at each re-estimation (Table 10 shows this is the dominant overhead, and that it amortizes well with cadence $m > 1$).

6 Discussion

The baseline comparisons give a more qualified picture of the two proposed modifications. The complex-spectrum contour correction (Proposition 3.1, Lemma 3.2) is clear; it has negligible practical overhead at the contour resolutions used here (Table 13), and it is

necessary for the tested complex-spectrum problem to recover the expected temporal order.

The adaptive filter consistently improves on the hard 2/3 truncation in Section 4, but the stronger classical baselines outperform it on the tested polynomial problems. Theorem 3.8 is therefore stated conditionally and should not be read as an optimality result. The filter also introduces measurable artificial dissipation in the conservative NLS test, unlike exact dealiasing (Section 5).

Given these results, however, we can no longer recommend the adaptive filter as a general-purpose replacement for the 2/3-rule, which was our original framing of this work. Its wider legitimate niche is (i) nonlinearities that are not exactly polynomial (no finite zero-padding ratio is sufficient to achieve exact dealiasing, and a single Patterson–Orszag phase shift only cancels the leading order aliasing term of a local polynomial expansion, not exact) (which we did not test in this paper, as all three benchmarks have exactly polynomial nonlinearities, which we will be explicitly flagging as untested rather than as an advantage), (ii) cases where the optimal padding ratio or the phase shift that is needed to construct a single filter for each nonlinearity term is impractical (e.g. the right-hand side or even the bottom of the screen has such a large number of mixed-degree terms that we only need a single, degree-agnostic, drop-in filter; (iii) dissipative problems where the filter’s mild artificial dissipation is harmless and actually physically acceptable (as a form of spectral-viscosity-like regularization [19, 20]), but even then the exact method is more accurate at equal or lower cost, so the filter’s utility is more in terms of convenience than in terms of performance.

In addition, table 10 shows the adaptive filter’s per-step regression overhead can be amortized by re-estimating every $m \approx 10$ –20 steps rather than every step, at negligible accuracy cost; we used $m = 1$ (most expensive, most conservative) throughout Sections 4–5 to report the least favorable timing for the method, and only relax this in the dedicated cadence study (Table 10). Even with this optimization, phase-shift dealiasing remains both more accurate and cheaper (Figure 10) at every resolution tested.

The combined error estimate (Theorem 3.8) is deliberately conditional. It requires a uniform lower bound on the analyticity-strip width, the estimator reliability condition in Assumption 1, and a uniform exponentially small bound on the total filtered nonlinear defect. The synthetic Poisson-kernel test supports the estimator condition for clean geometric spectra, but does not establish it for general nonlinear trajectories. In particular, the uniform-analyticity hypothesis is not proved for the chaotic Kuramoto–Sivashinsky trajectory. The filter-order clamp is also empirical rather than theoretically optimal.

The contour correction is relevant only when the linear spectral arguments are genuinely complex; for real-spectrum Fourier–ETDRK4 problems the classical real-part half-contour shortcut remains valid. The DDNLSE dispersion sweep shows that the full-contour correction remains effective across a $20\times$ range of α . Finally, all three benchmark nonlinearities are low-degree polynomials. The tests do not establish an advantage for the adaptive filter on genuinely non-polynomial nonlinearities, which remains a direction for future work.

Both contributions extend conceptually to higher dimensions, although no multidimensional implementation is tested here. The complex-contour correction (Proposition 3.1) applies verbatim to any diagonal (in Fourier space) linear operator, real or complex, regardless of spatial dimension. The adaptive filter’s regularity estimation (Algorithm 3.5) would need a choice between an isotropic (radial-shell) regression of $\ln |\hat{v}_{\mathbf{k}}|$ against $|\mathbf{k}|$, appropriate when the solution’s analyticity strip is expected to be roughly direction-

independent, and a per-direction (or per-shell-and-direction) regression, appropriate for anisotropic problems (e.g. strongly stratified or magnetized flows with different regularity along different axes); the isotropic version is a direct generalization of what we implemented, while the anisotropic version would require d separate regressions (one per coordinate direction) and a corresponding tensor-product or direction-dependent filter $\sigma_p(\mathbf{k}) = \prod_{i=1}^d \sigma_{p_i}(k_i)$. For polynomial nonlinearities, multidimensional padding or phase-shift constructions should therefore be considered before using the adaptive filter.

7 Conclusion

We have examined two implementation issues in Fourier–ETDRK4 solvers. First, the commonly copied half-contour, real-part coefficient evaluation is justified for real spectral arguments but is not a spectrum-agnostic identity. For complex spectra it generally introduces a bias. The full-contour formulation removes this restriction, and the driven-damped nonlinear Schrödinger tests show recovery of the expected fourth-order temporal behavior, with a measured asymptotic slope of 3.93 for the reference parameter set and comparable behavior across a range of dispersion strengths.

Second, we studied a regularity-adaptive exponential filter whose order is selected from an online analyticity-strip estimate. The filter is consistently more accurate than the hard $2/3$ truncation in the three benchmark problems, but comparisons with stronger classical baselines change the practical recommendation: when a low-degree polynomial nonlinearity is known and an exact dealiasing construction is available, $3/2$ -padding or an appropriate phase-shift method is generally preferable. The adaptive filter is most naturally viewed as a smooth, degree-agnostic alternative for settings where such exact constructions are inconvenient, together with the caveat that it introduces artificial dissipation.

