

Fast All-vs-All Conjunction Screening

Screening every catalog object against every other in seconds to minutes, for any propagator

Anthony “TJ” Koury III and Claude Opus 5.5

Koury: Edgesource, Space Data Network · tj@edgesource.com. Claude: Anthropic.

Technical whitepaper 1.9 | 11 October 2026

Numerical evidence cutoff: 11 October 2026

© Edgesource Corporation

Executive summary

Space Data Network (SDN) screens the full public catalog, 32,514 objects, all-vs-all over three days in 19 s with SGP4 on ordinary processors, and in 21 s with a GPU. With numerical high-precision orbit propagation (HPOP) under the full force model of section 5, it takes 28 CPU-hours, 1.8 hours on 16 dedicated cores, almost all of it propagation; search and refinement take 138 s on 2 module threads. The screen finds every pair that comes within 5 km, and reports each close approach’s time (TCA), miss distance, relative speed and maximum collision probability. The measurements were taken on one 28-core workstation, in Node.js and in a standard web browser. In WasmEdge, the runtime SDN nodes use, the SGP4 screen takes 23 s. Every run on one catalog and propagator reports the same conjunctions.

Five design choices produce that speed:

1. **A spatial grid proposes pairs; the module decides.** At every 60 s step, a uniform grid finds the pairs that could be close among the 528.6 million, on a GPU (WebGPU) or in the WebAssembly (WASM) conjunction module on CPU threads. The module re-tests each candidate in f64 and refines it. Every reported number comes from the module.
2. **Each trajectory bounds its own motion.** Between two samples, an object cannot stray from a straight line by more than a bound its trajectory source computes. A pair whose straight-line approach stays farther apart than the threshold plus both bounds cannot close, and is discarded without further work. The bound holds for any force model, maneuvers included.
3. **One propagator per run, in time windows.** A run screens either SGP4 element sets or one numerical propagator’s trajectories, never a mix. Long screens run as consecutive windows, so memory stays at one window’s trajectories while the next window is propagated.
4. **Load once, move records unchanged.** The catalog loads into the module once. A propagator’s output passes into the conjunction module as the records it emitted, without re-encoding.
5. **Refinement proves, then solves.** Where the range of a pair provably has one minimum, Newton’s method on the range rate finds it in about ten evaluations.

The screen reproduces the module’s single-call screen where both can run. On 4,000 objects over one day, all 1,068 conjunctions match, with TCA within 0.64 ms and miss distance within 5 mm; the single call takes 32.9 s and the windowed screen 0.44 s.

Against SOCRATES on the element sets SOCRATES itself used, TCA agrees within 0.5 ms and miss distance within 2.4 m for 95 % of cataloged pairs, the precision SOCRATES reports. Each event reports the covariance-free maximum probability. With an empirical model of element-set prediction error, events also carry a covariance-based probability. It is labeled calibrated only where held-out precise orbits confirm the covariance: LEO 600 to 800 km, 13,734 of the 294,544 three-day conjunctions. There, screening on the uncertainty region missed 0–0.4 % of true collisions, against 2.7–16 % for a probability threshold of $1e-4$, whose probabilities a large covariance dilutes (section 7). The model¹ adds 1.1 s to the 19 s screen (section 7).

Operators keep their most precise orbits and planned maneuvers private. Section 8 specifies **private screening**: two operators learn when their objects come close without exchanging trajectories, by computing distances on homomorphically encrypted positions.


- **Cost.** Its arithmetic is measured: 0.17 s and 8.7 MB per pair-day at 1 s steps.
- **What it reveals.** The paper shows how a naive design leaks distances, how invented trajectories can locate a hidden satellite, and which defenses stop that, with or without a public track for either satellite.
- **What it cannot hide.** Real close approaches reveal what safety requires.

1 The problem

An all-vs-all screen compares every object with every other: $n(n-1)/2$ pairs, 528.6 million for 32,514 objects. A three-day screen at 60 s steps evaluates each pair at 4,321 instants, about 2.3 trillion pair-steps. A close approach at 10 km/s relative speed lasts under a second inside a few kilometers, so a coarse test that misses it loses the event.

The module’s existing single-call screen (`screen_catalog`) uses a perigee/apogee prefilter and a k-d tree per time step. It holds every encounter of the window in WASM linear memory, at most 2 GiB. It runs out of memory above about 7,000 objects for a one-day window².

2 Pipeline



All physics runs in WASM modules; the host moves records and contains none.

The host code (JavaScript) moves records between the module and the GPU and contains no physics. The conjunction module, the SGP4 implementation, the HPOP propagator and the epoch-state converter are WASM modules. They are built with the SDN Module SDK and run unchanged in a browser, in the WasmEdge runtime and in its Docker image³.

Stage	Where	Precision	Decides
Coarse grid	WASM module, threaded	f64 computed, f32 stored	Sample states and bounds
Pair search	GPU, WebGPU compute	f32 with 0.01 km slack	Candidates only
Pair search without a GPU	WASM module, threaded	f64	Candidates only
Re-test and refinement	WASM module, threaded	f64	Every reported result

Stage	Where	Precision	Decides
Window merge	Host	Exact comparisons	Which window owns each TCA

3 The candidate test

Self-bounded motion

Let h be half the coarse step (30 s at 60 s steps). For a sample of an object at step k , with position $\mathbf{r}$ and velocity $\mathbf{v}$, the trajectory source supplies a bound D such that the true position stays within D of the straight line:

$$|\mathbf{r}(t_k + \tau) - \mathbf{r}_k - \mathbf{v}_k \tau| \leq D \quad \text{for all } |\tau| \leq h.$$

Two objects can then come within the threshold d during the interval only if the straight-line relative path does so within $d + D_1 + D_2$. The test computes the clamped closest approach of the relative straight line over $[-h, h]$ and compares it with that limit. A pair that fails the test at every step of an interval cannot close in it.

Each source kind computes D from what it knows:

Source	Bound	Basis
SGP4 element set	$\frac{1}{2} A h^2$, with $A = 1.05 \mu / r_{\min}^2$ and $r_{\min} = \mathbf{r} - \mathbf{v} h$, floored at 6,000 km	SGP4 motion is unpowered; the 5 % margin covers J2 (0.16 % at the surface) and drag
Chebyshev trajectory (PPE) from any propagator	$\frac{1}{2} A h^2 + (J + G) h + P$	A bounds the position series' second derivative using $ T_k'' \leq k^2(k^2 - 1)/3$. J and P sum velocity and position jumps where intervals meet. G is the gap between the velocity series and the position series' derivative

The trajectory bound uses only the coefficients, so it holds for any force model and for impulsive maneuvers that appear as interval jumps. A native test checks it against the trajectory itself every 0.1 s around samples taken every 7 s. The trajectory is a 7,000 km orbit interpolated as a propagator would export it, with a 5 m/s maneuver at an interval boundary. Tested at $h = 30$ s and $h = 90$ s, the bound was never exceeded, and away from the maneuver it was at most 1.34 times the true deviation².

