Vedaksha: A Clean-Room, Sub-Arcsecond Ephemeris Engine for Computational Vedic Astrology

Amit Gupta
ArthIQ Labs LLC
info@arthiq.net
ORCID: 0009-0008-9450-8134

2026-09-01

Technical report describing Vedaksha v9.1.0. Not peer reviewed.
Every measured figure herein is specific to that version.

Abstract

Vedaksha is a clean-room ephemeris and Vedic astrology (jyotish) computation engine written in Rust. It computes planetary positions from the analytical VSOP87A and ELP/MPP02 series and from the JPL DE440s SPK numerical kernel — validated against an independent JPL DE441 ephemeris served by the NASA/JPL Horizons system to a mean residual of 0.103 arcseconds (maximum 1.187 arcseconds) over 15,350 longitude comparisons across ten bodies spanning 1900–2025, re-measured for this paper, and cross-validated against two national-agency ephemerides from two different agencies (JPL DE440 and IMCCE INPOP21a) using a separate parity harness that consumes only the engine's published output. The analytical path, which needs no runtime data file, is separately characterised across 1550–2650 CE with the time scale removed from the comparison by addressing both sides in Terrestrial Time: mean residuals remain sub-arcsecond for the Sun, Moon, Mercury, and Venus across that entire span, while growing to several arcseconds for Uranus and Neptune at its extremes. The engine derives from these positions the constructs used in Vedic astrology — nakshatras, dashas, vargas (divisional charts), panchanga, shadbala, and muhurta — exposed through native Rust, Python, WebAssembly, and Model Context Protocol (MCP) interfaces. Every algorithm is derived directly from a cited primary source rather than from an existing implementation: the ephemeris and sidereal-conversion algorithms as shipped in the current release contain no source code, header, or documentation from the AGPL-licensed Swiss Ephemeris library, on which prior Vedic astrology software has commonly relied. (Three early sidereal-offset constants were briefly cross-referenced against a published Swiss Ephemeris header solely to identify a numeric value; all three were independently re-derived from cited primary sources in a subsequent clean-room pass, documented below.) This clean-room provenance is not merely asserted: it is documented per data source, with fetch dates and, where established, content hashes, in a public provenance ledger, and specific re-derivations carry dated, auditable clean-room evidence trails. Vedaksha is released under the Business Source License 1.1 (converting to Apache License 2.0 four years after each version's release) and is distributed via crates.io, PyPI, npm, GHCR, and GitHub.

1 Introduction

Software for computing Vedic astrology (jyotish) charts requires, as its foundation, a planetary ephemeris of at least sub-arcsecond precision. In practice, the dominant ephemeris used by existing Vedic astrology software — commercial and open source alike — is the Swiss Ephemeris library, which is dual-licensed under the GNU Affero General Public License or a separate commercial Swiss Ephemeris Professional license. A developer who wants to build or redistribute a Vedic astrology tool under a different license, or who wants to independently audit where a returned planetary position actually comes from, is left with two options: adopt Swiss Ephemeris's AGPL obligations (or pay for its commercial license), or hand-verify numerical tables against published references without any systematic provenance trail. Neither option gives a permissively-licensable, provenance-auditable alternative. Vedaksha was built to close that gap.

By "clean-room" we mean specifically this: every algorithm and constant in Vedaksha is derived from a primary source — a peer-reviewed analytical theory such as VSOP87A for the major planets or ELP/MPP02 for the Moon [Chapront and Francou, 2003], or a numerical kernel distributed by a national agency, such as the JPL DE440s SPK ephemeris — independently validated against a JPL DE441 ephemeris served by the NASA/JPL Horizons system, which is never itself ingested as a local kernel — and the algorithms and constants shipped in the current release contain no source code, header file, or documentation belonging to an existing proprietary or copyleft-licensed ephemeris implementation. The project's history is more instructive than a blanket denial: three early ayanamsha (sidereal-zodiac offset) constants were originally identified by reading a published header file of the AGPL-licensed Swiss Ephemeris library to find a numeric value — no source code was copied, and the numeric constants themselves are uncopyrightable facts — and all three were subsequently re-derived from independently cited primary sources in a 2026-08 clean-room pass that first removed the original values from the working tree, so the re-derivation could not see its own target. This claim is deliberately made verifiable rather than asserted on trust: `DATA_PROVENANCE.md`, committed at the root of the repository, lists every external data source the project ingests, together with its primary URL, license status, fetch date, and a content hash where one applies; and the `docs/audit/` directory carries dated, per-derivation clean-room records — including a firewall-enforcement transcript for the 2026-08-17 re-derivation of the ayanamsha (sidereal-zodiac offset) constants — that document what was and was not consulted at the time each derivation was made public. These records are public alongside the code, not retrospective summaries, so a downstream user can check the provenance claim against the same evidence this paper cites.

The remainder of this paper is organized as follows: Section 2 describes the ephemeris pipeline — the analytical and numerical position computations, reference-frame handling, ayanamsha derivation, and clean-room methodology; Section 3 reports independent numerical validation against a public ephemeris oracle; Section 4 describes the software's internal property-graph representation, its distribution channels, interfaces, and licensing; and Section 5 concludes. This paper treats the ephemeris foundation in depth; the derived jyotish (Vedic astrology) computation layer built on top of it — nakshatras, dashas, vargas, panchanga, shadbala, muhurta, and the rest — is described only at the interface level, as the set of MCP tools it exposes in Section 4, and is not given a separate methods treatment here.