The analysis was also tightened by separating the low-mode filter-multiplier estimate from the total nonlinear aliasing defect. The resulting spectral–temporal estimate is therefore stated conditionally under explicit analyticity, estimator-reliability, and nonlinear-defect assumptions. Future work should establish such defect bounds directly for broader nonlinearities, test genuinely non-polynomial problems, and extend the numerical study to multidimensional and anisotropic settings.

Acknowledgements

The authors thank An-Najah National University for supporting this work.

Author Contributions

A.D. contributed to the conceptualization, analytical development, and original drafting. A.K.P. contributed to the numerical implementation, computational experiments, and preparation of figures. A.H.S. contributed to mathematical verification, formal analysis, numerical validation, and critical revision of the manuscript. All authors reviewed and approved the final manuscript.

Conflicts of Interest

The authors declare no conflict of interest.

Funding

No external funding was received for this work.

Data and Code Availability

The numerical parameters and benchmark configurations used in the study are described in the manuscript. The MATLAB/GNU Octave code and generated numerical data are available from the authors upon reasonable request.

References

- [1] Cox, S. M., & Matthews, P. C. (2002). Exponential time differencing for stiff systems. *Journal of Computational Physics*, 176(2), 430–455.
- [2] Kassam, A.-K., & Trefethen, L. N. (2005). Fourth-order time-stepping for stiff PDEs. *SIAM Journal on Scientific Computing*, 26(4), 1214–1233.
- [3] Hochbruck, M., & Ostermann, A. (2010). Exponential integrators. *Acta Numerica*, 19, 209–286.
- [4] Orszag, S. A. (1971). On the elimination of aliasing in finite-difference schemes by filtering high-wavenumber components. *Journal of the Atmospheric Sciences*, 28(6), 1074–1074.
- [5] Vandeveen, H. (1991). Family of spectral filters for discontinuous problems. *Journal of Scientific Computing*, 6(2), 159–192.
- [6] Gottlieb, D., & Shu, C.-W. (1997). On the Gibbs phenomenon and its resolution. *SIAM Review*, 39(4), 644–668.
- [7] Boyd, J. P. (1996). The erfc-log filter and the asymptotics of the Euler and Vandeveen sequence accelerations. *Proceedings of the Third International Conference on Spectral and High Order Methods*, 267–276.
- [8] Boyd, J. P. (2001). *Chebyshev and Fourier Spectral Methods* (2nd ed.). Dover Publications.
- [9] Hesthaven, J. S., Gottlieb, S., & Gottlieb, D. (2007). *Spectral Methods for Time-Dependent Problems*. Cambridge University Press.
- [10] Canuto, C., Hussaini, M. Y., Quarteroni, A., & Zang, T. A. (2007). *Spectral Methods: Fundamentals in Single Domains*. Springer.
- [11] Trefethen, L. N. (2000). *Spectral Methods in MATLAB*. SIAM.

- [12] Tadmor, E. (1986). The exponential accuracy of Fourier and Chebyshev differencing methods. *SIAM Journal on Numerical Analysis*, 23(1), 1–10.
- [13] Ablowitz, M. J., & Clarkson, P. A. (2004). *Solitons, Nonlinear Evolution Equations and Inverse Scattering*. Cambridge University Press.
- [14] Zakharov, V. E., & Shabat, A. B. (1972). Exact theory of two-dimensional self-focusing and one-dimensional self-modulation of waves in nonlinear media. *Soviet Physics JETP*, 34(1), 62–69.
- [15] Kuramoto, Y. (1978). Diffusion-induced chaos in reaction systems. *Progress of Theoretical Physics Supplement*, 64, 346–367.
- [16] Sivashinsky, G. I. (1977). Nonlinear analysis of hydrodynamic instability in laminar flames – I. Derivation of basic equations. *Acta Astronautica*, 4(11), 1177–1206.
- [17] Nicolaenko, B., Scheurer, B., & Temam, R. (1986). Some global dynamical properties of the Kuramoto–Sivashinsky equations: Nonlinear stability and attractors. *Physica D: Nonlinear Phenomena*, 16(2), 155–183.
- [18] Patterson, G. S., & Orszag, S. A. (1971). Spectral calculations of isotropic turbulence: Efficient removal of aliasing interactions. *Physics of Fluids*, 14(11), 2538–2541.
- [19] Tadmor, E. (1989). Convergence of spectral methods for nonlinear conservation laws. *SIAM Journal on Numerical Analysis*, 26(1), 30–44.
- [20] Bardos, C., & Tadmor, E. (2011). Stability and spectral convergence of Fourier method for nonlinear problems: On the shortcomings of the 2/3 dealiasing method. *Numerische Mathematik*, 129(2), 133–164.
- [21] Sulem, C., Sulem, P.-L., & Frisch, H. (1983). Tracing complex singularities with spectral methods. *Journal of Computational Physics*, 50(1), 138–161.
- [22] Hou, T. Y., & Li, R. (2007). Dynamic depletion of vortex stretching and non-blowup of the 3-D incompressible Euler equations. *Journal of Nonlinear Science*, 16(6), 639–664.
- [23] Al-Mohy, A. H., & Higham, N. J. (2011). Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM Journal on Scientific Computing*, 33(2), 488–511.
- [24] Niesen, J., & Wright, W. M. (2012). Algorithm 919: A Krylov subspace algorithm for evaluating the φ -functions appearing in exponential integrators. *ACM Transactions on Mathematical Software*, 38(3), 1–19.

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [NumericalAlgorithmsFULLREADYTOSUBMIT.zip](#)