Spatial grid

Testing all 528.6 million pairs at every step is wasted work: almost all pairs are thousands of kilometres apart. At each step, an object's straight-line segment over $[-h, h]$, widened by $d/2 + D$, gives an axis-aligned box. Two objects can pass the straight-line test only if their boxes overlap, so a search over overlapping boxes drops nothing the test would pass.

The boxes go into a uniform grid whose cell is the 99th-percentile box edge. Each overlapping pair is counted once, in the cell holding the low corner of the two boxes' overlap. The few boxes that span four or more cells on an axis on the GPU, or five or more on CPU threads, such as fast perigee passes, are tested against every object. The same grid runs in two places:

- **On a GPU**, in three WebGPU compute passes per step. The first inserts each box into its cells, hashed into one lock-free list per slot. The second has each object walk its cells' lists. The third tests the large boxes.

- **Without a GPU**, in the WASM module on CPU threads: in a browser without WebGPU, in Node.js, or in WasmEdge on an SDN node or in Docker.

Either way, the grid proposes exactly the pairs an all-pairs test would. On the full catalog over three days, the GPU grid proposed the same 2,678,533 candidates as an all-pairs GPU kernel did.

SDN nodes run modules in WasmEdge compiled ahead of time. That needs SDN’s patched WasmEdge 0.16.4: the stock compiler mishandles offsets on atomic memory instructions, and with it the full-catalog search hung in its second window. Interpreted WasmEdge is about 75 times slower.

Two precisions, one decision

The GPU evaluates the test in f32. Positions near 7,000 km round to about 0.5 m in f32, so the GPU adds 0.01 km of slack and proposes a superset. The module repeats the test in f64 for every candidate. Over one day of the full catalog the GPU proposed 829,990 candidates and 828,910 passed the f64 test. Over three days the GPU proposed 2,678,533 and the CPU grid, which applies the f64 test directly, 2,675,123; both refine to the same conjunctions.

Refinement

Passing steps of a pair join into encounters. Each encounter is refined by the same TCA search the single-call screen uses, so both report the same TCA.

For most encounters the search first proves that the range has one minimum in its window, then solves for it directly. Let A bound both objects’ accelerations over the window; for SGP4 this is $1.05 \mu / r^2_{\min}$, the bound of the candidate test. With T the window’s half-width and $\Delta \mathbf{r}$, $\Delta \mathbf{v}$ the relative state at its middle:

$$|\Delta \dot{\mathbf{r}}| \geq |\Delta \mathbf{v}| - \varepsilon - AT, \quad |\Delta \mathbf{r}| \leq |\Delta \mathbf{r}_{\text{mid}}| + |\Delta \mathbf{v}|T + \frac{1}{2}AT^2.$$

ε (1 m/s) allows for a source’s velocity differing from the rate of its position. If $q = |\Delta \mathbf{r}| \max A / |\Delta \dot{\mathbf{r}}| \min^2 < 1$, the squared range is convex on the window. Its one minimum is then the root of the range rate $f = \Delta \mathbf{r} \cdot \Delta \mathbf{v}$, and Newton steps of $-f / |\Delta \mathbf{v}|^2$ converge by a factor q or better per step. About ten state evaluations reach the root to the resolution of a Julian date, about 47 μs .

Where the proof fails, for example a slow pair that stays close for minutes, the search samples the range every 5 s and refines each minimum it brackets with golden-section search.

Probability uses the Alfano maximum⁴. When a refine call would stage more than the module’s 16,384-event output limit, the host splits it in two.

4 One propagator per run

A run screens one propagator’s output. A resident index holds either mean elements, evaluated with the module’s SGP4, or trajectories from one generator, identified by the Chebyshev ephemeris (PPE) record’s EPHEMERIS_SOURCE. The module refuses an index that mixes them (**mixed-propagators**).

Chebyshev series per interval are the standard form for precomputed ephemerides. NASA’s SPICE toolkit stores them that way: SPK data types 2, 3, 14 and 20 hold Chebyshev coefficients for position, velocity or both¹⁵, and JPL distributes its planetary ephemerides as Chebyshev polynomials fit to positions and velocities in 32-day intervals^{16,17}. Deprit, Pickard and Poplarchek recommended Chebyshev approximation for compressing ephemerides because it minimizes the largest error over the interval¹⁸. A Chebyshev interpolant of degree n is within a factor $2 + (2/\pi) \log(n + 1)$ of the best polynomial approximation of that degree¹⁹. The same form serves spacecraft and catalog orbits: onboard ephemeris generation²⁰ and orbit information with continuous state and covariance at a predetermined accuracy²¹. Open-source readers include jplephem (MIT License,²²) and ANISE (Mozilla Public License 2.0,²³); NAIF’s toolkit is free to use under its own rules¹⁵.

Trajectories reach the module as their propagator exported them. The HPOP module’s export is a size-prefixed \$PRW record per object. The conjunction module’s index preparation takes those records on its `trajectories` port without re-encoding, and the request carries only each object’s identity. Trajectories in TT or TDB map to UTC linearly between their converted interval ends, using the vendored ERFA library⁵. TDB – TT is evaluated once per minute and interpolated, exact to about 1e-14 s.

5 Time windows

A long screen runs as consecutive windows; the merged result equals one screen of the whole span:

- A conjunction belongs to the window that holds its TCA.
- A TCA within the refinement tolerance of a shared edge is reported by both windows and kept once.

Window length trades memory for per-window overhead. A day of HPOP’s trajectories is about 240 KB per object, so HPOP screens in 1-hour windows: about 390 MB of trajectories for the full catalog, which the conjunction module indexes in 1.05 GiB of its 2 GiB. SGP4 windows reuse one index loaded at the start, so they are longer (6 hours).

Tests check four windows against a single screen for both propagators. The SGP4 case uses CelesTrak element sets that include an object SGP4 cannot propagate. The HPOP case uses crossing orbits integrated by HPOP, passed through the `trajectories` port^{2,33}.

The HPOP propagation farm

1. The epoch-state module converts each element set into a GCRF state at its epoch: SGP4 at zero elapsed time, then TEME to GCRF through ERFA. This is the OMM-to-numerical handoff of the companion paper¹, section 4.
2. HPOP resident instances, at most 1,024 objects each, integrate those states on worker threads with the force model below. Objects are dealt round-robin into about two instances per worker, so element sets of every age spread evenly.
3. Each window, every instance exports the window and the records go to the browser unchanged. The next window is exported while the current one is screened.

