

Finite-amplitude realization of prescribed Reynolds stress in three-dimensional viscous flow: projection, local forcing, and grid convergence

Jinlin Wang

wangjlin26@mail2.sysu.edu.cn

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

Research Article

Keywords: Reynolds-stress synthesis, incompressible flow, oscillatory pulses, projection, numerical verification

Posted Date: October 14th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11356942/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.

Finite-amplitude realization of prescribed Reynolds stress in three-dimensional viscous flow: projection, local forcing, and grid convergence

Jinlin Wang

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

13 October 2026

Abstract

Oscillatory velocity fields can realize a prescribed Reynolds stress in a formal high-frequency construction, but finite-amplitude delivery in a viscous solver requires the field to survive incompressibility projection, diffusion, nonlinear interaction, and interaction with a nonuniform background. We study these restrictions in a periodic three-dimensional incompressible flow. A naive Cholesky construction matches its target covariance before projection but suffers an order-one stress change afterwards. We give a six-mode solenoidal construction, characterize its realizable correlation tensors by three inequalities, and observe less than 0.01% projection-induced change for the tested constant targets. Isolated-pulse decay follows $1 - \exp(-2\nu_{\text{eff}}k^2t)$ with $\nu_{\text{eff}}/\nu = 0.9991$; appreciable self-interaction begins near $\varepsilon kt = 1$ in the tested amplitude range. For a Gaussian-vortex background, we additionally prescribe a spatially varying positive-definite target and its nonzero force $-\nabla \cdot T$. At $N = 128$, $t = 0.01$, and fixed spectral cutoff $K = 6$, the local stress and full-force errors are 0.316% and 1.796%. After evolution to $t = 1.9$, the induced three-component force field differs by 0.166% between $N = 128$ and 192 and by 0.0114% between $N = 192$ and 256 on a common low-mode grid. These results establish a verified numerical example of local stress and force delivery, while leaving general target fields, the unfiltered spectrum, and any implication for finite-time singularities open.

Keywords: Reynolds-stress synthesis; incompressible flow; oscillatory pulses; projection; numerical verification.

1 Introduction

The use of oscillatory velocity fields to supply a prescribed momentum flux is a recurring device in constructions for incompressible flows. In a recent preprint, [OpenAI \(2026\)](#) present a forced three-dimensional Navier–Stokes blow-up construction whose oscillatory correction supplies a required stress to a mean flow. Related smooth-forcing constructions for the three-dimensional Euler equations have been reported by [Alpöge & Buckmaster \(2026\)](#), building on the program of [Córdoba & Martínez-Zoroa \(2023\)](#). These papers address mathematical existence questions. We study a narrower numerical question: at finite amplitude and finite resolution, how accurately can a specified stress and its induced force be realized in a viscous incompressible computation?

The construction is perturbative. An approximate blow-up is assembled from a leading-order field and a hierarchy of corrections. The residual of the assembled fields,

$$R(u, p) = \partial_t u + (u \cdot \nabla)u - \nu \Delta u + \nabla p, \quad (1)$$

is treated as the external force. A central step in the cited Navier–Stokes preprint is the correction of a residual stress by high-frequency oscillations whose averaged quadratic products supply a momentum flux. The finite-amplitude computation below does not reproduce that paper’s space–time profiles, prove that its hypotheses hold numerically, or assess its blow-up claim. It isolates one constitutive step in a simple periodic box.

At finite resolution, a designed oscillation may be changed by the pressure projection, viscously attenuated before its stress acts on the background, or modified by nonlinear mode coupling. These effects need not share a single small parameter. We therefore separate algebraic realizability, projection, isolated-pulse evolution, and interaction with a nonuniform background. We use matched positive and negative pulse amplitudes and a no-pulse control to distinguish the even stress and force from the background and from odd cross terms.

Our principal new test prescribes a positive-definite stress that varies in space and therefore has a nonzero target divergence. This directly checks local stress and force delivery, rather than inferring it from a globally averaged tensor or a force created only after subsequent evolution. The target, norm, spectral cutoff, core region, and advancement gates were fixed before the higher-resolution experiments.

Section 2 describes the construction. Section 3 sets out the numerical and diagnostic protocol. Sections 4–7 report the measured constraints and local-delivery results. Section 8 assembles a candidate parameter window; section 9 states the evidential limits.

2 Problem formulation

2.1 Synthesis and the projection obstruction

Let T be a target Reynolds-stress tensor, symmetric and positive definite. We synthesise an oscillatory velocity field as

$$u'_i(\mathbf{x}) = \sqrt{2}\varepsilon \sum_{m=1}^3 L_{im} \cos(\mathbf{k}_m \cdot \mathbf{x} + \varphi_m), \quad T = LL^\top, \quad (2)$$

with L the lower Cholesky factor of T and integer wave vectors $\mathbf{k}_m$ so that the modes are exactly orthogonal in a periodic box. For distinct wave vectors the spatial average gives

$$\langle u'_i u'_j \rangle = 2\varepsilon^2 \sum_{m,n} L_{im} L_{jn} \langle \cos(\mathbf{k}_m \cdot \mathbf{x} + \varphi_m) \cos(\mathbf{k}_n \cdot \mathbf{x} + \varphi_n) \rangle = \varepsilon^2 \sum_m L_{im} L_{jm} = \varepsilon^2 T_{ij}, \quad (3)$$

so that (2) realises the target exactly.

