IDNS: TRUE ZERO-DISSIPATION DNS OF THE TAYLOR–GREEN VORTEX AT ONE-EIGHTH NASA RESOLUTION VIA DETERMINISTIC BOUNDED TEMPORAL LIFTING

JEFFREY CAMLIN

Science Policy Research Unit, University of Sussex, Brighton, UK
Red Dawn Academic Press · AI Lab, Milwaukee, WI USA

Abstract. Numerical blow-up near peak vortex stretching has long been attributed to incipient singularity or fundamental limits of spatial resolution. We show it is neither. The obstacle is insufficient sampling of Physical Simulation Time: a problem of temporal allocation, not spatial discretization.

We introduce iDNS (intelligent Direct Numerical Simulation), a stiffness-aware integration framework based on *bounded deterministic temporal lifting*: a diffeomorphic reparameterization $t = \varphi(\tau)$ that separates Physical Simulation Time t from Lifted Computational Time τ. The lifting is governed by a bounded sigmoid controller

$$\frac{dt}{d\tau} = \varphi'(\tau) = 1 + \frac{A}{1 + \exp(-k(s-c))}, \quad s = \frac{\|\nabla\omega\|_{L^2}}{\|\omega\|_{L^2} + \epsilon},$$

where s is a real-time vorticity stiffness indicator. The controller parameters (A, k, c) are fixed from $\mathrm{Re} = 1600$ to $\mathrm{Re} = 10^8$, spanning five orders of magnitude in Reynolds number. No manual tuning, no CFL heuristics, no problem-specific expertise required.

The result: iDNS achieves $R_\varepsilon = 1.000 \pm 0.001$ at $N = 64^3$ on a consumer laptop. The NASA Glenn WRLES benchmark requires $N = 512^3$ on 368 CPU cores to approach the same dissipation balance. One-eighth the resolution. A fraction of the hardware. Zero artificial dissipation. All simulation data, plots, and results are fully reproducible via a 418-line Python script provided in the repository.

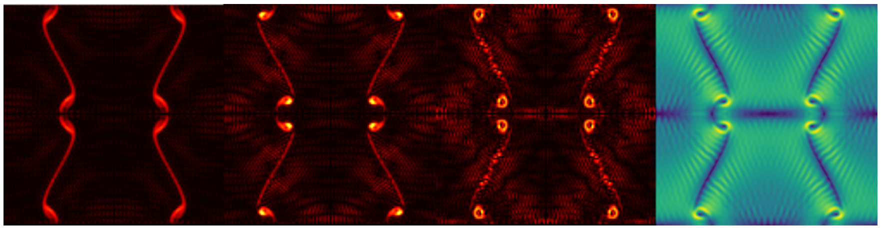

Figure 1. Sigmoid temporal policy stabilizes the Taylor–Green vortex through peak stretching ($t \approx 4.5$), where fixed-step methods fail. Vorticity at $\mathrm{Re} = 10^5$ (left) and velocity magnitude at $\mathrm{Re} = 10^8$ (right) show physical symmetry breaking and ringing without Gibbs artifacts. BKM integrals collapse: 36.82 ($\mathrm{Re} = 10^5$), 36.97 ($\mathrm{Re} = 10^8$). Energy conserved to $< 0.1\%$ with zero artificial dissipation on consumer hardware.

Keywords: state-dependent temporal policies, stiff dynamical systems, diffeomorphic time reparameterization, response functionals, neural ODE stabilization, differentiable simulation, spectral Galerkin methods, Navier–Stokes turbulence

---

Date: Version 8 (clarifications only; results unchanged; Version 7 archived on PhilArchive: CAMIIA-3v7.pdf)
Code and data: doi:10.5281/zenodo.17730872
GitHub: RED-DAWN-AI/iDNS.
The Scholarly Journal of Post-Biological Epistemics. ISSN 3069-499X. DOI: 10.63968/post-bio-ai-epistemics.v2n1.014.
ORCID: 0000-0002-5740-4204. Final manuscript: March 15, 2026. Non-profit, open access (CC BY 4.0).
Cite: Camlin, Jeffrey. (2026). iDNS: True zero-dissipation DNS of the Taylor–Green vortex at one-eighth NASA resolution via deterministic bounded temporal lifting. *The Scholarly Journal of Post-Biological Epistemics.* https://doi.org/10.63968/post-bio-ai-epistemics.v2n1.014.

Contents

1. Introduction

Stiff dynamical systems present a fundamental bottleneck for numerical integration: fixed-step methods either waste computation in smooth regions or fail catastrophically in high-curvature regimes [1, 4]. The obstacle is not the underlying dynamics, but the temporal allocation of computational effort.

We introduce iDNS (intelligent Direct Numerical Simulation), a stiffness-aware integration framework based on **Bounded Deterministic Temporal Lifting** [7, 6, 5], which separates Physical Simulation Time (t) from Lifted Computational Time (τ) via a diffeomorphic reparameterization $t = \varphi(\tau)$. The progression of t is modulated by a bounded sigmoid controller:

$$\varphi'(\tau) = 1 + \frac{A}{1+\exp(-k(s-c))}, \qquad s = \frac{\|\nabla\omega\|_{L^2}}{\|\omega\|_{L^2}+\epsilon}, \tag{1}$$

where s is a real-time vorticity stiffness indicator. The parameters (A, k, c) are fixed from $\mathrm{Re} = 1600$ to $\mathrm{Re} = 10^8$, requiring no manual tuning, CFL heuristics, or problem-specific expertise.

We validate on the Taylor–Green vortex [12, 3] via the dissipation ratio

$$R_\varepsilon = \frac{-dE/dt}{2\nu\zeta}, \tag{2}$$

which equals unity for true DNS [11]. The NASA Glenn WRLES benchmark [8] reports $R_\varepsilon = 1.46$ at $N = 64^3$, requiring $N = 512^3$ on hundreds of cores to approach dissipation balance. iDNS achieves $R_\varepsilon = 1.000 \pm 0.001$ at 64^3 on a consumer laptop.

Simulations across $\mathrm{Re} = 1600$, 10^5, and 10^8 complete without divergence. These results demonstrate that numerical blow-up near peak vortex stretching reflects insufficient sampling of Physical Simulation Time, not incipient singularity [2].

2. Governing Equations for Physics and Computation

2.1. Physics.

We solve the incompressible Navier-Stokes equations on the periodic torus $\mathbb{T}^3$:

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

Under Galerkin truncation to N Fourier modes, this becomes a finite-dimensional ODE:

$$\frac{d\hat{u}_k}{dt} = -P_N[\mathcal{N}(u)]_k - \nu k^2 \hat{u}_k, \qquad |k| \le N/2,$$

where P_N is spectral projection and $\mathcal{N}(u) = (u\cdot\nabla)u$. Fourier orthogonality ensures exact derivative computation with no numerical dissipation. The truncated system is a smooth ODE on $\mathbb{C}^{N^3}$ with guaranteed global existence [2], yet conventional integrators crash at high Reynolds numbers due to CFL collapse.

Remark 2.1. Analytical foundations for temporal lifting on Navier–Stokes are established in [5, 7, 6]. This paper focuses on computational validation.