HPOP	Setting
Gravity	EGM2008 to degree and order 70, in Earth-fixed axes from IERS Earth orientation ^{24,26}
Sun, Moon and planets	Sun, Moon, Venus, Mars and Jupiter from JPL DE440s ²⁵
Drag	NRLMSISE-00 ²⁷
Radiation pressure	Cannonball with a conical Earth shadow; Earth albedo and infrared, Knocke model on 15° elements ²⁸
Tides and relativity	IERS 2010 solid tides and relativistic terms; FES2004 ocean tides to degree and order 30 ^{26,29}
Earth orientation and space weather	As released at the screen’s start: IERS finals2000A rows to the day before, the last one held; GFZ Kp, ap and F10.7, then NOAA SWPC’s latest daily report with its 3-day forecast, the last day held ³⁰
Drag coefficient	$Cd \cdot A/m = 2B^*/\rho_0$ from each element set’s B^* , with SGP4’s reference density $\rho_0 = 2.461 \times 10^{-5}$ XKMPER kg/m ² per Earth radius, XKMPER = 6,378.135 km ³¹
Reentry	HPOP ends a trajectory at 100 km geodetic altitude; an object is screened through the last full window before its reentry
Assumptions	$Cd \cdot A/m = 0.01$ m ² /kg for the 6,069 element sets with $B^* \leq 0$; $Cr \cdot A/m = 0.02$ m ² /kg for all 32,514 objects
Integrator	Runge–Kutta–Fehlberg 7(8), relative tolerance 10^{-12} , steps of at most 300 s, restarted only where the space weather changes ³²

HPOP	Setting
Export	Chebyshev series of degree 12 on 300 s segments of TT, fitted to the integrator’s dense output; each segment states its largest difference from the integrator’s accepted states

The tolerance, the segment length and the resident integration were measured³³:

- **Tolerance.** Against 10^{-14} , the integration’s own error after three days is 0.1 m at 400 km, 1.2 m at 700 km sun-synchronous and 0.18 m for Molniya. At 10^{-11} it is 6 to 9 m in LEO.
- **Segment length.** A degree-12 series resolves about four cycles per segment. 300 s keeps the 70×70 field’s shortest LEO terms, with periods near 80 s; over 600 s, the terms above degree 35 leave centimetre fit errors. Over four days the exported series stay within 0.31 mm of the integrator’s dense output in LEO and 0.18 mm for Molniya.
- **One integration.** A resident object takes the same steps as one HPOP request for that object. From a day before its epoch to three days after, the two agree within 0.05 mm in LEO, 9 μ m for Molniya and 9 nm for GPS and geostationary orbits.

HPOP integrates each object from its element epoch, so the first window carries that catch-up: 55,578 object-days from the element epochs to the screen’s start, against 97,542 for the three days screened. HPOP drops cached intervals that end before the requested window, so an instance holds about one window per object.

6 Results

All runs used the Space-Track GP catalog of 30 September 2026, 32,514 objects⁶. Screens started 2026-10-01T00:00Z, with a 5 km threshold and 60 s coarse steps, on a Mac Studio with 28 cores and 27 module threads. GPU runs used headless Chrome with Metal for WebGPU. The workstation was shared with other work (1-minute load average 15–65 on 28 cores), so times are upper bounds.

Full catalog, three days

Propagator	Pair search	End to end	Sampling and search	Refine	Conjunctions
SGP4, 12×6 h	CPU, Node.js	19.1 s	9.7 s	6.6 s	294,544
SGP4	CPU, WasmEdge (AOT, SDN patches)	22.9 s	11.9 s	8.7 s	294,544
SGP4	GPU	21.3 s	6.1 + 2.3 s	7.2 s	294,544
HPOP, 72×1 h	CPU, Node.js	6,383 s	97.7 s	40.1 s	283,700

- For each propagator, every run reports the same conjunctions with the same TCA and miss distance.
- **SGP4’s** remaining time is sampling: $4,321 \text{ steps} \times 32,514 \text{ SGP4 evaluations}$. On the GPU path the module samples (6.1 s) and the GPU searches (2.3 s).
- **HPOP’s** time is propagation: 99,929 CPU-seconds on 8 worker threads, 0.65 s per object-day, 36 % of it the catch-up from element epochs to the screen’s start. The table’s end-to-end time is that CPU time on 16 dedicated cores plus the 137.7 s of search and refinement on 2 module threads; on the shared workstation (1-minute load average median 44, maximum 82) the screen took $15,210 \text{ s}$ ³³.
- The two propagators report different conjunctions. Of HPOP’s 283,700, 39,425 (14 %) are also in the SGP4 screen, the same pair with TCAs within 10 minutes; on those, TCAs differ by a median of 1.2 s and miss distances by 691 m. The share falls from 22 % of the first day’s conjunctions to 9 % of the third’s. Neither result is a measure of accuracy (section 7).

Step size

The candidate test is exhaustive at any step, so the step size trades sampling against candidates. SGP4, three days:

Step	Candidates	CPU search	CPU end to end	GPU end to end	Conjunctions
5 s	4,748,277	92.1 s	102.0 s	98.6 s	294,544
30 s	1,414,694	17.8 s	29.9 s		294,544
60 s	2,675,123	9.7 s	19.1 s	21.3 s	294,544
120 s	17,116,089	10.0 s	39.4 s		294,544

Every step size reports the same conjunctions, with miss distances within 2 mm. TCAs differ by up to 173 ms, but only on flat minima, where the miss distance is the same. A 5 s screen samples 12 times as often and finds nothing more. Shorter steps cost sampling; longer steps give larger bounds and more candidates. 60 s is near the optimum.

How the time came down

Three days, SGP4, CPU search in Node.js, same host:

Change	End to end	Refine
Dense TCA scan around every minimum	131.4 s	114.3 s
Unimodal proof, golden-section search	63.9 s	45.6 s
Newton steps on the range rate	55.3 s	38.4 s
The same solve for each fast encounter's first pass	34.7 s	17.6 s
SGP4 state read without a lock; element sets converted once	25.2 s	12.7 s
Refinement work shared in small chunks; events built in parallel	19.1 s	6.6 s

Every row reports the same conjunctions. On the GPU path, replacing the all-pairs kernel with the grid cut the GPU search from 8.5 s to 2.3 s.

The first version also sent the whole catalog with every call. Loading it once into the module cut one day of the full catalog from 266.7 s to 55.4 s, with the same 77,667 conjunctions.

