Real-Time Space Traffic Management for 50,000 Satellites: A Spectral Digital Twin

The entire LEO catalog. Updated in two minutes.

Tam\{}\{}'as Nagy, Ph.D.

tnagyphd@gmail.com

Draft

Abstract

The low Earth orbit environment is projected to host over 50,000 active satellites by 2028, generating approximately 125,000 relevant conjunction pairs per day. Current conjunction assessment operates in overnight batch mode, producing alerts hours after the screening window closes. We propose a **spectral digital twin** architecture in which every tracked object carries a compact spectral state — nominal trajectory, N spectral uncertainty coefficients, a non-Gaussian correction factor γ, and an analytical sensitivity $\partial P_c/\partial(\Delta v)$ — enabling full-catalog conjunction screening in **2 minutes** (125,000 pairs × 0.001 s/pair). The architecture comprises four operational layers: catalog maintenance, pairwise screening, threshold alerting, and optimal maneuver response. We demonstrate that the spectral digital twin enables capabilities impossible with current systems: intraday risk updates, pre-launch conjunction assessment, multi-operator conflict resolution with analytical gradients, and formally verified P_c computation (Lean 4). The system is deterministic, auditable, and reproducible — properties increasingly demanded by emerging space traffic management regulations (FAA, ESA, UN COPUOS). We estimate the development cost at \$5–10M with annual industry-wide savings of \$50–100M from optimized maneuver coordination, reduced false alerts, and extended satellite lifetime. The spectral digital twin integrates the entire spectral space situational awareness pipeline: non-Gaussian P_c (Nagy, 2026d), analytical maneuver optimization (Nagy, 2026, Maneuver), and cascade risk monitoring (Nagy, 2026, Kessler) into a unified operational system.

---

1. Introduction

1.1 The Scale Problem

As of March 2026, the U.S. Space Command tracks 36,000 cataloged objects in Earth orbit. SpaceX has deployed 6,500+ Starlink satellites with a license for 42,000. Amazon Kuiper has begun deployment toward 3,236 satellites. OneWeb operates 630 with plans for 6,372. By 2028, conservative estimates project 50,000+ active satellites in LEO, with an additional 100,000+ tracked debris objects.

The conjunction screening problem scales quadratically. For n objects, the number of potential conjunction pairs is $\binom{n}{2} \approx n^2/2$. After geometric filtering (objects in non-intersecting orbits are discarded), approximately 125,000 pairs per day require detailed P_c computation.

1.2 Current Operational Architecture

The 18th Space Defense Squadron (18 SDS) provides conjunction data messages (CDMs) to satellite operators through the Space Safety Portal. The workflow:

1. **Observation**: Radar and optical sensors track objects, producing state vectors with covariance.
2. **Screening**: Batch process identifies close approaches within a 7-day prediction window.
3. **CDM generation**: For each close approach, compute P_c using 2D Gaussian (Akella and Alfriend, 2000).
4. **Dissemination**: CDMs distributed to operators (delay: 4–12 hours from observation).
5. **Operator response**: Manual review, maneuver planning (trial-and-error), execution.

The latency from observation to operator action is typically **12–36 hours**. For LEO objects with orbital periods of 90 minutes, this means conjunction geometry can change significantly between screening and action.

1.3 Why Real-Time Matters

Several operational scenarios demand faster-than-batch conjunction assessment:

- **Launch conjunction**: A rocket is on the pad. The launch window opens in 2 hours. Does the trajectory create conjunctions with existing objects? Current answer time: 4–8 hours (too slow; launches proceed with incomplete screening).
- **Maneuver deconfliction**: SpaceX plans to maneuver Starlink-2847 to avoid debris. Does the post-maneuver orbit create a conjunction with OneWeb-445? Current process: phone call, email exchange, manual re-screening. Time: hours to days.
- **Breakup response**: A satellite fragments (explosion, collision). Fragments are detected within hours. Which existing satellites are at immediate risk? Current process: emergency screening, manual prioritization. Time: 12–24 hours.
- **Continuous monitoring**: Rather than screening once per day, track P_c evolution in near-real-time as observations update covariances. Currently impossible at catalog scale.