In a velocity–pressure solver, however, the initial field is projected onto the solenoidal subspace by the pressure Poisson equation solved during the first time step. For a mode with wave vector $\mathbf{k}$ the projection removes the component of the amplitude vector along $\mathbf{k}$,

$$\mathbf{a}^{(k)} \longrightarrow \mathbf{a}_\perp^{(k)} \equiv \mathbf{a}^{(k)} - (\hat{\mathbf{k}}_k \cdot \mathbf{a}^{(k)}) \hat{\mathbf{k}}_k, \quad \mathbf{a}^{(k)} = (L_{1k}, L_{2k}, L_{3k})^\top, \quad (4)$$

and the realised stress is $\varepsilon^2 \sum_k \mathbf{a}_\perp^{(k)} \otimes \mathbf{a}_\perp^{(k)}$, which is not $\varepsilon^2 T$ in general. Table 1 shows the effect for the reference target $T = \begin{pmatrix} 1 & 0.6 & 0 \\ 0.6 & 1 & 0 \\ 0 & 0 & 0.4 \end{pmatrix}$ with $\mathbf{k}_k \propto (1, 0, 1), (0, 1, 1), (1, 1, 0)$: the diagonal entries are rescaled by 0.25, 0.52 and 2.03. The projection is therefore not a small correction but a leading-order modification, and a synthesis that ignores it measures a property of the projection rather than of the flow.

2.2 A divergence-free pulse basis

To remove the obstruction we construct the pulses so that $\nabla \cdot \mathbf{u}' \equiv 0$ by construction. Take three wave vectors aligned with the coordinate axes, $\mathbf{k} \in \{k\hat{\mathbf{x}}, k\hat{\mathbf{y}}, k\hat{\mathbf{z}}\}$, and two orthogonal

	s_{11}	s_{22}	s_{33}	s_{12}	s_{13}	s_{23}
design, $\varepsilon^2 T$	10.00	10.00	4.00	6.00	0	0
measured before projection	10.00	10.00	3.97	6.00	0	0
measured after projection	2.50	5.20	8.07	3.00	-2.50	-4.60
prediction (4)	2.50	5.20	8.10	3.00	-2.50	-4.60

Table 1: Reynolds-stress components in units of 10^{-5} for the Cholesky synthesis (2) with $\varepsilon = 10^{-2}$, measured in a 64^3 periodic box. The tabulated pre-projection components are within 0.75% of the design entries, while the post-projection components agree with the modewise prediction (4) to the displayed precision.

polarisations for each, one with a sine and one with a cosine phase; the polarisation of a mode necessarily lies in the plane normal to its wave vector, so $\nabla \cdot \mathbf{u}' \equiv 0$ identically.

The essential point is a pairing rule. The wave-vector families along z , y , and x contribute, respectively, to the (xx, xy, yy) , (xx, xz, zz) , and (yy, yz, zz) blocks of T . Because each off-diagonal entry is controlled by exactly one block, and because a 2×2 symmetric positive-semidefinite block is realised exactly by its two orthogonal polarisations (with amplitudes $\sqrt{2\mu_i}$ for eigenvalues μ_i), the synthesis reduces to distributing the diagonal entries among three blocks.

Writing the distribution as parameters $p, q, r \in [0, 1]$ and the pairwise correlations as $\rho_{ij} = T_{ij}/\sqrt{T_{ii}T_{jj}}$, the three blocks are positive semidefinite if and only if

$$\boxed{pq \geq \rho_{xy}^2, \quad (1-p)r \geq \rho_{xz}^2, \quad (1-q)(1-r) \geq \rho_{yz}^2.} \quad (5)$$

This is the additional restriction that the divergence-free constraint imposes on the positive-semidefinite cone. It involves the correlations only, and is therefore independent of the anisotropy magnitude.

The construction is exact: for three generic targets, including the reference T above, the realised stress matches the target with a relative error of 2.2×10^{-16} and a divergence residual of zero. In the solver the projection then leaves the realised stress unchanged to within 0.01 per cent, against the factors of up to four in table 1.

2.3 Contraction of the realisable set

Conditions (5) do not exhaust the positive-semidefinite cone. Sampling 2×10^4 random correlation matrices with unit diagonal, the minimal six-mode basis is realisable for 44.5 per cent of them, and the feasibility rate collapses as the correlations strengthen (table 2). A perfectly correlated (rank-one) target is infeasible. This is a property of the minimal basis and not a fundamental limit: adding wave-vector families enlarges the realisable set, and determining how many families are required to cover a given T is itself a well-posed question that we leave open.