Agreement with the single-call screen

Objects, window	Single call	Windowed, CPU search	Conjunctions	Largest difference
2,000, 0.1 day	1.1 s	0.08 s	23 of 23	TCA 0.16 ms, miss 0.6 mm
4,000, 1 day	32.9 s	0.44 s	1,068 of 1,068	TCA 0.64 ms, miss 5 mm

On the 2,000-object case the CPU search and the GPU found the same 424 candidates and 23 conjunctions, in Node.js, WasmEdge and a browser. Differences from the single call come from refinement starting from different coarse brackets, and sit inside the module's parity tolerances of 10 ms and 10 cm.

Validation findings

7 Uncertainty and probability of collision

The companion paper¹ sets the conditions this screen respects, in its sections 4, 6 and 11:

- An OMM supplies no covariance. The initial uncertainty of an OMM-seeded HPOP trajectory, including this paper’s HPOP runs, is unknown: no validated representation of it exists.
- A probability of collision requires relative-state uncertainty with independent calibration and stated cross-correlation, geometry and hard-body-radius assumptions.

Every event reports the Alfano maximum: the largest probability any covariance size could give for the reported miss distance and combined hard-body radius⁴. It needs no covariance and serves as an upper bound. Each object’s radius comes with its basis: supplied, half the catalog size, from the radar cross section, or the request default. No uncertainty is invented: an event carries covariance, and a covariance-based probability, only when a source supplied covariance or the empirical model below applies, and it states which. The work below follows the companion paper’s three parts⁸:

1. Initial covariance
2. Propagation
3. Calibration on held-out evidence

Independent truth

Calibration needs reference orbits that SGP4 did not produce. SDN’s reference-states module turns precise orbits into GCRF states at stated times, each with its stated uncertainty as covariance¹⁰:

- IGS final GPS orbits;
- ILRS orbits for LAGEOS-1/2, ETALON-1/2, Ajisai, Starlette, Stella, LARETS, WESTPAC, LARES and LARES-2;
- Sentinel-1C/1D precise orbits;
- Swarm A/B/C precise orbits.

Their frames are converted with IERS Earth orientation parameters. For 2026-08-02 to 15, this gives 48 objects. The conversion matches an independent pyerfa computation within 1 mm, and a Vallado textbook case within its printed precision.

An empirical error model, and its calibration

A model of SGP4 prediction error was built from 1.89 million element sets (2026-07-12 to 08-08), stratified by orbit regime and prediction age. It uses differences between consecutive element sets of the same object. Those differences are not truth, and the reference orbits show by how much:

- From 600 km up, under half a day, they understate the true along-track error 4 to 11 times, because consecutive fits share their error.
- Below 600 km they overstate it, by 10 to 90 times beyond a day.

The calibration gate scales each stratum’s covariance on one week of reference orbits (08-02 to 08), then tests it on the next (08-09 to 15). It compares each error’s Mahalanobis distance with χ^2 on 3 degrees of freedom, and every sample counts. A stratum is CALIBRATED only when all of these hold:

- its 1σ and 2σ containment are within 5 points of nominal;
- no more than 1 % falls outside 3σ ;
- it has at least 30 samples from at least 3 objects.

LEO 600 to 800 km, prediction age	Samples (objects)	Inside 1σ / 2σ / 3σ	Label
0 to 0.5 day	14,968 (4)	0.642 / 0.954 / 0.997	CALIBRATED
0.5 to 1 day	14,876 (4)	0.627 / 0.947 / 0.996	FAILED
1 to 2 days	31,327 (4)	0.647 / 0.943 / 0.997	CALIBRATED
2 to 3 days	31,805 (4)	0.636 / 0.922 / 0.994	CALIBRATED
3 to 5 days	58,981 (4)	0.701 / 0.960 / 0.997	CALIBRATED
5 to 7 days	49,187 (4)	0.728 / 0.965 / 0.991	CALIBRATED
Nominal		0.683 / 0.954 / 0.997	

Every other stratum is UNCALIBRATED:

- **GPS and the other MEO strata fail**, for example 0.37 / 0.70 / 0.89 for LAGEOS-1/2 and LARES-2 within half a day. GPS element sets carry per-satellite along-track offsets of -1.6 to $+3.8$ km.
- **The other LEO bands have only one or two reference objects.**
- **GEO, eccentric and beyond-GEO orbits have no independent truth here.**

The labels hold for these objects in this week.

Covariance probability on screened events

With the scaled model, element-set screens give each event's objects the covariance of their stratum at the TCA's prediction age. The event then reports the selected method's probability, Foster's by default⁹. Each event is labeled:

- the covariance as synthesized from the empirical model;
- the two objects' errors as independent;
- calibrated only when both objects' strata passed the gate.

On the three-day full catalog the screen found the same 294,544 conjunctions:

- 289,455 received a covariance probability;
- 13,734 of those are calibrated;
- refinement took 7.8 s instead of 7.0 s, and the run 19.6 s instead of 18.5 s.

Events outside the model keep the Alfano maximum. These are a negative or too-old prediction age, or a stratum the model lacks.

Propagated and fitted covariance

- **Propagation.** HPOP resident instances accept an initial covariance P_0 per object and return $P(t) = \Phi P_0 \Phi^T + Q$, with the state transition matrix of the full force model. Q is a declared white-acceleration process noise, inertial or radial-transverse-normal, given as a spectral density and discretization interval. It is recorded in the trajectory (SDS 1.232.0). The state transition matrix matches finite differences, and Q matches its closed form.
- **Orbit determination.** Fits publish their formal covariance as an OCM. The OCM lists the noise models, weighting, observations used, process noise, a priori and convergence criteria, with biases neither estimated nor considered and no consider parameters. It is marked Uncalibrated.
- **HPOP covariance, calibrated in part.** It was checked the same way as the empirical model, for this paper's HPOP screen product: an element-set epoch state with P_0 measured at the epoch, then HPOP with the full force model of section 5 and the space weather released on the element set's day. A white-acceleration Q was fitted by maximum likelihood on the fit week; a regime kept $Q = 0$ where P_0 alone passed more fit-week strata³³.
 - **LEO 600 to 800 km, 0 to 0.5 day and 1 to 3 days: CALIBRATED with P_0 alone.** P_0 carried through HPOP's state transition matrix predicts the along-track growth, 1.3 to 17 km, within 11 %.
 - **LEO 600 to 800 km, 0.5 to 1 day: FAILED.** 75 % of the samples fall inside 1σ , 7 points above nominal.
 - **LEO 600 to 800 km beyond 3 days: FAILED.**
 - **GPS: FAILED at every age.** The normal error stays within 0.15 km for a week, while the along-track error grows to 36 km RMS at 5 to 7 days, against 16 km predicted. Started from the independent reference state instead, the same force model's along-track error after 6 days is 0.2 km (median), and 1.0 km with the radiation-pressure coefficient doubled; started from the element set it is 8.4 km. The along-track error is the epoch state's.
 - **Other bands: INSUFFICIENT**, except MEO 2,000 to 10,000 km within half a day: FAILED.
- **Orbit-determination covariance is not calibrated.** Covariance supplied in an OEM or OCM enters conjunction assessment as supplied covariance and keeps the source's calibration label.