2 Methods

2.1 Analytical ephemeris pipeline

Vedaksha computes planetary positions for Mercury through Neptune from the VSOP87A analytical theory [Bretagnon and Francou, 1988] and lunar positions from the ELP/MPP02 lunar theory [Chapront and Francou, 2003]. Both are evaluated as truncated trigonometric series against coefficient tables generated once from the IMCCE primary distributions and committed to the repository as generated Rust source, so this analytical pipeline needs no external data file at runtime. As measured directly (`find crates/vedaksha-ephem-core/src/analytical/`

`coefficients/ -type f | wc -l`), the committed coefficient tree currently holds 331 files, consistent with the 331-file, roughly 30,000-retained-term count recorded in `MAINTENANCE.md` §10. VSOP87A and ELP/MPP02 are evaluated in the dynamical equinox and ecliptic of J2000.0 as their respective authors defined it, using the corresponding IAU 1976 obliquity value of 84381.448 arcseconds rather than the IAU 2006 value used elsewhere in the engine, because rotating either series's output by the IAU 2006 value instead would introduce a 0.042 arcsecond frame error against the theory being evaluated.

2.2 Numerical ephemeris pipeline

As a second, independently implemented path behind the same `EphemerisProvider` trait, Vedaksha reads the JPL DE440s SPK kernel distributed by NASA's Navigation and Ancillary Information Facility [NAIF/JPL, 2021] through its own NAIF DAF/SPK binary-format parser and evaluates the kernel's Chebyshev polynomial position segments (SPK data type 2) directly, without depending on an existing SPICE toolkit implementation. Read from the kernel's own segment headers, DE440s covers 1850-01-01 to 2150-01-01; requests outside that span are rejected rather than extrapolated. The analytical pipeline above requires no data file and is used where one cannot be shipped, such as WebAssembly and edge/Cloudflare Workers deployments, while the SPK pipeline is offered as an installable sub-arcsecond ephemeris tier in the Python binding. The two pipelines are cross-checked against each other and, independently of both, against apparent geocentric ecliptic positions served by the NASA/JPL Horizons system, which serves the DE441 ephemeris rather than DE440s and is therefore a genuinely separate numerical source: a public web oracle used only to validate the two shipped pipelines' output, never a kernel Vedaksha reads locally.

2.3 Reference frame handling

Positions are converted from the J2000.0 mean equatorial frame to the mean equatorial frame of date by precession, then to the true equatorial frame of date by nutation, before rotation into the ecliptic of date. Precession uses the IAU 2006 model in its P03 four-angle Fukushima–Williams parameterization of Capitaine, Wallace & Chapront [Capitaine et al., 2003], composed as the rotation product $R_x(-\epsilon_A)\,R_z(-\bar{\psi})\,R_x(\bar{\phi})\,R_z(\bar{\gamma})$, in which — per `MAINTENANCE.md` §10 — the rotation order is the entire content of the expression: reversing it changes the result by a negligible 0.014 milliarcsecond at J2000.0 but by 0.56 arcsecond at 499 CE, an error invisible to any test that samples only the 1900–2100 window. This composition is checked, in the test `precession_matrix_matches_erfa_across_six_millennia`, against the IAU SOFA board's ERFA library (BSD-3; used here only as a numerical verification oracle, never as a source of shipped position data) via its `eraEqec06` routine at four epochs spanning 6100 years. Nutation in longitude and obliquity is evaluated from the IAU 2000B truncated 77-term series.

2.4 Sidereal (ayanamsha) conversion

Vedic astrology charts are cast against the sidereal zodiac, which requires subtracting an ayanamsha — the angular offset between the tropical and sidereal zero points — from a tropical longitude. The quantity subtracted is the *true* ayanamsha, the mean (precession-only) value plus nutation in longitude, rather than the mean value alone: the tropical longitudes it is applied to are apparent positions and therefore already carry nutation, so rotating them by a mean-only offset would leave an uncancelled nutation term of up to roughly 17–20 arcseconds in every sidereal position. Both quantities are exposed under explicitly named functions, since an engine that is silent about which one it returns invites comparison against the wrong column. Vedaksha implements eleven ayanamsha systems, each derived forward from a named primary source rather than reverse-engineered from another implementation's output. The normative derivation record is `docs/audit/2026-08-17-ayanamsha-cleanroom/spec.md`, with its machine-readable inputs in the same directory's `derivation-inputs.json`: the Indian official (Lahiri / Chitra-paksha) ayanamsha from *The Indian Astronomical Ephemeris 2022*; Fagan–Bradley from Fagan & Firebrace's *Primer of Sidereal Astrology*; Krishnamurti (KP) from Krishnamurti's *Krishnamurti Padhdhati Vol-I*; Raman from B. V. Raman's *A Manual of Hindu Astrology* (1935); Surya Siddhanta from the Surya Siddhanta itself (Ch. 3, vv. 9–12); Yukteshwar (also spelled Yukteswar) from Swami Sri Yukteswar's *The Holy Science* (1894); the star-anchored Revati-paksha and True Chitra ayanamshas, fixed respectively to ζ Piscium and Spica; Pushya-paksha from Narasimha Rao's published proposal; True Mula (Chandra Hari) from Chandra Hari's paper in the *Indian Journal of History of Science* 33(4), 1998; and a Galactic Centre ayanamsha anchored at 0° Sagittarius, whose primary is the VLBI absolute-astrometry determination of Sagittarius A*'s position and proper motion in the ICRF3 frame by Gordon, de Witt & Jacobs [Gordon et al., 2023] — a separate, more recent radio measurement of Sgr A* itself, and not part of the van Leeuwen/Hipparcos re-reduction used below for the star-anchored systems, since Sgr A* is not a Hipparcos catalogue star. Star positions for the star-anchored systems are taken from the van Leeuwen (2007) re-reduction of the Hipparcos catalogue [van Leeuwen, 2007] via VizieR, cross-checked against the original 1997 Hipparcos catalogue. Each ayanamsha is independently implemented twice — once in the Rust engine and once in a Python reference script that regenerates a committed fixture from the same input file — and the two agree to 5.7×10^{-14} degrees over 192 comparisons spanning 6000 years.

