

Spectral Vorticity Bridge for Turbulent Flows

A mode-level bridge between vortices, energy flux, and cascade direction

A build-once spectral operator view for turbulence diagnostics in incompressible flow

Tamas Nagy, Ph.D.

tnagyphd@gmail.com

Draft

Executive Summary (Non-Technical)

Turbulence is usually described as chaotic and hard to predict in detail, but in practice engineers and scientists rarely need exact trajectories. They need reliable statements about **where energy goes across scales, when flow becomes unstable, and which coherent vortex structures matter for transport and dissipation.**

The central gap is that vortex-language and spectral-language are often used separately. Vortex diagnostics are geometric and local; spectral diagnostics are modal and global. This paper opens one unified object: a **Spectral Vorticity Bridge** that links both descriptions through explicit mode-level identities and flux bounds.

The intended contribution is not a claim to solve 3D turbulence in full generality. The claim is narrower and operationally stronger: for incompressible flow classes of interest, one reusable spectral operator representation can support triad-level transfer signs, band-limited flux control, and interpretable vortex-coherence diagnostics.

If this bridge works, it gives a build-once, query-many-times turbulence workflow similar in spirit to existing spectral programs in this repository: construct the right operator once, then query multiple downstream quantities (transfer, cascade direction, coherence markers, and stress-relevant indicators) without re-deriving each quantity from scratch.

This paper does **not** claim a full Navier-Stokes regularity proof, and it does **not** claim to replace DNS/LES. It aims to provide a rigorous spectral control layer that can improve theorem-driven understanding and practical turbulence diagnostics.

Abstract

We introduce a Spectral Vorticity Bridge for incompressible turbulent flow, designed to connect vortex-centered and mode-centered descriptions through a shared operator framework. The bridge has three layers. Layer 1 establishes triad-level transfer identities and sign-structured interaction lemmas in Fourier space. Layer 2 derives band-limited flux bounds that separate direct and inverse transfer regimes under explicit assumptions. Layer 3 defines vortex-coherence spectral diagnostics that can be queried from the same operator object and linked to intermittency proxies.

The goal is a reusable law-level representation: once the spectral operator and projection surfaces are fixed, transfer and cascade observables become repeated linear-algebra and bounded nonlinear

queries rather than isolated one-off derivations. We position the framework as a rigorous intermediary between pure geometric vortex narratives and purely statistical spectral scaling narratives.

The paper is a foundational step, not a millennium claim: we do not prove full 3D global regularity. Instead, we provide a theorem-oriented bridge architecture intended to support future regularity criteria, intermittency analysis, and physically interpretable turbulence diagnostics across simulation and reduced-order settings.

1. Why This Direction

1.1 The Object Mismatch

Fluid dynamics practice switches between two powerful but weakly coupled stories:

- Vorticity geometry: tubes, sheets, stretching, reconnection.
- Spectral transfer: triads, flux, cascade, scaling bands.

The mismatch is not conceptual only. It blocks theorem transport: statements proven in one language are often hard to reuse in the other.

1.2 The Bridge Thesis

The thesis of this paper is:

A single spectral-operator representation can serve as a bridge object between vorticity geometry and cascade flux diagnostics, with reusable theorem-level consequences.

1.3 Non-Claims

- No claim of full 3D Navier-Stokes regularity.
- No claim that intermittency is fully solved.
- No claim that LES/DNS can be replaced.

The objective is a mathematically explicit bridge layer that can be grown theorem by theorem.

2. Program Structure (L1-L3)

L1. Triad Transfer Core (foundational)

- Define mode interaction form for incompressible flow.
- Prove sign-structured transfer lemmas for admissible triads.
- Identify forbidden/neutral triad families under structural constraints.

L2. Flux Band Control (intermediate)

- Define shell/band flux operator $\Pi(k)$.
- Derive explicit bounds under regularity and forcing assumptions.
- Classify conditions for direct vs inverse dominance.