1.4 Our Contribution

We propose the **spectral digital twin**: a compact, updatable representation of every tracked object that enables full-catalog conjunction screening in 2 minutes, with analytical maneuver optimization, non-Gaussian P_c, and formal verification. The system integrates the spectral framework papers (Nagy, 2026a–d) into a coherent operational architecture.

---

2. The Spectral Object State

2.1 State Representation

Each tracked object o_j in the spectral digital twin carries a **spectral state vector**:

$$\mathcal{S}_j = \left(\mathbf{x}_j(t),\ \{A_k^{(j)}\}_{k=0}^{N-1},\ \gamma_j,\ \nabla_{\Delta v} P_c^{(j)}\right) \tag{1}$$

Component	Description	Size	Update frequency
$\mathbf{x}_j(t)$	Nominal trajectory (RK78 ephemeris)	6 floats/epoch	Every observation
$\{A_k^{(j)}\}$	Spectral uncertainty coefficients	$N \leq 64$ floats	Every observation
γ_j	Non-Gaussian correction factor	1 float	Precomputed lookup
$\nabla_{\Delta v} P_c^{(j)}$	Maneuver sensitivity	3 floats/conjunction	On demand

Table 1. Spectral state components. Total storage per object: $ $600 bytes (vs. $ $10 KB for full covariance history in current systems).

2.2 Spectral Coefficient Update

When a new observation arrives for object o_j, the spectral coefficients update via:

$$A_k^{(j)}(t_{\text{new}}) = e^{M_j \Delta t} \, A_k^{(j)}(t_{\text{old}}) + \delta A_k^{(\text{obs})} \tag{2}$$

where $\delta A_k^{(\text{obs})}$ incorporates the measurement update (analogous to the Kalman gain, but in spectral space). The propagation $e^{M_j \Delta t}$ is a matrix-vector multiply: $O(N^2)$ operations, $< 1\ \mu$s.

2.3 The Non-Gaussian Correction Factor

The correction factor γ_j encodes how much the non-Gaussian spectral P_c differs from the Gaussian 2D-Pc:

$$\gamma_j = \frac{P_c^{\text{spectral}}}{P_c^{\text{Gaussian}}} \tag{3}$$

From Nagy (2026d): γ depends primarily on orbital altitude, inclination, and propagation time. A precomputed lookup table $\gamma(h, i, \Delta t)$ provides the correction factor without recomputation. Typical values: $\gamma \in [1.0, 3.5]$ (always ≥ 1, since non-Gaussianity increases tail probabilities).

---

3. The Four-Layer Architecture

3.1 Layer 1: Catalog Maintenance

The catalog is a database of spectral states $\{\mathcal{S}_j\}_{j=1}^{N_{\text{cat}}}$, continuously updated from observation sources:

- **Space Surveillance Network** (SSN): Radar and optical observations, 200,000+ observations/day.
- **Commercial providers** (LeoLabs, ExoAnalytic): Supplementary tracking data.
- **Operator-provided ephemerides**: High-accuracy predictions from satellite operators.

Each observation triggers a spectral state update (Eq. 2). The catalog supports 150,000+ objects with spectral states, requiring $ $90 MB of storage.

3.2 Layer 2: Pairwise Screening

For each pair (o_i, o_j) that passes geometric pre-filtering:

$$P_c^{(i,j)} = \gamma_{i,j} \cdot \sum_{k=0}^{N-1} A_k^{(i,j)}(T_{\text{TCA}})\, G_k^{(R_{i,j})} \tag{4}$$

Computation per pair: 0.001 seconds (from Nagy, 2026d). For 125,000 daily pairs:

$$T_{\text{screen}} = 125{,}000 \times 0.001 \text{ s} = 125 \text{ s} \approx 2 \text{ minutes} \tag{5}$$

This enables **multiple screenings per day** — a fundamental capability upgrade from overnight batch processing.