2.5 Clean-room methodology

Every source consulted for the pipelines above was fetched directly from its issuing scientific agency, with no intermediate copyleft-licensed codebase in the provenance chain: the VSOP87A and ELP/MPP02 coefficient data were fetched from IMCCE's and IMCCE/SYRTE's own FTP distributions respectively, and the DE440s SPK kernel was fetched from NAIF's own distribution server. `DATA_PROVENANCE.md`, committed at the repository root, records for each of these the fetch date, license status, and, where established, a content hash verified by the fetch script itself: the DE440s kernel and the JPL Horizons oracle each carry a pinned and verified SHA256, while the VSOP87A ledger entry instead records that a SHA256 ledger for that source is an open follow-up, since its IMCCE FTP primary is currently treated as byte-stable rather than independently hash-pinned. No Swiss Ephemeris source code, header file, or documentation was consulted at any point in deriving the VSOP87A, ELP/MPP02, or DE440s pipelines described in this section. The project's one documented instance of consulting a Swiss Ephemeris header — to identify three early ayanamsha numeric constants, all three of which were subsequently independently re-derived from the primary sources listed above — is confined to the sidereal-conversion layer and is discussed in Section 1 and in `DATA_PROVENANCE.md`.

3 Validation

3.1 Oracle construction

Vedaksha's numerical pipeline (`SpkReader`, reading the locally-shipped DE440s kernel) is validated against a committed oracle of 24,350 apparent geocentric ecliptic longitudes: 10 bodies (Sun, Moon, Mercury through Pluto) at 30-day steps from 1900-01-01 to 2100-01-01, fetched from the NASA/JPL Horizons system [Giorgini, 2015]. Horizons serves the DE441 ephemeris, a kernel Vedaksha never reads locally, so the comparison is against a genuinely independent numerical integration of solar-system dynamics rather than a restatement of Vedaksha's own input; it exercises the DAF/SPK Chebyshev evaluation and the whole apparent-place pipeline (light-time, aberration, precession, nutation) against an external source. Reference values are astronomical facts served by a public-domain US Government work, so using them as the project's independent reference raises no copyleft question against the Business Source License 1.1 position. The fixture, `tests/oracle_jpl/reference_positions.json`, was fetched 2026-07-16 and is committed and SHA256-pinned; regenerating it re-fetches from Horizons and compares hashes rather than trusting a stale copy.

One property of this construction is worth making explicit, because it strengthens what the residuals mean. Both sides of the comparison are addressed in UT: Vedaksha's `apparent_position` takes a UT1 Julian Day and converts to Terrestrial Time internally via its own ΔT implementation, while Horizons observer tables read the input as UT and apply their own. No ΔT correction is applied when building the fixture, and neither side is told what the other used. The agreement reported below is therefore between two pipelines that independently chose a time-scale conversion as well as independently computing positions — not between two position computations handed a common, externally supplied ΔT. That makes the sub-arcsecond figure a joint statement about the ephemeris and the time-scale handling together, which is the harder result; it is also precisely why the residuals diverge past 2026, where the two ΔT implementations stop reading a shared measured record and begin extrapolating separately (Section 3.4).

3.2 Re-measured residuals

The test `compare_against_reference` in `crates/vedaksha-ephem-core/tests/oracle_comparison.rs` was re-run for this paper on 2026-08-28 (`cargo test -p vedaksha-ephem-core --test oracle_comparison -- --nocapture`); Table 1 reports the residuals it measured and printed, split at 2026-01-01 — the date past which `delta_t.rs` stops interpolating measured IERS ΔT values and starts extrapolating them. The 2026–2100 mean is not printed by the test itself (which asserts only a per-point ceiling for that era); it was computed for this paper directly from the same run's 24,350-row per-comparison output, and the resulting max independently reproduces the value the test itself printed for that era (n follows by construction: 24,350 total minus 15,350 measured-era comparisons, with zero skipped).

Era	n	Mean residual (arcsec)	Max residual (arcsec)
1900–2025 (measured ΔT)	15,350	0.103	1.187 (Uranus, 1948-06-18)
2026–2100 (predicted ΔT)	9,000	2.200	44.912 (Moon, 2099-02-06)

Table 1: Residuals between Vedaksha's `SpkReader` (DE440s) and JPL Horizons (DE441), 24,350 comparisons total, re-measured 2026-08-28. The same test run also prints an unsplit, whole-fixture mean of 0.878 arcseconds across all 24,350 comparisons, which a reader can verify against the two per-era figures above by the comparison-weighted arithmetic $(15{,}350 \times 0.103 + 9{,}000 \times 2.200)/24{,}350 \approx 0.878$.

