

Interplanetary Time Attestation Under Light-Time Delay

Heime Jorgen

`heime.jorgen@proton.me`

kenosian <https://orcid.org/0009-0006-3356-9369>

Research Article

Keywords: space, propagation, ttts, golay

Posted Date: September 23rd, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11126395/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: The authors declare no competing interests.

Article

Interplanetary Time Attestation Under Light-Time Delay

Relativistic TCB Propagation, GNSS-Denied Positioning,
and Byzantine-Resilient Peer Consensus

Companion to. *Proof-of-Time: Byzantine-Resilient Temporal Ordering in Untrusted Networks*, SSRN abstract 6467220. This paper extends the temporal-ordering model to non-terrestrial links; it does not replace the core wire protocol.

Specification. TTTPS draft-06 and the companion `draft-helmprotocol-ttpps-deepspace-00`. The deep-space document is an informational profile intended for TIPTOP and DTN discussion.

Evidence status. Results reported here are source-level, deterministic simulation, and unit-test evidence. No interplanetary flight measurement or live deep-space mesh is claimed.

Availability. Implementation evidence is drawn from the `opentt-server` deep-space modules and the shared Hydra-KV `PotTier` definition.

Heime Jorgen^{1,*}

¹*Independent research.*

**heime.jorgen@proton.me*

Abstract. Time ordering becomes a different engineering problem when the propagation delay is comparable to, or larger than, the protocol's ordinary tolerance. This paper presents a deep-space companion profile for Proof-of-Time: a delay-tolerant tier, relative Barycentric Coordinate Time propagation, propagation-aware Byzantine aggregation, deterministic ephemeris and pulsar navigation interfaces, and a confidence path that distinguishes marginal source spread from inter-source dependence. The wire record remains unchanged; physical context and verification evidence are carried out of band. We trace the profile to an implementation containing integer-preserving TCB arithmetic, relative gravitational and velocity corrections, a hysteretic tier selector, peer-health routing, and a four-pulsar least-squares navigation solver. The reported tests establish algorithmic and simulated properties, including the expected growth of the tolerance envelope with one-way light time and recovery of a known position in a noise-free synthetic navigation fixture. They do not establish flight performance, live X-ray-detector operation, or an actually connected interplanetary mesh. The result is therefore a reproducible design and implementation profile, not a claim of deployed space-time accuracy.

Key result. Deep-space operation is treated as a verification context rather than a new timestamp format: local issuance, relative TCB propagation, OWLT accounting, peer aggregation, and evidence capture remain separate. A kernel or protocol result is not promoted beyond the assumptions and provenance that produced it.

Introduction

Interplanetary time attestation is not an attempt to make a remote clock local. It is an attempt to state, preserve, and verify what can be known about an event when the communication path itself has a large and variable light-time. The distinction matters for autonomous spacecraft, delay-tolerant relays, and any future system that must continue ordering events while disconnected from a terrestrial reference.

The contributions of this companion are fourfold. First, it defines a relative TCB propagation path in which OWLT, range-rate, gravitational potential, and velocity terms are explicit. Second, it gives a positioning hierarchy for GNSS-denied operation and records source authority rather than assuming it. Third, it combines light-time-adjusted peer projection with a Byzantine median and low-peer failure states. Fourth, it places Shannon, von-Neumann, and D3 signals in separate confidence layers so that confidence is not confused with cryptographic integrity or physical truth.

A long link is not a fast link with a larger timeout.

The original Proof-of-Time construction orders events using signed, hash-linked records and an ensemble of sources. Its ordinary tiers assume that a node can compare observations within a bounded operational window. A deep-space link violates that assumption without invalidating the ordering problem itself.

This companion makes the separation explicit. The core record and its cryptographic verification remain the protocol boundary. Propagation physics, position sources, peer health, and confidence are inputs to the verification decision and its evidence, not new fields silently inserted into the wire record. This preserves interoperability while allowing an implementation to choose a tier whose tolerance reflects the link.

The profile is informational. It describes an implementable path and its assumptions; it does not claim that a terrestrial implementation has measured an interplanetary link.

Why terrestrial assumptions fail

In a terrestrial setting, the delay between issuing and receiving a timestamp is often treated as a perturbation around a shared civil-time reference. In an interplanetary setting, the delay can dominate the tolerance itself. A node cannot determine from arrival time alone whether a remote event was recent, old, or generated during a disconnected interval. Retransmission makes this more

subtle: a correctly signed record can arrive after a newer record and can still be valid as a historical observation.

The absence of GNSS creates a second separation. Position is no longer a free input to a clock discipline algorithm. It is an observation with a source, uncertainty, epoch, coordinate frame, and authority status. A system that uses an unlabelled position estimate can produce a numerically precise correction whose physical applicability is unknown.

Problem definition

Let an observation o contain a local epoch, a peer report, a position context, and a provenance chain. The verifier must determine whether o can be compared with another observation o' under a declared coordinate-time model. This requires more than evaluating a signature:

$$\begin{aligned} \text{Admit}(o, o') = & \text{Integrity}(o) \wedge \text{Context}(o, o') \\ & \wedge \text{DelayBound}(o, o') \wedge \text{PeerPolicy}(o, o'). \end{aligned} \quad (1)$$

If any required term is unavailable, the result is HOLD or UNVERIFIABLE, depending on whether the missing evidence is a policy input or an authority measurement. This paper does not treat an unavailable term as zero.

