

Geometry and finite-time delivery of prescribed Reynolds stresses by transverse waves

Jinlin Wang

wangjlin26@mail2.sysu.edu.cn

Sun Yat-sen University <https://orcid.org/0009-0001-1409-2597>

Research Article

Keywords: Reynolds stress, incompressible flow, transverse waves, realisability, filtered force, direct numerical simulation

Posted Date: October 20th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11409458/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Geometry and finite-time delivery of prescribed Reynolds stresses by transverse waves

Jinlin Wang^{1†} and Yi Heng^{2‡}

¹School of Systems Science and Engineering, Sun Yat-sen University, Guangzhou 510275, China

²School of Computer Science and Engineering, Sun Yat-sen University, Guangzhou 510006, China

Corresponding author: Yi Heng, hengyi@mail.sysu.edu.cn

Positive Reynolds-stress tensors need not be deliverable by an incompressible velocity field or retain their prescribed force under evolution. We construct periodic transverse waves and give an exact realisability criterion for a fixed three-axis basis; positive definiteness alone is insufficient. Band-limited envelopes give explicit bounds on projection error, filtered stress and force, and conditional finite-time evolution. A six-mode pulse prescribes a nonconstant stress on a Gaussian-vortex background. At $t = 0.01$, its three-dimensional filtered stress and full-force errors relative to the initial target are 0.316 % and 1.8 % across $N = 128$ –512. At $t = 0.1$, they rise to 2.54–2.55 % and 18.0 %. The $K = 6$ low-mode force-field difference at $t = 1.9$ decreases across successive grid pairs from 0.166 to 0.0114 and 0.00231 %. An independent spectral-element solver reproduces the isolated-wave cross-stress trend; additional amplitudes verify its leading ε^2 coherence loss. Matched background scans identify strain-driven evolution of the odd pulse as the largest aligned source of the early force departure. Fourfold amplitude variation changes the $t = 0.1$ force error by at most 1.23×10^{-4} percentage points. For the sampled carrier, viscosity, amplitude, filter and core, $\Omega_0 t$ separates all pass/fail cases at 1, 5 and 10 % full-force error tolerances, yielding an empirical operating window.

Keywords: Reynolds stress; incompressible flow; transverse waves; realisability; filtered force; direct numerical simulation.

1. Introduction

A Reynolds stress is more than a prescribed covariance. Its divergence enters the mean momentum equation, and its components describe how fluctuating motion transports momentum between directions (Reynolds 1895; Launder *et al.* 1975). In turbulence modelling, positive semidefiniteness is the first physical constraint on the stress, but it leaves open how a particular velocity field could realise that tensor. Classical work

[†]Jinlin Wang: ORCID <https://orcid.org/0009-0001-1409-2597>.

[‡]Yi Heng: ORCID <https://orcid.org/0000-0001-6894-6976>.

established positive-semidefinite and invariant constraints on Reynolds stresses (Lumley 1979; Schumann 1977; Pope 2000); later frame-indifferent and process constraints (Wang 1997; Mishra & Girimaji 2014) and geometric representations of anisotropy (Banerjee *et al.* 2007) refine that account. Our contribution is not the algebraic positivity condition itself. For the three coordinate axes, the admissible tensor cone is the factor-width-two cone studied in matrix theory (Boman *et al.* 2005). We express that geometry as a transverse-wave allocation criterion, then examine the projection error and finite-time stress and force supplied by the resulting velocity field.

The question matters whenever an imposed fluctuation is expected to do dynamical work. Synthetic turbulent inflows are commonly designed to match moments and correlations, but a match at insertion need not survive downstream adjustment. Recycling approaches (Lund *et al.* 1998), digital filtering (Klein *et al.* 2003), synthetic eddies (Jarrin *et al.* 2006), Fourier and related constructions (di Mare *et al.* 2006), and efficient random-field methods (Xie & Castro 2008) all address this practical difficulty from different directions. Comparative reviews (Tabor & Baba-Ahmadi 2010; Wu 2017; Dhamankar *et al.* 2018) make clear that prescribing one-point statistics is only part of an inflow problem: spatial structure, solenoidality and compatibility with the evolution matter as well. The present work does not propose a new boundary condition. It uses the inflow literature to motivate a controlled periodic initial-value experiment in which the target tensor, generated field and resulting force can be compared without an inlet adjustment region.

Several studies show why the distinction is consequential. A priori and a posteriori comparisons of inflow fields expose the rapid loss of inadequately resolved synthetic content (Keating *et al.* 2004). Corrections to synthetic-eddy normalisation improve recovery of inhomogeneous statistics (Skillen *et al.* 2016; Schau *et al.* 2022); a finite-mode Fourier generator can also carry a sizeable sampling bias even when its prescribed covariance is correct in expectation (Patterson *et al.* 2021). Divergence-free inflow formulations reduce spurious pressure responses (Kim *et al.* 2013), while later covariance-based and spectrum-controlled constructions make solenoidality an explicit design condition (Dreze *et al.* 2023; Cai *et al.* 2023; Fyhn & Fyhn 2026). Volume forcing has likewise been used to seed fluctuations within a hybrid flow calculation (Davidson & Billson 2006). These examples suggest a common issue: the stress realised by a velocity field should be assessed after the incompressible constraint has acted, and its divergence should be measured separately from the stress itself.

A second difficulty is temporal. For a Fourier mode with wavenumber k , viscosity damps velocity as $e^{-\nu k^2 t}$, hence its quadratic stress as $e^{-2\nu k^2 t}$. Nonlinear interactions can redistribute that stress among components even when the total kinetic-energy decay remains close to the viscous prediction. Direct numerical simulation is useful precisely because it resolves these mechanisms rather than replacing them with a closure (Kim *et al.* 1987; Moin & Mahesh 1998). But a small difference between two grids is not, by itself, a bound on target error, and a fitted exponential rate for one stress component need not represent a material viscosity. We therefore distinguish the analytical target, the field initially delivered after projection, and the field later produced by the Navier–Stokes evolution.

Pressure projection is the immediate obstacle. The Helmholtz–Leray projector removes the longitudinal part of each Fourier amplitude. Projection methods have a long numerical history (Chorin 1968; Brown *et al.* 2001); for stress synthesis the relevant point is algebraic: projecting a generic Cholesky factor changes its covariance at leading order. A transverse basis avoids that change for constant amplitudes, but imposes a geometric restriction on the set of tensors that the chosen directions can realise. Spatial modulation restores a nonzero

stress divergence while reintroducing a calculable projection defect. Filtering then defines a resolved stress and force; the distinction between filtered momentum flux and an unfiltered covariance is familiar from large-eddy analysis (Leonard 1975; Germano *et al.* 1991; Lilly 1992). Our use of a sharp low-mode filter is diagnostic rather than a subgrid model.

Oscillatory stresses also appear in mathematical constructions for incompressible equations. The *h*-principle viewpoint (De Lellis & Székelyhidi 2012), non-uniqueness constructions for Euler (Daneri & Székelyhidi 2017; Isett 2018), and weak-solution results for Navier–Stokes (Buckmaster & Vicol 2019) use rapidly varying fields to alter coarse momentum fluxes. Our finite-resolution, smooth-solution calculations neither reproduce those schemes nor bear on singularity formation. They isolate one step shared by many constructions: what a specified oscillatory field actually supplies as stress and force in a viscous solver.

We pursue three connected claims. First, an explicit allocation criterion for a fixed three-axis transverse basis determines whether a constant positive tensor belongs to the established factor-width-two cone and supplies infeasibility certificates. Secondly, envelope bandwidth and carrier separation give checkable projection, filtered-force and conditional finite-time bounds. Thirdly, matched numerical experiments quantify initial delivery and subsequent dynamics for a prescribed nonconstant stress. Matched 128^3 – 512^3 vortex triplets measure local three-dimensional stress and force; a separate 512^3 isolated-wave scan probes the componentwise decay that a trace-based fit misses. Section 2 defines the construction and diagnostics, §3 examines isolated waves, §4 gives the nonuniform-background results, and §5 discusses the finite-time window and limits of inference. Proofs and quantitative estimates are collected in appendix A.

2. Wave construction and measurement

2.1. Transverse covariance and its geometry

Vector and tensor symbols are written without boldface throughout; component indices distinguish, for example, the velocity vector u from its scalar component u_i , and the stress tensor R from R_{ij} . The position x , polarisation b_j and wavevector κ_j in (2.1) are vectors. On the periodic torus $\mathbb{T}^3 = (\mathbb{R}/2\pi\mathbb{Z})^3$, let v be a finite sum of oscillatory velocities. Distinct non-resonant carriers, each with a transverse polarisation $b_j \cdot \kappa_j = 0$, generate the low-frequency covariance

$$v(x) = \sqrt{2} \varepsilon \sum_j a_j(x) b_j \cos(\kappa_j \cdot x + \phi_j), \quad T_v(x) = \varepsilon^2 \sum_j a_j(x)^2 b_j b_j^\top, \quad (2.1)$$

provided that the envelope and carrier bands are separated as specified in appendix A. Equation (2.1) specifies the constructed covariance before flow evolution. For constant envelopes the field is exactly solenoidal. In contrast, applying a Cholesky factor L of an arbitrary positive tensor $T = LL^\top$ to generic carrier directions produces amplitudes b_j with longitudinal components. The Leray projector $\mathcal{P}(\kappa) = I - \kappa\kappa^\top/|\kappa|^2$ then replaces each b_j by $\mathcal{P}(\kappa_j)b_j$, changing the covariance before physical evolution begins.

For the three fixed carrier directions e_x, e_y, e_z , each wave-family covariance is a positive semidefinite block supported on the plane normal to its carrier. Write $d_i = T_{ii} > 0$ and $\rho_{ij} = T_{ij}/\sqrt{d_i d_j}$. Appendix A proves that this basis realises T if and only if there exist $p, q, r \in [0, 1]$ such that

$$pq \geq \rho_{xy}^2, \quad (1-p)r \geq \rho_{xz}^2, \quad (1-q)(1-r) \geq \rho_{yz}^2. \quad (2.2)$$

The criterion is exact for this fixed basis. In matrix terminology, its cone consists of sums of positive-semidefinite blocks supported on two coordinates, equivalently matrices of factor width at most two (Boman *et al.* 2005). The allocation variables make that existing cone constructive for these transverse carriers; the restriction does not apply to arbitrary periodic wave-vector sets. It gives a simple counterexample to the assumption that positive definiteness suffices: writing $\iota = (1, 1, 1)^\top$, the equicorrelation tensor $T_\rho = (1 - \rho)I + \rho\iota\iota^\top$ is positive definite for $0 \leq \rho < 1$, whereas the fixed three-axis basis is feasible only when $\rho \leq 1/2$. When directions may depend on the tensor, almost every positive tensor needs at least three lattice directions; a pointwise six-direction construction exists without a bound on their wavenumbers.

The numerical reference tensor is

$$T_0 = \begin{pmatrix} 1 & 0.6 & 0 \\ 0.6 & 1 & 0 \\ 0 & 0 & 0.4 \end{pmatrix}. \quad (2.3)$$

A six-mode transverse synthesis realises $\varepsilon^2 T_0$. For its particular phases and amplitudes, the isolated initial field reduces, up to stored coefficient rounding, to

$$\frac{v_0}{\varepsilon} = \begin{pmatrix} \sqrt{6/5} \cos kz + \sqrt{4/5} \cos ky \\ \sqrt{6/5} \cos kz + \sqrt{4/5} \cos kx \\ \sqrt{2/5}(\sin kx + \sin ky) \end{pmatrix}. \quad (2.4)$$

2.2. Local stress, force and errors

We use a spherical Fourier cut-off G_K on the full three-dimensional field and define

$$R_{ij}^K(u) = G_K(u_i u_j) - G_K u_i G_K u_j, \quad F_i^K(u) = -\partial_j R_{ij}^K(u). \quad (2.5)$$

This is a deterministic filtered momentum flux. The sharp Fourier filter is not positivity preserving, so R^K need not be pointwise positive semidefinite even when the prescribed tensor is positive. The filter is fixed at $K = 6$ for primary comparisons, with $K = 4, 8$ as sensitivity checks. Let $u^+(t), u^-(t), u^0(t)$ be solutions started from $\mathcal{P}(U+v), \mathcal{P}(U-v), \mathcal{P}U$, respectively. For each observable Q , the even pulse-induced part is $Q^{\text{ev}}(t) = \{Q(u^+(t)) + Q(u^-(t))\}/2 - Q(u^0(t))$. This cancels the odd-in-amplitude contribution. In the target and intergrid errors below, $Q_N(t)$ denotes this matched even observable, not the total-flow observable. For $Q \in \{R^K, F^K\}$, define $\mathcal{E}_{\text{target}}^Q(t; N) = \|Q_N(t) - Q_{\text{tar}}\|_{2,C} / \|Q_{\text{tar}}\|_{2,C}$, where $r = [(x - \pi)^2 + (y - \pi)^2]^{1/2}$, $C = \{0.3 \leq r \leq 1.2, z \in [0, 2\pi)\}$, the tensor norm is Frobenius, and Q_{tar} is the initial prescribed field. For two grids transferred to the common 128^3 Fourier grid, define $\mathcal{E}_{\text{grid}}^Q(N_1, N_2; t) = \|Q_{N_1}(t) - Q_{N_2}(t)\|_{2,C} / \|Q_{N_2}(t)\|_{2,C}$, with N_2 the finer grid. The intergrid definition also applies to the even low-mode velocity response. The projected-force target error uses the same definition with $Q = \mathcal{P}F^K$ and the projected initial target in the denominator.

For band-limited envelopes of width q_j , appendix A gives the initial projection estimate

$$\|\mathcal{P}v - v\|_2 \leq \varepsilon \sum_j \frac{q_j}{|\kappa_j| - q_j} |b_j| \|a_j\|_2, \quad |\kappa_j| > q_j. \quad (2.6)$$

The same appendix bounds the difference of filtered forces by the velocity difference times a known factor depending on K and field norms. A conditional smooth-solution estimate adds an explicit viscous term and a nonlinear finite-time term. These are absolute bounds; converting them to relative force errors also requires a lower bound on the target-force

	R_{xx}	R_{yy}	R_{zz}	R_{xy}	R_{xz}	R_{yz}
Designed	10.00	10.00	4.00	6.00	0	0
Before projection	10.00	10.00	3.97	6.00	0	0
After projection	2.50	5.20	8.07	3.00	-2.50	-4.60
Modewise prediction	2.50	5.20	8.10	3.00	-2.50	-4.60

Table 1. Constant-target Cholesky control at $N = 64$, $\varepsilon = 10^{-2}$; entries are in units of 10^{-5} . The small pre-projection discrepancy in R_{zz} reflects stored coefficient precision. For R_{zz} , the measured post-projection value differs from the modewise prediction by 0.37 %; the remaining components agree at the displayed precision.

norm. Appendix A evaluates the projection estimate for the actual nonconstant pulse and separates its velocity correction from the much looser general filtered-force bound.