The 1900–2025 figures — mean 0.103 arcseconds, max 1.187 arcseconds over 15,350 comparisons — are what the sub-arcsecond accuracy claim in this paper's abstract rests on, since ΔT over that span is the measured IERS record rather than a prediction. These figures differ marginally from the mean 0.106 / max 1.184 arcseconds recorded when the fixture's README was first written on

2026-07-16, because a since-released fix to Vedaksha's observer-position computation (not to the reference fixture itself) changed the computed positions being compared in the interim.

3.3 Cross-validation against a second agency

The residuals above measure the numerical (SPK) path against a single reference lineage: JPL, via the Horizons service. Agreement with one agency's ephemeris leaves open whether the engine agrees with the underlying dynamics or merely with that agency's realization of them. It also leaves the *analytical* path — the VSOP87A/ELP-MPP02 pipeline that requires no runtime data file, and which is what makes the WebAssembly and browser distributions possible — unmeasured against any external reference in the account above.

Both gaps were addressed by a measurement performed in a separate parity-testing project that consumes Vedaksha only through its published interfaces and holds its reference kernels locally rather than querying any service. That measurement compared the analytical path's tropical longitudes at 2,000 instants for nine bodies against two ephemerides produced by two different agencies:

Reference	Agency	n	Mean (arcsec)	Max (arcsec)
JPL DE440 (local kernel)	NASA/JPL	13,874	0.130	2.940 (Moon)
IMCCE INPOP21a (local kernel)	IMCCE	11,732	0.115	1.270 (Jupiter)

Table 2: Cross-validation of the *analytical* ephemeris path against two national-agency ephemerides. Counts are evaluable comparisons, which differ between references because each kernel's own coverage and body set differ.

Two things follow. First, the analytical path agrees with both references at a level comparable to the numerical path's agreement with Horizons (0.103 arcseconds), despite evaluating truncated trigonometric series rather than reading a numerically integrated kernel — so the data-file-free pipeline is not the weaker of the two in practice. Second, and more importantly for the claim this paper makes, the agreement is not specific to JPL: an ephemeris independently produced by IMCCE places the same bodies within the same tolerance. The residual against INPOP21a is in fact marginally smaller than that against DE440.

The provenance of these figures should be stated exactly. They were produced by a separate parity harness with its own toolchain and its own copies of the reference kernels, but that harness is maintained by the same authors as the engine. What is independent here is the methodology and the references, not the party: the harness reads only Vedaksha's published output and never its internals, and the reference ephemerides are produced by national agencies and are not under our control. The measurement itself is ours. The harness is public at `https://github.com/arthiqlabs/vedaksha-parity`, so the comparison can be re-run by anyone now: these figures are reproducible in fact, not merely independently derived.

3.4 Component-level validation of the SPK reader

The comparisons above are end-to-end: each exercises an entire pipeline, with the apparent-place machinery (light-time, aberration, precession, nutation) and a UT-to-TT conversion layered on top of the position computation itself. A disagreement anywhere in that stack appears as a longitude residual, which makes those figures the right measure of what a caller receives but a blunt instrument for locating error.

One component admits a sharper test. Vedaksha's SPK reader — the DAF/SPK binary parser and Chebyshev-segment evaluator written for this project rather than taken from an existing SPICE toolkit — is reachable directly, through an entry point that takes a TDB Julian Day and returns a raw state vector with no apparent-place processing and no time-scale conversion applied. Two implementations reading comparable kernels at the same TDB instant should therefore agree to within their shared numerical precision, with nothing else in the comparison at all.

Measured by the same parity harness, against two kernels' own raw segment evaluation: for nine of the ten bodies — the Sun, Mercury, Venus, Mars, Jupiter, Saturn, Uranus, Neptune and Pluto — agreement is sub-metre, and mostly sub-millimetre, against both references. The barycentre-target convention (Mars, Jupiter, Saturn, Uranus, Neptune and Pluto addressed as barycentres, matching the segments DE440s actually stores) was verified against both kernels before any residual was interpreted.

The tenth body carries a convention that is worth stating precisely, because it is not evident from the returned values and a caller who assumes otherwise will be wrong by thousands of kilometres. For the Moon, this entry point returns a position relative to the *Earth–Moon barycentre* — neither heliocentric nor referred to Earth's own body centre. The external measurement established this by explicit subtraction against a reference kernel's own EMB segment at several widely separated instants, rather than assuming it. Once that convention is applied, one of the two references agrees with Vedaksha to the same sub-millimetre precision as every other body; against the other, a small consistent residual remains, of the order of the Earth-to-EMB offset. That residual is reported here as measured and is not explained: accounting for it would require inspecting the internals of one implementation or the other, which this project's clean-room position does not permit, and it does not affect the apparent-position figures reported above, which are computed through a different path.

This result is independent of the longitude comparisons rather than a restatement of them. It validates a component the paper claims as independently implemented, in isolation, with the time-scale question removed entirely from the comparison.

3.5 Bit-level reproducibility is architecture-scoped; accuracy is not

Vedaksha evaluates both trigonometric series through a SIMD kernel, and that kernel does not produce identical bit patterns on every architecture: its four-lane double-precision operations map to a single 256-bit AVX2 register on x86-64 but to a pair of 128-bit NEON registers on aarch64, and the two paths round differently. The project maintains fixed-point digests over its raw series output for regression purposes, and those digests genuinely differ between the two architectures — they are pinned per-architecture rather than as a single value, and a build on an architecture with no pinned value fails loudly rather than silently skipping the check.