Contribution boundaries

The contribution is a compositional profile and its implementation evidence. It is not a new astronomical timescale, an alternative to DSN navigation, or a claim that a Byzantine median removes relativistic or sensor error. The profile is useful precisely because those responsibilities remain visible: the physical layer estimates, the peer layer aggregates, the confidence layer qualifies, and the cryptographic layer attests what was actually observed.

The remainder of the paper follows that responsibility split. Sections 3–7 define the physical and peer model; Sections 8–9 define confidence and D3; Sections 10–12 describe the data flow, implementation, and experiments; and Sections 13–16 make the threat, failure, status, and reproducibility boundaries explicit. The final relationship section identifies which claims belong to TTTPS, Confidence-00, P2/D3, or K-Lean.

Relation to Prior Work

Proof-of-Time supplies the signed, hash-linked temporal-attestation primitive and source-ensemble model. TTTPS supplies the core wire and tier vocabulary. This paper adds only the non-terrestrial profile: propagation context, position authority, peer projection, delayed-link failure semantics, and their reproducible evidence. K-Lean supplies a separate formal-contract and certificate layer. Confidence-00 supplies the confidence and AdaptiveSwitch terminology, while the P2/D3 research supplies the higher-order dependence motivation. CCSDS, DSN, JPL ephemerides, and

pulsar catalogues provide external technical context, not deployment evidence from this work. The core temporal and delay-tolerant references are cited as [1, 2, 6], while the navigation reference artifacts are treated as [4, 5].

Proof-of-Time and the TTTPS core

The Proof-of-Time paper establishes the motivation for attesting temporal ordering when a channel controller can benefit from misordering events. Its contribution is the temporal-attestation primitive: signed records, linked history, source synthesis, and verification of the resulting record. TTTPS turns that primitive into a protocol context with wire fields, tiers, freshness, and integrity procedures. Neither work assumes that the physical propagation context is absent; rather, the core protocol does not need to encode every deployment-specific physical quantity in its compact record.

The present paper therefore extends the deep-space material introduced in the TTTPS draft-06 discussion while adding a fuller implementation and evidence boundary. A deep-space verifier must know how a timestamp was propagated and under which epoch and source assumptions it is comparable. Those facts are carried as verification context and evidence around the record. This avoids two opposite errors: treating the core wire record as a complete physical navigation solution, or treating a physical navigation estimate as a cryptographic attestation.

Formal verification and K-Lean

K-Lean is relevant as a model for machine-checkable contracts: a theorem that compiles under a pinned toolchain is evidence about the stated formal object. It does not certify that a prose spacecraft claim and a Lean declaration have the same identity. That distinction is also present here. A compiled TCB function, a passing signature check, and a physically applicable coordinate frame are different predicates. The paper retains the K-Lean separation between a formal certificate and the claim to which an operator attaches it.

Confidence-00 and P2/D3

Confidence-00 provides the Shannon/G-Score, von-Neumann correlation arm, AdaptiveSwitch, hysteresis, and insufficient-knowledge vocabulary. This paper uses those concepts only where delayed, sparse, or cross-round peer evidence requires a confidence decision. It does not repeat their full derivations.

The P2/D3 work contributes the observation that pairwise agreement can leave higher-order dependence unresolved. D3 is consequently placed above the physical and pairwise confidence layers as an informative diagnostic. It is not imported as a replacement for the TTTPS record, TCB propagation, XNAV, or Byzantine median.

Space-network context

CCSDS and DSN practice supplies the operational vocabulary for spacecraft time, navigation, and delayed links. JPL ephemerides and the ATNF pulsar catalogue supply reference data for the deterministic implementation tests. These sources define context and comparison material; they do not turn the reported synthetic tests into flight measurements. The distinction is central to the paper’s evidence policy.

Table 1 — What each cited layer contributes to this companion.

Layer	Contribution here
Proof-of-Time	temporal attestation and source-ordering problem
TTTPS	wire, integrity, tier, and freshness boundary
K-Lean	formal-contract and certificate discipline
Confidence-00	confidence and adaptive state vocabulary
P2/D3	higher-order dependence diagnostic
CCSDS/DSN/JPLspace	context and reference data

System Model and Assumptions

Actors and observations

An issuer creates a local event record, a relay forwards it, and an optional ground anchor supplies a reference context. A peer observation contains an identity, local epoch, range, range-rate, coordinate frame, and provenance. The relay is not trusted to improve the meaning of the observation. Every context field is retained even when the timestamp signature verifies.

Network and fault model

The channel is delay-tolerant rather than synchronously reliable. Loss and reordering are ordinary cases; duplication and replay are security cases. A Byzantine peer may report an out-of-bound value, withhold a value, or collude with another peer. A common sensor or ephemeris source remains a correlation risk even when software identities differ.