2.3. Numerical setting

The unforced incompressible equations are

$$\partial_t u + (u \cdot \nabla)u = -\nabla p + \nu \Delta u, \quad \nabla \cdot u = 0,$$

with periodic boundaries and pressure divided by the constant density. The experiments use XCOMPACT3D, with sixth-order compact spatial derivatives, a three-stage Runge–Kutta advance and an FFT pressure solve (Lele 1992; Laizet & Lamballais 2009; Bartholomew *et al.* 2020; Rolfo *et al.* 2023). The domain is $(2\pi)^3$. The primary vortex runs use $\nu = 10^{-5}$, $\Delta t = 0.01$, carrier $k = 24$, amplitude $\varepsilon = 10^{-3}$, and $N = 128, 192, 256, 512$ points in each direction. The freely evolving Gaussian vortex has initial azimuthal velocity $U_\theta(r) = \Omega_0 r \exp[-r^2/(2\sigma_0^2)]$, with $\Omega_0 = 0.5$ and $\sigma_0 = 0.6$. The Gaussian formula is sampled about (π, π) and periodically repeated by the numerical discretisation; it is a confined background, not an exact smooth periodic Gaussian solution. The matched no-pulse run supplies its evolved reference for the filtered diagnostics. The explicit subgrid model and background-maintenance forcing are disabled. All local diagnostics start from the three-dimensional Cartesian fields. Four-grid target errors, intergrid differences and background scans use an identical mask on a common 128^3 Fourier grid. Older native-grid quadratures of the same physical core are identified where retained. The isolated 512^3 runs use the same ν , ε and carrier shape, with $k = 20, 24, 28, 32, 40$; they have zero background and 20 diagnostic times from 0.1 to 2. For $K \leq 8$, filtered algebra can be evaluated on an intermediate 40^3 Fourier grid: products of retained modes have highest wavenumber 16, below its Nyquist limit 20. The result is then evaluated on the common 128^3 grid for core integration. Velocity products are first formed on each native DNS grid, before this restriction; the intermediate grid is not used to form the original high-frequency products.

2.4. Projection, reference backgrounds and verification

The projection obstruction can be measured independently of subsequent dynamics. For a Cholesky synthesis of T_0 using the oblique carriers proportional to $(1, 0, 1)$, $(0, 1, 1)$ and $(1, 1, 0)$, the componentwise prediction after Leray projection captures the large stress changes in table 1. The displayed R_{zz} values, 8.07 and 8.10, differ by 0.37 % relative to the prediction; the other components agree at the displayed precision. By contrast, the transverse six-mode construction reproduces its constant covariance to relative error 2.2×10^{-16} in the Fourier construction, has zero divergence there, and retains its covariance to within 0.01 % after the solver projection. Thus the first numerical comparison tests an algebraic property of the construction, rather than a fitted flow response.

Unit-diagonal target	Fixed-axis realisable?	Certificate
I	yes	$p = q = r = 1/2$
$T_{xy} = 0.6, T_{xz} = T_{yz} = 0$	yes	$p = q = 0.6, r = 1/2$
$T_\rho, \rho = 1/2$	yes	$p = q = r = 1/2$
$T_\rho, \rho = 0.6$	no	$\text{tr}(HT_\rho) < 0$

Table 2. Exact examples for the fixed three-axis cone. Here $T_\rho = (1 - \rho)I + \rho u u^\top$. The last row is positive definite, so positivity alone is insufficient for this specified basis.

The restriction caused by the fixed transverse basis is equally concrete. Table 2 illustrates criterion (2.2) with the equicorrelation family. The dual matrix $H = 2I - u u^\top$ is nonnegative on every coordinate-plane positive semidefinite block, while $\text{tr}(HT_\rho) = 3 - 6\rho$. A negative value therefore certifies failure for this basis even though T_ρ remains positive definite when $\rho < 1$. The exact proof and the distinction between fixed directions and tensor-dependent lattice directions are given in appendix A.

For the Gaussian-vortex constant-target case, the unfiltered fluctuation is formed after subtracting the analytic viscous evolution of the initial swirl. Spatially varying targets are analysed from matched $+\varepsilon$, $-\varepsilon$ and zero-pulse runs using the filtered Cartesian definition (2.5). The cylindrical radial profiles reported below are a separate azimuthally averaged diagnostic; the three-dimensional filtered force is evaluated directly from Cartesian fields.

Two checks calibrate the evolution independently of the target comparison. For the isolated $k = 6$ pulse, the fitted slopes of log total fluctuation energy are -0.0721 and -0.3604 at $\nu = 10^{-3}$ and 5×10^{-3} , respectively, versus the single-carrier predictions -0.0720 and -0.3600 . Centred discretisation also preserves the imposed reflection symmetry to residuals of order 10^{-15} in the archived symmetry check. These controls support the synthesis and time integration at their tested parameters; neither establishes convergence of a later local force.

3. Isolated-wave delivery and componentwise dynamics

Projection can cause an order-one failure before time integration. For the tensor (2.3), a generic Cholesky synthesis on oblique carriers yields post-projection diagonal stresses equal to approximately 0.25, 0.52 and 2.03 times their designed components; the transverse six-mode field instead retains its constant covariance to within 0.01 % in the solver. Figure 1 contrasts this static obstruction with the errors introduced by slow envelope modulation. The exact criterion (2.2) explains why the transverse construction works for (2.3) but cannot be extended by the same three directions to every positive tensor.

For a single carrier under linear diffusion, the stress ratio is $e^{-2\nu k^2 t}$. This identifies a delivery time t_d through $2\nu k^2 t_d \ll 1$; it also supplies an end-to-end calibration. At $k = 6$, the measured logarithmic energy slopes are -0.0721 and -0.3604 for $\nu = 10^{-3}$ and 5×10^{-3} , against -0.0720 and -0.3600 analytically. At finite amplitude the six-mode field interacts with itself. For the archived 64^3 , $k = 6$, $\nu = 10^{-5}$ isolated runs, the total energy ratio at $t = 1.8$ decreases monotonically across $\varepsilon = 0.05, 0.1, 0.2, 0.4$, with excess loss relative to single-rate diffusion increasing in magnitude from 5.2×10^{-5} to 5.748×10^{-3} . These ratios describe the whole-field energy, including new modes; they are not a measurement of a universal nonlinear threshold.

Delivery of prescribed Reynolds stresses




Figure 1. Projection-only sensitivity of the transverse-wave synthesis to modulation wavenumber and carrier scale. The plotted errors are offline field comparisons, separate from the time-evolved vortex results and the 512³ coherence scan.

ε	εkt	$E(t)/E_{\text{design}}(0)$	change from initial energy	deviation from linear diffusion
0.05	0.54	0.998653	-0.001347	-0.0000520
0.10	1.08	0.998484	-0.001516	-0.0002204
0.20	2.16	0.997646	-0.002354	-0.0010586
0.40	4.32	0.992957	-0.007043	-0.0057481

Table 3. Isolated $N = 64$ pulse at $k = 6$, $\nu = 10^{-5}$ and $t = 1.8$. Here $E = \frac{1}{2}\langle |u|^2 \rangle$ and $E_{\text{design}}(0) = 1.2\varepsilon^2$, so their ratio equals the whole-field stress trace divided by $2.4\varepsilon^2$. It includes all generated Fourier modes. The fourth column is $E(t)/E_{\text{design}}(0) - 1$; the fifth subtracts the linear-diffusion comparator $e^{-2\nu k^2 t} = 0.99870484$ from the same energy ratio. The deviations are dimensionless, not percentages.

3.1. Viscous window and finite-amplitude energy transfer

For a single isolated carrier, the delivery error in its box-averaged stress is $1 - e^{-2\nu k^2 t}$. Thus $2\nu k^2 t \ll 1$ is a necessary short-time condition for retaining the initial covariance, independently of how accurately that covariance was synthesised. In the documented zero-background resolution scan at $(k, \nu, N, \varepsilon) = (20, 10^{-3}, 128, 10^{-3})$, an exponential stress-decay fit gave $\nu_{\text{eff}}/\nu = 0.9998$, i.e. $100(\nu_{\text{eff}}/\nu - 1) = -0.02\%$ relative to molecular viscosity. A straight-line fit to the curved error trace would not estimate the initial rate: for $2\nu k^2 = 0.072$, the endpoint secant between $t = 0.15$ and 2 divided by the initial slope is 0.92620. The exponential reference is used throughout.

The four archived 64³, $k = 6$, $\nu = 10^{-5}$ amplitude runs expose the limitation of a single-rate comparator. Table 3 gives the final stored stage at $t = 1.8$; figure 2 shows the ten stored times. The total energy decays monotonically, consistent with the unforced periodic energy balance, while its departure from $e^{-2\nu k^2 t}$ becomes more negative as amplitude grows. In particular, doubling ε from 0.1 to 0.2 increases the magnitude of that departure from 2.204×10^{-4} to 1.0586×10^{-3} . Energy transferred into new Fourier modes has different viscous rates, so the departure is not by itself a measure of a new viscosity.

An earlier excess rate, expressed as $\nu_{\text{eff}} - \nu$, was initially attributed to the spatial scheme. At fixed physical wavenumber its fitted value was 1.878×10^{-6} on $N = 64$ and 1.891×10^{-6} on $N = 128$, with no resolved decrease on refinement. Across $\varepsilon = 10^{-3}, 3 \times 10^{-3}, 10^{-2}, 3 \times 10^{-2}$, the ratio of that excess to ε^2 lay between 1.876×10^{-2} and 1.891×10^{-2} . This combination points to an amplitude-dependent interaction, not a




Figure 2. Archived isolated-pulse amplitude sweep at $N = 64$. Points are final-stage diagnostics at steps 20, 40, $\dots$, 200; the ordinate is $100[E(t)/E_{\text{design}}(0) - e^{-2\nu k^2 t}]$. Lines join repeated observations within a deterministic run.

k	$R_{xy}/\sqrt{R_{xx}R_{yy}}$ at $t = 2$	g_{tr}/g_0	g_{xy}/g_0	c
20	0.599812569	1.000224	1.020718	0.195698
24	0.599732608	1.000318	1.020627	0.194206
28	0.599639979	1.000429	1.020514	0.192409
32	0.599535688	1.000557	1.020388	0.190383
40	0.599296112	1.000855	1.020086	0.185591

Table 4. Isolated 512^3 six-mode pulse at $\varepsilon = 10^{-3}$, $\nu = 10^{-5}$. Rates are ordinary least-squares slopes with a free intercept of $-\log R_i$ at the 20 times $t = 0.1, \dots, 2.0$, using one sample at the end of the final Runge–Kutta stage of each three-stage step; c is fitted through the origin in $-\log[(R_{xy}/R_{xx})/0.6] = c\varepsilon^2 k^2 t^2$. The initial ratio rounds to 0.6000 in printed diagnostics but is not exactly that value.

demonstrated discretisation viscosity. For the identified $N = 128, k = 20$ case, $k\Delta x = 2\pi k/N = 0.982$, or 6.4 points per wavelength. The tested second-order central scheme at its separately reported resolution differed by -0.72% and then -0.18% on doubling resolution. These calibrations motivate, but do not prove as a uniform theorem, the empirical resolution condition used in §5.

The 512^3 scan resolves a subtler effect at $\varepsilon = 10^{-3}$. At each saved physical time, only the stress record from the final Runge–Kutta stage is retained. Let g_i be the slope from an ordinary least-squares fit with a free intercept to $-\log R_i(t)$ at the selected final-stage times, and $g_0 = 2\nu k^2$. The trace-based ratios in table 4 differ from unity by only 0.02–0.09%, whereas an exponential fit to R_{xy} yields $g_{xy}/g_0 \approx 1.020$ at all five wavenumbers. This does not indicate a separate shear viscosity: the fitted interval matters. At $k = 24$, fitting $-\log R_{xy}$ by that same method at all saved endpoints in $0.1 \leq t \leq 1.0$ and $1.0 \leq t \leq 2.0$ gives $100(g_{xy}/g_0 - 1) = 1.0857\%$ (10 times) and 2.9379% (11 times), respectively.

For the continuum projected Euler equation with initial field (2.4), the Fourier calculation in appendix A.5 gives $R'_{xx}(0) = R'_{xy}(0) = R''_{xx}(0) = 0$ and $R''_{xy}(0) = -(6/25)\varepsilon^4 k^2$. Consequently,

$$-\log \left[\frac{R_{xy}(t)/R_{xx}(t)}{3/5} \right] = \frac{1}{5}\varepsilon^2 k^2 t^2 + \text{higher-order terms.} \quad (3.1)$$

ε	g_{xx}/g_0	g_{xy}/g_0	c	$Q(2)$
0.0005	1.000084	1.005162	0.194306	0.000111474
0.0010	1.000318	1.020627	0.194206	0.000445804
0.0020	1.001271	1.082351	0.193889	0.001779213

Table 5. Amplitude control for the isolated 512³, $k = 24$ pulse, using the same 20 final-stage diagnostic times as table 4. Here $Q(t) = -\log[(R_{xy}/R_{xx})/0.6]$; the common linear reference is $g_0 = 2\nu k^2$. Each run was integrated for 200 steps.

The fitted c lies about 2–7 % below the inviscid leading coefficient $1/5$ in (3.1), and decreases monotonically with k . At fixed time, both $\nu k^2 t$ and the finite-amplitude scale $\varepsilon k t$ increase with k , so viscous and higher-order finite-time terms can contribute to this systematic deficit; discretisation is not excluded. The approximately fixed 2 % apparent rate follows because both the viscous reference and leading quadratic correction contain k^2 . The independent amplitude control at $k = 24$ makes this interpretation testable. Table 5 compares $\varepsilon = 0.0005, 0.001, 0.002$, with the wave shape, grid, viscosity, time step, output schedule and solver build held fixed. Relative to $\varepsilon = 0.001$, $g_{xy}/g_0 - 1$ changes by factors 0.2503 and 3.9925, close to the quadratic predictions 0.25 and 4. The direct endpoint coherence loss $Q(2)$ changes by factors 0.2501 and 3.9910. A log–log fit of the extra apparent rate across the three amplitudes has exponent 1.9979. The fitted coefficients $c = 0.1939$ – 0.1943 vary by less than 0.3 % across the triplet. Thus the amplitude dependence favours the continuum nonlinear explanation over a fixed additional linear viscosity at this parameter point. The remaining difference from $c = 1/5$ includes unresolved finite-time, viscous and numerical contributions.