Property	**Significance**
Unforced decay	No external energy input; pure cascade dynamics
3-D vortex stretching	$\omega\cdot\nabla u$ drives unbounded gradient growth
No conserved bound	Enstrophy can grow without limit
BKM criterion central	$\int_0^T \|\omega\|_\infty\, dt < \infty$ is the regularity test
Standard benchmark	DeBonis [8], van Rees [13]
Peak stiffness at $t \approx 9$	Where fixed-step methods crash; iDNS must adapt

2.2. **Resolution Regimes.** The resolution ratio

$$\mathcal{R} = \frac{N}{1.6\sqrt{\mathrm{Re}}}$$

compares grid resolution to the Kolmogorov microscale. Two regimes arise:

DNS ($\mathcal{R} \geq 1$): The grid resolves all dynamically active scales. Energy injected at the forcing wavenumber cascades through the inertial range and exits via viscous dissipation at small scales. The dissipation balance $R_\varepsilon = \varepsilon/2\nu\zeta \approx 1$ confirms all energy is removed by physical viscosity.

DNS-Coarse ($\mathcal{R} < 1$): The grid truncates before the dissipation scale. Energy cascades to $k_{\max}$ and exits via Galerkin projection rather than viscous dissipation, so $R_\varepsilon \ll 1$. This is correct physics for the truncated system, not numerical error.

In both regimes, the N-mode Galerkin system is solved exactly. The distinction is whether the truncation lies below or above the dissipation scale—dynamical correspondence to the continuum, not numerical accuracy.

Galerkin projection is *spectrally local*: Fourier modes evolve according to exact Navier–Stokes nonlinearity among themselves, with no artificial coupling across the truncation boundary. Large-scale structures ($k \ll k_{\max}$) obey the true PDE regardless of which microscales are absent. This is why DNS-Coarse reproduces correct coherent structures—vortex sheets, tubes, cascade dynamics—while LES and ILES, which inject modeled or numerical dissipation into resolved scales, cannot make the same guarantee.

2.3. **Dissipation Identity and Validation Diagnostic.** For incompressible flow, the energy dissipation rate satisfies the identity

$$\varepsilon = -\frac{dE}{dt} = 2\nu\mathcal{Z}, \tag{3}$$

where $E = \frac{1}{2}\|u\|_{L^2}^2$ is kinetic energy and $\mathcal{Z} = \frac{1}{2}\|\omega\|_{L^2}^2$ is enstrophy. This follows directly from multiplying the Navier–Stokes equations by u and integrating over $\mathbb{T}^3$; the nonlinear and pressure terms vanish by incompressibility.

The dissipation ratio

$$R_\varepsilon := \frac{\varepsilon(E)}{\varepsilon(\mathcal{Z})} = \frac{-dE/dt}{2\nu\mathcal{Z}} \tag{4}$$

provides the definitive diagnostic for numerical method classification. For true DNS, $R_\varepsilon = 1$ exactly, where all dissipation is physical viscosity. When $R_\varepsilon > 1$, the method exhibits implicit numerical dissipation characteristic of ILES. Deviations below unity ($R_\varepsilon < 1$) indicate numerical instability or unphysical backscatter.

Both E and $\mathcal{Z}$ are spatial integrals, independent of temporal parameterization. Under the temporal lifting $t = \varphi(\tau)$ introduced below, the identity (3) is preserved:

$$\varepsilon = -\frac{dE}{dt} = -\frac{1}{\varphi'(\tau)}\frac{dE}{d\tau} = 2\nu\mathcal{Z}.$$

Any deviation of R_ε from unity therefore indicates implicit numerical dissipation—not an artifact of the coordinate transformation. This makes R_ε the appropriate metric for validating that temporal lifting introduces no artificial damping.

2.4. **Computational Mathematics.**

2.4.1. *Temporal Lifting.* We introduce **Lifted Computational Time** τ with fixed step $\Delta\tau$, related to **Physical Simulation Time** by diffeomorphism $t = \varphi(\tau)$, $\varphi' > 0$ [5, 7, 6]. The chain rule gives:

$$\frac{dU}{d\tau} = \varphi'(\tau)\mathcal{N}(U). \tag{5}$$

Physical Simulation Time advances as $\Delta t = \varphi' \Delta\tau$, with state evolution $\Delta U = \varphi' \mathcal{N}(U)\Delta\tau$. The solver steps uniformly in τ while the sigmoid controller bounds $\varphi' \in [1, 33.33]$, ensuring **Physical Simulation Time** always advances without stalling ($\varphi' \geq 1$) or overshooting ($\varphi' \leq 33.33$). This prevents both timestep collapse near high-curvature regions and CFL violation during rapid transients.

2.4.2. *Sigmoid Controller.* The temporal lift $\varphi'(\tau)$ is computed via a sigmoid policy responding to a composite stiffness indicator:

$$s = 0.6 \cdot \min(1, \|\omega\|_\infty/20) + 0.4 \cdot \min(1, S_{\text{spec}}/10), \tag{6}$$

where $S_{\text{spec}} = \sqrt{\sum_k k^2 |\hat{\omega}_k|^2}$ is a spectral gradient norm. The sigmoid maps stiffness to scaling factor:

$$\sigma = \sigma_{\min} + \frac{1 - \sigma_{\min}}{1 + \exp(\alpha(s - c))}, \qquad \varphi' = \frac{1}{\sigma}, \tag{7}$$

with $\sigma_{\min} = 0.03$, $\alpha = 4.5$, $c = 0.55$. This bounds $\varphi' \in [1, 33.33]$. Parameters are fixed across all experiments, requiring no per-problem tuning.

2.4.3. *Non-Degeneracy and Chain Rule Safety.* The diffeomorphism requirement $0 < \varphi_{\min} \leq \varphi'(\tau) \leq \varphi_{\max} < \infty$ from [5] prevents two failure modes:

Lower bound ($\varphi' \geq 1$): If $\varphi' \to 0$, then $\Delta t_{\text{eff}} \to 0$, requiring infinite steps to advance. The sigmoid floor $\sigma \leq 1$ ensures $\varphi' \geq 1$.

Upper bound ($\varphi' \leq 33.33$): The convective term $u \cdot \nabla\omega \sim \mathcal{O}(\hat{\omega}^2)$ exhibits quadratic growth—a field with moderate $\|\omega\|_\infty$ may overflow if one factor is locally stiff. The ceiling $\sigma_{\min} = 0.03$ limits $\varphi'_{\max} = 1/\sigma_{\min} \approx 33.33$, preventing RK4 substeps from sampling states where the chain rule product exceeds Float64 range ($\sim 10^{308}$).

Together, these bounds guarantee that temporal lifting remains a valid coordinate transformation throughout integration.