3.3 Layer 3: Threshold Alerting

Conjunctions are classified by P_c and time-to-TCA:

Alert Level	P_c Threshold	Time-to-TCA	Action
Green	$< 10^{-5}$	Any	Monitor
Yellow	10^{-5} – 10^{-4}	> 48 h	Plan maneuver
Orange	10^{-4} – 10^{-3}	> 12 h	Execute maneuver
Red	$> 10^{-3}$	Any	Emergency maneuver
Red (imminent)	$> 10^{-4}$	< 6 h	Immediate action

Table 2. Alert classification. The spectral system's fast update rate (2 min vs. 12 h) means alerts can be issued at the Yellow level with high confidence, giving operators more reaction time.

3.4 Layer 4: Optimal Maneuver Response

When a conjunction exceeds the action threshold, the spectral maneuver optimizer (Nagy, 2026, Maneuver) computes:

$$\Delta v^* = \arg\min_{|\Delta v|} \text{ s.t. } P_c(\Delta v) < \epsilon \tag{6}$$

using the analytical gradient $\partial P_c / \partial(\Delta v)$ from the spectral representation. Convergence: 3–5 iterations, $< 10\ \mu$s. Post-maneuver full-catalog screening: 2 minutes.

---

4. Multi-Operator Coordination

4.1 The Coordination Problem

Multiple operators sharing LEO must coordinate maneuvers. Without coordination:

- SpaceX maneuvers Starlink-2847 to avoid debris fragment D-12045.
- The post-maneuver orbit creates a new conjunction with OneWeb-445.
- OneWeb, unaware, does not maneuver.
- Collision risk increases rather than decreases.

Current coordination: bilateral phone calls and email exchanges between operators. Latency: hours to days. Scalability: impossible at 50,000+ satellites.

4.2 Spectral Conflict Resolution

The spectral digital twin enables automated multi-operator conflict resolution:

1. **Shared P_c computation**: All operators use the same spectral P_c (deterministic, reproducible). No disputes about risk level.
2. **Joint optimization**: For a conjunction involving satellites from different operators:

$$\min_{\Delta v_A, \Delta v_B} \; w_A |\Delta v_A|^2 + w_B |\Delta v_B|^2 \quad \text{s.t.} \quad P_c(\Delta v_A, \Delta v_B) < \epsilon \tag{7}$$

where w_A, w_B reflect operator costs (fuel budget, mission priority). The analytical gradient extends to joint maneuvers: $\partial P_c / \partial(\Delta v_A, \Delta v_B)$.

3. **Cascading deconfliction**: After computing $(\Delta v_A^*, \Delta v_B^*)$, screen both post-maneuver orbits against the full catalog (4 minutes total). If new conjunctions arise, iterate with additional operators.

4.3 The Coordination Protocol

$$\text{Detect} \;\rightarrow\; \text{Notify} \;\rightarrow\; \text{Compute joint } \Delta v^* \;\rightarrow\; \text{Screen} \;\rightarrow\; \text{Execute}$$

Total time from detection to verified maneuver plan: < 10 minutes (vs. hours/days currently). This enables **same-orbit maneuver coordination** — both operators can act within the same orbital pass.

* * *

5. Regulatory Compliance and Auditability

5.1 The Regulatory Landscape

Space traffic management regulation is accelerating:

- **FAA** (USA): Requires conjunction assessment for launch licenses. Considering mandatory collision avoidance for licensed operators.
- **FCC** (USA): Requires orbital debris mitigation plans for spectrum licenses. SpaceX's Starlink license includes avoidance commitments.

- **ESA**: Proposed Space Traffic Management framework (2024) calls for standardized P_c computation and mandatory avoidance thresholds.
- **UN COPUOS**: Long-Term Sustainability Guidelines (2019) recommend conjunction assessment and collision avoidance for all operators.

5.2 Properties Required by Regulators

Property	Current (Gaussian MC)	Spectral Digital Twin
Deterministic	No (MC sampling)	Yes
Reproducible	Seed-dependent	Exact
Auditable	Opaque (statistics)	Inspectable coefficients
Formally verified	No	Yes (Lean 4)
Non-Gaussian tails	No (Gaussian only)	Yes
Real-time	No (batch)	Yes (2 min)
Standardized	De facto (2D-Pc)	Proposed standard

Table 3. Regulatory compliance comparison. The spectral method satisfies every property that emerging STM regulations demand.