A matched viscosity scan further tests this interpretation at $k = 24$, $\varepsilon = 10^{-3}$ and $N = 512$. Using the same 20 final-stage samples at $t = 0.1, \dots, 2.0$, we fit $Q(t) = -\log[(R_{xy}/R_{xx})/0.6]$ through the origin against $\varepsilon^2 k^2 t^2$. The full-window coefficient increases from 0.194206 at $\nu = 10^{-5}$ to 0.196624 and 0.197859 at $\nu = 5 \times 10^{-6}$ and 2.5×10^{-6} , respectively; $Q(2)$ increases from 4.4580×10^{-4} to 4.5227×10^{-4} and 4.5555×10^{-4} . The deficits $1/5 - c$ are 0.005794, 0.003376 and 0.002141: successive viscosity halvings reduce them by factors 0.583 and 0.634. Thus the measured finite-window coefficient moves monotonically toward the independently derived inviscid short-time value $1/5$ as viscosity decreases. This is positive evidence that viscosity contributes to the deficit, alongside the spatial trend from the independent Nek5000 calculation in table 6. The three-point scan at one spatial resolution does not separate viscosity from finite-window and discretisation effects or determine a viscosity-extrapolated coefficient.

To test whether the componentwise drift is specific to the compact finite-difference solver, we repeated the isolated $k = 24$, $\varepsilon = 10^{-3}$ pulse with Nek5000 v19.0 (Fischer *et al.* 2008). The initial velocity was evaluated from (2.4) on a triply periodic spectral-element mesh, independently of the Xcompact3d fields. We used degree-seven elements (Patera 1984), 12-point nonlinear overintegration, a P_N – P_N pressure formulation and BDF2/OIFS stepping. The 20³- and 25³-element meshes reproduce the designed initial covariance to round-off. With eight Gauss–Lobatto–Legendre nodes per degree-seven element, these meshes store 160 and 200 element-local values along a coordinate line; continuous, periodic assembly leaves 140 and 175 distinct nodal positions per direction, respectively. These nonuniform nodal distributions are not equivalent to uniform 160³ or 200³ finite-difference grids. Table 6 applies the same 20 physical sampling times and fit definitions as table 4. Both Nek5000 meshes recover a 2.06–2.07 % excess fitted rate in R_{xy} , while the trace-based rate

Solver and distinct periodic nodes	Δt	g_{tr}/g_0	g_{xy}/g_0	c	$Q(2)$
Xcompact3d, 512 ³ points	0.010	1.000318	1.020627	0.194206	0.000445804
Nek5000, 140 ³ nodes	0.010	1.000290	1.020645	0.194829	0.000446885
Nek5000, 140 ³ nodes	0.005	1.000290	1.020645	0.194833	0.000446890
Nek5000, 175 ³ nodes	0.010	1.000318	1.020688	0.194963	0.000447193

Table 6. Independent solver check for the unforced, zero-background six-mode pulse at $k = 24$, $\varepsilon = 10^{-3}$, $\nu = 10^{-5}$, $t \leq 2$. Nek5000 uses degree-seven elements with 12-point overintegration: 20³ and 25³ elements yield the stated counts of distinct periodic Gauss–Lobatto–Legendre nodes. Rates use ordinary least-squares slopes with a free intercept at $t = 0.1, \dots, 2.0$, and c is fitted through the origin in $Q(t) = c\varepsilon^2 k^2 t^2$, where $Q = -\log[(R_{\text{xy}}/R_{\text{xx}})/0.6]$. Each Nek row is a volume-weighted, mean-subtracted stress; each Xcompact3d row uses the stress sampled at the end of the final Runge–Kutta stage.

stays close to g_0 . Halving the 20³-element time step changes c by 3.8×10^{-6} ; increasing the element count changes it by 1.34×10^{-4} . The Nek coefficient rises from 0.194829 on the 20³-element mesh to 0.194963 on the 25³-element mesh, approaching the inviscid short-time coefficient $1/5$. Across these meshes its distance from the Xcompact3d value 0.194206 increases from 6.23×10^{-4} to 7.57×10^{-4} . Nek refinement therefore moves c toward $1/5$ but away from this fixed-grid Xcompact3d result; the largest solver difference is below 8×10^{-4} . The compared discretisations use different nodal distributions and projection and time-stepping procedures, so the sign of their difference cannot yet be assigned to a particular numerical error. Two Nek meshes cannot establish a continuum limit or separate spatial and temporal effects. The independent discretisation corroborates the isolated coherence effect; the separate nonuniform-background force check is reported in §4.4.

3.2. Envelope projection and a factorial interaction control

For a solenoidal carrier $\mathbf{v}$ multiplied by a scalar envelope a , the initial correction is $\mathcal{P}(a\mathbf{v}) - a\mathbf{v} = -\nabla \Delta^{-1}(\mathbf{v} \cdot \nabla a)$. The band-limited estimate (2.6) suggests a velocity correction of order q/k , but it does not assign the same power to the filtered force, since taking a divergence can cancel leading contributions. In a projection-only Fourier calculation for the scalar three-dimensional envelope, $q = 1, 2, 3$ and $k = 12, 18, 24, 36$, descriptive slopes at fixed q were -1.04 to -1.13 for stress and -1.92 to -1.96 for full force. Repeating $q = 1$, $k = 12, 24, 36$ on $N = 256$ reproduced the $N = 128$ construction values to floating-point precision. The measured stress slopes are consistent with the leading $O(q/k)$ velocity correction in (2.6) and the first-order perturbation of a quadratic stress. Since (2.6) is an upper estimate, the slopes do not establish a matching lower bound or a universal exponent. The approximately -2 force slopes remain descriptive of this construction: the general force estimate does not establish the extra cancellation. The curves in figure 1 therefore describe this construction, not an asymptotic law for all envelopes.

To separate initial projection and linear sideband diffusion from subsequent deformation, define a continuous linear reference mode by mode:

$$\widehat{u}_{\text{lin}}(\kappa, t) = e^{-\nu|\kappa|^2 t} \mathcal{P}(\kappa) \widehat{u}_{\text{raw}}(\kappa, 0). \quad (3.2)$$

This comparator retains the envelope sidebands, unlike a single-rate exponential. Its difference from a numerical solution can still contain discrete projection and diffusion error, so it cannot be called a pure nonlinear residual.

k	q	ν	$E_{\text{lin}}^F(\varepsilon=0.005) (\%)$	$E_{\text{lin}}^F(\varepsilon=0.4) (\%)$	$E_{\text{tar}}^F(\varepsilon=0.4) (\%)$
12	1	10^{-5}	0.002262	9.102	9.209
12	1	10^{-3}	0.002146	8.994	12.76
12	2	10^{-5}	0.003402	11.12	11.77
12	2	10^{-3}	0.003254	11.13	15.91
24	1	10^{-5}	0.008996	13.13	13.23
24	1	10^{-3}	0.01499	18.96	34.32
24	2	10^{-5}	0.01470	15.92	16.24
24	2	10^{-3}	0.01134	22.46	38.29

Table 7. All 16 factorial combinations at $t = 0.2$, grouped into eight rows by amplitude. E_{lin}^F uses the continuous linear-reference force norm, whereas $E_{\text{tar}}^F = \mathcal{E}_{\text{target}}^{F^K}(0.2; 128)$ uses the initial prescribed force norm. The arguments identify the pulse amplitude ε ; every column is evaluated at $t = 0.2$, with $K = 6$ and the same core.

For the factorial comparison at $t = 0.2$, define $E_{\text{lin}}^F(\varepsilon) = \|F_{\text{num}}^{K,\text{ev}} - F_{\text{lin}}^{K,\text{ev}}\|_{2,C} / \|F_{\text{lin}}^{K,\text{ev}}\|_{2,C}$, and let $E_{\text{tar}}^F(\varepsilon) = \mathcal{E}_{\text{target}}^{F^K}(0.2; 128)$ for that same amplitude and target. The linear-reference and initial-target denominators therefore remain distinct. An isolated 2^4 design varied $k \in \{12, 24\}$, envelope wavenumber $q \in \{1, 2\}$, $\varepsilon \in \{0.005, 0.4\}$ and $\nu \in \{10^{-5}, 10^{-3}\}$ on $N = 128$ with $\Delta t = 0.005$. All 16 matched triplets were retained at $t = 0.05, 0.1, 0.2$. At $t = 0.2$, the full-force difference from (3.2) is only 0.002146–0.01499 % at the small amplitude, but 8.994–22.46 % at the large one (table 7, figure 3). Even after sideband diffusion is represented, the high-amplitude cases therefore undergo a substantial deformation. Their precise nonlinear/discrete split needs a separate high-amplitude resolution check: at fixed N , changing k also changes $k\Delta x$, and two parameter levels cannot identify a general boundary.

4. A nonuniform target in a vortex background

4.1. Constant target: global covariance, production and radial response

The first Gaussian-vortex campaign imposed the spatially constant $\varepsilon^2 T_0$. It is useful as a background-flow control, but its target divergence vanishes. After subtracting the analytic viscous evolution of the initial vortex, we formed the unfiltered volume-averaged stress and production

$$\overline{P} = -\langle u'_i u'_j S_{ij} \rangle_{\Omega} = \langle -2u'_r u'_\theta S_{r\theta} \rangle_{\Omega}. \quad (4.1)$$

The matched positive-pulse, negative-pulse and no-pulse triplet extends this constant-target comparison through $N = 512$. Table 8 shows a negligible no-pulse baseline and close agreement across the four grids in global stress and production. At $N = 128$, the global target error is 0.1085 % at $t = 0.1$ using the Euclidean norm of the six volume-averaged stress components. At $N = 256$, the six-component global stress differs from its constant initial target by 0.1086 % at $t = 0.1$ and 1.649 % at $t = 2$; the corresponding $N = 512$ errors are 0.1086 % and 1.649 %. The late-time target differences include physical evolution.

The unfiltered cylindrical-coordinate diagnostic azimuthally averaged velocity products and then differentiated their radial profile. It uses a different stress operator from the three-dimensional filtered Cartesian force in (2.5). The even $+\varepsilon, -\varepsilon, 0$ response is needed because a single pulse sign can carry a larger odd component. Table 9 gives its force and mean azimuthal-velocity response over $0.2 \leq r \leq 1.2$. The $N = 512$ even response profiles differ




Figure 3. Isolated factorial experiment at $t = 0.05, 0.1, 0.2$, comparing the full induced force with the sideband-resolved linear reference (3.2). A common $K = 6$ and core mask are used; lines connect times from one deterministic triplet.

Grid	No-pulse baseline / target	$\bar{P}(t = 2)$	Stress difference / $N = 256$
$N = 128$	1.44×10^{-10}	1.58403×10^{-8}	2.875×10^{-5}
$N = 192$	1.08×10^{-10}	1.58309×10^{-8}	1.883×10^{-6}
$N = 256$	1.09×10^{-10}	1.58302×10^{-8}	reference
$N = 512$	2.15×10^{-10}	1.58301×10^{-8}	6.89×10^{-7}

Table 8. Constant-target analytic-background vortex at $k = 24$, $\varepsilon = 10^{-3}$, $\nu = 10^{-5}$, $\sigma_0 = 0.6$, $\Omega_0 = 0.5$. The last column is the relative Euclidean distance between the six unfiltered, volume-averaged stress components at $t = 2$, with $N = 256$ as denominator; the $N = 512$ value is 6.89×10^{-7} , or $6.89 \times 10^{-5} \%$. Production is the volume average (4.1).

from $N = 256$ by 0.599 % in count-weighted relative L^2 , while the force profiles differ by 7.09 %; both grids use the same 64 bin centres. The $N = 128$ and 192 force differences of 9.20 and 8.03 % also use $N = 256$ as their single reference. These are not successive grid-pair differences, errors against a nonzero prescribed force, or evidence that the radial force has converged. The smaller response-profile difference likewise cannot by itself establish an asymptotic convergence rate. The pointwise radial angular-momentum residual was 5–9 % of the force RMS. A shell-integrated check at $t = 1.9$, using the total radial angular-momentum flux at the shell endpoints, gave a residual below 1.7 % of integrated absolute tendency on all three grids. Coarsening the $N = 256$ profile from 64 to 32 bins changed the pointwise residual from 8.95 to 5.82 %, while 16 bins increased it to 13.48 %. These

Grid	Force RMS	Response RMS	Force difference	Response difference
$N = 128$	1.174×10^{-6}	1.172×10^{-6}	9.20 %	5.07 %
$N = 192$	1.180×10^{-6}	1.173×10^{-6}	8.03 %	3.36 %
$N = 256$	1.182×10^{-6}	1.174×10^{-6}	reference	reference
$N = 512$	1.180×10^{-6}	1.173×10^{-6}	7.09 %	0.599 %

Table 9. Constant-target cylindrical-coordinate diagnostic at $t = 2$. Entries are count-weighted radial-profile RMS values or relative differences from $N = 256$ over $0.2 \leq r \leq 1.2$, after matched-sign decomposition. The force is formed by differentiating unfiltered binned products, not by the Cartesian filtered operation (2.5). Both $N = 256$ and $N = 512$ use 64 common radial bins. These constant-target $N = 512$ diagnostics are radial-profile comparisons; the separate nonconstant-target triplets provide the three-dimensional Cartesian comparisons reported below.

changes identify binning and differentiation as material to the radial budget and do not imply pointwise closure. We therefore use this binned cylindrical operator to document diagnostic sensitivity, not to support the fixed-band three-dimensional force-convergence claim.

Full three-dimensional snapshots from this constant-target family allow the definitions to be compared without a fresh simulation. At $K = 6$, the filtered core stress differs from the *constant initial* target by 2.573, 2.579 and 2.573 % at $t = 0.1$ on $N = 128, 192, 256$. The projected Cartesian momentum-budget residual at $t = 1.9$ is 0.218, 0.0494 and 0.0434 % of induced force RMS. The common-grid full-force difference at that time falls from 0.1895 % between $N = 128$ and 192 to 0.01306 % between 192 and 256 for $K = 6$; the corresponding pairs are 0.1266/0.00913 % at $K = 4$ and 0.2291/0.01522 % at $K = 8$. The force magnitude itself changes with K .