3 Numerical platform and verification

3.1 Solver and case

All simulations use XCOMPACT3D (Bartholomew et al., 2020; Laizet & Lamballais, 2009), a three-dimensional incompressible solver with sixth-order compact spatial derivatives, a three-stage Runge–Kutta time advance, and an FFT-based pressure Poisson solver. A case module imposes the background and pulse as initial conditions and evaluates in-situ diagnostics. The experiments reported here use a periodic box of side 2π . The primary vortex and spatially varying-target cases have $\nu = 10^{-5}$, $\Delta t = 0.01$, $N \in \{128, 192, 256\}$ grid points per direction, $k = 24$, $\varepsilon = 10^{-3}$, $\sigma_0 = 0.6$, and $\Omega_0 = 0.5$; they run for 200 steps and save all three velocity

$\max \rho_{ij} $	realisable fraction
[0, 0.5)	100.0 %
[0.5, 0.7)	83.1 %
[0.7, 0.9)	37.8 %
[0.9, 1.0]	6.4 %
rank one	0 %
all	44.5 %

Table 2: Feasibility of the minimal divergence-free basis (three axis-aligned wave vectors, two polarisations each) as a function of the strongest pairwise correlation, from 2×10^4 random samples.

components and pressure every ten steps. The explicit subgrid-stress branch is disabled in these inputs (`'ilesmod=0'`); the unused `'jles'` field in the input file does not activate it. Background-maintenance forcing is also disabled (`'hold_mean_run=0'`), so the vortex evolves freely after initialization.

Two mean flows are available in the case module: a Gaussian vortex column $u_\theta(r) = \Omega_0 r \exp(-r^2/2\sigma_0^2)$ and a plane shear layer. All parameters — ε , k , σ_0 , Ω_0 , the target tensor and the diagnostic cadence — are read at run time from a namelist, so that parameter scans require no recompilation. For the three-dimensional local diagnostic we use the fixed spherical spectral cutoff G_K with $|\mathbf{k}| \leq K$ and define

$$R_{ij} = G_K(u_i u_j) - G_K(u_i) G_K(u_j), \quad F_i = -\partial_j R_{ij}. \quad (6)$$

The induced even component of any field Q is $(Q_{+\varepsilon} + Q_{-\varepsilon})/2 - Q_0$, where the three runs share all parameters except the pulse sign. Unless stated otherwise, field norms use the three-dimensional core $0.3 \leq r \leq 1.2$ over all z , the tensor Frobenius norm for stress, and the three-component Euclidean norm for force and velocity. $K = 6$ is the primary cutoff; $K = 4, 8$ are sensitivity checks. Time points from one deterministic run are not treated as independent statistical replicates.

3.2 Mean-flow decomposition

The fluctuation field needs a specified reference background. The original radial-bin estimate of the mean is delicate in two respects that we record because each can silently corrupt the result. The later Gaussian-vortex experiment instead subtracts the analytic viscous evolution of its initial background. It was designed after the original no-pulse baseline failed the predeclared one-per-cent threshold and is analysed as a separate family.

First, the azimuthal average of the *Cartesian* components of a pure swirl vanishes identically, $\langle -\Omega r \sin \theta \rangle_\theta = \langle \Omega r \cos \theta \rangle_\theta = 0$. A radial-bin average of Cartesian components therefore does not remove a swirling mean flow at all; the resulting contamination reaches two orders of magnitude above the signal. The average must be taken on the cylindrical components (u_r, u_θ, u_z) and transformed back.

Second, bins at small radius contain fewer than one grid point, since an annulus of width Δr at radius r has area $\sim 2\pi r \Delta r$. If empty bins are left at zero, the interpolation used to evaluate the mean selects zero, and the contamination *grows* with the number of bins. A nearest-valid-bin fill restores monotone convergence; with 64 bins the mean is recovered to 0.57 per cent.

For the plane shear a horizontal average can remove pulse components that are uniform on horizontal planes. The corrected periodic-shear control instead subtracts the analytically decaying sinusoidal background. The original tanh profile is discontinuous across the periodic boundary, so its frozen-background fluctuation growth is not used as a physical result (section 6).

3.3 Verification

The platform is verified by three independent checks, each of which the conclusions of this paper depend on.

Exact synthesis. With the mean flow switched off, the first diagnostic call returns the stress of the unprojected initial condition; it reproduces the design value $\varepsilon^2 T$ to five significant figures (table 1). With the divergence-free basis the same value is returned after the projection as before it.

Viscous decay against theory. For an isolated pulse ($\Omega_0 = 0$) the fluctuation energy must decay as $\exp(-2\nu k^2 t)$ exactly. Fitting $\langle u'^2 \rangle(t)$ for $k = 6$ gives $d \ln \langle u'^2 \rangle / dt = -0.0721$ at $\nu = 10^{-3}$, against the analytic value -0.0720 , and -0.3604 at $\nu = 5 \times 10^{-3}$, against -0.3600 . The agreement, 0.1–0.2 per cent, is an end-to-end test of the synthesis, the diagnostic and the time integration together.

Symmetry. The initial condition is odd in z about $z = 0$ and about $z = L_z/2$; centred discretisations commute with this symmetry and the residual remains at 10^{-15} throughout. The in-situ azimuthal spectrum (appendix A) returns an identically zero sine component for $m = 0$, as it must.

4 Delivery fidelity and the viscous window

The quantity that the synthesis step must control is not the pulse energy but the *fidelity* with which the prescribed stress is delivered,

$$\eta(t) = \frac{\|\langle \mathbf{u}' \otimes \mathbf{u}' \rangle(t) - \varepsilon^2 T\|_F}{\|\varepsilon^2 T\|_F}. \quad (7)$$

For an isolated pulse whose modes all decay at the same rate νk^2 , $\eta = 1 - \exp(-2\nu_{\text{eff}} k^2 t)$, so that the delivered stress is faithful provided

$$\boxed{\eta \simeq 2\nu_{\text{eff}} k^2 t \ll 1 \quad \Longleftrightarrow \quad t_{\text{deliver}} \ll \frac{1}{2\nu_{\text{eff}} k^2}}. \quad (8)$$

Fitting (7) with the mean flow switched off returns $\nu_{\text{eff}}/\nu = 0.9991$, confirming both the form and the coefficient. (8) is the operational constraint on the delivery time.