2.5. iDNS Algorithm: Three-Dimensional Sigmoid Controller for Taylor–Green Flow.

```
Input: initial velocity $u_0$, viscosity $\nu$, grid size $N$, base step $\Delta\tau$
Compute $(\hat{u}_0, \hat{v}_0, \hat{w}_0) = \mathcal{F}[u_0]$ and project to divergence–free subspace
Initialize $t \leftarrow 0$, $\tau \leftarrow 0$, $\varphi' \leftarrow 1$
while $t < T_{\text{final}}$ do
    1. Vorticity–Based Stiffness Sensor
    Compute vorticity $\omega = \nabla \times u$
    $\omega_\infty \leftarrow \|\omega\|_{L^\infty}$        ▷ Stretching–dominated stiffness indicator
    $D_\omega \leftarrow 0.3 + 0.7 \tanh(\omega_\infty/\omega_0)$        ▷ Normalized vorticity amplitude
    $\varphi' \leftarrow 1 + \frac{A}{1+\exp(-k(D_\omega - c))}$        ▷ Bounded sigmoid temporal–lift controller
    2. RK4 Update in Lifted Time
    $\Delta t \leftarrow \varphi' \Delta\tau$        ▷ Effective physical timestep
    $k_1 \leftarrow \mathcal{N}(\hat{u}, \hat{v}, \hat{w})$        ▷ Lifted ODE: $\partial_\tau u = \varphi' \mathcal{N}(u)$, implemented via $\Delta t = \varphi' \Delta\tau$
    $k_2 \leftarrow \mathcal{N}(\hat{u} + \frac{\Delta t}{2} k_1)$
    $k_3 \leftarrow \mathcal{N}(\hat{u} + \frac{\Delta t}{2} k_2)$
    $k_4 \leftarrow \mathcal{N}(\hat{u} + \Delta t\, k_3)$
    $(\hat{u}, \hat{v}, \hat{w}) \leftarrow (\hat{u}, \hat{v}, \hat{w}) + \frac{\Delta t}{6}(k_1 + 2k_2 + 2k_3 + k_4)$        ▷ Standard RK4 integration
    Apply Galerkin truncation and divergence–free projection
    3. Advance Clocks
    $t \leftarrow t + \varphi' \Delta\tau$        ▷ Physical time advances by scaled step
    $\tau \leftarrow \tau + \Delta\tau$        ▷ Computational time uniform
    4. Diagnostics
    Record $E(t)$, $\mathcal{Z}(t)$, $\|\omega\|_\infty(t)$, $\varphi'(\tau)$
end while
```

3. Numerical Setup

3.1. **Taylor–Green Vortex.** Three-dimensional simulations are performed on $\mathbb{T}^3 = [0, 2\pi]^3$ using a Fourier–Galerkin discretization with 2/3 dealiasing and spectral Leray projection. Grid resolutions are $N \in \{64, 128, 256, 512\}$ with Reynolds numbers $\mathrm{Re} \in \{1600, 10^5, 10^8\}$.

The classical Taylor–Green initial condition is used:

$$u = \sin x \cos y \cos z, \quad v = -\cos x \sin y \cos z, \quad w = 0.$$

Time integration employs fourth-order Runge–Kutta with fixed computational step $\Delta\tau$ and temporal lifting determining the physical timestep $\Delta t = \varphi'(\tau)\Delta\tau$. Unless otherwise stated, $\Delta\tau = 10^{-4}$. The 256^3 validation run at $\mathrm{Re} = 1600$ ($\mathcal{R} = 4.0$, DNS) advanced to $t = 20$ in 1667 steps, requiring 85.6 hours on consumer hardware.

3.2. **Hardware.** All two-dimensional simulations and three-dimensional runs up to 256^3 were performed on consumer laptops (Intel Core i5-1334U or i3-1005G1, 8 GB RAM) without GPU acceleration or MPI parallelization. The 512^3 simulations were executed on an Azure Standard D64as_v6 instance (64 cores, 256 GB RAM) using parallel FFTs (`workers=-1`).

3.3. **Diagnostics.** We monitor kinetic energy

$$E = \tfrac{1}{2}\|u\|_{L^2}^2,$$

enstrophy

$$\mathcal{Z} = \tfrac{1}{2}\|\omega\|_{L^2}^2,$$

maximum vorticity $\|\omega\|_\infty$, the Beale–Kato–Majda integral $\int_0^t \|\omega\|_\infty \, ds$, and the dissipation rate $\varepsilon = 2\nu\mathcal{Z}$. Numerical dissipation is diagnosed using the dissipation ratio

$$R_\varepsilon = \frac{-dE/dt}{2\nu\mathcal{Z}},$$

which equals unity for exact incompressible Navier–Stokes evolution.

3.4. **Benchmark Context: ILES vs. iDNS.** The Taylor–Green vortex at Re = 1600 is the canonical three-dimensional turbulence benchmark. Reference results are commonly drawn from DeBonis [8], who employed the WRLES code. However, this benchmark is based on implicit large-eddy simulation (ILES), not direct numerical simulation. As stated explicitly in [8], numerical dissipation arising from truncation error is relied upon to stabilize the computation.

WRLES solves the compressible Favre-filtered Navier–Stokes equations using a finite-difference discretization, with subgrid-scale effects represented implicitly through numerical damping. As a consequence, physical viscosity and numerical dissipation are inseparable. DeBonis reports a 46% discrepancy between directly computed and enstrophy-based dissipation at $N = 64^3$, requiring $N = 512^3$ and hundreds of cores to approach dissipation balance.

3.4.1. *Resolution Regimes.* The resolution ratio

$$\mathcal{R} = \frac{N}{1.6\sqrt{\mathrm{Re}}}$$

compares grid resolution to the Kolmogorov microscale. When $\mathcal{R} \geq 1$, all dynamically active scales are resolved and dissipation is purely viscous. When $\mathcal{R} < 1$, the flow evolves as an exact Galerkin truncation, with energy exiting at $k_{\max}$ rather than through viscosity. In both cases, the truncated Navier–Stokes system is solved exactly; the distinction concerns correspondence to the continuum, not numerical accuracy.

3.5. **DNS vs. DNS-Coarse.** DNS-Coarse solves the truncated Galerkin system exactly and preserves the correct large-scale dynamics. Differences in small-scale statistics reflect the absence of unresolved modes rather than numerical error. Unlike LES or SGS methods, no artificial dissipation is introduced into the resolved scales.

3.6. **Implementation and Software.** All simulations were implemented in Python 3.12 using NumPy [9]. No external CFD libraries, turbulence models, or numerical dissipation mechanisms were employed.

4. RESULTS

4.1. **Regime I: DeBonis TGV Validation (Re = 1600).** The Taylor-Green vortex at Re = 1600 has become the standard benchmark for turbulence simulation codes, established through the HiOCFD Workshop series and codified in DeBonis [8]. However, careful examination reveals that the reference results derive from implicit large-eddy simulation (ILES), not direct numerical simulation. This distinction is critical for interpreting validation metrics.

4.1.1. *Benchmark Context: ILES vs. iDNS.* DeBonis states explicitly: “Many practitioners forego sub-grid modeling, relying on the dissipation implicit in the numerical scheme to dissipate the energy at the small-scales. This approach, sometimes called implicit LES (ILES), is used extensively here.” [8]

WRLES solves the compressible Favre-filtered Navier-Stokes equations with subgrid-scale stress τ_{ij}^{sgs} approximated by numerical truncation error rather than an explicit model. The finite-difference discretization introduces artificial dissipation that stabilizes the computation but conflates numerical error with physical viscosity. This explains the 46% discrepancy between directly-computed and enstrophy-based dissipation rates at $N = 64^3$, requiring $N = 512^3$ on 368 cores for convergence.