Reconstructing the unfiltered binned radial force from the archived three-dimensional fields reproduces the original logs within 1.5×10^{-9} relative L^2 error. At $t = 1.9$ and $K = 6$, applying radial binning and differentiation to the *filtered* tensor differs from radial averaging of its direct three-dimensional Cartesian divergence by 3.57, 1.79 and 3.28 % at $N = 128, 192, 256$, respectively. Each of these differences is divided by the L^2 norm of the radially averaged direct three-dimensional force over $0.3 \leq r \leq 1.2$; in particular, the 1.79 % at $N = 192$ is an operator-comparison discrepancy, not the $t = 0.01$ initial-target force error in table 10. The filtered binned force differs from the original unfiltered radial force by 28.6, 26.8 and 28.4 % across those grids. The dominant discrepancy between the two force series is thus their stress definition: unfiltered velocity products relative to an analytic background versus a filtered covariance. Since the prescribed constant tensor has zero divergence, this campaign cannot by itself establish delivery of a prescribed nonzero local force.

4.2. Prescribed nonconstant stress and its local force

To require a nonzero prescribed force, we modulate the six-mode field by $a(x) = 1 + 0.25 \cos x$. The initial target is

$$T_{\star}(x) = \varepsilon^2 a(x)^2 T_0, \quad F_{\star i}(x) = -\partial_j T_{\star ij} = 0.5 \varepsilon^2 a(x) \sin x (T_0)_{i1}. \quad (4.2)$$

Since $a \geq 0.75$ and the eigenvalues of T_0 are 0.4, 0.4, 1.6, this tensor remains positive definite throughout the box. Modulation makes the raw velocity slightly non-solenoidal, so the first projection is included in the delivery error. The target itself is fixed before evolution and is never reset to a projected or later-time field.

N	stress error (%)	full-force error (%)	projected-force error (%)
128	0.3162	1.796	3.238
192	0.3157	1.793	3.234
256	0.3156	1.793	3.234
512	0.3156	1.793	3.234

Table 10. Target errors $100\mathcal{E}_{\text{target}}^{R^K}$, $100\mathcal{E}_{\text{target}}^{F^K}$ and $100\mathcal{E}_{\text{target}}^{\mathcal{P}F^K}$, respectively, at $t = 0.01$ against the prescribed initial nonconstant tensor and its divergence; $K = 6$, $\Delta t = 0.01$, and the identical core mask on the common 128^3 Fourier comparison grid. Each row comes from a matched $+\varepsilon, -\varepsilon, 0$ Gaussian-vortex triplet. The projected-force error is evaluated after applying the Leray projector to both computed and initial target forces, and is normalised by the projected target norm. Multiplication by $r_\star = 0.5144958$ converts that column to units of the full target-force norm; it is not directly comparable with the full-force column before this conversion. The isolated zero-background cases are reported separately.

The projection-only Fourier construction at $t = 0$, evaluated with the same $K = 6$ filter and core metric as the simulations, has stress and full-force target errors of 0.1843 and 0.1240 %, respectively; appendix A gives the associated projection-bound evaluation. The first saved simulation time is $t = 0.01$. Table 10 shows that the stress and full-force errors against (4.2) are almost unchanged by grid refinement. The force error is larger than the stress error because a divergence amplifies spatial differences. The projected-force diagnostic is a third, separately normalised quantity. For this particular target, the two force denominators have an exact relation. Writing $F_\star = c(x)(1, 0.6, 0)^\top$, all its nonzero Fourier modes are parallel to the x -axis, so $\mathcal{P}F_\star = c(x)(0, 0.6, 0)^\top$ and

$$r_\star := \frac{\|\mathcal{P}F_\star\|_{2,C}}{\|F_\star\|_{2,C}} = \frac{0.6}{\sqrt{1.36}} = 0.5144958.$$

Multiplying a projected-force target error by $r_\star$ therefore expresses its absolute error in units of the full initial target-force norm, without changing either diagnostic. These are early delivery measurements; they do not predict a small error against the same *initial* target at later times.

At $t = 0.01$, the zero-background control gives stress and full-force target errors of 0.183 and 0.136 %, compared with 0.316 and 1.796 % in the Gaussian-vortex case. Here $2\nu k^2 t = 1.152 \times 10^{-4}$, $\varepsilon k t = 2.4 \times 10^{-4}$, and $\Omega_0 t = 0.005$. The paired and termwise analysis below shows that the background mainly alters the evolving odd-pulse quadratic force; its direct cross force is much smaller.

An independent linear reference quantifies the early isolated error. Starting from the projected initial pulse, we multiply each velocity Fourier coefficient by $\exp(-\nu|\kappa|^2 t)$, retaining the envelope sidebands, and then form the quadratic stress, apply G_6 , and evaluate the force with the same common-core norm. This modewise Stokes calculation predicts full-force target errors of 0.135138 % at $t = 0.01$ and 0.237048 % at $t = 0.1$; the matched isolated $N = 256$ DNS gives 0.135150 % and 0.237090 %, respectively. Thus the initial projection defect and its subsequent linear viscous evolution account quantitatively for the isolated error at these parameters. The calculation adds the error fields before taking their norm; adding the scalar projection error and $1 - \exp(-2\nu k^2 t)$ would neglect their relative directions and the distinct sideband decay rates.

By $t = 0.1$, the stress errors against the initial target on the common 128^3 core are 2.5459, 2.5432, 2.5430 and 2.5429 % at $N = 128, 192, 256, 512$; the corresponding full-force errors are 18.0427, 18.0197, 18.0181 and 18.0178 %. They include the evolving background and pulse. At $N = 512$, the projected-force target errors are 3.234 % at $t = 0.01$ and 32.695 %

t	isolated error	vortex error	$\Delta F^K(O)$	cross force
0.01	0.136	1.796	1.790	0.00155
0.1	0.238	18.043	18.052	0.05299

Table 11. Paired 128^3 , $K = 6$ nonconstant-target runs with identical inputs apart from $\Omega_0 = 0$ (isolated) or 0.5 (vortex). Entries are vector RMS values in per cent of the initial target-force RMS on the common core. Here $\Delta F^K(O)$ is the vortex-minus-isolated odd-quadratic force. The last two columns are component magnitudes, not additive portions of the error.

at $t = 0.1$, each divided by $\|\mathcal{P}F_\star\|_{2,C}$. In units of the full target norm, these are 1.664 and 16.821 %, compared with full-force errors of 1.793 and 18.018 %. At $t = 0.1$, projection reduces the vortex error norm by only 6.64 %. For the isolated $N = 256$ control at the same time, the separately normalised projected error is 0.18440 %, or 0.09487 % of the full target norm, compared with a full-force error of 0.23709 %: a norm reduction of 59.99 %. The vortex departure therefore largely survives projection, whereas the isolated error is much more sensitive to it. These are reductions of core norms, not orthogonal error fractions: the periodic Leray projector is not orthogonal in the restricted core inner product. Moreover, a gradient correction to the initial velocity need not produce a purely gradient correction to the quadratic stress divergence. A matched-background comparison and the projected odd-velocity budget below separate the direct cross force from the background-induced evolution of the pulse. Figure 4 shows the early field and its target. The global constant-target metric instead uses the six unfiltered volume-averaged stress components.

4.3. Background-induced change of the local force

The matched triplets permit an exact decomposition of the filtered force at each saved time. Let $B = u^0$ be the no-pulse velocity, $E = (u^+ + u^-)/2 - B$ the even velocity response and $O = (u^+ - u^-)/2$ the odd response. Applying (2.5) to $u^\pm = B + E \pm O$ gives

$$F_i^{K,\text{ev}} = F_i^K(O) + F_i^K(E) - \partial_j \{G_K(B_i E_j + E_i B_j) - (G_K B_i)(G_K E_j) - (G_K E_i)(G_K B_j)\}. \quad (4.3)$$

The three terms are the odd-pulse quadratic force, even-pulse quadratic force and direct background–even-response cross force. The products in each term are evaluated on the native grid before the $K = 6$ restriction. The following comparisons use the common 128^3 Fourier lift and $0.3 \leq r \leq 1.2$ core; every percentage in table 11 is divided by the norm of the *initial full Cartesian target force*. The term norms need not sum because the vector fields have different directions.

At $t = 0.1$, the direct cross force is only 0.0530 % of the initial target, whereas the odd-pulse quadratic force differs between the two backgrounds by 18.052 %. The latter accounts for almost the entire background-dependent change in the full force. At $N = 256$, the same cross and even-quadratic norms are 0.05298 and 0.000260 %, while the total target error is 18.0181 % under this common-grid metric. An older native-grid quadrature gives 18.070 % on the same physical core. The four-grid comparison above now uses the common-grid quadrature throughout; the force definition and initial-target denominator are unchanged. The early discrepancy therefore does not arise primarily from an instantaneous background–pulse cross stress. Rather, the background changes the odd pulse and hence its quadratic stress. In particular, the $t = 0.01$ discrepancy cannot mainly be assigned to the static envelope-projection defect: the paired isolated error is 0.136 %, close to the 0.124 % offline projection-only force error.

N	observed change	strain	rotation	advection	closure
128	16.2639	+87.424	+24.776	-12.295	0.1806
256	16.2417	+87.520	+24.844	-12.320	0.00313

Table 12. Integrated sources of the paired odd-quadratic filtered-force change over $0.01 \leq t \leq 0.1$. The observed-change column is the vector RMS in per cent of the initial full target-force RMS on the common 128^3 core. Strain, rotation and advection are signed projections of the integrated source fields onto that observed change, divided by its squared norm and expressed in per cent; they are not additive RMS shares. Closure is the RMS difference between the summed sources and observed change, divided by the observed-change RMS, in per cent. Both grids use the same $K = 6$ filter and time quadrature.

To identify how the background changes the odd pulse, subtract the two pulse equations and apply the periodic Leray projector $\mathcal{P}$:

$$\partial_t O = -\mathcal{P}\{(B \cdot \nabla)O + (O \cdot \nabla)B + (E \cdot \nabla)O + (O \cdot \nabla)E\} + \nu \Delta O. \quad (4.4)$$

We split $(O \cdot \nabla)B = [S(B) + A(B)]O$, with S and A the symmetric and antisymmetric parts of ∇B . For each of two matched grids, $N = 128$ and 256 , we saved ten velocity snapshots from $t = 0.01$ to 0.1 in both the vortex and isolated cases. At each time, we projected the terms on the right-hand side of (4.4), passed them through the bilinear derivative of $F^K(O)$, and integrated the vortex-minus-isolated force-rate fields by the trapezoidal rule with spacing 0.01 . The velocity products were formed before the $K = 6$ filter, as in the force diagnostic.

Table 12 identifies background strain as the largest aligned source at both grids; background rotation adds to the change while advection opposes part of it. The viscous signed projection is -0.0463% at both grids, and each even-response term contributes less than 0.0001% . The three dominant source projections agree within 0.10 percentage points across grids. The smaller closure residual at $N = 256$ is consistent with reduced spatial-derivative error; it does not by itself separate spatial and temporal errors. This split does not uniquely identify wavevector refraction as a separate pathway and has not been repeated as a time-resolved source budget at $N = 512$. The exact algebraic decomposition (4.3) is also stable on the late $N = 512$ fields. At $t = 1.8$ and 2.0 , the $N = 256$ – 512 common-core differences in the background cross force are 0.00365 and 0.00366% relative to the 512^3 term; the odd-quadratic differences are 0.00222 and 0.00227% . Adding all three terms reconstructs the independently archived 512^3 filtered force to better than $7 \times 10^{-9}\%$ of that force at either time. These late checks validate the decomposition and its fixed-band spatial consistency, but do not extend the early $N = 128$ and 256 temporal source fractions to the later vortex evolution.

4.4. Independent-solver check of the local force

The nonuniform-target result was repeated with Nek5000 on the Gaussian-vortex background, using independent matched $+\varepsilon, -\varepsilon, 0$ triplets. The initial pulse was evaluated directly from the continuum Leray projection of $a(x)$ times the six-mode carrier; no Xcompact3d snapshot was transferred. Nek5000 used degree-seven periodic spectral elements, with 20^3 and 25^3 elements, 12-point nonlinear overintegration allocation, and BDF2/OIFS stepping at $\Delta t = 0.01$. The meshes have 140 and 175 distinct, nonuniform periodic GLL nodal positions per axis, respectively. Each full Nek velocity field was interpolated to the common uniform 128^3 comparison grid before products were formed. We then applied the same spherical $K = 6$ filter, Cartesian stress divergence, even-pulse subtraction and three-dimensional core norm used for the Xcompact3d triplet. This

Time	Solver	Stress/target	Force/target	Stress/Xcompact3d	Force/Xcompact3d
0.01	Xcompact3d, 128^3	0.3162	1.7957	—	—
	Nek5000, 20^3 elements	0.3246	1.7999	0.1117	0.1157
	Nek5000, 25^3 elements	0.3168	1.8145	0.0245	0.0376
0.1	Xcompact3d, 128^3	2.5459	18.0427	—	—
	Nek5000, 20^3 elements	2.5491	18.0668	0.1298	0.5265
	Nek5000, 25^3 elements	2.5561	18.1265	0.0322	0.1831

Table 13. Independent three-dimensional nonconstant-target comparison in per cent. The first two metrics are $100\mathcal{E}_{\text{target}}^{R^K}$ and $100\mathcal{E}_{\text{target}}^{F^K}$, with the prescribed initial stress and force as denominator; the last two use the evolved Xcompact3d field at the same time. Each solver row represents the matched even response from positive, negative and no-pulse runs. Nek fields are first interpolated to a uniform 128^3 grid; the same Fourier filter and core norm are then applied to both solvers.

interpolation precedes multiplication for Nek5000, unlike the native-product route used for Xcompact3d, so the solver differences include diagnostic-transfer error as well as evolution error. Table 13 compares fields at fixed physical times.