A caution on measurement is worth recording. A linear fit of $\eta(t)$ over a finite interval returns a slope depressed by the curvature of the exponential: over $t \in [0.15, 2]$ with $2\nu k^2 = 0.072$ the average slope is smaller than the initial slope by a factor 0.9306, and a linear fit returns exactly that factor. The exponential form must be used whenever a rate is to be extracted.

5 Onset of self-interaction

At finite amplitude the pulse is not an exact solution: the quadratic term transfers energy to harmonics and to a mean, which modifies the decay of the primary modes. The natural parameter is $\varepsilon k t$, the number of nonlinear times elapsed. Table 3 reports the deviation of $\langle u'^2 \rangle$ from its initial value over $t = 1.8$ at $Re = 10^5$, where the viscous contribution is negligible.

The smallest-amplitude case already differs from the viscous energy-ratio baseline by 4.6×10^{-4} . At $\varepsilon k t = 1.08$ and 2.16 the excess is of order 10^{-2} , while at $\varepsilon k t = 4.32$ the pulse energy is 21.9% below its initial value. These cases locate the onset of appreciable self-interaction near $\varepsilon k t = 1$ and demonstrate substantial loss by $\varepsilon k t \approx 4$, consistent with the nonlinear time scale $(\varepsilon k)^{-1}$; they do not define a universal sharp threshold.

ε	εkt	$\langle u'^2 \rangle$ ratio	deviation	excess over viscous
0.05	0.54	0.999166	-8.3×10^{-4}	$+4.6 \times 10^{-4}$
0.10	1.08	1.008489	$+8.5 \times 10^{-3}$	$+9.8 \times 10^{-3}$
0.20	2.16	1.002818	$+2.8 \times 10^{-3}$	$+4.1 \times 10^{-3}$
0.40	4.32	0.781363	-2.2×10^{-1}	-2.2×10^{-1}

Table 3: Deviation of the fluctuation energy from its initial value over $t = 1.8$ for an isolated divergence-free pulse at $k = 6$, $\nu = 10^{-5}$. The viscous baseline is -1.30×10^{-3} .

grid	baseline / target	$\langle P \rangle$ at $t = 2$	stress diff. / $N = 256$
$N = 128$	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

Table 4: Analytic-background Gaussian vortex at $\nu = 10^{-5}$, $k = 24$, $\varepsilon = 10^{-3}$, $\sigma_0 = 0.6$, and $\Omega_0 = 0.5$. Stress difference is the relative Euclidean norm of the six-component global stress at $t = 2$. Production is $\langle -2R_{r\theta}S_{r\theta} \rangle$.

6 Nonuniform background: global stress and production

The synthesis step is meant to operate in a spatially varying mean flow. We measure the volume-averaged production term

$$P = -\langle u'_i u'_j \rangle S_{ij}, \quad (9)$$

For the Gaussian vortex we subtract the analytically diffusing background before forming the stress and use its analytic $S_{r\theta}$. This diagnostic was introduced after the original radial-bin method failed the predeclared no-pulse baseline limit: at $N = 128$ its residual stress was 2.728 per cent of the designed stress norm, above the 1 per cent limit. The repaired analytic-background family is therefore reported separately.

The $N = 256$ stress at $t = 0.1$ differs from its designed tensor by 0.1086 per cent in norm. At $t = 2$ this difference is 1.649 per cent and includes physical viscous and background-flow evolution. The matched zero-mean controls give zero $\langle P \rangle$ at $t = 2$ on all three grids. These results establish a convergent global covariance and volume-averaged production at one parameter point. We additionally output radial profiles of cylindrical stress components and their axisymmetric radial divergence, and compare the azimuthal mean-velocity response with a matched no-pulse run. Because a single pulse sign contains a substantial odd-in-amplitude response on some grids, we report the even component from the $+\varepsilon$ and $-\varepsilon$ pair.

The paired decomposition materially changes the response comparison: the single-sign response differs strongly between $N = 192$ and $N = 256$, whereas its even component differs by 3.36 per cent. The radial angular-momentum budget residual is 5–9 per cent of the force

grid	force RMS	response RMS	force diff.	response diff.
$N = 128$	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

Table 5: Core-weighted RMS of the even-in- ε radial force and mean azimuthal-velocity response at $t = 2$, and weighted profile differences from $N = 256$ over $0.2 \leq r \leq 1.2$. The force is the axisymmetric radial average, not the full three-dimensional $-\nabla \cdot \mathbf{R}(\mathbf{x})$.

RMS over the three grids. The $N = 128$ half-time-step comparison changes the force profile by less than 10^{-5} in relative L^2 . These checks support a grid-consistent radial diagnostic at this parameter point, but do not establish the full spatial tensor divergence or exact local delivery of a prescribed force field.