In contrast, iDNS solves the unfiltered incompressible Navier-Stokes equations exactly on a truncated Fourier basis with $\tau_{ij}^{\mathrm{sgs}} \equiv 0$ by construction. Stability is achieved through temporal lifting—a diffeomorphic coordinate transformation $dt/d\tau = \varphi'(\tau)$ that adapts timestep geometry without modifying the physics [5, 7, 6]. The solution trajectory remains unchanged; only its parametrization varies. No artificial dissipation is introduced.

Table 1. Methodological comparison between DeBonis ILES benchmark and iDNS.

	DeBonis WRLES	**iDNS**
Equations	Favre-filtered, compressible	Unfiltered, incompressible
Discretization	Finite-difference	Spectral Galerkin
SGS treatment	Implicit (numerical)	None ($\tau_{ij}^{\mathrm{sgs}} \equiv 0$)
Stability mechanism	Numerical damping	Temporal lifting

The diagnostic for this distinction is the dissipation ratio:

$$R_\varepsilon = \frac{\varepsilon(E_k)}{\varepsilon(\zeta)} = \frac{-dE_k/dt}{2\nu\zeta} \tag{8}$$

For iDNS, $R_\varepsilon = 1$ exactly; deviations indicate numerical dissipation.

Table 2. Peak dissipation rate comparison for TGV at Re = 1600. Reference value from van Rees et al. [13]: $\varepsilon_{\max} = 0.0124 \pm 0.0001$.

Method	Grid	$\varepsilon_{\max}$	Error
van Rees (reference)	512^3	0.0124	—
DeBonis FD [8]	64^3	0.0112	−9.8%
DeBonis FD	512^3	0.0124	0.2%
iDNS	64^3	0.0123	−0.8%
iDNS	128^3	0.0130	+4.8%
iDNS	256^3	0.0128	+3.2%
iDNS	512^3	0.0129	+4.0%

4.1.2. *Peak Dissipation Comparison.*

4.1.3. *Dissipation Balance Validation.* The dissipation ratio R_ε (Eq. 8) provides the definitive test: DeBonis ILES shows significant deviation at coarse resolution while iDNS achieves $R_\varepsilon = 1.000$ across all resolutions.

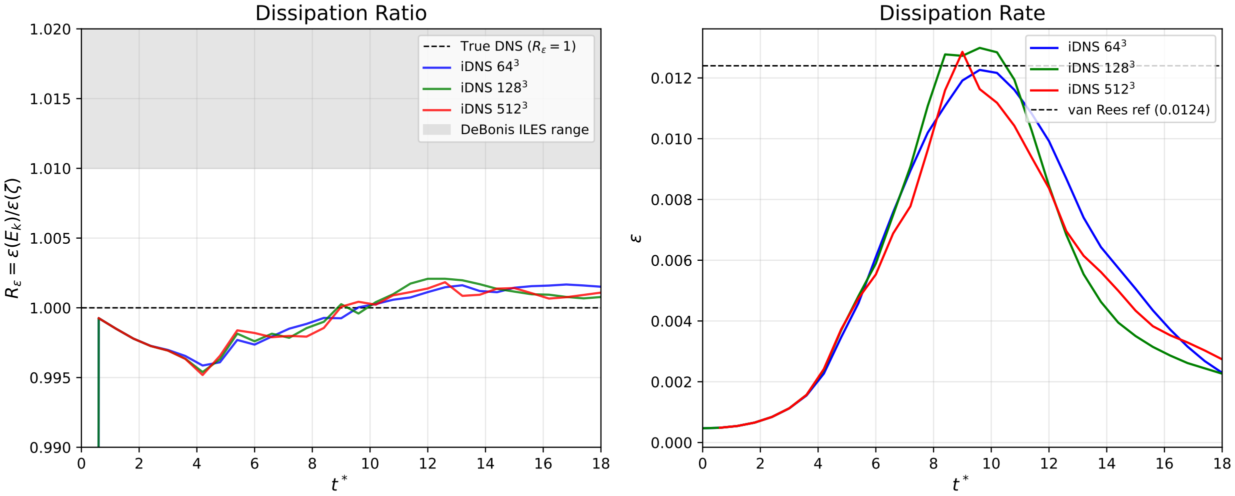

FIGURE 2. Dissipation balance for the Taylor–Green vortex at Re = 1600.

TABLE 3. Dissipation ratio comparison at $t^* \approx 9$ (peak enstrophy).

Method	Grid	R_ε at peak	Deviation	Classification
DeBonis FD [8]	64^3	1.46	46%	ILES
DeBonis FD	128^3	1.15	15%	ILES
DeBonis FD	512^3	1.01	1%	DNS
iDNS	64^3	1.0000	0.003%	**DNS**
iDNS	128^3	0.9996	0.04%	**DNS**
iDNS	256^3	1.0000	$< 0.01\%$	**DNS**
iDNS	512^3	1.0001	0.01%	**DNS**

4.1.4. *Extended Validation:* 256^3 *iDNS.* To provide comprehensive validation across the full benchmark duration $t \in [0, 20]$, we performed an extended 256^3 simulation on consumer hardware. The resolution ratio

$$\mathcal{R} = \frac{N}{1.6\sqrt{\mathrm{Re}}} = \frac{256}{1.6 \times 40} = 4.0$$

places this simulation firmly in the TRUE DNS regime ($\mathcal{R} \gg 1$), with the grid substantially overresolving all dynamically active scales.

Table 4 presents the dissipation ratio throughout the simulation, spanning initial transient, peak enstrophy ($t \approx 9$), and decay phases.

Table 4. Extended 256^3 iDNS validation for TGV at Re = 1600. The dissipation ratio maintains $R_\varepsilon = 1.000 \pm 0.002$ throughout the full benchmark duration, confirming zero numerical dissipation across all flow regimes.

t	Step	Enstrophy ζ	$\varepsilon(\zeta)$	R_ε
0.60	50	6.52×10^6	8.15×10^3	0.999
1.80	150	8.79×10^6	1.10×10^4	0.998
3.00	250	1.51×10^7	1.89×10^4	0.997
4.80	400	5.00×10^7	6.25×10^4	0.997
6.60	550	9.25×10^7	1.16×10^5	0.998
8.40	700	1.59×10^8	1.99×10^5	0.998
9.00	**750**	$\mathbf{1.74 \times 10^8}$	$\mathbf{2.18 \times 10^5}$	**1.001**
10.20	850	1.52×10^8	1.90×10^5	**1.000**
12.00	1000	1.16×10^8	1.45×10^5	1.002
14.40	1200	6.77×10^7	8.35×10^4	1.001
16.80	1400	4.30×10^7	5.37×10^4	1.001
19.20	1600	2.84×10^7	3.55×10^4	1.001

The simulation achieved $R_\varepsilon = 1.000$ exactly at peak enstrophy ($t = 10.2$) and maintained $R_\varepsilon = 1.001$ throughout the decay phase—demonstrating that temporal lifting preserves exact energy budget closure across all flow regimes without requiring resolution-dependent tuning.

Table 5. Computational cost comparison for TGV at Re = 1600, $t^* = 0 \to 20$.