At $t = 0$, the refined Nek stress and force errors against the initial target are 0.1820 % and 0.1234 %, close to the offline projection values 0.1843 % and 0.1240 %. The refined Nek and Xcompact3d *evolved field differences* are 0.0245 % in stress and 0.0376 % in force at $t = 0.01$, and 0.0322 % and 0.1831 % at $t = 0.1$. These differences are far below the errors of either solver relative to the unchanged initial target, corroborating the interpretation that the target departure reflects background-driven evolution. Halving the time step on the 20^3 -element mesh changes stress and force by 0.00133 % and 0.0127 %, respectively, at $t = 0.1$, whereas changing from 20^3 to 25^3 elements changes them by 0.1316 % and 0.6266 %. The force is more sensitive to spatial resolution than to this time-step change; two Nek meshes do not certify continuum convergence. These runs reproduce the early nonuniform-background local-force calculation across solvers. Separately, the matched Xcompact3d nonconstant-target triplets now supply 512^3 three-dimensional Cartesian stress and force comparisons through $t = 1.9$, described below; the older constant-target 512^3 archives contain global and radial diagnostics. Nek5000 was compared only through $t = 0.1$, so the late-time Cartesian result has no independent-solver counterpart.

Table 14 makes the comparison bases explicit. In particular, the $N = 128$ global constant-target error of 0.1085 %, 2.573 % filtered constant-target core error and 2.546 % filtered nonconstant-target core error all concern $t = 0.1$, but apply different operators or target families. The 0.316 % and 1.796 % early values are at $t = 0.01$, whereas the 0.166/0.0114 % pair in table 15 is an intergrid comparison at $t = 1.9$.

The matched 512^3 vortex triplet provides a direct fixed-band check beyond the target-error comparison. At $K = 6$, the $N = 256$ –512 relative core field differences are 0.0000111 % for stress and 0.0000433 % for full Cartesian force at $t = 0.01$, rising to 0.0000592 % and 0.000414 % at $t = 0.1$. The corresponding even velocity-response differences are 0.000486 % and 0.000633 %. In every early-time observable and at $K = 4, 6, 8$, the $N = 256$ –512 difference is smaller than the $N = 192$ –256 difference. Products were formed on each native grid before Fourier filtering and comparison on the common core. The exceptionally small $t = 0.01$ differences chiefly confirm consistent representation of the initial construction and first saved state; they alone do not demonstrate convergence of an evolved force. The later comparisons at $t = 0.1$ and 1.9 address finite-time grid consistency at each specified cutoff, not the unfiltered force.

Target	Time	Quantity and reference	$N = 128$ value
Constant	0.1	Global six-component stress / initial tensor	0.1085 %
Constant	0.1	Filtered core stress / initial tensor	2.573 %
Nonconstant	0.01	Filtered core stress / initial tensor	0.316 %
Nonconstant	0.1	Filtered core stress / initial tensor	2.546 %
Nonconstant	0.01	Full core force / initial force	1.796 %
Nonconstant	0.1	Full core force / initial force	18.043 %

Table 14. Map of different target errors. The global stress uses the Euclidean norm of six volume-averaged components; filtered stress uses core-volume Frobenius RMS at $K = 6$; full force uses the three-component core-volume RMS. The filtered-stress and full-force rows report $100\mathcal{E}_{\text{target}}^{R^K}$ and $100\mathcal{E}_{\text{target}}^{F^K}$, respectively. Every target comparator is the stated *initial* analytic field.

The evolved $K = 6$ force is grid consistent at fixed low modes. At $t = 1.9$, the relative three-dimensional field differences for successive grid pairs $N = 128$ –192, 192–256 and 256–512 are 0.0925, 0.00620 and 0.00124 % for stress; 0.166, 0.0114 and 0.00231 % for full force; and 0.151, 0.0125 and 0.00270 % for the even low-mode velocity response. These values are $\mathcal{E}_{\text{grid}}^Q$ for the stated observable and grid pair. The same decreasing ordering holds at $K = 4, 8$, where each successive grid-pair difference also decreases for all three observables.

The nonuniform refinement ratios allow a more quantitative spatial check. We first recover absolute differences $D_{ij} = \|Q_{N_i} - Q_{N_j}\|_{2,C}$ by multiplying each relative difference by its own finer-grid norm. Under a single leading spatial error model $Q_N = Q_\infty + CN^{-p}$, an effective order solves

$$\frac{D_{12}}{D_{23}} = \frac{N_1^{-p} - N_2^{-p}}{N_2^{-p} - N_3^{-p}}.$$

For the triplets (128, 192, 256) and (192, 256, 512), respectively, the resulting orders are 6.43 and 6.19 for stress, 6.37 and 6.16 for force, and 5.88 and 5.95 for the even velocity response. These effective orders, 5.88–6.43, are consistent with the nominal sixth-order compact spatial discretisation for the evolved fixed-band fields at $t = 1.9$. They rely on the leading-error model at fixed time step and cutoff and do not establish the order of the complete space–time workflow or convergence of the unfiltered fields.

At the original output half-width $h = 0.1$, the Leray-projected momentum-budget residual at $t = 1.9$ is 0.197, 0.0429 and 0.0371 % of the induced-force RMS at $N = 128, 192, 256$, respectively; an independent $N = 512$ low-band audit gives 0.0364 % under the common-core $K = 6, h = 0.1$ definition. A separate $N = 128$ output-cadence control gives 0.1872 % at $h = 0.02$; corresponding refined-cadence values have not been established on the other grids, including $N = 512$. Figure 5 shows the fixed-band intergrid field comparisons. For this small pulse the no-pulse force is roughly three orders of magnitude larger than the even induced force, which makes the matched control essential.

The constant-target radial profiles in table 9 and the nonconstant-target three-dimensional comparisons in table 15 use different target families, stress operators, times and denominators. In particular, all radial-profile differences use $N = 256$ as a single reference, whereas the Cartesian comparisons use successive grid pairs; the radial percentages cannot be read as a convergence sequence or as competing estimates of the filtered force. The temporal controls and additional target families below test whether the fixed-band conclusions survive changes in output cadence and target geometry.

Delivery of prescribed Reynolds stresses




Figure 4. Early stress and force delivery for the nonconstant target (4.2). The line profiles use $N = 128$ at $t = 0.01$; panel (c) includes $N = 128, 192, 256, 512$ and reports the three target-error metrics in table 10, with the same initial targets, low-mode cut-off and core norm.




Figure 5. Common-grid comparisons of the nonconstant-target stress, force and even velocity response at $t = 1.9$ for $K = 4, 6, 8$, showing all three successive grid pairs $N = 128-192, 192-256$ and $256-512$. The ordinates are $100\mathcal{E}_{\text{grid}}^Q$ for each stated observable; the $K = 6$ values are also given in table 15. Each denominator is the finer-grid field in its pair.

4.5. Temporal controls and multidirectional targets

The early-delivery comparison uses products formed on each native grid before spectral filtering. For common-grid evaluation, the filtered algebra is represented on a 40^3 Fourier grid; at the largest diagnostic cutoff $K = 8$, a product of retained low modes reaches wavenumber 16, below the grid Nyquist wavenumber 20. An independent native-grid implementation changed the baseline early full force by only 2.0×10^{-10} relatively. This numerical agreement concerns the evaluation of the chosen filtered diagnostic.

The time-integration step and the saved-output spacing were varied separately at $N = 128$. With the nonconstant one-dimensional envelope, step halvings from $\Delta t = 0.01$ to 0.005 and then 0.0025 changed the $t = 0.1$ force by 0.00882 and 0.00441 %, respectively, and the even velocity response by 0.295 and 0.147 %. At $t = 0.01$ the response differences were 3.20 and 1.58 %, when its denominator was still small. The approximate halving does not verify the nominal third-order time advance for the full initialisation and projection workflow; its early response sensitivity remains unresolved.




Figure 6. Nonconstant-target time evolution and temporal controls at $K = 6$. Central tendencies use times $t - h$ and $t + h$. The short-time and late-time checks vary Δt and output half-width separately; each curve contains repeated diagnostics of a deterministic matched triplet.

Central differences used to estimate the momentum tendency have their own half-width h . At $t = 0.1$, reducing h from 0.08 to 0.005 reduced the projected residual from approximately 0.0913 to 0.0333 % of the induced-force RMS, with little subsequent change on reducing Δt . A later matched triplet with output every 0.02 reached $t = 2$. At $t = 1.9$, reducing h from 0.1 to 0.02 at fixed $\Delta t = 0.01$ changed the residual from 0.1972 to 0.1872 %; the tendency-vector change was 0.03501 % of the induced force. At $\Delta t = 0.005$, $h = 0.02$ the residual remained 0.1872 %. Halving Δt changed the late force and even response fields by 0.004025 and 0.007864 %, respectively. The two controls constrain different operations and subtraction of their residual norms would not constitute an exact error budget (figure 6).

The $K = 6$ fixed-band field comparisons at $t = 1.9$ are given in table 15 and illustrated alongside the $K = 4, 8$ comparisons in figure 5. It compares evolved *fields* and divides

Induced field, $K = 6, t = 1.9$	$N = 128\text{--}192$	$N = 192\text{--}256$	$N = 256\text{--}512$
Filtered stress tensor	0.0925 %	0.00620 %	0.00124 %
Full three-component force	0.166 %	0.0114 %	0.00231 %
Even low-mode velocity response	0.151 %	0.0125 %	0.00270 %

Table 15. Relative core-volume L^2 field differences $100\mathcal{E}_{\text{grid}}^Q$ for the nonconstant target, with $Q = R^K, F^K$ or the even low-mode velocity response in the three rows. Each field is transferred to the same 128^3 Fourier grid before applying an identical three-dimensional core mask. The second grid in each pair supplies the denominator.




Figure 7. Filter-cutoff sensitivity and fixed-cutoff grid differences for the nonconstant Gaussian-vortex target. (a) Change in the $N = 256$ core RMS of the matched even filtered stress and full three-component force relative to $K = 6$, at $t = 0.1$ (dashed, open markers) and $t = 1.9$ (solid, filled markers). (b) Relative core-volume L^2 difference between $N = 192$ and $N = 256$ fields at $t = 1.9$, mapped to a common 128^3 Fourier grid; the finer-grid norm is the denominator. Products are formed on the native DNS grids before filtering. The curves across K represent different filtered observables, whereas panel (b) compares grids at fixed K . The runs are deterministic matched triplets, so no statistical error bars are implied.

by the finer-grid norm, whereas table 10 divides by the initial target norm. The $K = 4, 8$ sensitivity cutoffs preserve the ordering, but their force magnitudes differ, so this is a fixed-band result.

To separate filter dependence from grid resolution, we reprocessed the archived matched triplets at $K = 4, 6, 8, 10, 12$ and $t = 0.1, 1.9$, forming velocity products on each native DNS grid before filtering. The 128^3 common comparison grid resolves products through wavenumber $2K = 24$ at the largest cutoff. For this cutoff-only scan, target errors use native-grid sampling of the physical core, whereas the intergrid field differences below use the common 128^3 mask. At $N = 256, t = 0.1$, the native-grid stress and full-force errors relative to the initial target change from 2.5466 and 18.0703 % at $K = 6$ to 2.5462 and 18.0931 % at $K = 12$. By $t = 1.9$, the induced force RMS is 16.9 % larger at $K = 12$ than at $K = 6$, and 20.0 % smaller at $K = 4$ (figure 7a). Thus the force magnitude remains sensitive to the specified filter. At each fixed K , the $N = 192\text{--}256$ relative field differences at $t = 1.9$ stay below 0.0082 % for stress and 0.0147 % for full force (figure 7b). These comparisons support grid consistency for each stated filtered observable; they do not establish a cutoff-independent force or convergence of the unfiltered field.

The one-dimensional modulation is not the only positive target tested. Two scalar envelopes are $a_{2D} = 1 + (\cos x + \cos y)/8$ and $a_{3D} = 1 + (\cos x + \cos y + \cos z)/12$. With $T_\star = \varepsilon^2 a^2 T_0$, both remain above $a = 0.75$, and the three-dimensional case has a

Target	N	Stress (%)	Full force (%)	Projected force (%)
Scalar 2D	128	0.2828	2.570	4.204
Scalar 3D	128	0.2969	4.122	7.349
Rotating axes	128	0.3777	4.380	8.019
Rotating axes	192	0.3766	4.367	7.999
Rotating axes	256	0.3766	4.366	7.998

Table 16. Errors against each prescribed initial tensor and force at $t = 0.01$, $K = 6$, $\Delta t = 0.005$, $\varepsilon = 10^{-3}$, in the Gaussian-vortex background. The same three-dimensional core norm is used for all entries.

nonzero third force component. A separate target varies its principal axes:

$$T_{\text{rot}} = \varepsilon^2 \sum_{p=1}^3 b_p^2 B_p, \quad (b_1, b_2, b_3) = (1 + 0.2 \cos x, 1 + 0.2 \cos y, 1 + 0.2 \cos z), \quad (4.5)$$

where the nonzero coordinate-plane blocks are

$$B_1|_{xy} = \begin{pmatrix} 0.6 & 0.25 \\ 0.25 & 0.4 \end{pmatrix}, \quad B_2|_{xz} = \begin{pmatrix} 0.4 & 0.15 \\ 0.15 & 0.2 \end{pmatrix}, \quad B_3|_{yz} = \begin{pmatrix} 0.6 & 0.1 \\ 0.1 & 0.2 \end{pmatrix}. \quad (4.6)$$

The target in (4.5) uses the coordinate-plane blocks in (4.6), each realised by two transverse polarisations. Because $b_p \geq 0.8$, the minimum eigenvalue of $T_{\text{rot}}/\varepsilon^2$ is at least 0.23143; the sampled dominant-eigenvector direction varies by up to 27.34° . Its pre-projection covariance is the prescribed tensor, up to input rounding, and its target force is $0.4\varepsilon^2 \sum_p b_p \sin x_p (B_p)_{ip}$.

Table 16 gives the early target errors with $\Delta t = 0.005$. These matched vortex triplets use $\mathcal{E}_{\text{target}}^{R^K}$, $\mathcal{E}_{\text{target}}^{F^K}$ and $\mathcal{E}_{\text{target}}^{PF^K}$ for stress, full force and projected force, respectively. The three $N = 128$ targets were followed to $t = 0.2$. The rotating target has additional early $N = 192, 256$ checks: its full-force differences are 0.01336 % for 128–192 and 0.0009475 % for 192–256 on the common grid. Figure 8 shows all three force components without axisymmetric averaging. Larger-grid rotating-target evolution to $t = 2$ has not been established.