We audited the budget after observing this residual. The stored radial moment $R_{r\theta}^{\text{bin}} = \overline{u_r(u_\theta - U_{\text{bg}})}^r$ gives the total radial angular momentum flux as $R_{r\theta}^{\text{bin}} + U_{\text{bg}}\overline{u_r}^r$. Integrating the azimuthal equation over $0.3 \leq r \leq 1.2$ uses this flux only at the shell endpoints and avoids differentiating a binned profile at every radius. At $t = 1.9$, the shell-integrated residual is below 1.7 per cent of the integrated absolute tendency on each of the three grids across four local polynomial interpolation stencils. Coarsening the archived $N = 256$ profile from 64 to 32 radial bins changes the pointwise residual from 8.95 to 5.82 per cent, while further coarsening to 16 bins increases it to 13.48 per cent. Thus radial binning and differentiation materially affect the pointwise budget; the shell result is a separate, weaker closure check and does not establish pointwise balance.

We subsequently saved all three velocity components and pressure at 20 times for matched positive-pulse, negative-pulse and no-pulse cases at $N = 128, 192, 256$. For a periodic spectral filter G_K with $|\mathbf{k}| \leq K$, we computed the full Cartesian tensor $R_{ij} = G_K(u_i u_j) - G_K(u_i)G_K(u_j)$ and its divergence $F_i = -\partial_j R_{ij}$, without imposing axisymmetry. At $t = 0.1$, the core-volume Frobenius RMS difference between the even-in-amplitude stress and $\varepsilon^2 T$ is 2.573, 2.579 and 2.573 per cent at the three grids, respectively, for $K = 6$ over $0.3 \leq r \leq 1.2$ and all z . At $t = 1.9$, the Leray-projected residual of the filtered Cartesian momentum budget is 0.218, 0.0494 and 0.0434 per cent of the full three-component force RMS. At $N = 256$, this ratio decreases from 0.166 to 0.0434 per cent across $t = 0.5, 1.0, 1.5, 1.9$. These core-volume RMS checks include the saved pressure in the unprojected budget; they do not bound the maximum pointwise residual.

To compare the spatial force rather than only its norm, we transfer the filtered force to the same 128^3 Fourier grid and use an identical core mask. At $t = 1.9$, the relative three-component L^2 difference falls from 0.1895 per cent between $N = 128$ and 192 to 0.01306 per cent between $N = 192$ and 256 for $K = 6$. The corresponding pairs are 0.1266 and 0.00913 per cent for $K = 4$, and 0.2291 and 0.01522 per cent for $K = 8$. This supports convergence of the force field at the selected low-mode filter scales and parameter point; the force magnitude changes appreciably with K , so the result does not cover the unfiltered full spectrum.

Reconstructing the old radial diagnostic from unfiltered three-dimensional snapshots reproduces its archived force to relative L^2 accuracy better than 1.5×10^{-9} . Applying the old radial bin and difference formula to the $K = 6$ filtered tensor differs from the radial average of its direct three-dimensional divergence by 3.57, 1.79 and 3.28 per cent across the grids, but differs from the old radial force by 28.6, 26.8 and 28.4 per cent. The main distinction is therefore the stress definition: the old quantity averages products of velocity relative to the analytic background over all resolved modes, whereas the new tensor is a filtered covariance. The designed tensor is spatially constant, so its target divergence is zero; the evolved nonzero force is not evidence of delivery of a prescribed nonzero local force field.

We therefore made a separate test with a prescribed nonconstant local target. The six-mode pulse was multiplied at initialisation by $a(x) = 1 + 0.25 \cos x$ on the 2π -periodic domain, with $k = 24$ and $\varepsilon = 10^{-3}$. The target was fixed before simulation as

$$\begin{aligned} \mathbf{T}_{\text{target}}(x) &= \varepsilon^2 a(x)^2 \mathbf{T}_0, \\ \mathbf{T}_0 &= \begin{pmatrix} 1 & 0.6 & 0 \\ 0.6 & 1 & 0 \\ 0 & 0 & 0.4 \end{pmatrix}, \\ F_i^{\text{target}} &= -\partial_j T_{ij}^{\text{target}} = 0.5 \varepsilon^2 a(x) \sin x (T_0)_{i1}. \end{aligned} \tag{10}$$

Since $a(x) \geq 0.75$ and the eigenvalues of $\mathbf{T}_0$ are 0.4, 0.4, 1.6, the target is positive definite

everywhere and its force is nonzero. Multiplication by $a(x)$ makes the raw pulse no longer exactly solenoidal; the solver’s pressure projection is part of the construction error, not a reason to redefine the target after the fact.

Before extending the grids, we specified a primary filter cutoff $K = 6$, an annular core $0.3 \leq r \leq 1.2$ over all z , and separate thresholds for the stress, full force, and Leray-projected force at $t = 0.01$. The $N = 128$ isolated-pulse errors were 0.183, 0.136 and 0.081 per cent, below the respective 5, 10 and 10 per cent limits. With the Gaussian-vortex background $\Omega_0 = 0.5$, $\sigma_0 = 0.6$, they were 0.316, 1.796 and 3.238 per cent, below the 8, 15 and 15 per cent limits. The projected three-dimensional momentum-budget residual at $t = 0.1$ was 0.0335 per cent of the induced-force RMS, below the predeclared 2 per cent limit. Each gate used matched positive-pulse, negative-pulse and no-pulse cases, with 20 saved three-dimensional times per case.