5.3 Formal Verification as Regulatory Asset

The spectral P_c computation chain is verified in Lean 4 (Nagy, 2026d): the density representation, propagation, integration, and sensitivity formulas are machine-checked mathematical proofs. This provides a level of assurance unprecedented in space safety computations:

- The P_c formula is **provably correct** (not just tested on examples).
- Regulators can inspect the proof chain from axioms to final result.
- Software implementations can be validated against the formal specification.
- No other conjunction assessment method offers comparable assurance.

---

6. Performance Analysis

6.1 Full System Timing

Operation	Time	Frequency
Catalog update (per observation)	1 μs	200,000/day
Full-catalog screening (125K pairs)	2 min	12×/day
Single conjunction P_c	1 ms	On demand
Maneuver optimization	10 μs	Per alert
Post-maneuver screening	2 min	Per maneuver
Cascade risk update ($\mathcal{K}_I$)	15 min	Daily

Table 4. Operational timing budget. Total daily compute: < 30 min for 12 screening cycles, plus on-demand maneuver optimization.

6.2 Comparison with Current Systems

Metric	Current (18 SDS/CARA)	Spectral Digital Twin	Improvement
Screening latency	4–12 hours	2 minutes	120–360×
Screening frequency	1/day	12+/day	12×
P_c model	Gaussian 2D	Non-Gaussian spectral	Exact tails
Maneuver planning	Trial-and-error (min)	Analytical gradient (μs)	$10^6\times$
Multi-operator coord.	Manual (hours)	Automated (10 min)	10–100×
Alert lead time	12–24 hours	Real-time	Continuous
Formal verification	None	Lean 4 proofs	Unique

Table 5. System-level comparison. The spectral digital twin represents a generational improvement in every operational metric.

---

7. Cost-Benefit Analysis

7.1 Development Costs

Component	Estimated Cost	Timeline
Spectral engine (core algorithms)	\$2M	12 months
Catalog integration (SSN interface)	\$1.5M	6 months
Multi-operator coordination layer	\$1.5M	9 months
Formal verification (Lean 4)	\$1M	6 months
Testing and validation	\$2M	12 months
Regulatory certification	\$1M	6 months
Total development	**\$9M**	**24 months**

Table 6. Development cost estimate.

7.2 Operational Savings

Savings Source	Annual Estimate	Basis
Fuel savings (optimal maneuvers)	\$20–30M	20% improvement × \$100M Starlink
False alert reduction	\$5–10M	30% fewer unnecessary maneuvers
Multi-operator coordination	\$10–20M	Automated vs. manual deconfliction
Satellite lifetime extension	\$15–40M	3–6 months × replacement cost
Total annual savings	**\$50–100M**	**Industry-wide**

Table 7. Operational savings estimate. ROI: 5–11× in year 1.

7.3 Value Beyond Cost Savings

The spectral digital twin enables capabilities that have no current-system equivalent:

- **Pre-launch safety certification**: Prove a new constellation does not degrade the orbital environment before launching a single satellite.
- **Insurance pricing**: Provide actuarially sound collision risk metrics for space insurance underwriters.
- **Debris removal prioritization**: Identify which debris objects contribute most to cascade risk ($\mathcal{K}_I$ sensitivity), guiding active debris removal missions.
- **Regulatory compliance automation**: Generate auditable, formally verified P_c reports for regulatory filings.

8. Discussion

8.1 Integration with Existing Infrastructure

The spectral digital twin is designed as an **augmentation** to existing systems, not a replacement:

- **Input**: Standard two-line elements (TLEs), conjunction data messages (CDMs), operator ephemerides — the same data feeds current systems consume.
- **Output**: Enhanced CDMs with spectral P_c, γ correction, and maneuver recommendations — backward-compatible with existing operator tools.
- **Transition**: Operators can run spectral screening in parallel with existing systems during a validation period, comparing results before switching.

8.2 The Path from Single-Pair to System-Level

The spectral framework represents a progression from micro to macro:

1. **Single pair** (Nagy, 2026d): Non-Gaussian P_c for one conjunction. The foundation.
2. **Single maneuver** (Nagy, 2026, Maneuver): Optimal avoidance for one conjunction. The gradient.
3. **Cascade risk** (Nagy, 2026, Kessler): Long-term sustainability for the population. The consequence.
4. **Full system** (this paper): Real-time management of the entire catalog. The integration.