The system contains issuing nodes, relay nodes, and optional ground anchors. Each node has local proper time, a position context, a peer identity, and a signing key. Links may delay, lose, reorder, duplicate, or partition messages. Peers may be Byzantine. The baseline robustness assumption is $f < N/2$; a peer count alone does not prove independent clocks, independent sensors, or independent custody.

The peer topology is represented as a directed peer graph $G = (V, E)$, where each edge records an admissible observation path, source and receive epochs, and relay or custody context. A ground anchor is an authority-bearing source, not automatically an independent clock. A JPL-

verified or otherwise authority-backed ephemeris can establish position-source provenance, but it does not by itself establish peer independence or current applicability.

The model records a JPL or other authority status, position-source class, epoch, coordinate frame, peer set, and uncertainty. Low-peer and unavailable authority regimes enter HOLD, holdover, or recovery rather than silently asserting a stronger result. These are applicability conditions, not claims that the environment always satisfies them.

Observation lifecycle

An observation enters as an authenticated payload and is first assigned an execution and correlation identifier. Normalisation then binds its epoch, units, frame, and source identity. Only after these fields are available may the verifier project the observation across a peer link. A delayed observation can remain valid as historical evidence while being ineligible for a current freshness decision; validity and current applicability are separate states.

Assumption ledger

The profile carries an assumption ledger with at least the following entries:

1. coordinate scale, frame, epoch, and unit convention;
2. position and ephemeris source plus authority status;
3. OWLT estimate, round-trip observation, and asymmetry assumption;
4. peer population, fault bound, and independence classification; and
5. clock, range, velocity, and model uncertainty.

An absent entry is not interpreted as a benign default. It prevents a strong PASS and produces HOLD or UNVERIFIABLE according to the missing evidence.

Low-peer regime

The low-peer regime is a first-class operating condition. If the peer graph cannot supply the population required by the configured policy, the node may retain a bounded holdover value but must expose its age and reduced confidence. When a new peer becomes available, recovery creates a new transition linked to the holdover interval. It does not claim that the disconnected interval was retrospectively synchronised.

Light-Time and Coordinate-Time Model

Coordinate frame and epoch

TCB, TDB, and TT are recorded as distinct coordinate-time choices. The context also records epoch, BCRS/GCRS frame, ephemeris revision, and unit system. Two values from different frames cannot be compared merely because their serialized timestamps have the same width.

TCB is the relativistic coordinate-time scale used for the barycentric propagation model. TDB is a scaled barycentric time argument used in astronomical ephemeris practice, while TT is the terrestrial reference scale used for Earth-centred timing. They are related time arguments, not aliases: the profile requires the conversion convention, reference epoch, and frame to be declared before values are compared. A TDB-labelled ephemeris value is not inserted into a TCB chain by relabelling its unit, and a TT observation is not treated as barycentric without the corresponding transformation context.

The paper does not replace a standards-defined relativistic time transformation. It specifies the engineering boundary at which that transformation enters the profile and requires the selected convention and revision to be retained as evidence.

One-way light time and tier selection

Let d be the estimated one-way distance and c the speed of light. The one-way light time is

$$t_{\text{owlt}} = d/c. \quad (2)$$

For the deep-space tier the engineering tolerance is represented as

$$T_{\text{dspace}} = T_{\text{base}} + 2t_{\text{owlt}}, \quad (3)$$

where the two-way term is an operational bound for an exchange, not a claim that propagation delay is an error in the local timestamp. This distinction prevents the protocol from hiding a physical delay inside the record's local Error Bound field.

The implementation selects the deep-space tier when the observed exchange no longer fits the fixed terrestrial tolerance. It uses hysteresis: a downgrade requires several consecutive observations within a safety margin. This avoids flapping when a link is near a tier boundary.

Table 2 — Illustrative propagation scales used by the implementation tests.

Path	OWLT	Round-trip envelope
Earth–Moon	1.282 s	$T_{\text{base}} + 2(1.282) \text{ s}$
Earth–Mars, minimum	182.13 s	$\approx 6.1 \text{ min plus base}$
Earth–Mars, maximum	1337.59 s	$\approx 44.6 \text{ min plus base}$
100 AU	49900.48 s	$\approx 27.7 \text{ h plus base}$
163 AU illustrative	$\approx 22.6 \text{ h}$	$\approx 45.2 \text{ h plus base}$

The 163 AU line is an order-of-magnitude test anchor, not live spacecraft telemetry. The numerical scales are recomputed from astronomical distance and c ; they are not measured protocol precision.

One-way, two-way, and asymmetric observations

An observation must identify whether it is one-way, two-way, or a composed exchange. A one-way record carries an emission epoch and a reception epoch; the receiver can estimate the propagation interval only when the relevant

position and coordinate-time models are available. A two-way exchange adds a return path and therefore exposes a round-trip interval, but it does not by itself identify the one-way split. Processing time, queueing, asymmetric links, and clock realization must not be silently absorbed into OWLT.