After these gates passed, we ran matched Gaussian-vortex triplets to $t = 2$ at $N = 128, 192, 256$, saving all three velocity components and pressure every ten steps. For an observable Q , its induced even component was defined by $(Q_+ + Q_-)/2 - Q_0$. At $t = 0.1$, the stress errors against the *initial* target were 2.546, 2.553 and 2.547 per cent at the three grids; the full-force errors were 18.043, 18.167 and 18.070 per cent. These later errors include background evolution, and the initial target was not assumed to remain stationary. At $K = 6, t = 1.9$, the induced-force RMS was $(1.1744, 1.1781, 1.1752) \times 10^{-6}$, respectively. The Leray-projected budget residual was 0.197, 0.0429 and 0.0371 per cent of that force. The un-subtracted no-pulse force is about 2.375×10^{-3} , so matched control subtraction is essential at this parameter point.

We transferred the $t = 1.9$ stress, force and even low-mode velocity response to the same 128^3 Fourier grid before comparing fields on an identical core mask. At $K = 6$, the relative L^2 differences between $N = 128$ and 192, then $N = 192$ and 256, were respectively 0.0925 and 0.00620 per cent for stress, 0.166 and 0.0114 per cent for the full three-component force, and 0.151 and 0.0125 per cent for response. The same ordering held at the preselected sensitivity cutoffs $K = 4, 8$.

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

Table 6: Relative core-volume L^2 field differences for the prescribed nonconstant-target experiment. Each field is first mapped to a common 128^3 Fourier grid; each denominator is the second grid in its pair.

This supplies numerical evidence for local delivery and subsequent grid-consistent evolution for this prescribed nonconstant target at fixed low-mode scales. It does not establish a general tensor-synthesis theorem, an unfiltered full-spectrum limit, or pointwise maximum-norm budget closure.

Single-factor $N = 128$ slices in k , σ_0 , ν , ε , and Ω_0 show changes in force and response, but are exploratory: they do not resolve parameter interactions or establish a joint feasibility boundary. For example, at $t = 2$ the core force RMS increases from 2.90×10^{-7} to 1.16×10^{-6} and 4.65×10^{-6} across $\varepsilon = 0.0005, 0.001, 0.002$, while it ranges from 5.86×10^{-7} to 2.80×10^{-6} across $\Omega_0 = 0.25$ to 1. These trends are single-grid observations, not scaling laws.

The original tanh plane-shear profile is discontinuous across the periodic boundary. Its fitted logarithmic fluctuation-energy growth decreases from 1.4772 to 0.9023 between $N = 128$ and $N = 256$, and at $N = 128$ the no-pulse control contributes 99.81 per cent of the pulse-run energy at $t = 1$. A smooth periodic sine shear, analysed against its diffusing background, has a near-zero fitted rate (-0.0007763) on all three grids. Since the mean profile changed, this is a

diagnostic control, not a convergence demonstration for the original tanh problem. We draw no physical-instability or quantitative-delivery inference from the original shear runs.

7 Decomposition of the apparent excess decay

A quantity that we initially took to be the numerical dissipation of the scheme — the excess decay rate of the pulse beyond the viscous prediction, $\nu_{\text{num}} = \nu_{\text{eff}} - \nu$ — is in fact neither numerical nor a dissipation. We report the two discriminating experiments because the distinction changes the operational constraints.

Doubling the resolution at fixed *physical* wavenumber leaves ν_{num} unchanged: 1.878×10^{-6} at $N = 64$ and 1.891×10^{-6} at $N = 128$, a fitted scaling exponent of -0.01 in Δx . A discretisation error would scale as a positive power — Δx^6 for the sixth-order compact scheme. Varying the amplitude instead gives $\nu_{\text{num}}/\varepsilon^2 = 1.876 \times 10^{-2}$, 1.889×10^{-2} , 1.891×10^{-2} and 1.889×10^{-2} for $\varepsilon = 10^{-3}$, 3×10^{-3} , 10^{-2} , 3×10^{-2} , constant to three significant figures. The excess decay is therefore the signature of a quadratic nonlinearity: it is the self-interaction of section 5, measured in the language of a linear decay.

Genuine numerical dissipation is much smaller, and separable. At $\varepsilon = 10^{-3}$, where nonlinearity is negligible, the effective viscosity agrees with the molecular value to within 0.02 per cent even at $k\Delta x \approx 0.98$ (six grid points per wavelength), and the deviation is *negative* — the discrete Laplacian underestimates $|k|^2$, so the pulse decays more slowly than continuum theory. A second-order central scheme at the same resolution deviates by -0.72 per cent, converging to -0.18 per cent when the resolution is doubled: exactly second order.

We draw two conclusions. First, the numerical error is not the binding constraint: $k\Delta x \lesssim 1.2$ suffices for sub-per-cent accuracy with the sixth-order scheme. Second, and more generally, any quantity identified as a numerical error should be subjected to both a grid-refinement and an amplitude -variation test before it is reported as one; here the two tests together reversed the interpretation.

8 The feasible wedge

Collecting the isolated-pulse results gives four candidate conditions for a jointly feasible nonuniform-background regime:

1. *scale separation*, $\epsilon_{\text{WKB}} \equiv 2\pi/(k\sigma_0) \ll 1$, so that the pulse wavelength is short compared with the mean-flow scale;
2. *viscous delivery*, $2\nu k^2 t < \eta_{\text{max}}$, from (8);
3. *weak nonlinearity*, $\epsilon k t \lesssim 0.5$, from section 5;
4. *resolution*, $k\Delta x \lesssim 1.2$, from section 7.