That is a real limitation on a reproducibility claim, and it is worth separating from a claim about accuracy, because the two do not follow from each other. To measure the second directly, the validation of Section 3.2 was re-run on both architectures at the same commit, against the same reference fixture:

The residuals are identical at the reported precision of one milliarcsecond, per body as well as in aggregate, while the underlying bit patterns are known to differ. The correct statement is therefore narrower than "reproducible" and stronger than nothing: results are bit-identical between builds of a given commit *on a given architecture*, and are equal to well below the accuracy floor across architectures. Nothing a caller consumes at the precision this engine claims depends on which of the two it runs on.

We report this because the distinction is easy to leave unexamined. A project asserting bit-

Architecture	n	Mean (arcsec)	Max (arcsec)
aarch64 (Apple Silicon)	13,815	0.239	1.896 (Neptune, 2007-02-09)
x86-64 (AVX2)	13,815	0.239	1.896 (Neptune, 2007-02-09)

Table 3: Analytical-path residuals measured on two architectures at one commit. Every per-body mean and maximum was also identical to the reported precision, and the worst case fell on the same body at the same date on both.

reproducibility without qualifying it by architecture would, on this evidence, be asserting something false; a project treating a per-architecture digest mismatch as an accuracy defect would be misreading the same evidence in the other direction.

3.6 Why residuals grow after 2026

The 2026–2100 residuals are markedly larger, and this is a time-scale disagreement rather than ephemeris error. Past 2025, both Horizons and Vedaksha's `delta_t.rs` must predict ΔT (the offset between Terrestrial Time and UT1) rather than read it from the measured record, and the two predictions diverge; that divergence appears as a longitude error proportional to each body's own angular rate, since the two sides are effectively evaluating the same position at slightly different UT1 times. This was re-verified for this paper, not assumed: at the single largest-error comparison, Moon at 2099-02-06 (44.912 arcseconds), the reference fixture's own longitude values were used to compute each body's angular rate near that date directly (a finite-difference of the committed `ref_longitude` values on either side of the date), and dividing the measured residual by that rate gives the implied time offset. For the Sun, Venus, and Mars — bodies slow enough that this 60-day-span finite difference resolves an unambiguous rate — the implied offsets are 67.0 s, 67.9 s, and 65.1 s respectively: a tight, freshly-computed cluster around ∼66–68 seconds, consistent with the ∼68-second ΔT divergence by 2099 already documented in `delta_t.rs` and `tests/oracle_jpl/README.md`. (The Moon and Mercury move too far within a single 30-day fixture step — the Moon by more than a full revolution, Mercury through a retrograde loop at this particular date — for the same finite-difference method to give a reliable instantaneous rate, so no implied-offset figure is reported for them individually here; the Moon's 44.912-arcsecond residual is nonetheless the dataset's largest precisely because its own angular rate — documented in `oracle_comparison.rs` as ∼0.64 arcseconds per second, by far the fastest of the ten bodies — converts the same few tens of seconds of ΔT disagreement into the largest longitude error of any body in the fixture.) ΔT prediction past the measured record is unpredictable in principle, so this divergence is expected to persist and is not treated as a defect; it is why `oracle_comparison.rs` asserts the two eras against separate error budgets rather than one bound covering both.

3.7 Accuracy across the full supported range, with ΔT eliminated

The entanglement described above — a residual outside the measured record being the sum of theory truncation error and ΔT divergence, in unknown proportion — was a limitation of the software's interface, not of the measurement. Engine version 9.1.0 adds `ecliptic_position_tt`, a TT-addressed analytical position entry point that performs no ΔT conversion at all, alongside the already TT-addressed raw SPK query. With both sides of a comparison addressed in TT, no ΔT model enters on either side, and the binding constraint becomes each kernel's stated coverage span rather than the looser window over which a ΔT model is considered reliable. The results below were produced by the same parity harness as the cross-validation reported earlier.

That the new path is genuinely conversion-free was verified before it was relied upon, and the check yielded a result of its own. Rather than take a ΔT value from an outside source, the implied ΔT was derived from Vedaksha alone, as the disagreement between the UT1-addressed and TT-addressed paths evaluated at the same instant, divided by each body's angular rate. The rate used was each body's true instantaneous speed rather than an average: Saturn was near stationary and retrograde at the test epoch, which an average-motion estimate misrepresents, and correcting this changed the result. Across seven bodies spanning three orders of magnitude of angular rate — from that near-stationary Saturn to the Moon at roughly 12 degrees per day — the implied ΔT values agree with one another to within 0.017 seconds, and coincide with the standard published value at J2000.0 of approximately 63.8 seconds, which is also what `delta_t.rs` tabulates. Two independent paths through the engine therefore differ by exactly one correctly-applied, correctly-signed ΔT and nothing else, at every angular rate the engine computes.

The sweep itself used DE440 only, and the exclusion of the second agency's ephemeris is worth stating because it is a coverage fact rather than a preference. INPOP21a's kernel covers 1897–2102, which does not extend meaningfully beyond the window the paper already reports, so it can contribute nothing to a distance-from-J2000 question; it was omitted rather than reported as though it corroborated a range it cannot reach. DE440's own coverage, 1550–2650 CE, was confirmed directly from the kernel rather than assumed, and bounds the sweep. Within it, 43 instants at 25-year spacing were compared, both sides addressed in TT, with the reference side computing an apparent place — light-time and aberration applied — against the same kernel, so that the comparison is like-for-like with Vedaksha's apparent ecliptic output rather than against raw geometry.