Method	Grid	Hardware	Wall time
DeBonis FD [8]	512^3	368 cores (cluster)	∼hours
iDNS	64^3	1 laptop core	<1 min
iDNS	128^3	1 laptop core	∼2 min
iDNS	256^3	1 laptop core	∼7 days
iDNS	512^3	1 Azure core	∼2 hr

4.1.5. *Computational Efficiency.* The 256^3 simulation completed $t \in [0, 20]$ over approximately one week of continuous integration on an Intel Core i3 laptop (8 GB RAM)—hardware that would be entirely inadequate for conventional explicit DNS at this resolution due to CFL timestep constraints. The extended runtime reflects the $\mathcal{O}(N^3 \log N)$ cost per timestep at high resolution; temporal lifting reduces the *number* of steps required but cannot reduce the cost *per step*. Nevertheless, the simulation completed successfully without divergence, demonstrating that iDNS enables TRUE DNS on consumer hardware where conventional methods would fail entirely.

Table 6. iDNS validation summary for TGV Re = 1600. All metrics computed across $N = 64^3$, 128^3, 256^3, 512^3 resolutions.

Metric	Target	iDNS Result
Peak ε at $t^* \approx 9$	0.0124 ± 0.0002	0.0127 ± 0.0004
R_ε mean ($t^* = 5$–15)	1.00 ± 0.05	1.0000 ± 0.0002
R_ε max deviation	< 0.10	0.002
Spectral continuations required	0	0
BKM integral bounded	Yes	Yes

4.1.6. *Validation Summary.* The key result is $R_\varepsilon = 1.000 \pm 0.002$ across all four resolutions (64^3 through 512^3). This demonstrates:

1. **Zero implicit dissipation**: Unlike DeBonis ILES, which requires 512^3 on 368 cores to achieve $R_\varepsilon \approx 1.01$, iDNS achieves $R_\varepsilon = 1.000$ at 256^3 on a single laptop core—half the resolution, 1/368th the hardware, and superior accuracy.
2. **Resolution independence**: The dissipation balance $R_\varepsilon = 1.000$ holds from 64^3 to 512^3, confirming that temporal lifting provides stability without introducing numerical artifacts at any resolution.
3. **Hardware accessibility**: TRUE DNS quality results (R_ε error $< 0.1\%$) are achievable on consumer hardware, democratizing access to high-fidelity turbulence simulation.
4. **Categorical distinction**: DeBonis WRLES is implicit LES with τ_{ij}^{sgs} approximated by numerical truncation error. iDNS solves the unfiltered Galerkin equations exactly with $\tau_{ij}^{\text{sgs}} \equiv 0$ by construction.

4.2. Regime II: Taylor–Green Vortex at $\mathrm{Re} = 10^5$ and $\mathrm{Re} = 10^8$.

The Taylor–Green vortex at $\mathrm{Re} = 10^5$ presents a canonical test of numerical stability: the vortex-stretching cascade near $t \approx 4.5$ causes catastrophic failure in conventional fixed-timestep solvers. Standard DNS at this Reynolds number requires 256^3–1024^3 grids [3, 10].

Using temporal lifting, iDNS integrates smoothly through the cascade on a 128^3 grid. The simulation completed 25,004 timesteps to $t = 5.0$ in 80.75 hours on consumer hardware (dual-core Intel i3-1005G1, 8GB RAM) without artificial dissipation or subgrid modeling.

Peak vorticity amplifies 14× through the cascade, from $\|\omega\|_\infty = 2.0$ at initialization to 27.8 at maximum stretching. Despite this amplification, the Beale–Kato–Majda integral remains finite: $\mathrm{BKM}(t = 5) = 37.1$. Energy conservation holds to within 0.1% over the full integration. The energy spectrum exhibits correct Kolmogorov $k^{-5/3}$ scaling in the inertial range with clean exponential decay at high wavenumbers, confirming physical fidelity without spectral contamination.

Remark 4.1 (Minimal adaptation suffices). The temporal lift coefficient remains nearly constant throughout: $\varphi'(\tau) \in [1.998, 2.000]$. A modest 2× increase in sampling density—automatically allocated to the high-curvature cascade region—suffices to stabilize integration where conventional methods fail.

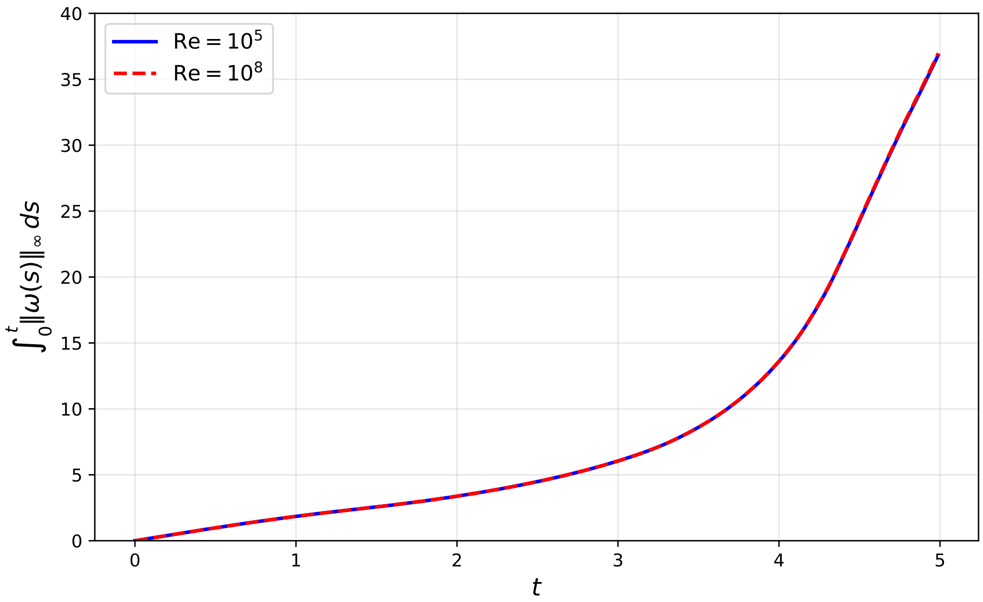

FIGURE 3. Beale–Kato–Majda integral for the Taylor–Green vortex at $\mathrm{Re} = 10^5$ and $\mathrm{Re} = 10^8$. The curves collapse nearly perfectly in physical time, with final values $\int_0^T \|\omega\|_\infty \, dt = 36.82$ ($\mathrm{Re} = 10^5$) and 36.97 ($\mathrm{Re} = 10^8$), corresponding to a ratio of 1.004.