Conditions 1 and 4 bound k from below and above and combine with the requirement that the mean flow be confined well inside the domain to give a pure grid constraint,

$$\boxed{k\sigma_0 \lesssim N/10}, \quad (11)$$

so that a scale separation $k\sigma_0$ requires $N \gtrsim 10 k\sigma_0$; $k\sigma_0 = 50$, for instance, requires approximately $N \geq 500$.

Within a given grid the window in k is $[4\pi/\sigma_0, \min(\sqrt{\eta_{\text{max}}/(2\nu t)}, 0.5/(\epsilon t), 1.2/\Delta x)]$, listed in table 7. The dominant bound is the viscous delivery one: decreasing ε by three orders of magnitude raises the nonlinear bound by the same factor and leaves the viscous bound

ν	Re	k_{fid}	k_{nl}	feasible window
10^{-5}	10^5	50.0	≥ 250	[16.8, 36.7]
10^{-4}	10^4	15.8	≥ 250	empty
10^{-3}	10^3	5.0	≥ 250	empty

Table 7: Predicted candidate wavenumber window for $\sigma_0 = 0.75$, $t = 2$, $\eta_{\text{max}} = 0.1$, $N = 192$ and $\varepsilon \leq 10^{-2}$. The lower bound $k_{\text{min}} = 16.8$ is set by scale separation ($\epsilon_{\text{WKB}} \leq 0.5$).

untouched. Within this low-amplitude slice, reducing viscosity widens the candidate window more effectively than further reducing amplitude.

For the listed parameters the predicted window is narrow and opens at $\nu = 10^{-5}$ among the three viscosity levels examined. We tested the isolated pulse at $\sigma_0 = 0.75$, $\nu = 10^{-5}$ ($Re = 10^5$), $k = 24$, $\varepsilon = 10^{-3}$ on a 192^3 grid (eight grid points per wavelength). At $t = 2$ the measured stress error is $\eta = 2.27$ per cent and $\nu_{\text{eff}}/\nu = 1.0006$; $\varepsilon kt = 0.048$ and $k\Delta x = 0.785$. The input has zero mean-flow amplitude, so $\epsilon_{\text{WKB}} = 0.349$ checks only a chosen scale ratio, not WKB delivery in a varying background. The subsequent Gaussian-vortex experiment tests one nonuniform-background point: at $k = 24$, $\sigma_0 = 0.6$, $\nu = 10^{-5}$, $\varepsilon = 10^{-3}$ and $t = 2$, the four ratios are $2\pi/(k\sigma_0) = 0.436$, $2\nu k^2 t = 0.0230$, $\varepsilon kt = 0.048$, and $k\Delta x = 0.785$ at $N = 192$. This point satisfies the stated numerical thresholds, but one point does not locate the joint boundary or establish an asymptotic WKB regime.

9 Conclusions

This study isolates the finite-amplitude realization of prescribed Reynolds stress in a viscous incompressible solver. The algebraic choice of pulse basis matters: pressure projection changes a naive Cholesky field at leading order, whereas the tested six-mode solenoidal field preserves its stress to within 0.01%. Its exact covariance construction applies to the subset of target tensors satisfying (5). For isolated pulses the measured decay agrees with the viscous exponential; nonlinear self-interaction becomes visible around $\varepsilon kt = 1$ in the tested range. These two mechanisms set different limits on delivery time and amplitude.

The prescribed nonconstant target in (10) provides the strongest local test. Its $N = 128$ early stress and full-force errors pass the predeclared thresholds, and the later induced three-dimensional stress, force, and response show substantially smaller $N = 192$ – 256 than $N = 128$ – 192 differences at three fixed low-mode cutoffs. The three-component filtered momentum budget has a 0.0371% projected residual relative to the induced-force RMS on $N = 256$ at $t = 1.9$. This is evidence for one locally prescribed stress and force in an evolving Gaussian-vortex background, not a general realizability or convergence theorem.

Negative controls constrain the interpretation. The original radial-bin diagnostic failed its no-pulse baseline gate; the uncorrected periodic plane-shear runs were dominated by their no-pulse control. The radial-bin force and the later filtered three-dimensional force use distinct stress definitions and must not be combined into one convergence sequence. The present local conclusions are restricted to a single tensor, one background parameter point, core-volume RMS norms, and $K = 4, 6, 8$. The unfiltered spectrum, other targets, and a joint feasibility boundary remain unmeasured.

For the selected $\sigma_0 = 0.75$, $t = 2$, $N = 192$ and low-amplitude parameters, the predicted window appears at $Re = 10^5$ but is empty at the two lower Reynolds numbers tabulated. This threshold depends on those choices; the Gaussian-vortex cases provide one nonuniform-flow point but do not locate the boundary. Whether the constructions of [OpenAI \(2026\)](#) and [Alpöge & Buckmaster \(2026\)](#) operate inside or outside this window cannot be decided without the explicit profiles of their constructions, which are not reproduced here.

The planned $N = 512$ parameter scan and reproduction with an independent solver have not yet been performed. They are needed to extend the parameter window and test code dependence; the present claims are confined to the completed simulations.