The reference side was addressed in TT throughout, with no TDB anywhere in the measurement, so the sub-millisecond TT–TDB distinction discussed below does not enter these figures. Table 4 gives the result.

Table 4: Analytical-path longitude residuals against DE440, by distance from J2000.0. Both sides addressed in TT, so no ΔT model enters on either side. Each cell is mean / max in arcseconds. Sample counts fall in the outer bands as the kernel's coverage margin is approached. The maxima are lower bounds on the true worst case, not measured worst cases: see the discussion of aliasing below.

Body	Distance from J2000.0 (years)					
	0–100 ($n = 8$)	100–200 ($n = 8$)	200–300 ($n = 8$)	300–400 ($n = 8$)	400–500 ($n = 6$)	500–600 ($n = 4$)
Sun	0.084/0.102	0.094/0.147	0.121/0.189	0.156/0.249	0.160/0.273	0.145/0.165
Moon	0.009/0.019	0.017/0.027	0.046/0.088	0.078/0.168	0.242/0.350	0.509/0.603
Mercury	0.084/0.119	0.103/0.157	0.124/0.246	0.138/0.227	0.177/0.353	0.151/0.225
Venus	0.080/0.102	0.091/0.136	0.115/0.186	0.135/0.221	0.196/0.358	0.109/0.144
Mars	0.065/0.113	0.115/0.302	0.221/0.646	0.366/0.758	1.259/3.794	2.299/6.277
Jupiter	0.178/0.372	0.324/0.755	0.450/0.639	0.413/0.828	0.476/0.870	0.531/0.591
Saturn	0.192/0.302	0.224/0.432	0.360/0.585	0.499/0.702	0.780/1.027	1.181/1.325
Uranus	0.539/1.215	0.582/2.270	1.399/3.132	1.895/4.587	3.209/6.583	5.484/10.145
Neptune	1.056/2.308	1.908/3.632	3.385/4.488	4.812/6.463	5.791/6.379	7.826/8.294

Truncation error grows with distance from J2000, and the growth is strongly body-dependent. In absolute terms it is largest for the outer planets: Neptune rises from 1.056 to 7.826 arcseconds in the mean, and Uranus from 0.539 to 5.484, roughly a factor of seven to ten across the sampled range. In relative terms the steepest growth belongs to the Moon, whose mean residual is the smallest of any body in the nearest band at 0.009 arcseconds and rises to 0.509 — a factor of almost sixty, though still sub-arcsecond at the farthest band. The Sun, Mercury, and Venus barely grow at all, remaining within roughly a factor of two across six centuries, and their apparent dip in the outermost band should not be read as improvement: with $n = 4$ that band cannot resolve a change of that size.

These figures are not directly comparable to the mean 0.239-arcsecond analytical residual reported earlier in this section, and the difference is not a discrepancy. That figure was measured against DE441 as served by Horizons, addressed in UT1, over 1,535 comparisons per body in 1900–2025; these are measured against DE440, addressed in TT, over at most eight instants per cell. The band figures rest on small samples and should be read as characterising the shape of the growth, not as competing precision claims.

The Mars anomaly, resolved: a synodic-period term the grid could not see

An earlier pass over this data showed Mars departing from the otherwise monotonic pattern, with its 400–500 and 500–600 year bands sitting well above a smooth extrapolation. That anomaly has since been resolved, and the explanation revises how the entire max column of Table 4 should be read.

Progressively finer sampling identified the cause. A five-year grid remained jagged and non-monotonic — the signature of aliasing rather than of a resolved signal — so a quarter-year scan was run across a single thirty-year window inside the anomalous band, 440 to 470 years forward from J2000, at 121 points. The residual there is not noise but a clean periodic oscillation: sharp peaks reaching about 4.5 arcseconds separated by broad troughs around 0.5 to 0.9 arcseconds. Measured directly from fourteen identified peaks, the mean peak-to-peak spacing is 2.154 years, with every one of the thirteen intervals falling between 2.00 and 2.25 years, a spread consistent with the scan's own grid resolution rather than with scatter. That matches the Mars–Earth synodic period, 779.94 days or approximately 2.135 years, to within about one percent. The identification rests on a published astronomical constant, not on any reference implementation's internals.

The consequence is that the original 25-year grid steps through roughly 11.7 synodic cycles at a time, and the five-year grid through roughly 2.3, so both sample the oscillation at an effectively arbitrary phase — which is precisely the aliasing both exhibited. This is also physically unsurprising: the geocentric angular error induced by a given heliocentric position error scales inversely with the Earth–Mars distance, so a residual governed by that distance should peak sharply near close approach and stay low between approaches.

Two consequences follow for the table, and the second is the more important. First, the band-level growth for Mars is real, not an artefact. Second, and applying to every body, *the maxima in Table 4 are lower bounds on the true worst case rather than measured worst cases*: a sparse grid rarely lands on a peak of a short-period term, so it systematically understates peak height. The other eight bodies were not tested for an analogous term in this cycle, and their maxima therefore carry the same caveat rather than a clean bill of health. Two questions remain open: whether the oscillation's amplitude envelope is itself smooth across the full 400–600 year range, only one thirty-year window having been densely scanned, and whether comparable short-period terms exist for other bodies at their own Earth-synodic periods — approximately 1.09 years for Jupiter and 1.035 for Saturn — which would be sampled even more coarsely by a 25-year grid.