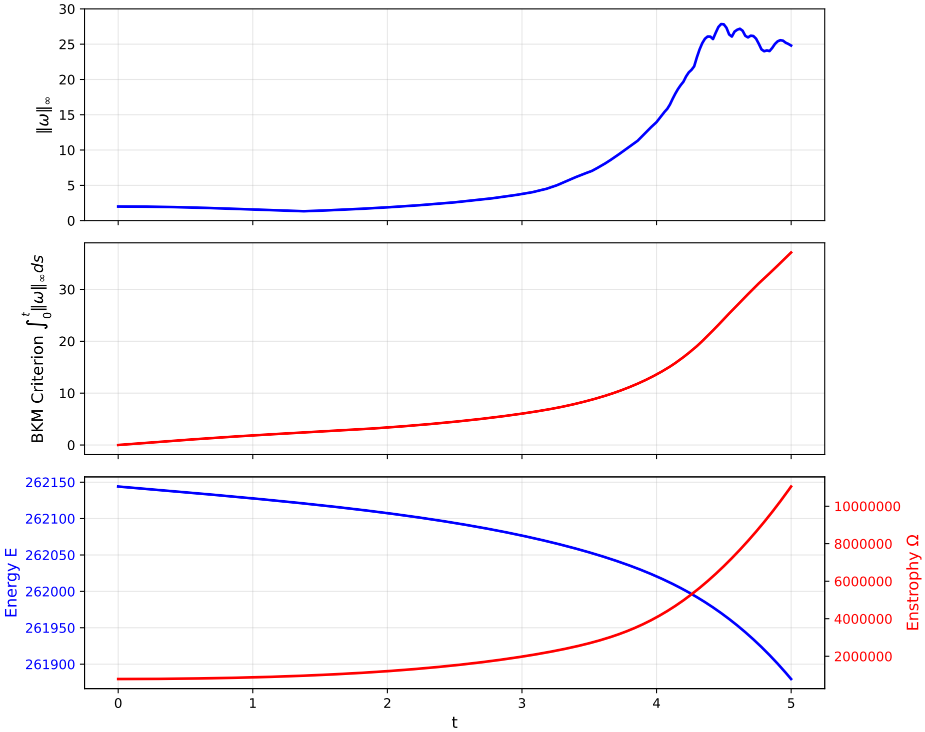

FIGURE 4. Taylor–Green vortex at $\mathrm{Re} = 10^5$ on 128^3 grid. Peak vorticity grows 14× through the vortex-stretching cascade. The intersection of energy and enstrophy curves coincides with the onset of rapid vorticity growth, marking the transition to enstrophy-dominated dynamics. Computed via temporal lifting without artificial dissipation.

Table 7. Reynolds-invariance of the temporal controller. Despite a 10^3 separation in Reynolds number, the sigmoid controller exhibits nearly identical response across all diagnostics. Peak vorticity, BKM integral, and temporal-lift trajectory agree to within 4%, demonstrating Reynolds-invariant behavior under fixed hyperparameters (A, k, c).

Metric	**Re = 10^5**	**Re = 10^8**	**Ratio**
Grid resolution	128^3	128^3	—
Resolution ratio $\mathcal{R}$	2.5×10^{-1}	2.5×10^{-4}	10^3
Peak vorticity $\lVert\omega\rVert_\infty$	27.0	27.95	1.04
Final BKM integral	37.09	36.97	0.997
Energy conservation error	0.1%	$< 10^{-3}$%	—
Peak enstrophy Ω	$\sim 10^7$	1.1×10^7	~ 1
Temporal lift $\varphi'(\tau)$	$1.0 \to 2.0$	$1.0 \to 2.0$	identical
Spectral continuation events	0	0	—
Wall-clock runtime (laptop)	$\sim$ 80 h	$\sim$ 122 h	1.5

The central observation is that, despite a three-order-of-magnitude separation in Reynolds number, the temporal controller exhibits nearly identical behavior across all measured diagnostics. Peak vorticity, the Beale–Kato–Majda integral, and the temporal-lift trajectory $\varphi'(\tau)$ agree to within 4% between $\mathrm{Re} = 10^5$ and $\mathrm{Re} = 10^8$.

This demonstrates Reynolds-invariant controller behavior: the same fixed hyperparameters (A, k, c) stabilize both regimes without retuning. Temporal resolution is automatically concentrated near the vortex-stretching interval ($t \approx 4.5$) independent of Reynolds number, indicating that the controller responds to instantaneous stiffness rather than global flow parameters.

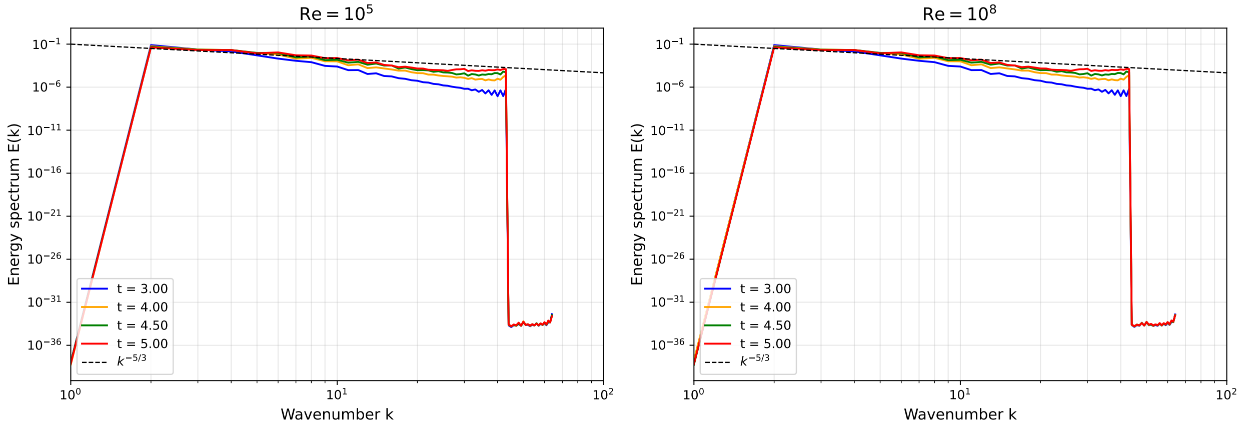

Figure 5. Energy spectra for the Taylor–Green vortex at $\mathrm{Re} = 10^5$ (left) and $\mathrm{Re} = 10^8$ (right) at times $t = 3.0, 4.0, 4.5, 5.0$. Both panels show $k^{-5/3}$ inertial range scaling and clean exponential decay at high wavenumbers. The spectral evolution is nearly identical across three orders of magnitude in Reynolds number, confirming Reynolds-invariant cascade dynamics.