Data and code availability

The numerical archive contains solver inputs, binary and input hashes, run-completion records, three-dimensional snapshots, and analysis scripts. The $N = 128$ – 256 varying-target results are recorded in `singularity-study/hpc-bundle/experiments/2026-10-13/` of the author’s project archive. That archive has not yet received a public repository identifier; the present arXiv source package contains the manuscript alone.

Computational assistance

Codex assisted with code changes, diagnostic processing, and manuscript drafting. The staged experimental-design and reporting procedure drew on the scientific-agent workflow of [Kassis et al. \(2026\)](#). Scientific responsibility and final verification remain with the named author.

A Azimuthal spectrum diagnostic

The in-situ diagnostic evaluates the azimuthal Fourier content of the fluctuation energy about the box axis by direct spectral projection,

$$E'(m) = \langle |\mathbf{u}'|^2 e^{-im\theta} \rangle, \quad \theta = \arctan[(y - y_c)/(x - x_c)], \quad (12)$$

accumulated pointwise so that no polar interpolation is required. For a pulse whose wave vectors are aligned with the axes, modes with $\mathbf{k} \parallel z$ carry $m = 0$ and those with $\mathbf{k} \parallel x$ or $\mathbf{k} \parallel y$ carry $m = \pm 1$; the measured spectrum is dominated by $m = 0$ with $m = 1$ next, as expected, and the $m = 0$ sine component vanishes identically. The diagnostic was verified on a short run and has not yet been used for a scientific conclusion.

B A note on the solver

The fourth-order options of the solver’s derivative switch (`ifirstder`, `isecondder` equal to 2 or 3) are documented in the input file but are not implemented: the coefficient routine aborts at run time with the message `Set of coefficients not ready yet`. Only the second-order central and sixth-order compact options are usable. Our cross-scheme verification in section 7 makes use of the two that work.

C Outstanding validation and reproducibility

The following items separate completed controls from measurements that remain necessary for a claim of local delivery or a joint feasibility boundary.

1. **Local delivery in a nonuniform mean flow.** The earlier Gaussian-vortex experiment has a constant target and a zero target divergence; its filtered three-dimensional force field converges at the selected low-mode cutoffs, while its original radial diagnostic uses a different stress definition. The subsequent experiment prescribes the nonconstant tensor and nonzero force in (10). Its $N = 128$ initial local stress and full force errors pass the predeclared thresholds, and its evolved stress, three-dimensional force and response agree

across three grids at $K = 4, 6, 8$. Wider-band or unfiltered checks and other target tensors remain necessary before asserting general local delivery.

2. **Plane-shear diagnostic interpretation.** The $N = 128, 192, 256$ runs and matched $N = 128$ no-pulse control are complete. The original non-periodic profile and frozen-mean diagnostic are confounded; a periodic-profile control diagnoses the issue. Further shear physics requires a separately designed periodic-background experiment.
3. **The wedge at $N = 512$.** Evaluation of (11) at $k\sigma_0 \approx 30$ requires $N \geq 512$. A scan of $k \in \{20, 24, 28, 32, 40\}$ at $\nu = 10^{-5}$, $\varepsilon = 10^{-3}$, $\sigma_0 = 0.75$ on a 512^3 grid, with the four-condition check reported for each point. Input files are provided in `hpc-bundle/inputs/wedge512/`; each point is analysed by `analysis/analyze.py`, which reports the four criteria directly.
4. **Independent reproduction.** The viscous window and the self-interaction threshold should be reproduced with a second code (spectral or a different finite-difference formulation). The pulse-synthesis tool is code-independent and can be reused directly.
5. **Reproduction of the construction’s own profiles.** If explicit profiles become available, the synthesis diagnostics can be applied to them without modification.

References

- Alpöge, L. and Buckmaster, T. 2026. Blowup for the Euler equations with smooth forcing. Preprint, <https://cims.nyu.edu/~tristanb/euler.pdf>.
- Bartholomew, P., Deskos, G., Frantz, R. A. S., Schuch, F. N., Lamballais, E. and Laizet, S. 2020. Xcompact3D: An open-source framework for solving turbulence problems on a Cartesian mesh. *SoftwareX* **12**, 100550. <https://doi.org/10.1016/j.softx.2020.100550>.
- Córdoba, D. and Martínez-Zoroa, L. 2023. Blow-up for the incompressible 3D-Euler equations with uniform $C^{1,1/2-\epsilon} \cap L^2$ force. arXiv:2309.08495. <https://arxiv.org/abs/2309.08495>.
- Kassis, T., Agarwal, V., He, Y., Patel, D. and Brueckner, A. M. 2026. Scientific Agent Skills: A Library of Procedural Knowledge for Research Agents. arXiv:2609.00065. <https://arxiv.org/abs/2609.00065>.
- Laizet, S. and Lamballais, E. 2009. High-order compact schemes for incompressible flows: a simple and efficient method with quasi-spectral accuracy. *Journal of Computational Physics* **228**, 5989–6015. <https://doi.org/10.1016/j.jcp.2009.05.010>.
- OpenAI. 2026. Finite time blowup for Navier–Stokes. Preprint. <https://cdn.openai.com/pdf/32d9f210-8b73-45e0-91bc-82a30aef8a9a/navier-stokes.pdf>.