3.8 Scope: what this validation does not cover

This oracle validates body longitudes only — not latitudes or distances, which are computed by the same pipeline but are not compared against the reference here.

The two analytical surfaces differ in which bodies they expose, and a reader citing either figure alone should know which. The `compute_natal_chart` tool serves nine bodies — the seven classical grahas from the Sun through Saturn, plus three lunar-node variants — and does not expose Uranus or Neptune at all, the Jyotish tradition it serves having no use for them. The TT-addressed `ecliptic_position_tt` entry point, by contrast, reaches every body the analytical theories model, Uranus and Neptune included, and the full-range sweep above accordingly covers bodies that no chart this engine produces will ever contain. Pluto is modelled by neither: it is absent from the analytical theories, and a request for it is rejected as an unknown body rather than reported as a date-range failure.

It also covers exactly ten physical bodies: the Sun, the Moon, and the eight planets through Pluto, at 2,435 dates each. **The lunar nodes are deliberately absent from it, and the sub-arcsecond claim above therefore does not extend to them.** This is a property of the reference rather than an oversight: a node is not an output of a numerical ephemeris integration but a derived geometric quantity — the intersection of the Moon's orbital plane with the ecliptic — so a DE-series or INPOP-series kernel cannot be queried for one at all, and computing a node requires an independent theory choice (a mean secular rate, or an osculating element carrying short-period terms) that no such kernel makes on the caller's behalf. Validating a node against a national-agency kernel is thus not merely unavailable but ill-posed.

Nodes are instead checked separately. The osculating node is compared against a small set of reference points obtained from Horizons in its orbital-elements mode — a different query type from the observer-table data underlying the main oracle — agreeing to within 1.7×10^{-4} degrees over the sampled dates.

The mean node admits a partial external check, and the distinction that makes it possible is worth drawing precisely, because it is easy to mistake for a contradiction of the ill-posedness argument above. What a kernel cannot supply is the definitional choice: which smoothing of the real, wobbling node counts as "the" mean node. What a kernel can supply, through its state vectors and standard orbital mechanics, is the osculating node of the actual orbit — a well-defined geometric quantity requiring no convention at all. And while the *difference* between a mean node and the physical osculating node is convention-dependent, it must in every convention be bounded and periodic, with no secular term: a smoothing of a quantity cannot drift away from it. The absence of secular drift is therefore testable without adopting anyone's definition, even though the value is not.

That test has been performed by the same parity harness as the cross-validation above. At 1,801 instants from 1900 to 2050, the Moon's Earth-relative state vector was read from DE440 and the osculating node computed from it directly, then expressed in two frames and regressed against time. In the properly precession-corrected mean ecliptic and equinox of date, the drift against Vedaksha's tropical mean node is -0.036 ± 2.107 arcseconds per year — within 0.05 standard errors of zero. The control, which repeats the identical computation in an uncorrected fixed J2000 ecliptic, returns -50.445 ± 2.109 arcseconds per year, within 0.3% of the standard general-precession rate and inconsistent with zero by more than twenty standard errors. The method therefore demonstrably detects a precession-scale frame error, and does not find one; the main result is smaller than the control by a factor of roughly 700. As a check that the pipeline measures the intended quantity, the residual spans roughly $\pm$7,000 arcseconds after correction, matching the known one-and-a-half to two degree libration of an osculating node about its mean.

This bounds a secular frame error in the mean node at roughly two arcseconds per year; it does not constrain the bounded periodic difference, and it is not an accuracy claim about the node's value. Node accuracy remains a weaker claim than the body-longitude figures reported above, and this paper does not assert sub-arcsecond agreement for any node. The true node is still covered only by internal consistency tests.

The oracle says nothing about the accuracy of house-cusp computation, which is not measured against any reference in this project. Two earlier test files, `comprehensive_comparison.rs` and `extended_comparison.rs`, once claimed that ground, but read fixtures generated from Swiss Ephemeris that had never existed in this repository's git history; both had therefore never executed a single assertion while reporting green, and were removed on 2026-07-16 rather than revived. Horizons does not serve house cusps, ayanamshas, nutation, or tithis, so reviving that coverage would have meant taking on a Swiss Ephemeris dependency — precisely what this project's clean-room position exists to avoid. Covering house cusps in the future would require a reference that is itself public-domain or derivable from first principles; that work has not been done, and this paper makes no accuracy claim about house-cusp computation as a result.

Finally, the validated window is narrower than the supported one. The engine accepts dates from roughly 2000 BCE to 3000 CE, and every accuracy figure in this section comes from 1900–2100, with the sub-arcsecond claim resting on the 1900–2025 portion where ΔT is read from the measured record. Extending validation across the full supported range is not simply a matter of running more comparisons, because two error sources become entangled outside the measured era: the truncation error of the analytical theories themselves, which grows with distance from J2000, and the divergence between two independent ΔT extrapolations, which is unbounded in principle. A residual measured out there is a sum of the two, and the paper would be unable to say in what proportion.