For nodes i and j , with range r_{ij} and light speed c , the nominal one-way light time is

$$t_{ij}^{\text{owlt}} = \frac{r_{ij}}{c}. \quad (4)$$

If e_i is the emission epoch and r_j is the receiving epoch, the observation model is

$$r_j = e_i + t_{ij}^{\text{owlt}} + \delta_{ij}^{\text{clock}} + \delta_{ij}^{\text{model}}. \quad (5)$$

The clock and model terms are not estimated from distance alone. They require the declared coordinate-time relation, source epochs, and uncertainty inputs; when those inputs are absent the record remains observable but cannot support a strong propagation correction. An exchange model may then be written as

$$t_{\text{rtt}} = t_{ij}^{\text{owlt}} + t_{ji}^{\text{owlt}} + t_{\text{proc}} + t_{\text{queue}} + \epsilon_{\text{obs}}, \quad (6)$$

where the non-propagation terms are retained as separately typed evidence. The profile therefore does not infer a symmetric one-way delay from a round-trip measurement unless symmetry is an explicit, versioned assumption.

Delay, freshness, and local error bounds

OWLT is a physical propagation constraint. It is not interchangeable with the local uncertainty attached to an issued timestamp. The latter describes the precision and model uncertainty at the issuing node, whereas OWLT describes when information can become available at another node. A delayed record can be internally valid yet no longer eligible for a freshness-sensitive decision.

Accordingly, the implementation keeps three concepts distinct: the timestamp and its local Error Bound, the propagation interval used to project the timestamp, and the freshness/deadline policy applied at receipt. In particular, the OWLT term is not inserted into the fixed-width local Error Bound field. The projected interval and its uncertainty are carried in the propagation context and linked to the same evidence record.

Tier selection and hysteresis

The terrestrial tiers remain useful for short links, while the deep-space profile is selected when the observed exchange cannot satisfy the fixed terrestrial tolerance. The selector is deliberately conservative: a tier change is based on an observed violation or a stable sequence of safe rounds, not on a single noisy sample. The implementation's hysteresis parameters are configuration evidence

and are frozen with the run; they are not silently recalibrated from an audit result. In the reference selector, a downgrade requires three consecutive safe rounds within the configured safety margin. This prevents flapping; it does not remove physical uncertainty.

Table 3 — Evidence interpretation for propagation quantities.

Quantity	Role	Evidence class
Range/ephemeris/geometric input		derived or referenced
OWLT	propagation interval	computed, not measured
RTT	exchange observation	observed when instrumented
Processing/queue implementation delay	delay	observed or unavailable
Freshness decision	policy outcome	deterministic decision

Relativistic TCB Propagation

Model and relative variables

For a chain $A_0, \dots, A_k$, each hop contributes an integer interval, a relative correction, and a declared residual. Nominal time is composed separately from uncertainty. The accumulated bound is a sum of declared hop and model terms, not a claim that a long link has zero physical error.

Let τ be local proper time, Φ the gravitational potential, and v the velocity in a chosen reference frame. The first-order correction used by the profile is

$$TCB \approx \tau \left(1 + \frac{\Phi}{c^2} + \frac{v^2}{2c^2} \right). \quad (7)$$

For a hop from node A to node B , the implementation uses relative quantities:

$$t_{B,TCB} = t_{A,TCB} + (\tau_B - \tau_A) \left(1 + \frac{\Phi_{rel}}{c^2} + \frac{v_{rel}^2}{2c^2} \right). \quad (8)$$

The relative velocity is derived from range-rate, the time derivative of the measured range, with Doppler available as a cross-check. Absolute velocity is therefore not required for every hop.

This expression is a first-order engineering model around the declared frame and epoch. It is not a replacement for a relativistic orbit determination solution. The sign convention for the potential and velocity terms is fixed by the selected coordinate-time convention and is recorded with the run; a value copied between TCB, TDB, and TT contexts without that convention is not an admissible conversion.

The implementation preserves the integer timestamp in *i128* and applies the small correction as a floating-point factor before rounding the correction back to integer nanoseconds. This is essential: converting an absolute

nanosecond epoch of order 10^{18} directly to *f64* would erase precision.

Integer-preserving correction path

The correction path is intentionally asymmetric. The integer epoch remains the authoritative value, while the relativistic factor contributes a bounded delta. In pseudocode, the computation has the form

$$\Delta t = \text{round}_{\text{ns}} \left(\Delta \tau \left(\frac{\Phi_{rel}}{c^2} + \frac{v_{rel}^2}{2c^2} \right) \right), \quad t' = t + \Delta \tau + \Delta t. \quad (9)$$

This keeps the large epoch out of the lossy representation and makes the rounding point explicit in the evidence record. The record also carries the frame, epoch, potential source, velocity source, and model revision; a correction without those fields is not independently reconstructable.

Hop composition and bounds

For a route with hops $h = 1, \dots, k$, the projected coordinate time is composed hop by hop. The declared uncertainty is conservatively bounded by

$$E_{\text{route}} \leq E_0 + \sum_{h=1}^k (E_{h,\text{clock}} + E_{h,\text{range}} + E_{h,\text{model}} + E_{h,\text{round}}). \quad (10)$$

This is an engineering bound under the stated assumptions. It is not a statistical confidence interval and it does not imply independence between hops. Correlated model error is retained as a separate correlation class so that the confidence layer cannot count the same uncertainty multiple times.