Agreement with SOCRATES

SOCRATES is CelesTrak’s public conjunction screen⁷. The replay used the top 300 rows of its 2 October 2026 maximum-probability report, from 179,398 rows. Each pair was reassessed from the element sets SOCRATES used for it, and the report and every input are hashed. The days since epoch agree to SOCRATES’s printed precision, which confirms the inputs are identical.

Group	Rows	TCA median / p95 / max	Miss distance median / p95 / max
Cataloged pairs, faster than 10 m/s	112	0.4 / 0.5 / 24 ms	0.38 / 2.4 / 6.6 m
Pairs with an object numbered 100000 or above	126	0.4 / 4.2 / 7.6 ms	0.58 / 20.8 / 61 m
One formation pair, 0 to 5 m/s	62	19 / 69 / 363 ms	0.26 / 0.47 / 0.51 m

SOCRATES prints range to 1 m and speed to 1 m/s, so the cataloged pairs agree to its reported precision. Every fast miss difference over 20 m involves an object numbered 100000 or above, mostly Starlink, and the published inputs do not show why.

Maximum probability is about nine times lower here at a 10 m combined radius. That matches SOCRATES using larger, object-specific radii, which it does not publish: a combined radius of 30 m (median) would reproduce its values. For the formation pair, a short-term encounter does not hold for either computation, so the probabilities are not comparable.

Screening rules compared

Probabilistic, bounded-set and admissible-set screening compare as follows, on cases built from real SGP4 errors against the held-out reference orbits:

- two different objects’ errors in one stratum make one encounter;
- each encounter is placed head-on and at a 90° crossing;
- synthetic true misses are added from 0 to 5 km;
- a collision is a true miss within 20 m.

The three rules alert as follows:

- **Probabilistic:** Foster’s probability $\geq 1\text{e-}4$.
- **Bounded set:** the 99.73 % relative ellipse comes within 20 m.
- **Possibility:** min-joined χ^2 possibility per object, alerting unless the necessity of no collision reaches 0.9973. Possibility and necessity are not probabilities.

The results, in the calibrated strata (10,000 cases a row):

Rule	Missed collisions, head-on	Missed collisions, crossing	False alerts at 1 km miss	False alerts at 2 km miss
Probabilistic	2.7 %	16 %	6 to 7 %	0 %
Bounded set	0.2 %	0.4 %	21 to 45 %	0 to 4.5 %
Possibility	0 %	0 %	69 to 89 %	4 to 24 %

- **Probabilistic.** Along-track uncertainty here is 0.6 to 3 km, so a true collision’s probability averages only 0.002 head-on and 0.0006 crossing. This is dilution, and it is what hides the collisions the probability rule misses.
- **Possibility.** It missed none in calibrated strata, at the cost of the most false alerts. In uncalibrated strata it missed up to 8 %, so it is only as sound as its declared uncertainty.
- **Disclosure.** The baselines and cases were chosen independently of the TEAG and ESPF authors’ implementations, and no TEAG or ESPF code was run.

8 Private screening

Operators hold back their best data:

- **National security.** Classified orbits reveal capability and coverage.
- **Commercial value.** Constellation geometry and station-keeping are trade secrets.
- **Maneuver intent.** A planned burn tells competitors and adversaries that something is changing.
- **Liability.** Sharing creates obligations.

Without that data, the screen in this paper sees catalog estimates where precise trajectories exist. Private screening lets two operators learn when their objects come close without either seeing the other's trajectory.

This section specifies the protocol, measures its arithmetic, and analyses what it reveals and how it can be abused¹⁴.

What is computed on ciphertext

A (the requester) and B (the responder) each hold a trajectory, sampled on a common time grid. A wants to know when B's object passes within R of its own. The protocol uses homomorphic encryption: B computes on A's encrypted positions without being able to read them^{11,12}.

1. **A encrypts its trajectory under its own key.**

- The quantities are x , y and z , in integer metres.
- 8,192 time steps are packed per ciphertext (BFV, $n = 8192$, 128-bit security).
- Three 60-bit plaintext moduli, recombined by the Chinese remainder theorem, hold the range, so no value wraps.

2. **B computes on ciphertext.**

- B squares A's coordinates itself: $\text{Enc}(|a|^2) = \text{Enc}(x)^2 + \text{Enc}(y)^2 + \text{Enc}(z)^2$, once per submission, shared by all of B's objects.
- Since $|a - b|^2 = |a|^2 - 2a \cdot b + |b|^2$, B then forms $\text{Enc}(|a - b|^2 - R'^2)$ from products of its own plaintext with A's ciphertext.
- A never supplies $|a|^2$. A false value would set the radius of every step's test, and the same value would pass the tube check (defense 2 below).

3. **Both learn the result.** A secure comparison tells A and B, for each step, whether B is within R' .

A pair can come within R between two samples. Testing against $R' = \sqrt{(R^2 + (v_{\text{max}} \Delta t / 2)^2)}$ catches every such approach. With $v_{\text{max}} = 15.5 \text{ km/s}$, R' is 9.2 km at 1 s steps and 77.7 km at 10 s steps.

Key custody.

- Only A ever decrypts, and no key leaves A.
- No third-party assessor is involved. A party holding the key that decrypts the result can also decrypt every input it receives under that key, so an assessor that decrypts adds trust without adding privacy.

Rules of the exchange

Fixed grid.

- Every encrypted trajectory uses a grid set by the protocol: UTC-aligned windows, a fixed step Δt , and 8,192 steps per encrypted block.
- Slot j always means $t_0 + j\Delta t$, so a requester cannot choose its own sample times.

One answer per pair and window.

- B answers once per (counterpart object, window).
- It checks every step of the window, never a subset.
- A accepts one digitally signed response per pair, so it cannot ask twice and keep the better answer.

Completeness audit. Two stages:

1. **Structure, before answering, without decrypting.** B checks:

- the window identifier and start epoch;
- the step;
- the number of encrypted blocks per coordinate ($\text{steps} \div 8,192$);
- the encryption parameters.