Each layer builds on the previous, with the spectral representation as the unifying mathematical structure.

8.3 Limitations and Future Work

- **Atmospheric density uncertainty**: The Fokker–Planck generator M depends on atmospheric density models (NRLMSISE-00, JB2008) that have 10–30% uncertainty during geomagnetic storms. The spectral framework handles this via an ensemble of generators $\{M_i\}$ with storm-conditional weighting, but validation during extreme space weather events remains future work.

- **Very low orbits ($<$ 300 km)**: Rapid orbital decay makes the spectral expansion less efficient (coefficients change rapidly). For these objects, the spectral approach offers less advantage over direct Monte Carlo.

- **International participation**: The spectral digital twin's value increases with the number of participating operators. A system with only partial catalog coverage provides partial benefit.

---

9. Conclusion

The spectral digital twin transforms space traffic management from a batch, Gaussian, manual process into a real-time, non-Gaussian, automated system. Full-catalog screening in 2 minutes, analytical maneuver optimization in microseconds, multi-operator coordination in minutes, and formally verified P_c computation represent a generational advance in operational capability. The system integrates four years of spectral method development — the Fenton distribution, universal risk representation, conjunction assessment, maneuver optimization, and cascade risk — into a unified architecture that meets the demands of 50,000+ satellite operations. With an estimated development cost of \$9M and annual industry savings of \$50–100M, the spectral digital twin offers compelling economics alongside unprecedented technical capability. As LEO becomes the most congested transportation domain in human history, the tools for managing it must evolve from artisanal to industrial. This paper proposes the industrial solution.

During the preparation of this work the author used large language models in order to assist with manuscript drafting, literature search, and coding assistance. After using these tools, the author reviewed and edited the content as needed and takes full responsibility for the content of the published article.

References

- Akella, M. R. and K. T. Alfriend (2000). Probability of collision between space objects. *JGCD*, 23(5). DOI: 10.2514/2.4611
- Alfriend, K. T. et al (1999). Probability of collision error analysis. *Space Debris*, 1(1). DOI: 10.1023/a:1010056509803
- European Space Agency (2024). Space Traffic Management: ESA Framework Proposal. *ESA/ADMIN/2024-003.*, 2024-003.
- Flohrer, T., Krag, H., and Klinkrad, H (2008). Assessment and categorization of TLE orbit errors for the US SSN catalogue. *Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference..*
- Hejduk, M. D., Plakalovic, D., and Newman, L (2023). Non-Gaussian effects in operational conjunction assessment. *NASA CARA Technical Report..*
- Kelso, T. S (2009). Analysis of the Iridium 33-Cosmos 2251 collision. *Advances in the Astronautical Sciences.*
- McKnight, D. S., Speaks, S., Macdonald, J., and Ebright, K (2021). Identifying the 50 statistically-most-concerning derelict objects in LEO. *Acta Astronautica.*
- SpaceX (2022). Starlink conjunction avoidance statistics. *FCC submission..*
- United Nations COPUOS (2019). "Guidelines for the Long-Term Sustainability of Outer Space Activities." *A/AC.105/C.1/L.366.* A/AC.105/C.1/L.366*.
- Weeden, B. and Samson, V (2022). Global counterspace capabilities: An open source assessment. *Secure World Foundation..*
- Nagy, T. (2026). The Fenton Distribution Solved. *Working paper.*
- Nagy, T. (2026). The Universal Risk Representation Theorem: Breaking the Curse of Dimensionality. *Zenodo.* DOI: 10.5281/zenodo.18910566
- Nagy, T. (2026). Spectral Kernel Risk: The Fourier–Gaussian Process Duality for Portfolio Loss Distributions. *Working paper.*
- Nagy, T. (2026). Optimal Collision Avoidance in One Microsecond: Analytical Maneuver Planning via the Spectral Fokker-Planck Generator. *Working paper.*
- Nagy, T. (2026). Spectral Cascade Risk: Quantifying the Kessler Syndrome via Fokker-Planck Fragment Propagation. *Working paper.*