Where an implementation has an evidence-backed decomposition of its errors, the bound may be tightened conditionally rather than by replacing the conservative expression globally:

$$E_{\text{route}} \leq E_0 + \sum_{h=1}^k E_{h,\text{bias/common}} + \sqrt{\sum_{h=1}^k (E_{h,\text{clock}}^2 + E_{h,\text{range}}^2 + E_{h,\text{jitter}}^2)}. \quad (11)$$

The root-sum-square term is admissible only for declared independent or appropriately uncorrelated random components. Unknown correlation, shared ephemeris error, and common-mode bias remain in the additive conservative class. RSS is therefore an evidence-conditioned refinement, not a guarantee that long routes avoid an UNVERIFIABLE outcome.

Consistency tests

The test matrix exercises round-trip conversion, sign convention, epoch ordering, integer preservation, and repeated-hop composition. The tests use frozen fixtures and assert both the projected value and the declared bound. They do not validate relativistic flight hardware or replace a higher-fidelity ephemeris. A failed frame or

epoch precondition yields an explicit UNVERIFIABLE state rather than a silently corrected timestamp.

The principal regression is an epoch of order 10^{18} nanoseconds with a small relativistic correction. The integer-preservation assertion checks that the baseline epoch is unchanged by floating-point conversion and that only the explicitly rounded correction is added. A second regression composes multiple hops and verifies that each hop's frame, source, and residual remain recoverable in the final route record.

GNSS-Denied Positioning

Fallback and position uncertainty

The fallback order is a deterministic selection policy over available, authenticated inputs. Each change records whether the prior source was unavailable, stale, excessive in residual, or outside coverage. Range error δr and range-rate error $\delta \dot{r}$ affect OWLT and the velocity correction and are carried into the declared bound.

The selector does not merge incompatible frames merely because their position vectors have the same dimension. Every candidate is normalised to a declared frame and epoch before comparison. A source can be useful for coarse context while remaining inadmissible for a strong time correction if its residual, age, or authority evidence exceeds policy.

Position sources

GNSS is a preference where available, not a deep-space assumption. The profile's position-source order is: GNSS; X-ray pulsar navigation (XNAV); neighbour ranging and multilateration; and an analytic ephemeris. Every result records which source supplied the position and whether its authority check was available.

Table 4 — Position-source roles under GNSS denial.

Source	Contribution	Primary limitation
GNSS	local position and time	unavailable beyond coverage
XNAV	absolute celestial constraint	detector and catalogue evidence
Ranging	relative geometry	peer visibility and geometry
Ephemeris	model-based prior	model residual and authority age

The XNAV solver uses four catalogue pulsars and solves the four-unknown linearised time-of-arrival system by normal equations with partial pivoting. The deterministic test reconstructs a known synthetic position with error below 10^{-3} m in the noise-free case. This is a numerical property of the solver and test fixture, not a claim about the accuracy of a flight detector. The detector boundary is intentionally an interface: a real X-ray instrument must supply observations, and no such hardware measurement is claimed here.

Ephemeris and authority

The analytic ephemeris follows the published approximate planetary-element method and is checked at a recorded J2000 reference against frozen Horizons state vectors. The comparison is a reproducible fixture, not a live network lookup at build time. A separate Horizons path can return an unverifiable status when its feature or network authority is unavailable; it must not fabricate a verified position.

The offline path is consequently a reference-vector check, not an assertion that the analytic approximation is an authoritative spacecraft navigation solution. The authority-backed path records request epoch, response revision, frame, and source identity. If those fields cannot be retained, the result is not promoted to an authoritative position estimate.

Precision guardrail for analytic ephemeris

An analytic ephemeris may be adequate for a coarse propagation context while being inadequate for a nanosecond-scale relativistic correction. The profile therefore permits a precision guardrail. Let δr_{ephem} be a declared position-error bound, T_{base} the applicable local precision budget, and η a dimensionless policy parameter. If

$$\frac{\delta r_{\text{ephem}}}{c} > \eta T_{\text{base}}, \quad (12)$$

the position estimate MUST NOT support the precision-sensitive correction. The correction is suppressed or downgraded according to implementation policy, and the state is recorded as UNVERIFIABLE_POSITION_PRECISION; it is not silently converted into an authoritative position or a generic rejection. Coarse OWLT context may remain available if its own error contract is satisfied. The guardrail record must include the error source, authority status, frame, epoch, T_{base} , η , and the resulting disposition. No numerical planetary residual is claimed here without a frozen, authority-backed reference vector.

Source selection and stale-data handling

Position is treated as an evidence-bearing input rather than an implicit constant. A source candidate is admissible only if its epoch, frame, validity interval, and authority metadata are present. A stale or out-of-coverage source is not silently preferred over a lower-ranked valid fallback. Each fallback transition records the reason and preserves the prior source in the provenance chain.

For a range observation, the first-order contribution to light-time error is

$$\delta t_{\text{owlt}} \simeq \frac{\delta r}{c}, \quad (13)$$

while range-rate error contributes to the velocity-dependent coordinate-time correction. These terms are not interchangeable: a position source can be adequate for a delay estimate while inadequate for a relativistic

velocity correction. The verification record exposes both fields.