A short or malformed submission is refused.

2. **Content, under encryption.** A missing or zero-filled step is still a valid ciphertext, so structure alone cannot reveal it. The physics check (defense 1 below) tests every step, and a blank or invented step does not obey gravity.

What can be verified under encryption.

- **Before answering:**
 - the structure;
 - that every step obeys gravity, and that the trajectory continues the object’s previous window (defense 1 below);
 - where the requester’s object has a track, that every step lies inside the tube around it (defense 2 below).

Each check is arithmetic on ciphertext followed by the secure comparison. B learns only pass or fail.

- **After the window:** a requester with a public track opens the window; a requester without one answers challenges to its alerts (defense 3 below).

What it costs

Measured with Microsoft SEAL 4.1.1 on one core of the shared workstation¹⁴:

Step	R'	A: encrypt and upload, per object-day	B: compute, per pair-day	B: response, per pair-day	A: decrypt, per pair-day
1 s	9.2 km	179 ms, 34.6 MB	169 ms	8.7 MB	163 ms
10 s	77.7 km	22 ms, 4.2 MB	21 ms	1.0 MB	20 ms

- **Correctness.** Over 86,400 one-second steps for eight pairs, every step matched the plaintext distance.
- **Detection.** Each designed 2 km encounter raised three to six alert steps. A control pair 60 km apart raised none at 1 s steps.
- **Scale.** One private object screened against 1,000 objects for three days at 1 s steps costs B about 510 core-seconds and 26 GB of responses. Private screening suits operator-to-operator subsets; the open screen covers the catalog.
- **Scope.** The measurement had A encrypt $|a|^2$ as a fourth quantity, and it covers the screening arithmetic only. B’s squaring of $|a|^2$, the checks of defenses 1 and 2, the secure comparison and noise flooding come on top. The squaring and the physics check run once per submission, shared by all of B’s objects. The physics check multiplies ciphertexts to depth 3 (4 with J2), which needs larger encryption parameters than the screen.

What the result reveals

The intended output is the set of steps at which B is within R', and nothing more.

A naive design leaks the distance.

- The obvious way to hide the distance is to multiply each step’s value by a fresh random factor and let A read only its sign. One step’s value then bounds the distance only loosely.
- Neighbouring steps, however, have almost the same distance. Intersecting the bounds of 64 neighbouring steps, A recovered the distance within a factor of 2 at 87 % of steps (1 s grid). That is enough for range-only orbit determination of B’s object.

- The output must therefore be one bit per step: B adds a random additive mask, and A and B run a two-party secure comparison that reveals only the sign, to both of them¹³.
- B's response must also be noise-flooded, so that its ciphertext noise carries nothing about b.

Real conjunctions reveal what safety requires. Each alert places the other object within R' of a known object at a known time. In this paper's 3-day SGP4 screen, 2,565 objects outside Starlink came within 5 km of a Starlink satellite:

- 1,312 of them did so three or more times;
- 669 did so ten or more times.

Kuiper, OneWeb and Qianfan had 875, 481 and 591 such objects, mostly with one to four conjunctions each. A private object in a crowded shell therefore gives its densest neighbour enough position fixes for a coarse orbit within days. Private screening hides an orbit from operators it does not approach. It cannot hide it from those it does.

Orbit guessing

The probing attack:

- **The query.** A requester submits invented trajectories, rather than its real object, to find B's object. Each step then answers whether B is within R' of a chosen point at a chosen time.
- **A chosen radius.** If A supplied $|a|^2$, a false value would make each step test any radius around any point, which locates B without a prior. B therefore squares A's coordinates itself (step 2 above).
- **The cost.** Low Earth orbit (200–2,000 km altitude) holds about 1.27×10^{12} km³, and a 9.2 km ball is about 3,300 km³. A first hit with no prior takes about 3.9×10^8 steps: 47,000 one-window queries. At the measured rates, B pays about 760 core-seconds and 39 GB of responses.
- **With a prior.** If the prober knows the shell (± 25 km at 700 km), it needs about 1,200 queries, costing B 20 core-seconds and 1 GB.
- **After a first hit,** orbital dynamics narrow the search quickly.
- **In reverse.** The responder can probe the requester through its reactions. A disclosure request or a maneuver after an alert reveals proximity to an invented trajectory.

Unbounded querying is therefore cheap enough for a determined adversary. In mutual screening both operators can attempt it: each is a requester for its own object and a responder for the other's.

Who enforces the defenses.

- Every answer about an object comes from its owner's node, computed with its owner's plaintext trajectory. Nothing that could reveal the object exists anywhere else.
- The requester's software, modified or not, holds only its own data and results locked under its own key.
- The defenses below therefore run on the node of the party being probed. Open source does not weaken them; it lets each party check what its counterpart's node computes.
- A central service computing on everyone's encrypted orbits would reverse this, which is one more reason the protocol has none.

The defenses:

1. **Check the physics before answering.** This needs no track.
 - Each query names a registered object of the requester. Each object submits one trajectory, continued from window to window: a window's first k steps repeat the previous window's last k , and must match. When keys change per window, a key-switching key from the requester lets the responder compare them.
 - With the acceleration taken from second differences, $s_j = a_{j+k} - 2a_j + a_{j-k} \approx \ddot{a}(t_j)(k\Delta t)^2$, point-mass gravity requires

$$\mathbf{s}_j \times \mathbf{a}_j = \mathbf{0}, \quad |\mathbf{s}_j|^2 |\mathbf{a}_j|^4 = \mu^2 (k\Delta t)^4.$$