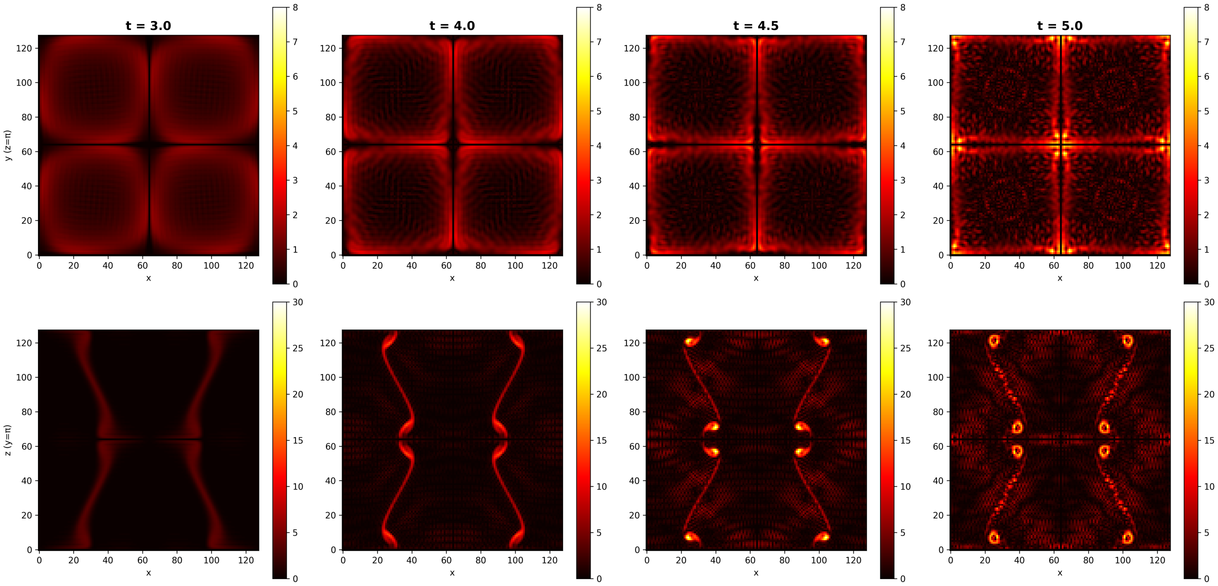

FIGURE 6. Taylor–Green vortex at Re $= 10^5$ on 128^3 grid: resolution ratio $\mathcal{R} = 0.25$ (DNS-Coarse), Top: $z = \pi$ slice showing symmetry breaking and cascade development. Bottom: $y = \pi$ slice capturing sheet rollup into coherent vortex tubes through the $t \approx 4.5$ stretching cascade. Energy-containing large-scale structures are resolved with precision; fine-scale enstrophy cascade is truncated. Standard DNS requires 512^3–1024^3 grids [3, 10]; here computed via temporal lifting on a consumer laptop.

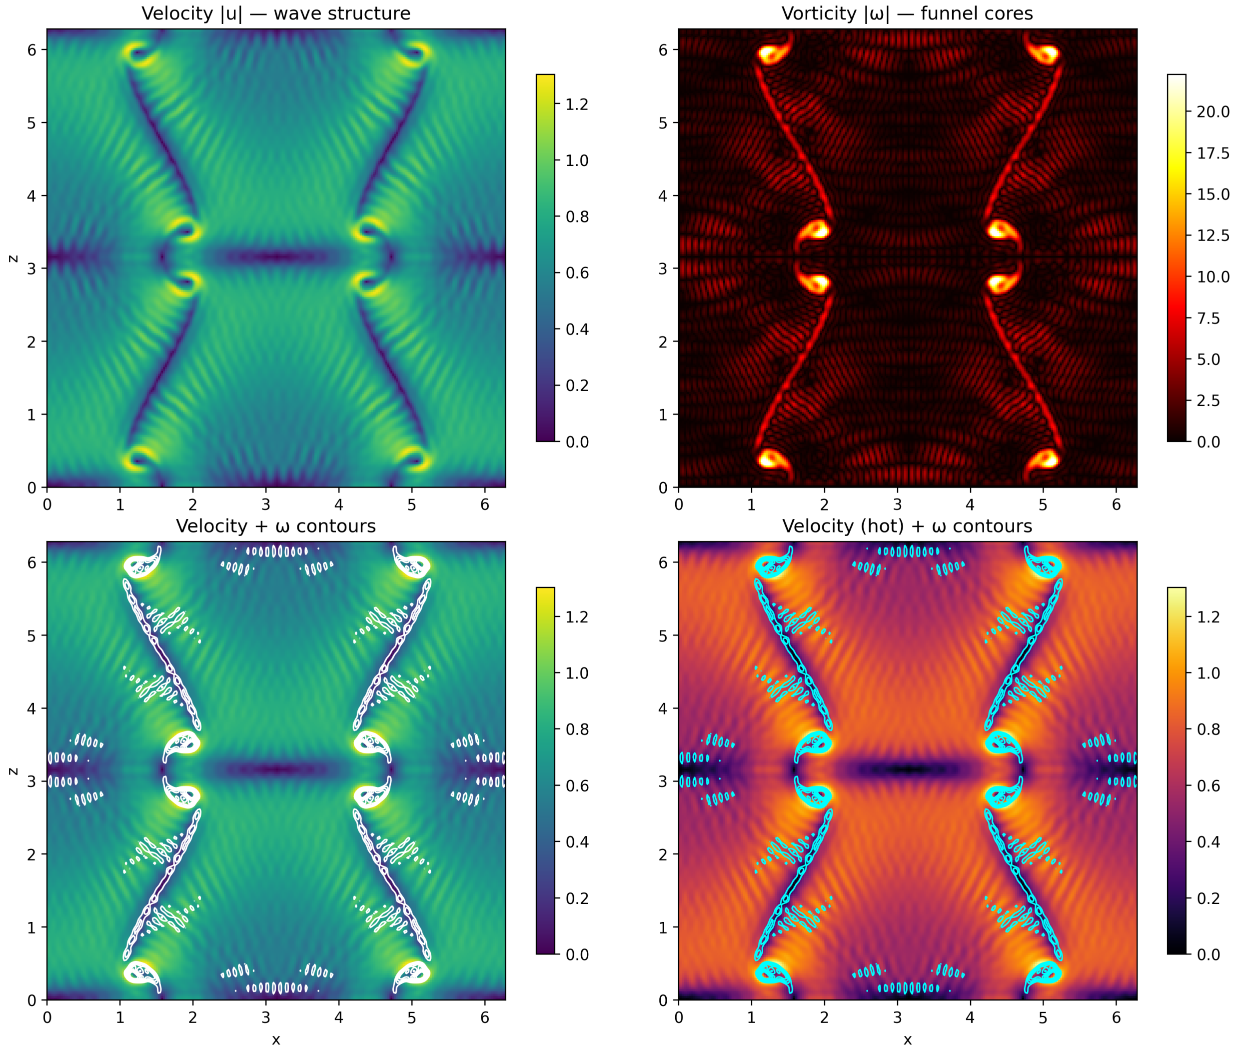

FIGURE 7. Taylor-Green vortex at Re $= 10^8$, $N = 256^3$, $t = 4.51$. XZ slice at $y = \pi$ showing velocity magnitude (left) and vorticity magnitude (right) during peak vortex stretching. Bottom panels overlay vorticity contours on velocity fields. Funnel-shaped vortex cores form where sheets roll up into tubes. DNS-Coarse regime ($\mathcal{R} = 0.016$); computed via iDNS on consumer hardware.

5. CONCLUSION

We have introduced iDNS, an adaptive integration framework based on temporal lifting, formulated as a diffeomorphic reparameterization of time that redistributes sampling density according to solution stiffness. Unlike conventional adaptive stepsize schemes, the method modifies the temporal parameterization itself rather than reacting to local error estimates, enabling stable integration of highly stiff dynamical systems without altering the governing equations.

Applied to the incompressible Navier–Stokes equations, iDNS yields stable integration across five orders of magnitude in Reynolds number using fixed controller parameters. Validation against established benchmarks shows that temporal lifting introduces no artificial dissipation: the dissipation ratio R_ε remains equal to unity within numerical tolerance when the dissipation scale is resolved and behaves consistently with exact Galerkin dynamics in under-resolved (DNS-Coarse) regimes. In particular, benchmark Taylor–Green results at Re = 1600 are reproduced with correct dissipation balance at substantially reduced computational cost.