XNAV as an interface contract

XNAV is specified at two distinct levels. The protocol contract defines the observation format, catalogue identity, measurement epoch, detector uncertainty, and solver output. The reference implementation provides a synthetic solver test. It does not claim that a spacecraft has acquired the necessary photon-arrival measurements, calibrated an X-ray detector, or achieved the synthetic fixture's error. Missing detector evidence produces **HOLD** or **UNVERIFIABLE**, depending on policy state.

Position-to-time admissibility

Before a position estimate can influence TCB propagation, the implementation checks frame compatibility, epoch validity, source authority, and residual bounds. If any required check is unavailable, the estimate remains available as context but cannot support a strong time-attestation verdict. GNSS denial is therefore handled by a typed fallback graph, not by treating an approximate position as equivalent to a measured one.

Peer Aggregation and Byzantine Resilience

Sparse and partitioned regimes

When the peer set falls below the required population, the implementation enters **HOLD** or **holdover** and retains the age of the last accepted observation. Recovery requires a new linked observation and does not rewrite the interval during which the system lacked sufficient evidence.

Peer TCB values are projected to a common hop context, checked with a light-time-adjusted comparison, and aggregated with a Byzantine median. The usual condition is explicit:

$$f < N/2, \quad (14)$$

where f is the number of faulty peers and N the participating peer count. The median does not remove physical ranging error, create peers, or prove peer independence. Those are separate evidence dimensions.

Dual-regime consensus and anchored holdover

The profile distinguishes an ensemble regime from a sparse-link regime. In the ensemble regime, $N \geq 3$ is evaluated under the declared honest-majority condition $f < N/2$ and the median path. In the sparse-link regime, $N < 3$ does not provide Byzantine-fault tolerance; in particular, $N = 2$ implies $f = 0$ under that condition. A deployment MAY instead define an Anchored Holdover Estimator (AHE) as a separate degraded profile. AHE is not a second Byzantine consensus algorithm and does not convert an untrusted peer into an honest anchor.

In AHE, the last authority-backed synchronization is retained with explicit age, clock-stability, aging-drift, and

anchor-provenance evidence. A general error budget is

$$E_{\text{holdover}}(t) = E_{\text{sync}} + \left| \int_{\tau_{\text{sync}}}^t \delta_{\text{drift}}(u) du \right| + \delta_{\text{aging}}(t). \quad (15)$$

The drift model, aging term, units, calibration interval, and confidence interpretation **MUST** be declared with the oscillator evidence. If the bound remains within the declared deep-space tolerance and the anchor is still admissible, the state may be recorded as **PASS_DEGRADED_HOLDOVER**; otherwise it is **HOLD** or **UNVERIFIABLE**. This is a proposed profile boundary in the present reference paper, not a claim that the reference implementation has live AHE or that sparse links can tolerate a Byzantine peer.

The peer-health path combines roster capacity, an RTT trend estimator, and a relay error-correction layer. Its routing score records the dimensions that are actually measured and leaves neutral values where a live signal is absent. This is preferable to reporting a complete score from fabricated telemetry.

What the median guarantees

For projected scalar time estimates $x_1, \dots, x_N$, the aggregate is the ordered median of the admissible peer values. Under an honest-majority assumption, arbitrary values from at most f faulty peers cannot move the median outside the interval determined by the honest population. The useful claim is therefore bounded-resilience under an explicit model, not exact truth recovery. The output remains limited by honest-clock dispersion, position error, propagation-model error, and any common source shared by the peers.

The condition $f < N/2$ is evaluated after rejecting stale, malformed, or duplicate identities and after recording the roster's independence classes. It is not enough to count messages. A Sybil-expanded roster or several nodes fed by one anchor may satisfy a numerical count while failing the intended independence assumption.

Projection, identity, and aggregation order

Aggregation is performed only after each peer observation has been projected to the receiving node's coordinate-time context. The projection record binds the peer identity, emission epoch, receive epoch, estimated OWLT, frame, and residual bound. A median over unprojected local readings would mix physical delay with clock disagreement and is therefore not an admissible fallback.

The roster is also an epistemic input. A larger count does not imply a larger independent population when several peers share an anchor, ephemeris, relay, or software image. The implementation records declared independence classes and collapses common-source groups before interpreting the $f < N/2$ condition.

Each peer contribution is consequently a typed observa-

tion with an identity, source class, issuance epoch, receive epoch, frame, OWLT estimate, projected value, and residual. The aggregation record retains both the complete admissible set and the values excluded by policy. This makes a later audit able to distinguish Byzantine filtering from simple loss of connectivity.

Partition, holdover, and recovery

During a partition, the last accepted value may support bounded holdover but does not become a fresh observation. The state transition records the time of loss, peer population, last accepted residual, and the policy that selected **HOLD**. Recovery requires a newly linked observation whose provenance and light-time context are independently checked. Historical holdover is never rewritten as if the link had remained available.