- The differences need only shifts and sums of ciphertexts, and both conditions are polynomials. The responder evaluates them with ciphertext products and tests them against a tolerance τ in the secure comparison. Multiplying through by $|a|^7$ keeps both polynomial with J2 included, at degree 16 instead of 6.
 - Around a declared burn, the check admits the declared Δv , drawn from a per-object budget.
 - The responder learns only pass or fail. Point clouds, steered sweeps and blank steps fail before any answer is computed.
 - The tolerance bounds the thrust a fake orbit can hide to about $\tau\mu/|a|^2$: with $\tau = 10^{-4}$ in low Earth orbit, 0.8 mm/s², about 70 m/s per day. Rounding to whole metres, the truncation of the differences and forces outside the check set how small τ can be.
 - A requester with K registered objects then learns what K satellites on its claimed orbits would.
2. **Answer only near the requester's track.** This applies where the requester's object has a track the requester does not control.
- **The track.** The public catalog's, or a certified private track: a tracker that already follows the object signs a copy of its track, encrypted under the requester's key. A certified track anchors a satellite missing from the public catalog without publishing it.
 - **The tube.** Under encryption, the responder tests that the submitted trajectory stays within D of the track. D is D_0 at the start and widens by $\Delta v_{\max} \cdot (t - t_b)$ after a burn declared at t_b , so planned maneuvers pass. The responder learns only pass or fail.
 - **The mask.** Wherever the responder's object b is farther than $D + R'$ from the track p, the answer is "far", whatever was submitted. An honest requester is within D of its track, so there $|a - b| \geq |b - p| - |a - p|$
 R' , and no honest answer changes. A masked step is a fresh encryption of a positive value, flooded like the rest, so the requester cannot tell it apart.
 - With a public track the responder computes the mask in plaintext, so it needs nothing from the requester. With a certified track it computes the mask under encryption, and a step is "close" only when both tests pass, in one secure comparison.
 - Probing with its real fleet, a requester reaches at most $D + R'$ around each of its satellites' tracks.
3. **Audit afterwards.**
- **A requester with a public track opens the window.** Each window is encrypted under a fresh key, and the responder keeps the ciphertexts it answered. After the window has passed, the requester hands over that window's key. The responder decrypts exactly what it answered: past positions, far less sensitive than planned maneuvers. The ciphertext is the commitment: it cannot be swapped afterwards, and no separate hash is needed. The responder checks every step and orbital dynamics, and agreement with independent tracking of the declared object, using the reference-orbit comparisons of section 7.
 - **A requester without one answers challenges.** Opening a window would expose its orbit. Instead, the responder may challenge any alert: the requester states, signed, its object's positions for a few minutes around it, which the responder already knows to within R' . The responder checks them against its own or a contracted tracker's observations, and an empty sky is a failure. This exposes a fake only where some tracker can see the stated position.
 - Queries, responses and alerts are digitally signed, so a failure is attributable evidence. Withholding the key or the positions counts as a failure.
4. **Stake and identity.** Only identities with stake or reputation may ask, and each object is registered, with stake, ahead of the windows it screens in. A failed check forfeits the stake and ends screening for that identity.
5. **Per-window budget.**
- An identity may ask no more questions per window than it has registered objects.
 - That bounds what a requester willing to lose its identity can extract before an audit catches it: one window of questions, inside the tubes of its own fleet or along the orbits it registered.
6. **Laplace noise with a safety offset.** Before the comparison, the responder adds noise drawn from a Laplace distribution of scale b to each step's distance. The noise is truncated at $\pm s$, and the test threshold is raised by s, so noise can add false alerts but never removes a real one.
- **It hides the precise position.** Two placements of the hidden object whose distances to the

probe differ by Δ change the odds of each answer by at most $e^{\Delta/b}$, and of k answers by $e^{k\Delta/b}$. This is differential privacy. With $b = 250$ m, ten answers barely separate positions 25 m apart.

- **Finding the boundary is expensive.** Locating the alert boundary to within δ takes about $(b/\delta)^2$ answers, more than the budget allows.
- **It does not hide coarse location.** An object 1,000 km away is never “close”. The noise complements the checks above; it does not replace them.
- **Repeats don’t help.** The noise is fixed per requester, object and step, so asking again does not average it away. Because of the truncation, the guarantee is approximate $((\varepsilon, \delta))$ differential privacy.
- **Cost.** The alert radius grows by s . With $s = 1.25$ km at $R = 5$ km, there are about 1.6 times as many alerts to resolve.

Direct authenticated streams protect integrity and metadata. The ciphertexts are protected by their owner’s key either way.

With these defenses in place, a prober learns what its real objects’ real close approaches reveal. A requester whose objects have no track learns what satellites on its claimed orbits would; a claimed orbit that no tracker follows cannot be told from a real hidden satellite until one of its alerts is challenged. If it is willing to lose its identity, a prober learns at most one more window of answers, within its own fleet’s tubes or along the orbits it registered.

9 Limits

Limit	Effect
SGP4 sampling	About half of a 19 s screen: 140 million SGP4 evaluations
HPOP propagation	Over 99 % of the HPOP screen, 0.65 CPU-seconds per object-day
Catch-up from element epochs	The first HPOP window waits for every object to reach the screen start: 36 % of the object-days
Memory per window	About 390 MB of trajectories per 1-hour HPOP window, indexed in 1.05 GiB of the module’s 2 GiB
HPOP coefficients	Drag from each element set’s B^* , $0.01 \text{ m}^2/\text{kg}$ where $B^* \leq 0$; one radiation-pressure coefficient for every object
One host	All timings from one shared 28-core workstation; the module gives the procedure for repeating them ²
Calibration coverage	Covariance calibrated only in LEO 600 to 800 km (empirical model; HPOP to 3 days, except 0.5 to 1 day), for 48 reference objects in one week
Probability inputs	Combined radius and covariance shape differ from SOCRATES’s unpublished ones

This paper reports computation speed, agreement between implementations and with SOCRATES on identical inputs, and covariance calibration where independent truth exists. It does not establish operational readiness or orbit accuracy; the companion paper¹ reports measurements of HPOP’s accuracy in its sections 7 and 10.

References

R1

Koury, A. Evidence-Supported ASO Catalog. Space Data Network technical whitepaper 2.0, 11 October 2026. Paper

R2

Edgesource. Conjunction assessment module: all-vs-all screening on a GPU or on CPU threads, time windows, TCA solve, parity and bound tests, and the benchmark procedure. Modules commit e743f2a279ab8754a55b509ccd770e3415aaf3e6; module 0.7.0, analysis/conjunction-assessment/dist/isomorphic/module.wasm, SHA-256 ca5b9898007b5067c88e51bb54de82004a5e0a6e114bdba59843677cd2b6c106. Files: analysis/conjunction-assessment/docs/gpu-all-vs-all.md, analysis/conjunction-assessment/docs/benchmark.m Private repository, available on request (tj@edgesource.com).

R3

Edgesource. Space Data Module SDK 0.8.26. Repository

R4

Alfano, S. Relating Position Uncertainty to Maximum Conjunction Probability. AAS 03-548, 2003. Paper

R5

ERFA (Essential Routines for Fundamental Astronomy), BSD-3, derived with permission from IAU SOFA. Vendored in the modules repository. Project

R6

Space-Track.org. General perturbations element sets, catalog snapshot of 30 September 2026. Service

R7

CelesTrak. SOCRATES conjunction screening service. Service

R8

Edgesource. Conjunction uncertainty program: reference states, GP prediction-error model, calibration gate, covariance probability on screened events, HPOP process noise, orbit-determination covariance, SOCRATES replay and screening evaluation. Modules commit 82a566fd2cafc062a43643c8855cc7fd371155cd. Files: analysis/reference-states/README.md, analysis/gp-error-model/docs/validation-2026-08.md, analysis/gp-error-model/docs/calibration-2026-08.md, analysis/conjunction-assessment/README.md, propagator/hpop/README.md, analysis/conjunction-assessment/docs/socrates-replay-2026-10-02.md, analysis/gp-error-model/docs/screening-evaluation-2026-08.md. Private repository, available on request (tj@edgesource.com).