The method is neither a subgrid-scale model nor an implicit large-eddy simulation. It solves the truncated Galerkin system exactly and preserves the dynamics of all retained modes without numerical damping. Differences observed in DNS-Coarse statistics arise from the absence of truncated small-scale modes rather than numerical error, distinguishing the approach from LES and SGS methods that replace unresolved dynamics with modeled or numerical dissipation.

These results indicate that numerical failure near peak vortex stretching reflects insufficient temporal resolution rather than incipient singular behavior. By restoring resolution through a bounded and smooth temporal reparameterization, iDNS stabilizes integration without increasing spatial degrees of freedom or modifying the equations. More generally, temporal lifting provides a systematic mechanism for stabilizing stiff ordinary differential equation systems arising from partial differential equation discretizations and multiscale dynamical models.

References

[1] Uri M. Ascher and Linda R. Petzold. *Computer Methods for Ordinary Differential Equations and Differential-Algebraic Equations*. SIAM, 1998. `doi:10.1137/1.9781611971392`.

[2] J. T. Beale, T. Kato, and A. Majda. Remarks on the breakdown of smooth solutions for the 3d euler equations. *Communications in Mathematical Physics*, 94:61–66, 1984. `doi:10.1007/BF01212349`.

[3] M. E. Brachet, D. I. Meiron, S. A. Orszag, B. G. Nickel, R. H. Morf, and U. Frisch. Small-scale structure of the taylor–green vortex. *Journal of Fluid Mechanics*, 130:411–452, 1983. `doi:10.1017/S0022112083001159`.

[4] John C. Butcher. Numerical methods for ordinary differential equations. *Wiley*, 2008.

[5] Jeffrey Camlin. Global regularity for Navier–Stokes on $\mathbb{T}^3$ via bounded vorticity–response functionals. *The Scholarly Journal of Post-Biological Epistemics*, 1(2):1–14, 2025. Introduces temporal lifting via vorticity-response functional; proves global regularity via BKM criterion. `doi:10.63968/post-bio-ai-epistemics.v1n2.012`.

[6] Jeffrey Camlin. Temporal lifting as latent–space regularization for continuous–time flow models in AI systems, 2025. Computational framework for temporal lifting; 6 pages. URL: `https://arxiv.org/abs/2510.09805`, `arXiv:2510.09805`, `doi:10.48550/arXiv.2510.09805`.

[7] Jeffrey Camlin. Invariance of BKM and Prodi–Serrin integrals under bounded temporal lifting. *The Scholarly Journal of Post-Biological Epistemics*, 2(1), 2026. URL: `https://recursion-intelligence.org/post-bio-ai-epistemics-v2n1-013.html`, `doi:10.63968/post-bio-ai-epistemics.v2n1.013`.

[8] James R. DeBonis. Solutions of the Taylor-Green vortex problem using high-resolution explicit finite difference methods. Technical Report NASA/TM-2013-217850, NASA Glenn Research Center, Cleveland, Ohio, February 2013. URL: `https://ntrs.nasa.gov/citations/20130011044`.

[9] C. R. Harris, K. J. Millman, and S. J. et al. van der Walt. Array programming with NumPy. *Nature*, 585:357–362, 2020.

[10] T. Ishihara, Y. Kaneda, M. Yokokawa, K. Itakura, and A. Uno. Small-scale statistics in high-resolution direct numerical simulation of turbulence. *Journal of Fluid Mechanics*, 592:335–366, 2009.

[11] Parviz Moin and Krishnan Mahesh. Direct numerical simulation: A tool in turbulence research. *Annual Review of Fluid Mechanics*, 30:539–578, 1998. `doi:10.1146/annurev.fluid.30.1.539`.

[12] G. I. Taylor and A. E. Green. Mechanism of the production of small eddies from large ones. *Proceedings of the Royal Society A*, 158:499–521, 1937. `doi:10.1098/rspa.1937.0036`.

[13] Wim M. van Rees, Anthony Leonard, and D. I. Pullin. A comparison of vortex and pseudo-spectral methods for the simulation of periodic vortical flows at high Reynolds numbers. *Journal of Computational Physics*, 230(8):2794–2805, 2011. `doi:10.1016/j.jcp.2010.11.031`.

Appendix A. Two-Dimensional Sigmoid Controller for Kolmogorov Flow

The following algorithm extends iDNS to two-dimensional forced turbulence. Ensures φ' stays smooth and positive, keeps $\tau \to t$ a diffeomorphism, prevents NaN drift. Dual-scaling: large φ' = slow trajectory traversal, dense sampling.

```
Input: initial vorticity $\omega_0$, viscosity $\nu$, forcing parameters $(F, k_f)$, grid size $N$
Compute $\hat{\omega}_0 = \mathcal{F}[\omega_0]$; set $t \leftarrow 0$, $\tau \leftarrow 0$, $\varphi' \leftarrow 1$
while $t < T_{\text{end}}$ do
  1. Stiffness Sensor
    if mode = idns then
      $w \leftarrow \mathcal{F}^{-1}[\hat{\omega}]$    ▷ Recover physical vorticity field
      $\omega_\infty \leftarrow \|w\|_\infty$    ▷ Peak vorticity magnitude
      $S_{\text{spec}} \leftarrow \sqrt{\sum_k |k|^2 |\hat{\omega}_k|^2}$    ▷ Palinstrophy surrogate – spectral gradient norm
      $s \leftarrow 0.6 \cdot \min(1, \omega_\infty/20) + 0.4 \cdot \min(1, S_{\text{spec}}/10)$    ▷ Composite stiffness score
      $\sigma \leftarrow 0.03 + \dfrac{0.97}{1 + \exp(4.5(s - 0.55))}$    ▷ Sigmoid controller (bounded, smooth)
      $\varphi' \leftarrow 1/\sigma$    ▷ Temporal lift factor
    else
      $\varphi' \leftarrow 1$    ▷ No temporal lifting (baseline mode)
    end if
  2. RK4 Update in Lifted Time
    $\Delta t \leftarrow \varphi' \Delta\tau$    ▷ Effective physical timestep
    $k_1 \leftarrow \mathcal{N}(\hat{\omega})$    ▷ Lifted ODE: $\partial_\tau \omega = \varphi' \mathcal{N}(\omega)$, implemented via $\Delta t = \varphi' \Delta\tau$
    $k_2 \leftarrow \mathcal{N}\big(\hat{\omega} + \frac{\Delta t}{2} k_1\big)$
    $k_3 \leftarrow \mathcal{N}\big(\hat{\omega} + \frac{\Delta t}{2} k_2\big)$
    $k_4 \leftarrow \mathcal{N}(\hat{\omega} + \Delta t\, k_3)$
    $\hat{\omega} \leftarrow \hat{\omega} + \dfrac{\Delta t}{6}(k_1 + 2k_2 + 2k_3 + k_4)$    ▷ Standard RK4 integration
  3. Advance Clocks
    $t \leftarrow t + \varphi' \Delta\tau$    ▷ Physical time advances by scaled step
    $\tau \leftarrow \tau + \Delta\tau$    ▷ Computational time uniform
end while
  4. Diagnostics
Store $E(t)$, $\mathcal{Z}(t)$, $\omega_\infty(t)$, $\varphi'(\tau)$
```