A separate plane-shear trial supplies a negative control on background definition. The original tanh shear profile is discontinuous across the periodic boundary; its fitted fluctuation-energy growth changed from 1.4772 to 0.9023 between $N = 128$ and 256. At $N = 128$, $t = 1$, the no-pulse case contributed 99.81 % of the pulse-run energy. A smooth periodic sine shear, measured relative to its analytically decaying background, instead gave a near-zero fitted rate of -0.0007763 across the three grids. Because the mean profile changed, this is a diagnostic control, not a convergence sequence for the original tanh problem. No instability or delivery claim is based on the latter runs.

5. Discussion and conclusions

The fixed-axis cone (2.2) establishes an algebraic boundary before a solver is run. Within that cone, a transverse field can avoid the leading projection error of a generic Cholesky construction. Slow modulation then produces a specified nonzero local stress divergence with an initial error controlled by the envelope-to-carrier ratio (2.6). Its subsequent delivery depends separately on viscous time, nonlinear interaction and resolution. The 512^3 coherence result illustrates why a single v_{eff} inferred from energy or diagonal stress cannot characterise every tensor component: the cross component has a quadratic-in-time correction already in the continuum projected equation.

Delivery of prescribed Reynolds stresses




Figure 8. Rotating target at $N = 128$, $t = 0.01$, $K = 6$: prescribed, delivered and difference fields for all three force components on $y = \pi$. Dotted lines delimit the two core bands; difference panels have their own stated scales.

Table 17 states the measured delivery window for the tested Gaussian-vortex target. All four grids reproduce the initial stress to within 0.317% and its full Cartesian force to within 1.80% at $t = 0.01$. By $t = 0.1$, the force error relative to that *same initial target* is at least 18.018%, despite the much smaller intergrid field differences. The paired control and table 18 attribute that loss predominantly to background-dependent evolution of the odd-pulse stress, rather than to the direct cross stress or the tested pulse-amplitude range. The early comparison establishes delivery and its subsequent loss for this target. The strength–time scan in table 19 makes the delivery window operational: within its 90 sampled cells, $\Omega_0 t$ separates passing and failing full-force errors at the 1, 5 and 10% tolerances. This empirical ordering is specific to the sampled setup and does not turn the separate carrier screens into sufficient force-error conditions.

A matched background–amplitude matrix measures the background-dependent force change at the same early comparison time. We vary $\Omega_0 = 0, 0.25, 0.5$ and $\varepsilon = 5 \times 10^{-4}, 10^{-3}, 2 \times 10^{-3}$ at $N = 128$ and 256, retaining the target and diagnostic definition of table 17. The no-pulse control is shared across amplitudes at fixed Ω_0 .

Time	$2\nu k^2 t$	$\varepsilon k t$	Stress/initial target (%)	Full force/initial target (%)
0.01	1.152×10^{-4}	2.4×10^{-4}	0.3156–0.3162	1.793–1.796
0.1	1.152×10^{-3}	2.4×10^{-3}	2.543–2.546	18.018–18.043

Table 17. Measured $K = 6$ core-volume target-error ranges $100\mathcal{E}_{\text{target}}^{R^K}$ and $100\mathcal{E}_{\text{target}}^{F^K}$ over the matched $N = 128, 192, 256, 512$ nonconstant-target Gaussian-vortex triplets, with $k = 24$, $\nu = 10^{-5}$, $\varepsilon = 10^{-3}$, $\sigma_0 = 0.6$, $\mathcal{Q}_0 = 0.5$ and $\Delta t = 0.01$. All four rows underlying these ranges use the common 128^3 core quadrature. The target and denominator remain the initial prescribed tensor and its full Cartesian divergence at both times. The background-amplitude products $\mathcal{Q}_0 t$ are 0.005 and 0.05, respectively. At both times $2\pi/(k\sigma_0) = 0.436$, while $k\Delta x$ ranges from 1.178 at $N = 128$ to 0.295 at $N = 512$. Ranges describe the four grids, not uncertainty intervals or intergrid errors.

$\mathcal{Q}_0$	ε	$N = 128$	$N = 256$
0	5×10^{-4}	0.237707	0.237066
0	10^{-3}	0.237731	0.237090
0	2×10^{-3}	0.237829	0.237189
0.25	5×10^{-4}	8.994806	8.982443
0.25	10^{-3}	8.994781	8.982418
0.25	2×10^{-3}	8.994683	8.982320
0.5	5×10^{-4}	18.042696	18.018096
0.5	10^{-3}	18.042672	18.018073
0.5	2×10^{-3}	18.042579	18.017979

Table 18. Full three-dimensional Cartesian $K = 6$ force error $100\mathcal{E}_{\text{target}}^{F^K}(0.1; N)$, in per cent of the initial analytic target-force RMS for each ε . All cases use $k = 24$, $\sigma_0 = 0.6$, $\nu = 10^{-5}$, $\Delta t = 0.01$, matched positive/negative/no-pulse runs and the common 128^3 core. Native-grid velocity products precede filtering. The rows are parameter samples, not uncertainty estimates.

At $N = 256$, increasing $\mathcal{Q}_0$ from zero to 0.25 and 0.5 raises the $\varepsilon = 10^{-3}$ full-force error from 0.2371 to 8.9824 and 18.0181 %. By contrast, the largest within-row error span over the fourfold amplitude range is only 1.23×10^{-4} percentage points. After division of the force fields by ε^2 , their largest difference from the $\varepsilon = 10^{-3}$ reference is 2.49×10^{-4} % of the initial target norm. At $N = 256$, the separately normalised projected-force errors at $t = 0.1$ and $\varepsilon = 10^{-3}$ are 0.1844, 16.3287 and 32.6954 % as $\mathcal{Q}_0$ increases. The maximum $N = 128$ –256 difference among the nine evolved full-force fields is 0.0310 % relative to the $N = 256$ field, on the same core.

The background-induced change of the $N = 256$ force field is 8.9896 and 18.0271 % of the initial target at $\mathcal{Q}_0 = 0.25$ and 0.5. These magnitudes are nearly proportional to $\mathcal{Q}_0$ over the sampled range, but the field is not exactly linear: the second difference between the three background strengths has RMS 0.4912 % of the initial target, or 2.725 % of the $\mathcal{Q}_0 = 0.5$ background response. Thus, at $t = 0.1$, the full-force target error is nearly proportional to background strength but changes by at most 1.23×10^{-4} percentage points under a fourfold change of pulse amplitude. The decomposition (4.3) in table 11 and the integrated sources in table 12 explain this contrast: background strain changes the odd pulse and hence its quadratic stress, while the direct cross force and the isolated envelope-projection error are small. This identifies the dominant measured origin of the approximately 18 % baseline departure. At $(\mathcal{Q}_0, \varepsilon) = (0.5, 10^{-3})$, the $N = 256$ common-core value 18.0181 % and the independently computed $N = 512$ value 18.0178 % both round to 18.018 %; the latter grid

Force-error budget	$t = 0.01$	$t = 0.05$	$t = 0.1$
1 %	0.25/0.5	0.05/0.1	0.025/0.05
5 %	0.5/–	0.25/0.5	0.125/0.15
10 %	0.5/–	0.5/–	0.25/0.5

Table 19. Empirical strength–time force-error tolerance brackets at $k = 24$, $\nu = 10^{-5}$, $\varepsilon = 10^{-3}$, $\sigma_0 = 0.6$, $\Delta t = 0.01$ and spherical $K = 6$. Each entry lists the largest passing strength followed by the first failing strength among the nine sampled values, separated by a slash. The dash means that the largest sampled value, 0.5, passes, with no stronger case tested. Full Cartesian force errors are vector RMS differences from the initial target, divided by its RMS on the common 128^3 Fourier-lifted core $0.3 \leq r \leq 1.2$; native-grid velocity products precede filtering. All three members of each triplet are saved at ten times. These brackets do not assign a crossing point between sampled strengths.

has not been run over the entire parameter matrix. The strength–time scan below locates empirical tolerance brackets for this configuration.

To locate a practical force-error window, we extend the matched positive/negative/control scan to nine strengths $0 \leq \Omega_0 \leq 0.5$ at $\varepsilon = 10^{-3}$, saving all three velocity fields every 0.01 through $t = 0.1$ on $N = 128$ and 256. The force and denominator are those of table 17; in particular, each comparison is with the *fixed initial* three-dimensional Cartesian target, not with a time-evolved or projected reference. Table 19 reports the largest passing and first failing *sampled* strength at three representative times. We classify a cell by $\max_{N \in \{128, 256\}} \mathcal{E}_{\text{target}}^{FK}(t; N)$, the larger of its two target errors, so every stated pass holds on both grids.

The ten-time grid also resolves latest-passing and first-failing times at each strength. At $\Omega_0 = 0.05$, for example, a 1 % budget passes at $t = 0.05$ but fails at 0.06; at $\Omega_0 = 0.15$, a 5 % budget passes at 0.09 but fails at 0.10. Across all 90 sampled strength–time cells, the product $\Omega_0 t$ orders the force-error classifications without overlap. A descriptive ordinary least-squares fit with a free intercept and equal weight for each of the 90 $N = 256$ cells gives

$$100\mathcal{E}_{\text{target}}^{FK}(t; 256) \simeq 0.0430 + 357.6 \Omega_0 t,$$

where the fitted ordinate is expressed in per cent. The RMS and maximum absolute residuals are 0.056 and 0.194 percentage points. This fit summarises the approximate background-exposure scaling over the sampled box; its constant intercept does not reproduce the time-dependent zero-background error. The correlated time samples are not independent replicates, and the fit does not replace the measured tolerance brackets. The largest passing and smallest failing products are 0.0025/0.0030 for a 1 % tolerance, 0.0135/0.0140 for 5 %, and 0.025/0.030 for 10 %. These three measured brackets are an operational summary of the force-delivery window for this fixed carrier, viscosity, amplitude, filter, core and sampled time range. The nonzero error at $\Omega_0 = 0$, the finite sampling gaps and the absence of tests outside this parameter box preclude a universal threshold or a sufficient condition at unsampled points.

Four dimensionless comparisons help screen candidate parameters: $2\pi/(k\sigma_0) \lesssim 0.5$ for carrier separation from the background scale, $2\nu k^2 t \lesssim \eta_{\text{max}}$ for viscous loss, $\varepsilon k t \lesssim 0.5$ for provisional pulse interaction, and $k\Delta x \lesssim 1.2$ for the tested compact discretisation. In a 2π box the last bound means $k \lesssim 1.2N/(2\pi)$. These are diagnostic screening conditions, not sufficient conditions for delivery to the fixed initial force: the tested $t = 0.1$ vortex satisfies all four with $\eta_{\text{max}} = 0.1$ but has an approximately 18 % force error. Background deformation is absent from these four screens. The measured $\Omega_0 t$ brackets in table 19

ν	k_{fid}	k_{nl}	k_{grid}	Screened interval
10^{-5}	50.0	250	36.7	[16.8, 36.7]
10^{-4}	15.8	250	36.7	empty
10^{-3}	5.0	250	36.7	empty

Table 20. Illustrative carrier-wavenumber screening interval for $\sigma_0 = 0.75$, $t = 2$, $N = 192$, $\varepsilon = 10^{-3}$ and $\eta_{\text{max}} = 0.1$. The lower edge is set by the stated scale-separation threshold. The nonlinearity and grid thresholds are empirical diagnostics for the tested family, so the interval is a design aid rather than a proven feasible domain.

provide the complementary force-error tolerance test for the sampled Gaussian-vortex setup. Table 20 separately screens carrier choices before that test.

For an illustrative carrier-screening calculation, let $\eta_{\text{max}} = 0.1$, $\sigma_0 = 0.75$, $t = 2$, $N = 192$ and $\varepsilon = 10^{-3}$. The scale-separation threshold $2\pi/(k\sigma_0) \leq 0.5$ requires $k \geq 16.8$, while the three proposed upper bounds are

$$k_{\text{fid}} = \sqrt{\frac{\eta_{\text{max}}}{2\nu t}}, \quad k_{\text{nl}} = \frac{0.5}{\varepsilon t}, \quad k_{\text{grid}} = \frac{1.2N}{2\pi}. \quad (5.1)$$

The continuous cutoffs in (5.1) are displayed in table 20; a particular computation must also choose an admissible integer carrier. Passing this screen does not imply a prescribed force-error tolerance is met, as table 17 shows for the Gaussian-vortex target. At $\nu = 10^{-5}$ the grid bound, not the viscous or provisional nonlinear bound, sets the upper edge. At $\nu = 10^{-4}$ and 10^{-3} , the viscous upper edge lies below the scale-separation lower edge: the screened carrier window closes entirely on the higher-viscosity side at the specified $t = 2$, $N = 192$ and $\sigma_0 = 0.75$. Increasing ε to 10^{-2} would move the nonlinear edge to 25 and contract the first-row screened interval accordingly.

The isolated $N = 192$, $k = 24$, $\sigma_0 = 0.75$, $\nu = 10^{-5}$, $\varepsilon = 10^{-3}$ check yielded a 2.27 % stress error at $t = 2$ and $\nu_{\text{eff}}/\nu = 1.0006$. Its mean-flow amplitude was zero, so the chosen ratio $2\pi/(k\sigma_0) = 0.349$ does not test delivery in a varying background. The later Gaussian-vortex case uses a different $\sigma_0 = 0.6$ and demonstrates that passing the four screens alone does not retain a small force error against the unchanged initial target. In a 2π box the grid threshold also gives $N \geq 2\pi k/1.2$, rounded up to an admissible mesh; $k\sigma_0 = 30$ and 50 at $\sigma_0 = 0.75$ require $N \geq 210$ and $N \geq 350$, respectively, under this empirical rule. Vortex confinement inside the periodic box is a separate consideration.