R9

Foster, J. L. and Estes, H. S. A Parametric Analysis of Orbital Debris Collision Probability and Maneuver Rate for Space Vehicles. NASA JSC-25898, 1992.

R10

Precise orbit products: IGS final orbits (IGS), ILRS analysis-centre orbits (ILRS), Copernicus Sentinel-1 precise orbit ephemerides (ESA), Swarm precise orbits (ESA), and IERS EOP 20 C04 (IERS).

R11

Microsoft Research. Microsoft SEAL 4.1.1, homomorphic encryption library. Repository

R12

Fan, J. and Vercauteren, F. Somewhat Practical Fully Homomorphic Encryption. IACR Cryptology ePrint Archive 2012/144. Paper

R13

Damgård, I., Geisler, M. and Krøigaard, M. Efficient and Secure Comparison for On-Line Auctions. ACISP 2007, LNCS 4586.

R14

Edgesource. Private screening: protocol, SEAL benchmark, leakage and exposure measurements, exchange rules and defenses. Modules commit 22db2691620599403885cd1353677011e08f3d22. Files: `analysis/conjunction-assessment/docs/private-screening.md`, `analysis/conjunction-assessment/bench/private-screening`. Private repository, available on request (tj@edgesource.com).

R15

NASA Navigation and Ancillary Information Facility (NAIF), Jet Propulsion Laboratory. SPK Required Reading, sections “Type 2: Chebyshev (position only)”, “Type 3: Chebyshev (position and velocity)”, “Type 14: Chebyshev Polynomials — Unequal Time Steps” and “Type 20: Chebyshev (velocity only)”. Document; toolkit terms: SPICE rules

R16

Jet Propulsion Laboratory, Solar System Dynamics. Planetary and Lunar Ephemerides: export information. Page

R17

Folkner, W. M., Williams, J. G., Boggs, D. H., Park, R. S. and Kuchynka, P. The Planetary and Lunar Ephemerides DE430 and DE431. IPN Progress Report 42-196, 15 February 2014. Paper

R18

Deprit, A., Pickard, H. and Poplarchek, W. Compression of Ephemerides by Discrete Chebyshev Approximations. Navigation 26(1):1–11, 1979. doi:10.1002/j.2161-4296.1979.tb01350.x

R19

Trefethen, L. N. Approximation Theory and Approximation Practice, Extended Edition. SIAM, 2019, chapter 16 (Chebyshev projections and interpolants are near-best). doi:10.1137/1.9781611975949

R20

Song, M.-S., Park, S.-Y., Kim, Y. and Yim, J. R. Development of Kinematic Ephemeris Generator for Korea Pathfinder Lunar Orbiter (KPLO). Journal of Astronomy and Space Sciences 37(3):199–208, 2020. doi:10.5140/jass.2020.37.3.199

R21

Braun, V. Providing Orbit Information With Predetermined Bounded Accuracy. Thesis, 2017. doi:10.5281/zenodo.228164

R22

Rhodes, B. `jplephem`: reads JPL planetary ephemerides (SPK types 2 and 3). MIT License. Repository

R23

Nyx Space. ANISE: SPICE kernels, including Chebyshev SPK types 2, 3, 14 and 20. Mozilla Public License 2.0. Repository

R24

Pavlis, N. K., Holmes, S. A., Kenyon, S. C. and Factor, J. K. The development and evaluation of the Earth Gravitational Model 2008 (EGM2008). *Journal of Geophysical Research* 117, B04406, 2012. doi:10.1029/2011JB008916

R25

Park, R. S., Folkner, W. M., Williams, J. G. and Boggs, D. H. The JPL Planetary and Lunar Ephemerides DE440 and DE441. *The Astronomical Journal* 161(3):105, 2021. doi:10.3847/1538-3881/abd414

R26

Petit, G. and Luzum, B., editors. IERS Conventions (2010). IERS Technical Note 36. Technical note

R27

Picone, J. M., Hedin, A. E., Drob, D. P. and Aikin, A. C. NRLMSISE-00 empirical model of the atmosphere: Statistical comparisons and scientific issues. *Journal of Geophysical Research* 107(A12), 1468, 2002. doi:10.1029/2002JA009430

R28

Knocke, P. C., Ries, J. C. and Tapley, B. D. Earth Radiation Pressure Effects on Satellites. AIAA/AAS Astrodynamics Conference, AIAA 88-4292, 1988. doi:10.2514/6.1988-4292

R29

Lyard, F., Lefevre, F., Letellier, T. and Francis, O. Modelling the global ocean tides: modern insights from FES2004. *Ocean Dynamics* 56:394–415, 2006. doi:10.1007/s10236-006-0086-x

R30

Earth orientation and space weather: IERS Rapid Service/Prediction Centre finals2000A (IERS); GFZ Kp, ap and F10.7 (Matzka, J., Stolle, C., Yamazaki, Y., Bronkalla, O. and Morschhauser, A. The geomagnetic Kp index and derived indices of geomagnetic activity. *Space Weather* 19, e2020SW002641, 2021, doi:10.1029/2020SW002641); NOAA Space Weather Prediction Center, Report and Forecast of Solar and Geophysical Activity (SWPC).

R31

Hoots, F. R. and Roehrich, R. L. Models for Propagation of NORAD Element Sets. Spacetrack Report No. 3, 1980, sections 8 and 9. Report

R32

Fehlberg, E. Classical Fifth-, Sixth-, Seventh-, and Eighth-Order Runge-Kutta Formulas with Stepsize Control. NASA Technical Report R-287, 1968. Report

Edgesource. HPOP on the resident path: force model, integrator and export, reentry, equivalence and interpolation tests, catalog cost, the propagation farm, the full-catalog screen (with the conjunction module of R2's commit) and the HPOP covariance calibration. Modules commit e743f2a279ab8754a55b509ccd770e3415aaf3e6. Files: propagator/hpop/README.md, propagator/hpop/catalog.mjs, analysis/conjunction-assessment/docs/gpu-all-vs-all.md, analysis/conjunction-assessment/docs/hpop-full-force-screen-2026-10.md, analysis/gp-error-model/docs/hpop-calibration-2026-08.md. Private repository, available on request (tj@edgesource.com).