Sybil, replay, and quorum transitions

Identity authentication prevents an unknown sender from entering the roster, but it does not by itself prove that authenticated peers are independent. A replayed observation is rejected or held according to its epoch and nonce policy; it is never counted as a new peer contribution. Quorum gain and loss are recorded as state transitions with the roster hash and observation window. This prevents a transient membership change from being mistaken for an improvement in clock evidence.

The recovery path is conservative: after a partition, a node may continue in bounded holdover only within the declared age and drift limits. Once those limits are exceeded, the state remains non-promotable until a fresh linked observation is available. A successful recovery appends a new transition and preserves the partition interval for replay.

This append-only linked recovery chain is the software meaning of a “self-healing D-chain” in the companion design: it reconnects evidence and state after a disruption without claiming physical self-repair, removal of light-time delay, or retrospective synchronisation of the disconnected interval.

Confidence Gating

The marginal Shannon score answers how spread the source weights are. A von-Neumann-style arm evaluates correlation from aligned residual history. The two are complementary: a uniform source distribution can still contain correlated sources.

Let p_i be the normalized contribution weight of source i . The marginal entropy is

$$H(p) = -\sum_{i=1}^N p_i \log p_i, \quad G(p) = \frac{H(p)}{\log N}, \quad (16)$$

for $N > 1$, with the usual limiting convention for zero weights. G measures spread of declared contributions;

it does not prove that the sources are independent or physically correct. A single apparently diverse anchor can therefore produce a high marginal score while retaining a low evidential diversity.

For aligned residual vectors z_i , the correlation arm constructs a normalized positive semidefinite density matrix ρ from the declared residual structure and evaluates

$$S_{\text{VN}}(\rho) = -\text{Tr}(\rho \log \rho). \quad (17)$$

The construction is meaningful only when samples share a declared epoch window, frame, and alignment rule. Missing alignment is not interpreted as zero correlation; it yields an insufficient-evidence state.

The higher-order D3 diagnostic belongs above these layers. It is a residual between the observed joint structure and the maximum-entropy structure consistent with the available pairwise marginals. Its intended interpretation is:

$$\begin{aligned} &\text{pairwise-compatible but not pairwise-explainable} \\ &\implies \text{D3 residual.} \end{aligned} \quad (18)$$

D3 is not entropy, and it is not a replacement for OWLT, TCB, ephemeris, XNAV, or Byzantine median. It is not an existing production deep-space gate, not an automatic replacement for the VN confidence arm, and not a base conformance requirement. It must not automatically reject a record. In the current implementation D3 is not claimed to be wired as a production deep-space admission gate; it is an informative shadow diagnostic for unresolved higher-order source dependence, requiring aligned samples and an explicit source-independence classification.

Epistemic entropy for sparse deep-space evidence

The deep-space profile extends the VN point operator with an epistemic operator bundle. The purpose is not to claim that the bundle is a quantum state of the spacecraft; the density-operator representation is a bounded mathematical carrier for competing interpretations of the same telemetry. For observations $C = \{(A_k, b_k)\}_k$, define

$$\mathcal{B}(C) = \{R = R^\dagger \succeq 0 : \text{Tr } R = 1, \text{Tr}(RA_k) \leq b_k \forall k\}. \quad (19)$$

The Epi result retains both the entropy interval and the identity ambiguity:

$$\begin{aligned} \text{EpiEnt}(\mathcal{B}) &= [\inf_{R \in \mathcal{B}} S_{\text{VN}}(R), \\ &\sup_{R \in \mathcal{B}} S_{\text{VN}}(R)], \\ \Delta_{\text{id}}(\mathcal{B}) &= \sup_{R, R' \in \mathcal{B}} \frac{1}{2} \|R - R'\|_1. \end{aligned} \quad (20)$$

The exact lower endpoint is not silently presented as a real-time result: the operational certificate uses the pu-

rity inequality

$$S_{\text{VN}}(R) \geq -\log_2 \text{Tr}(R^2),$$

$$\underline{S}_{\text{cert}} = -\log_2 \left(\sup_{R \in \mathcal{B}} \text{Tr}(R^2) \right). \quad (21)$$

and records whether this is a certified bound or an exact endpoint.

The deep-space binding is deliberately physical. The input envelope carries TCB residuals, OWLT error, ephemeris error, alignment, applicability, and authenticated provenance. Missing alignment or an inapplicable position context is therefore an epistemic hold, even when a numerically tidy point operator would produce a high *G*-Score. A sparse peer set is not padded with synthetic votes: after provenance collapse, $N_{\text{eff}} < 3$ selects the anchored holdover regime and preserves an explicit drift and age bound. Conversely, a shared ephemeris is retained as common provenance and is not counted as independent support.

$$N_{\text{eff}} = |\{g(i) : i \in I_{\text{admitted}}\}|,$$

$$N_{\text{eff}} < N_{\text{min}} \implies \text{AHE}(E_{\text{sync}}, \sigma_{\text{clock}}, \gamma_{\text{drift}}),$$

$$\neg \text{aligned} \vee \neg \text{applicable} \implies \text{HOLD}_{\text{Epi}}. \quad (22)$$

The Epi branch is attached after admission, provenance collapse, freshness, quorum, median, and VN qualification. It can withhold a new PoT commit by returning an epistemic hold, but it does not rewrite the PoT wire format or turn D3 into a normative gate. This is the minimal composition that closes the pre-Epi gap: a clean VN point estimate is no longer sufficient when the physical context is missing, the peer population is sparse, or multiple provenance-compatible interpretations remain admissible.