The evidence has distinct scopes. Fourier and cone calculations establish initial realizability for specified wave families. Constant-target vortex runs show close agreement in global covariance and volume-averaged production, while their radial profiles expose sensitivity to diagnostic definition and do not certify radial-force convergence. The 128^3 – 512^3 nonconstant-target triplets show accurate early local delivery, but their approximately 18 % force error at $t = 0.1$ marks a loss of agreement with the fixed initial target. Decreasing differences in the evolved *low-mode* fields through $t = 1.9$ establish a separate grid-consistency result; the $t = 0.01$ comparison mainly checks the common initial construction. For multidirectional targets, only the rotating family has an early three-grid comparison; none has a late-time three-grid result. The joint background–amplitude scan in table 18 maps the fixed $t = 0.1$ force error over a narrow parameter box: it varies strongly with Ω_0 and barely with ε . Table 19 extends that result to ten early times and locates sampled 1, 5 and 10 % full-force tolerance brackets at $N = 128$ and 256 . The isolated factorial and 512^3 amplitude scans probe different nonlinear effects; in particular, the observed ε^2 coherence

law has not been tested in the Gaussian-vortex background. Appendix A evaluates the initial projection bound at the nonconstant-target parameters; the general filter-stability constants remain too large to certify the measured core errors. Its conditional finite-time estimate has conservative Sobolev constants and does not sharply certify the high-frequency runs. Unfiltered convergence, a late-time independent-solver check of the nonuniform local force, and tolerance boundaries outside the sampled $(k, \nu, \varepsilon, \Omega_0, t)$ box remain open. No finite-time singularity claim follows from these smooth-solution computations.

Declaration of interests

The authors declare no competing interests.

Funding

This research was supported by the National Key Research and Development Program of China (No. 2025YFB3003401).

Data availability

The data supporting this study, including input files, analysis scripts and compact filtered summaries, will be made available by the corresponding author on reasonable request.

CRediT authorship contribution statement

Jinlin Wang: Writing – original draft, Validation, Software, Methodology. Yi Heng: Writing – review & editing, Supervision, Project administration, Funding acquisition.

Computational assistance

Codex (OpenAI) assisted with numerical workflow development, diagnostic processing and manuscript editing. The authors are responsible for verifying the analyses and the manuscript content.

Appendix A. Proofs and quantitative bounds

The following statements concern continuum fields on the periodic torus $\mathbb{T}^3 = (\mathbb{R}/2\pi\mathbb{Z})^3$. They separate exact algebraic realisability from the error incurred by modulation, incompressibility projection and evolution. All norms in this section use normalised Lebesgue measure on the torus.

A.1. The fixed three-axis cone

For a finite collection $\mathcal{N}$ of nonzero lattice directions, define

$$C_{\mathcal{N}} = \left\{ \sum_{n \in \mathcal{N}} B_n : B_n = B_n^{\top} \succeq 0, \quad B_{nn} = 0 \right\}. \quad (\text{A } 1)$$

For the three coordinate axes, (A 1) is the classical factor-width-two cone: each block is a sum of rank-one terms with at most two nonzero coordinates (Boman *et al.* 2005). The next theorem gives an explicit three-dimensional allocation form for that cone. Each B_n is the covariance of at most two polarisations transverse to n . Distinct carrier frequencies can be assigned to these polarisations; the low-frequency covariance is then the sum in (A 1).

Theorem 1 (exact fixed-axis criterion). Let T be symmetric positive definite, write $d_i = T_{ii} > 0$ and $\rho_{ij} = T_{ij}/\sqrt{d_i d_j}$. For $\mathcal{N} = \{e_x, e_y, e_z\}$, $T \in C_{\mathcal{N}}$ if and only if there are $p, q, r \in [0, 1]$ satisfying

$$pq \geq \rho_{xy}^2, \quad (1-p)r \geq \rho_{xz}^2, \quad (1-q)(1-r) \geq \rho_{yz}^2. \quad (\text{A } 2)$$

Proof. The z -transverse block must supply the sole xy covariance. Write its two diagonal entries as pd_x, qd_y . The y -transverse block must then have xx entry $(1-p)d_x$; write its zz entry as rd_z . The remaining x -transverse block has $yy = (1-q)d_y$ and $zz = (1-r)d_z$. All diagonal allocations are nonnegative exactly when $p, q, r \in [0, 1]$. The three determinants are nonnegative exactly when (A 2) holds. Thus every representation obeys these conditions. Conversely, the three blocks with the specified diagonal entries and off-diagonal entries T_{xy}, T_{xz}, T_{yz} are positive semidefinite and sum to T . Diagonalising each block produces at most two transverse polarisations. $\square$

The fixed-axis set has a useful certificate form. Let $P_n = I - nn^\top/|n|^2$. The cone in (A 1) is closed: in a convergent sequence of sums, positivity bounds the trace of each summand by the trace of the sum, so a subsequence of every summand converges. Finite-dimensional conic duality therefore gives

$$T \in C_N \iff \text{tr}(HT) \geq 0 \quad \text{for every symmetric } H \text{ with } P_n H P_n|_{n^\perp} \succeq 0 \text{ for all } n \in N. \quad (\text{A } 3)$$

Indeed, nonnegativity of $\text{tr}(HB_n)$ on each plane-positive block is equivalent to the displayed restriction on H ; the bipolar theorem then yields (A 3).

This criterion gives an analytic replacement for a sampled feasibility rate. In the equicorrelation family $T_\rho = (1-\rho)I + \rho u u^\top$, where $0 \leq \rho < 1$, the fixed three-axis construction is feasible if and only if $\rho \leq 1/2$. For sufficiency take $p = q = r = 1/2$. For necessity take $H = 2I - u u^\top$: every coordinate-plane restriction of H is positive semidefinite, whereas $\text{tr}(HT_\rho) = 3 - 6\rho$. Consequently $T_{0.6}$ is strictly positive definite but lies outside the fixed-axis cone. This does not obstruct constructions using other wave vectors.

A.2. How many lattice directions are needed?

Theorem 2 (a generic lower bound and a pointwise upper bound). A positive-definite 3×3 tensor needs at least two transverse lattice directions. With the directions allowed to depend on T , almost every such tensor needs at least three. Every positive-definite T admits a pointwise representation as a sum of at most six transverse rank-one blocks, with a possibly different nonzero integer lattice direction for each block. The almost-everywhere statement uses six-dimensional Lebesgue measure on the symmetric matrices; the six-direction claim does not bound the lengths of the integer vectors.

Proof. One transverse block has rank at most two. For two distinct directions, the linear space of symmetric matrices supported in either transverse plane has dimension three for each plane, and their intersection is the one-dimensional space supported on the planes' common line. Their sum therefore has dimension five. Each pair of lattice directions imposes a nontrivial linear equation on T . The countable union of these five-dimensional sets has measure zero.

For the upper bound, normalise T to trace one. Rank-one projectors onto rational directions are dense among all rank-one projectors. Their convex hull has closure equal to the trace-one positive-semidefinite set. A positive-definite trace-one T lies in its relative interior, and the relative interior of a finite-dimensional convex set equals the relative interior of its closure. Hence T is a finite convex combination of rational rank-one projectors. Carathéodory's theorem in the five-dimensional trace-one affine space reduces the combination to at most six terms. For every rational polarisation b , a nonzero integer $n \perp b$ exists, making its rank-one block transverse. Rescaling restores $\text{tr } T$. $\square$

Theorem 2 is pointwise. A smooth field $T(x)$ requires additional work to choose one finite set of carriers for every x , construct smooth nonnegative coefficients, and separate the Fourier supports. It also gives no claim that three directions suffice for every positive-definite tensor. General positive-stress decompositions have precedents in geometric-lemma and Mikado-flow constructions (Daneri & Székelyhidi 2017); the present fixed-axis and filtered force statements concern a more specific class of waves.

A.3. A quantitative projection and filtered-force bound

Let G_K be the orthogonal Fourier projector onto $\{m \in \mathbb{Z}^3 : |m| \leq K\}$, where $K \geq 1$, and set

$$R_K(w) = G_K(w \otimes w) - G_K w \otimes G_K w, \quad F_K(w) = -\nabla \cdot R_K(w). \quad (\text{A } 4)$$

The filter and force in (A 4) are the operators used in the following stability estimate. Define $C_K = \sqrt{\#\{m \in \mathbb{Z}^3 : |m| \leq K\}}$. Fourier Cauchy–Schwarz gives $\|G_K w\|_\infty \leq C_K \|w\|_2$.

Lemma 3 (filter stability). For bounded vector fields w, z on the torus,

$$\|R_K(w) - R_K(z)\|_2 \leq (\|w\|_\infty + \|z\|_\infty + C_K(\|w\|_2 + \|z\|_2))\|w - z\|_2, \quad (\text{A } 5)$$

$$\|F_K(w) - F_K(z)\|_2 \leq K(\|w\|_\infty + \|z\|_\infty + 2C_K(\|w\|_2 + \|z\|_2))\|w - z\|_2. \quad (\text{A } 6)$$

Proof. Expand each quadratic difference as $(w - z) \otimes w + z \otimes (w - z)$. The first term in R_K is an L^2 contraction of G_K . For the second, use $\|G_K w\|_\infty \leq C_K \|w\|_2$. The divergence of the first filtered term has Fourier support

Delivery of prescribed Reynolds stresses

within K , and that of the product of two filtered fields within $2K$. Applying these Bernstein multiplier bounds yields (A 5)–(A 6). $\square$

Consider finitely many real modulated waves

$$v(x) = \sqrt{2}\varepsilon \sum_{j=1}^J a_j(x) b_j \cos(\kappa_j \cdot x + \phi_j), \quad b_j \cdot \kappa_j = 0, \quad \text{supp } \widehat{a}_j \subseteq \{|\ell| \leq q_j\}. \quad (\text{A } 7)$$

Assume all envelopes are real and every individual carrier and every nonconstant self- or cross-product carrier remains outside the cutoff K after the envelope bandwidth is added. More explicitly it suffices that $|\kappa_j| - q_j > K$, $2|\kappa_j| - 2q_j > K$, and $|\kappa_i \pm \kappa_j| - q_i - q_j > K$ for $i \neq j$, while $2q_j \leq K$. The original paired six-mode basis is also covered: two waves may share the same carrier if their phases differ by $\pi/2$; their zero-frequency cross term then vanishes exactly, while the double-carrier term must remain above K . Under either form of sufficient separation,

$$G_K v = 0, \quad R_K(v) = T_v = \varepsilon^2 \sum_j a_j^2 b_j b_j^\top, \quad F_K(v) = -\nabla \cdot T_v. \quad (\text{A } 8)$$

Theorem 4 (initial projection and force error). Let $\mathcal{P}$ be the Leray projector and $u_0 = \mathcal{P}v$. For (A 7), with $|\kappa_j| > q_j$,

$$D_{\text{proj}} := \|u_0 - v\|_2 \leq \varepsilon \sum_j \frac{q_j}{|\kappa_j| - q_j} |b_j| \|a_j\|_2. \quad (\text{A } 9)$$

If the nonresonance conditions above hold, the stress and force errors against the prescribed T_v and its divergence are bounded by (A 5) and (A 6) with $w = u_0$, $z = v$, $\|w - z\|_2 = D_{\text{proj}}$.

Proof. At Fourier frequency $\kappa_j + \ell$, transversality gives

$$(I - \mathcal{P}(\kappa_j + \ell))b_j = \frac{(\ell \cdot b_j)(\kappa_j + \ell)}{|\kappa_j + \ell|^2}.$$

Its norm is at most $q_j |b_j| / (|\kappa_j| - q_j)$. The same estimate holds for the negative carrier. Parseval and the separation of positive and negative carrier supports give the single-wave estimate; the triangle inequality gives (A 9). The final assertion follows from (A 8) and Lemma 3. $\square$

These are absolute bounds. A relative force estimate requires a stated positive lower bound for $\|\nabla \cdot T_v\|_2$. The factor K in (A 6) also makes clear why a velocity $q_j/|\kappa_j|$ estimate alone does not prove the empirical force slopes in figure 1.

Evaluation for the nonconstant-target pulse. For $a(x) = 1 + \eta \cos x$, $\eta = 1/4$, the two in-phase z -carrier polarisations in (2.4) must be grouped into one vector-valued wave when applying the nonresonance statement. The resulting five carrier groups have

$$\sum_j |b_j| = \sqrt{6/5} + 2\sqrt{2/5} + 2\sqrt{1/5}.$$

At $k = 24$, $q = 1$, and $\varepsilon = 10^{-3}$, (A 9) gives

$$D_{\text{proj}} \leq \frac{\varepsilon}{k-1} \left(\sqrt{6/5} + 2\sqrt{2/5} + 2\sqrt{1/5} \right) \sqrt{1 + \eta^2/2} = 1.4371 \times 10^{-4}.$$

The constant part of a is already solenoidal; applying the same argument only to $a - 1$ lowers this estimate to 2.5016×10^{-5} . For the idealised coefficients in (2.4), only the x -polarised y - and z -carrier sidebands have $b_j \cdot \nabla a \neq 0$. Their orthogonal Fourier supports give the exact continuum value

$$D_{\text{proj}} = \frac{\varepsilon \eta}{\sqrt{2(k^2 + 1)}} = 7.3593 \times 10^{-6}.$$

Thus the q/k velocity scale is realised by this field, while the unrefined triangle bound leaves substantial avoidable slack.

The projection-only Fourier calculation at $t = 0$ uses the same $K = 6$ filter and $N = 128$ core mask as the vortex comparisons, but no time evolution or background flow. It gives 0.1843 % stress and 0.1240 % full-force error against the prescribed initial target. This direct calculation is more informative for this specific finite Fourier field than the general stability estimates. For completeness, the core occupies a fraction $f_C = 0.1076660$ of the comparison grid; its target stress and force norms are 1.082167×10^{-6} and 2.659842×10^{-7} . Using $\|E\|_{2,C} \leq f_C^{-1/2} \|E\|_2$, $C_6 = \sqrt{925}$, $\|v\|_2 = 1.573213 \times 10^{-3}$, $\|v\|_\infty \leq 3.856647 \times 10^{-3}$, and the eight projection sidebands to bound $\|u_0\|_\infty \leq \|v\|_\infty + \sqrt{8} D_{\text{proj}}$, (A 5)–(A 6) with the unrefined D_{proj} yield relative core bounds of 42.0 and 1.97×10^3 for stress and force (as dimensionless ratios, not percentages). Even inserting the exact

velocity correction leaves these general relative bounds at 2.14 and 100.7. The force bound 100.7 exceeds the directly computed projection-only force-error ratio 0.001240 by about 8.1×10^4 . The general estimate admits every Fourier mode through K when bounding products, whereas this prescribed field has only a finite set of carrier and modulation sidebands. Coefficientwise filtering and differentiation give the stated $t = 0$ error directly, but do not yet furnish a general sharp bound for the core norm or the evolved vortex field at $t = 0.01$.