L3. Vortex-Coherence to Intermittency Proxy (advanced)

- Define spectral coherence markers linked to vorticity concentration.
 - Prove bridge lemmas from coherence marker to burst-risk proxy.
 - Validate interpretability against controlled numerical regimes.
-

3. First Theorem Candidates

1. **Triad Compatibility Lemma:** admissible incompressible triads satisfy constrained transfer geometry.
2. **Signed Transfer Envelope:** transfer contribution per triad admits upper and lower envelopes by mode amplitudes and angle factors.
3. **Band Flux Bound:** shell flux $\Pi(k)$ is bounded by a computable spectral norm product over neighboring bands.
4. **Coherence-Flux Link:** elevated spectral coherence marker implies elevated short-horizon flux burst indicator under bounded forcing.

These are intended as the first Lean-friendly theorem chain, not the final capstone.

4. Implementation and Verification Plan

- **Math object:** spectral operator + projection family + flux functional.
 - **Formal layer:** Lean theorem chain for L1 and selected L2 statements.
 - **Numerical layer:** reproducible synthetic and benchmark flow tests to evaluate bridge diagnostics.
 - **Paper layer:** theorem statements in main text, proof mechanics and extended checks in appendix.
-

5. Immediate Next Steps

1. Fix baseline function spaces and boundary assumptions for the first theorem chain.
2. Freeze L1 theorem list (6-10 lemmas/theorems) as a bounded gym package.
3. Build one minimal numerical harness for triad transfer sanity checks.
4. Draft the first bridge figure: vortex view vs spectral view vs bridge object.

6. Benchmark and KPI Protocol

To make the bridge usable beyond theorem statements, the paper adopts a fixed benchmark protocol.

Benchmark families

- Controlled synthetic incompressible fields with known triad signatures.

- One fixed DNS/LES family for external realism checks.

Mandatory KPI table

- burst-event precision/recall,
- false alarm rate,
- early-warning lead time,
- per-query runtime cost,
- rerun stability.

Baseline comparison rule

Every experiment must include at least one classical baseline indicator and report:

- quality delta,
- cost delta,
- and uncertainty/variance of the estimate.

Initial KPI smoke result (synthetic)

We include a first bounded smoke benchmark to verify that the bridge scoring surface is operational and measurable.

- script: `examples/spectral_vorticity_turbulence_benchmark.py`
- artifact: `output/spectral_vorticity_turbulence/benchmark_smoke.json`
- setup: 24 synthetic incompressible latent-stress trials, series length 420

Mean KPI deltas (bridge vs baseline):

KPI	Baseline	Bridge	Delta
Precision	0.502	0.772	+0.270
Recall	0.107	0.168	+0.060
False alarm rate	0.063	0.028	-0.035
Lead time	0.849	0.898	+0.049
Query time (ms)	1.136	0.033	-1.102

This smoke result is not the main scientific claim. It is a reproducible operational sanity check confirming that the bridge framework can be measured by the declared KPI protocol and can show bounded quality/cost gains in a controlled setting.

External adapter result (DNS/LES-style path)

To avoid overfitting the protocol to internal synthetic generation, we also expose an external CSV adapter path:

- `python3 examples/spectral_vorticity_turbulence_benchmark.py --input-csv <path>`
- required input columns: `energy`, `coeff`, `residual`, `coherence`
- optional: `burst` (otherwise proxy labels are generated)

Adapter smoke run used here:

- `python3 examples/spectral_vorticity_turbulence_benchmark.py --write-dns-proxy-sample`
- `python3 examples/spectral_vorticity_turbulence_benchmark.py --input-csv output/spectral_vorticity_`

Result artifact:

- `output/spectral_vorticity_turbulence/benchmark_external.json`

Observed adapter deltas (single-run smoke):

KPI	Delta (bridge-baseline)
Precision gain	+0.206
Recall gain	+0.054
False alarm reduction	+0.024
Lead-time gain	+0.208
Query-time delta (ms)	-12.895

Real external file ingest run:

- command: `python3 examples/spectral_vorticity_turbulence_benchmark.py --input-csv evaluation/data`
- note: this file is external and non-proxy, but not turbulence-native; the adapter auto-infers bounded energy/coeff/residual/coherence channels from variable numeric columns.
- artifact copy: `output/spectral_vorticity_turbulence/benchmark_external_ml2_s1.json`

Observed deltas (single-run):

KPI	Delta (bridge-baseline)
Precision gain	+0.0034
Recall gain	+0.0027
False alarm reduction	+0.0003
Lead-time gain	+0.0137
Query-time delta (ms)	-0.0218

This still does not constitute external physical DNS/LES validation, but it confirms that the benchmark pipeline accepts genuine external CSV files beyond the synthetic proxy path while preserving the KPI contract.

Real DNS run (Lee-Moser channel DNS profile):

- provenance: Lee, M. and Moser, R. D., *J. Fluid Mech.* 774 (2015), 395-415
- source file: `output/spectral_vorticity_turbulence/LM_Channel_5200_mean_prof.dat`
- converted ingest: `output/spectral_vorticity_turbulence/LM_Channel_5200_ingest.csv`
- artifact: `output/spectral_vorticity_turbulence/benchmark_external_lm5200.json`

Observed deltas (single-run):

KPI	Delta (bridge-baseline)
Precision gain	+0.0000
Recall gain	+0.0000

KPI	Delta (bridge-baseline)
False alarm reduction	+0.0000
Lead-time gain	+0.0000
Query-time delta (ms)	-0.0627

Interpretation: this is a valid real DNS ingest and runtime confirmation, but the mean-profile dataset is not event-rich enough for burst KPI discrimination. Field-resolved DNS/LES time-series inputs are still needed for strong practical quality claims on detection metrics.

7. Build-Once Query-Many Pipeline

The target practical interface is:

1. Build spectral operator once on the bounded domain.
2. Query repeatedly for:
 - triad stress map,
 - shell flux risk score,
 - coherence-to-burst alert.
3. Export machine-readable diagnostics for downstream CFD monitoring.

This pipeline framing is central to practical impact: build once, query many.

8. Working Positioning

- **Scientific value:** a rigorous bridge between two major turbulence languages.
 - **Practical value:** reusable diagnostics from one operator object.
 - **Publication value:** a clear foundational paper that can spawn follow-up regularity/intermittency papers.
-

9. Practical Question, Value Surface, and Scope Boundaries

9.1 The practical question

The practical question is not “solve turbulence exactly.” The practical question is:

- when does the flow enter a higher-risk regime,
- where does that regime emerge in scale/region space,
- and how early and reliably can we alert on that transition.

This makes the bridge a risk-diagnostics layer rather than a full replacement for high-fidelity flow solvers.

9.2 Why this matters even in non-event periods

A useful monitoring layer must add value both when turbulence risk is present and when it is absent:

- in stable periods: suppress false alarms and certify stable operation;
- in transition periods: provide earlier warning and interpretable stress maps;
- in all periods: support operator confidence with bounded, repeatable indicators.

9.3 Highest expected business impact lane

The strongest near-term business lane is rotating machinery and turbomachinery operations (compressors, turbines), where instability episodes are expensive and early warning has direct operational value.

Priority order for impact exploration:

1. turbomachinery operations (highest immediate economic leverage),
2. CFD design-loop acceleration (reduced model iteration cost),
3. process-industry mixing/flow-monitoring diagnostics.

9.4 Hurricane forecasting boundary

This framework should be positioned as a turbulence/stress feature layer, not as a replacement for full numerical weather prediction models. A valid claim in hurricane applications would be “complementary diagnostic value,” and only after strict out-of-sample backtesting against established baselines.

9.5 Financial-market analogy boundary

The turbulence language may be used as a structural analogy for market regime transitions (intermittency/coherence bursts), but not as a literal fluid-dynamics equivalence. Any finance application should be framed as a risk-signal architecture validated by market-specific out-of-sample metrics.