Separating them is possible by addressing both sides in a uniform time scale, so that no ΔT conversion enters the comparison on either side and the remaining difference is attributable to the theories alone. Until engine version 9.1.0 Vedaksha did not support this on the analytical path: every published entry point accepted UT1 and performed the conversion internally, and the ΔT value was not itself exposed, so an external party could not compensate for it either. That gap has since been closed by a TT-addressed position entry point, and the resulting full-range measurement is reported in Section 3.7. The sub-arcsecond claim in this paper's abstract nonetheless remains scoped to 1900–2025, where ΔT is read from the measured record; the full-range sweep characterises how truncation error grows outside that window, and does not extend the headline figure to it.

4 Software Architecture and Availability

Internally, a computed chart is represented not as a flat response struct but as a property graph: `crates/vedaksha-graph/src/ontology.rs` defines nine node types (chart, planet, sign, house, nakshatra, pada, geometric aspect pattern such as a Grand Trine or T-Square, dasha period, and fixed star) and twelve edge types connecting them (a planet is placed in a sign, occupies a house, aspects another planet, rules a sign, disposits another planet, is conjunct a fixed star or in a nakshatra, participates in a pattern, or is the lord of a dasha period, among others). This representation is exposed, not merely internal: the `emit_graph` MCP tool serializes the same graph to Neo4j Cypher, SurrealDB SurrealQL, JSON-LD, plain JSON, or retrieval-augmented-generation (RAG) embedding text, so a downstream agent or graph database can query a chart's facts directly rather than parsing a bespoke, chart-specific response schema.

The MCP surface is the primary agent-facing interface: 17 tools (`tools/mcp-tools.json`, engine version 9.0.0), covering natal chart computation, dasha periods, karakas, combustion, shadbala, divisional charts (vargas), the graph emitter above, transits and transit search, muhurta search, ashtakavarga, gochara, panchanga, drishti, bhavas, synastry, and composite charts, served over both stdio and HTTP with bearer-token authentication.

Vedaksha is distributed across four language surfaces and one container registry, all currently at engine version 9.0.0:

Surface	Install	Notes
crates.io	`cargo add vedaksha`	7 published crates (workspace also contains an npm-only WebAssembly crate)
PyPI	`pip install vedaksha`	Engine via WebAssembly, `py3-none-any`, Python $\geq$ 3.9, no Rust toolchain
npm	`npm install vedaksha-wasm`	Browser and edge, no data files
GHCR	`docker run ghcr.io/arthiqlabs/vedaksha-mcp`	Multi-arch (amd64 + arm64) MCP server image
GitHub	`https://github.com/arthiqlabs/vedaksha`	Source, issue tracker, `docs/audit/` clean-room evidence

Table 5: Distribution surfaces, engine version 9.0.0.

Vedaksha is released under the Business Source License 1.1. Per the repository's `LICENSE` file: the Change Date is four years from each version's release date, after which that version converts to the Apache License, Version 2.0; the Additional Use Grant permits any non-commercial use, with a one-time USD $500 fee per organization for commercial use, covering unlimited products and seats perpetually and applying to the licensed version and every version released afterward.

Software DOI: [to be assigned via Zenodo archival, pending].

5 Conclusion

Vedaksha computes planetary positions from the analytical VSOP87A and ELP/MPP02 series and from the JPL DE440s numerical kernel, and, re-measured for this paper against an independent NASA/JPL Horizons DE441 reference over 15,350 comparisons spanning 1900–2025, the numerical pipeline agrees with that independent source to a mean residual of 0.103 arcseconds (maximum 1.187 arcseconds); the larger residuals reported past 2026 were traced to ΔT prediction divergence rather than ephemeris error, and house-cusp accuracy remains explicitly out of scope for this validation. Every ephemeris and sidereal-conversion algorithm in the current release was derived from a cited primary source with no code, header, or documentation carried over from the AGPL-licensed Swiss Ephemeris library on which prior Vedic astrology software has commonly relied, and this clean-room provenance is documented per data source — with fetch dates and, where established, content hashes — rather than merely asserted. A reader wishing to verify the provenance claim independently, rather than take it on the authority of this paper, can do so from the public repository at `https://github.com/arthiqlabs/vedaksha`, which carries `DATA_PROVENANCE.md` and the dated clean-room evidence trails under `docs/audit/` at the repository root.

References

P. Bretagnon and G. Francou. Planetary theories in rectangular and spherical variables: VSOP87 solution. *Astronomy & Astrophysics*, 202:309, 1988.

N. Capitaine, P. T. Wallace, and J. Chapront. Expressions for IAU 2000 precession quantities. *Astronomy & Astrophysics*, 412:567–586, 2003.

J. Chapront and G. Francou. The lunar theory ELP revisited. Introduction of new planetary perturbations. *Astronomy & Astrophysics*, 404:735–742, 2003. doi: 10.1051/0004-6361:20030529.

J. D. Giorgini. Status of the JPL Horizons ephemeris system. In *IAU General Assembly*, volume 29, 2015.

D. Gordon, A. de Witt, and C. S. Jacobs. Position and proper motion of Sagittarius A* in the ICRF3 frame from VLBI absolute astrometry. *The Astronomical Journal*, 165:49, 2023. doi: 10.3847/1538-3881/aca65b.

NAIF/JPL. JPL DE440s SPK kernel (numerical planetary and lunar ephemeris). `https://naif.jpl.nasa.gov/pub/naif/generic_kernels/spk/planets/de440s.bsp`, 2021.

F. van Leeuwen. Validation of the new Hipparcos reduction. *Astronomy & Astrophysics*, 474:653–664, 2007.