A.4. An explicit conditional finite-time estimate

The constant in the next theorem depends only on s and the chosen normalisation of the torus. One way to certify it from Fourier sums is to put $w_m = (1 + |m|^2)^{s/2}$ and define

$$E_s^2 = \sum_{m \in \mathbb{Z}^3} w_m^{-2}, \quad B_s^2 = \sum_{m \in \mathbb{Z}^3} |m|^2 w_m^{-2}, \quad D_s = \sup_{\ell, m \in \mathbb{Z}^3} \frac{|m| |w_{m+\ell} - w_m|}{|\ell| w_m + |m| w_\ell},$$

with a zero numerator and denominator interpreted as zero. For $s > 5/2$, the two sums converge and the mean-value inequality for w_m makes $D_s < \infty$. Fourier Cauchy–Schwarz gives the Sobolev embedding constant E_s ; the transport commutator estimate uses the displayed weight difference and two applications of $\sum |m| |\widehat{u}_m| \leq B_s \|u\|_{H^s}$. The commutator norm is bounded by $2D_s B_s M^2$, while the L^2 nonlinear term and embedding are bounded with E_s . Thus the explicit choice $A_s = E_s + 4D_s B_s$ is admissible below. No optimised numerical value of this deliberately conservative constant is claimed.

Theorem 5 (viscous and nonlinear evolution). Let $s > 5/2$, $u_0 = \mathcal{P}v \in H^{s+2}$, and let $u(t)$ be the smooth unforced periodic Navier–Stokes solution with viscosity $\nu > 0$ and initial field u_0 . Choose a Sobolev constant A_s large enough that, for a smooth solenoidal solution with $M = \|u\|_{H^s}$,

$$\frac{dM}{dt} \leq A_s M^2, \quad \|\mathcal{P}\nabla \cdot (u \otimes u)\|_2 \leq A_s M^2, \quad \|u\|_\infty \leq A_s M.$$

Write $M_0 = \|u_0\|_{H^s}$ and suppose $0 \leq t < 1/(A_s M_0)$. Then

$$\|u(t) - v\|_2 \leq D(t) := D_{\text{proj}} + \nu t \|\Delta u_0\|_2 + \frac{A_s M_0^2 t}{1 - A_s M_0 t}, \quad (\text{A } 10)$$

$$\|F_K(u(t)) + \nabla \cdot T_\nu\|_2 \leq K \left[\frac{A_s M_0}{1 - A_s M_0 t} + \|v\|_\infty + 2C_K(M_0 + \|v\|_2) \right] D(t), \quad (\text{A } 11)$$

provided (A 8) holds. Replacing the bracket in (A 11) by the bracket of (A 5) gives the corresponding stress bound.

Proof. The standard H^s energy and product inequalities, using $\nabla \cdot u = 0$, give $M(t) \leq M_0/(1 - A_s M_0 t)$. Duhamel’s formula relative to the heat semigroup gives

$$u(t) - e^{\nu t \Delta} u_0 = - \int_0^t e^{\nu(t-\tau)\Delta} \mathcal{P}\nabla \cdot (u \otimes u)(\tau) d\tau.$$

The L^2 contraction of the heat semigroup bounds this difference by $A_s M_0^2 t/(1 - A_s M_0 t)$. Also $\|e^{\nu t \Delta} u_0 - u_0\|_2 \leq \nu t \|\Delta u_0\|_2$. Adding the projection difference proves (A 10). The L^2 energy inequality gives $\|u(t)\|_2 \leq \|u_0\|_2 \leq M_0$. Apply Lemma 3 to $u(t)$ and v , and use (A 8). $\square$

For a solenoidal background U , apply Theorem 5 separately to the three raw initial fields $U + v$, $U - v$ and U . The quadratic identity

$$\frac{1}{2} \{F_K(U + v) + F_K(U - v) - 2F_K(U)\} = F_K(v)$$

shows that half the sum of the three individual error bounds (with the background bound counted twice) controls the matched even-in-amplitude diagnostic. This valid estimate can be loose: an $O(\varepsilon^2)$ sharp even-response estimate needs a second-variation analysis of the flow map. Theorem 5 is a finite-time smooth-solution bound and does not assert global regularity, a useful bound at $\varepsilon k t = O(1)$, or discretisation error control.

A.5. Short-time cross-stress coefficient

For the isolated Euler initial field v_0 in (2.4), define

$$a_0 = -\mathcal{P}[(v_0 \cdot \nabla)v_0], \quad b_0 = -\mathcal{P}[(a_0 \cdot \nabla)v_0 + (v_0 \cdot \nabla)a_0].$$

These are $\partial_t u(0)$ and $\partial_t^2 u(0)$. With $R_{ij} = \langle u_i u_j \rangle$, differentiation gives

$$R''_{ij}(0) = \langle (b_0)_i (v_0)_j + 2(a_0)_i (a_0)_j + (v_0)_i (b_0)_j \rangle.$$

Fourier multiplication of the six signed carrier frequencies gives $\langle (a_0)_x^2 \rangle = (9/25)\varepsilon^4 k^2$ and $\langle (v_0)_x (b_0)_x \rangle = -(9/25)\varepsilon^4 k^2$, so $R''_{xx}(0) = 0$. For the cross component, $\langle (a_0)_x (a_0)_y \rangle = 0$ and each of $\langle (b_0)_x (v_0)_y \rangle$ and $\langle (v_0)_x (b_0)_y \rangle$ equals $-(3/25)\varepsilon^4 k^2$. The first derivatives vanish by Fourier support. Consequently $R_{xy}(0) = (3/5)\varepsilon^2$ and $R''_{xy}(0) = -(6/25)\varepsilon^4 k^2$ give

$$-\log \left[\frac{R_{xy}(t)/R_{xx}(t)}{3/5} \right] = \frac{1}{5}\varepsilon^2 k^2 t^2 + o(t^2), \quad t \rightarrow 0,$$

at fixed ε and k . This is the inviscid local-in-time coefficient; a viscous fit over a finite interval need not equal it.

REFERENCES

- BANERJEE, S., KRAHL, R., DURST, F. & ZENGER, CH. 2007 Presentation of anisotropy properties of turbulence, invariants versus eigenvalue approaches. *Journal of Turbulence* **8**, N32.
- BARTHOLOMEW, PAUL, DESKOS, GEORGIOS, FRANTZ, RICARDO A.S., SCHUCH, FELIPE N., LAMBALLAIS, ERIC & LAIZET, SYLVAIN 2020 Xcompact3D: an open-source framework for solving turbulence problems on a Cartesian mesh. *SoftwareX* **12**, 100550.
- BOMAN, ERIK G., CHEN, DORON, PAREKH, OJAS & TOLEDO, SIVAN 2005 On factor width and symmetric H-matrices. *Linear Algebra and its Applications* **405** (1–3), 239–248.
- BROWN, DAVID L., CORTEZ, RICARDO & MINION, MICHAEL L. 2001 Accurate projection methods for the incompressible Navier–Stokes equations. *Journal of Computational Physics* **168**, 464–499.
- BUCKMASTER, TRISTAN & VICOL, VLAD 2019 Nonuniqueness of weak solutions to the Navier–Stokes equation. *Annals of Mathematics* **189**, 101–144.
- CAI, YUNZHU, WAN, JIAWEI & KAREEM, AHSAN 2023 A new divergence-free synthetic eddy method for generating homogeneous isotropic turbulence with a prescribed energy spectrum. *Computers & Fluids* **253**, 105788.
- CHORIN, ALEXANDRE JOEL 1968 Numerical solution of the Navier–Stokes equations. *Mathematics of Computation* **22**, 745–762.
- DANERI, SARA & SZÉKELYHIDI, LÁSZLÓ 2017 Non-uniqueness and h -principle for Hölder-continuous weak solutions of the Euler equations. *Archive for Rational Mechanics and Analysis* **224**, 471–514.
- DAVIDSON, L. & BILLSON, M. 2006 Hybrid LES-RANS using synthesized turbulent fluctuations for forcing in the interface region. *International Journal of Heat and Fluid Flow* **27**, 1028–1042.
- DE LELLIS, CAMILLO & SZÉKELYHIDI, LÁSZLÓ 2012 The h -principle and the equations of fluid dynamics. *Bulletin of the American Mathematical Society* **49**, 347–375.
- DHAMANKAR, NITIN S., BLAISDELL, GREGORY A. & LYRINTZIS, ANASTASIOS S. 2018 Overview of turbulent inflow boundary conditions for large-eddy simulations. *AIAA Journal* **56**, 1317–1334.
- DREZE, YANN, HAO, MUTING & DI MARE, LUCA 2023 Divergence-free turbulent inflow data from realistic covariance tensor. *Physics of Fluids* **35**, 025120.
- FISCHER, PAUL F., LOTTES, JAMES W. & KERKEMEIER, STEFAN G. 2008 Nek5000 web page. Official solver website: <https://nek5000.mcs.anl.gov/>.
- FYHN, EIRIK HOLM & FYHN, HURSANAY 2026 turbulentSyntheticEddyInlet: a self-correcting, divergence-free synthetic-eddy boundary condition for OpenFOAM. *SoftwareX* **35**, 102907.
- GERMANO, MASSIMO, PIOMELLI, UGO, MOIN, PARVIZ & CABOT, WILLIAM H. 1991 A dynamic subgrid-scale eddy viscosity model. *Physics of Fluids A: Fluid Dynamics* **3**, 1760–1765.
- ISETT, PHILIP 2018 A proof of Onsager’s conjecture. *Annals of Mathematics* **188**, 871–963.
- JARRIN, N., BENHAMADOU, S., LAURENCE, D. & PROSSER, R. 2006 A synthetic-eddy-method for generating inflow conditions for large-eddy simulations. *International Journal of Heat and Fluid Flow* **27**, 585–593.
- KEATING, ANTHONY, PIOMELLI, UGO, BALARAS, ELIAS & KALTENBACH, HANS-JAKOB 2004 A priori and a posteriori tests of inflow conditions for large-eddy simulation. *Physics of Fluids* **16**, 4696–4712.
- KIM, JOHN, MOIN, PARVIZ & MOSER, ROBERT 1987 Turbulence statistics in fully developed channel flow at low Reynolds number. *Journal of Fluid Mechanics* **177**, 133–166.
- KIM, YUSIK, CASTRO, IAN P. & XIE, ZHENG-TONG 2013 Divergence-free turbulence inflow conditions for large-eddy simulations with incompressible flow solvers. *Computers & Fluids* **84**, 56–68.
- KLEIN, M., SADIKI, A. & JANICKA, J. 2003 A digital filter based generation of inflow data for spatially developing direct numerical or large eddy simulations. *Journal of Computational Physics* **186**, 652–665.
- LAIZET, SYLVAIN & LAMBALLAIS, ERIC 2009 High-order compact schemes for incompressible flows: A simple and efficient method with quasi-spectral accuracy. *Journal of Computational Physics* **228**, 5989–6015.
- LAUNDER, B. E., REECE, G. J. & RODI, W. 1975 Progress in the development of a Reynolds-stress turbulence closure. *Journal of Fluid Mechanics* **68**, 537–566.

- LELE, SANJIVA K. 1992 Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics* **103**, 16–42.
- LEONARD, A. 1975 Energy cascade in large-eddy simulations of turbulent fluid flows. *Advances in Geophysics* pp. 237–248.
- LILLY, D. K. 1992 A proposed modification of the Germano subgrid-scale closure method. *Physics of Fluids A: Fluid Dynamics* **4**, 633–635.
- LUMLEY, JOHN L. 1979 Computational modeling of turbulent flows. In *Advances in Applied Mechanics*, vol. 18, pp. 123–176. Elsevier.
- LUND, THOMAS S., WU, XIAOHUA & SQUIRES, KYLE D. 1998 Generation of turbulent inflow data for spatially-developing boundary layer simulations. *Journal of Computational Physics* **140**, 233–258.
- DI MARE, L., KLEIN, M., JONES, W. P. & JANICKA, J. 2006 Synthetic turbulence inflow conditions for large-eddy simulation. *Physics of Fluids* **18**, 025107.
- MISHRA, AASHWIN A. & GIRIMAJI, SHARATH S. 2014 On the realizability of pressure–strain closures. *Journal of Fluid Mechanics* **755**, 535–560.
- MOIN, PARVIZ & MAHESH, KRISHNAN 1998 Direct numerical simulation: A tool in turbulence research. *Annual Review of Fluid Mechanics* **30**, 539–578.
- PATERA, ANTHONY T. 1984 A spectral element method for fluid dynamics: laminar flow in a channel expansion. *Journal of Computational Physics* **54** (3), 468–488.
- PATTERSON, J. W., BALIN, R. & JANSEN, K. E. 2021 Assessing and improving the accuracy of synthetic turbulence generation. *Journal of Fluid Mechanics* **906**, R1.
- POPE, STEPHEN B. 2000 *Turbulent Flows*. Cambridge University Press.
- REYNOLDS, OSBORNE 1895 IV. On the dynamical theory of incompressible viscous fluids and the determination of the criterion. *Philosophical Transactions of the Royal Society of London. (A.)* pp. 123–164.
- ROLFO, STEFANO, FLAGEUL, CÉDRIC, BARTHOLOMEW, PAUL, SPIGA, FILIPPO & LAIZET, SYLVAIN 2023 The 2DECOMP&FFT library: an update with new CPU/GPU capabilities. *Journal of Open Source Software* **8**, 5813.
- SCHAU, KYLE A., JOHNSON, CHELSEA, MULLER, JULIA & OEFELEIN, JOSEPH C. 2022 An ensemble synthetic eddy method for accurate treatment of inhomogeneous turbulence. *Computers & Fluids* **248**, 105671.
- SCHUMANN, U. 1977 Realizability of Reynolds-stress turbulence models. *The Physics of Fluids* **20**, 721–725.
- SKILLEN, A., REVELL, A. & CRAFT, T. 2016 Accuracy and efficiency improvements in synthetic eddy methods. *International Journal of Heat and Fluid Flow* **62**, 386–394.
- TABOR, G.R. & BABA-AHMADI, M.H. 2010 Inlet conditions for large eddy simulation: A review. *Computers & Fluids* **39**, 553–567.
- WANG, LIQIU 1997 Frame-indifferent and positive-definite Reynolds stress–strain relation. *Journal of Fluid Mechanics* **352**, 341–358.
- WU, XIAOHUA 2017 Inflow turbulence generation methods. *Annual Review of Fluid Mechanics* **49**, 23–49.
- XIE, ZHENG-TONG & CASTRO, IAN P. 2008 Efficient generation of inflow conditions for large eddy simulation of street-scale flows. *Flow, Turbulence and Combustion* **81**, 449–470.