Table 5 — Pre-Epi deep-space gaps and their Epi-aware disposition.

Gap	Epi-aware disposition	Evidence class
Unbound ephemeris OWLT	precision guardrail; coarse or unverifiable context, no false nanosecond correction	tested policy
Sparse peers	provenance collapse and AHE; no fabricated quorum	Rust/sim.
Missing alignment	explicit Epi hold; no zero-correlation inference	Rust test
Shared source	common provenance retained; N_{eff} reduced	hybrid test
Exact entropy infimum	purity certificate recorded	computational contract
D3 calibration gap	informative shadow only; no automatic rejection	offline shadow

Independent-source accounting

Source independence is a provenance property as well as a statistical one. The confidence layer records shared anchors, ephemerides, relays, software images, and observation windows before interpreting entropy or correlation. Sources with a common origin are collapsed into a dependency group for the independent-support calculation. The nominal peer count and the collapsed independent count are both retained so that a confidence result cannot hide the dependency reduction.

Confidence, integrity, and applicability

Confidence describes the state of the available evidence; integrity describes whether the record and its transition are internally protected; applicability describes whether the evidence supports the particular claim and scope under evaluation. A high confidence score cannot repair a missing provenance link, and a cryptographically intact record cannot establish that it applies to a different physical regime. The gate therefore retains these dimensions as separate fields and emits **HOLD** when a required dimension is absent. The explicit “InsufficientKnowledge” state denotes evidence too incomplete to distinguish safe acceptance from a proved incompatibility; operationally it maps to conservative **HOLD** or **UNVERIFIABLE** according to the missing authority or policy input.

Adaptive transitions

The AdaptiveSwitch changes operating mode only on an explicit evidence transition. It may enter conservative holdover when source diversity, residual history, or applicability falls below policy, and may return to active aggregation after the required stable observations. Each transition references its predecessor state and source artifacts. This makes recovery auditable and prevents a transient confidence spike from silently promoting a previously unresolved record.

Decision policy

Confidence values qualify a transition; they do not replace the structural or epistemic predicates. A strong score with missing applicability remains **HOLD**, while a low score is not by itself a contradiction. The final state combines confidence, integrity, provenance, and scope under the frozen policy version. Every change of mode records the pre-state, the measured inputs, the policy revision, and the resulting post-state.

Higher-Order Dependence and D3

Calibration boundary

The absolute-value D3 attempt and signed calibration are retained as separate generations. The signed result demonstrates calibration on its synthetic design; it does not establish that every physical collusion pattern is identifiable. A residual can request more evidence or a conservative transition, but cannot alone establish a physical

contradiction.

Why pairwise correlation is insufficient

Pairwise correlation describes the two-source projections of a joint distribution. It does not determine the joint distribution in general. For binary variables, let X_1 and X_2 be independent fair bits and define $X_3 = X_1 \oplus X_2$. Each pair is marginally independent, while the parity constraint is exact. A verifier that inspects only pairwise residuals can therefore miss a three-way dependency.

The D3 diagnostic compares the observed joint structure P with the maximum-entropy distribution $P^{(2)}$ satisfying the measured pairwise marginals. One residual form is

$$D_3 = D(P \parallel P^{(2)}), \quad (23)$$

where D is a declared divergence or calibrated signed statistic. The quantity is meaningful only relative to a sampling window, estimator, regularisation rule, and source-dependency class recorded in the artifact.

The von-Neumann arm uses pairwise residual structure. It cannot distinguish all joint distributions with the same pairwise marginals. A parity construction with $X_3 = X_1 \oplus X_2$ is a compact counterexample: every pair can look independent while the joint relation is not. D3 measures the residual between the observed joint structure and the maximum-entropy structure compatible with the observed pairwise marginals.

The recorded absolute-value D3 candidate failed its calibration target ($R^2 \approx 0.133$ and 0.060). A signed third-order cumulant diagnostic reached quadratic $R^2 = 0.994027$ and 0.975254 in an isolated synthetic calibration, with bootstrap lower bounds 0.991293 and 0.917285 . These numbers are calibration evidence only. D3 is not light-time correction, XNAV, Byzantine median, or a base conformance requirement; it remains an informative shadow diagnostic and cannot by itself produce a hard rejection.

Signed calibration and operational use

The signed calibration retains the direction of the third-order residual before fitting the relationship used by the diagnostic. The calibration protocol freezes the synthetic generator, sample count, noise model, seed, estimator, and fit family before observing the result. The reported R^2 and bootstrap intervals therefore describe that fixture and do not constitute a universal dependence detector.

In operation, D3 is an evidence request and routing signal. A stable residual may lower confidence, trigger additional aligned rounds, or preserve HOLD; it cannot establish that a peer is Byzantine and cannot turn a formal or physical mismatch into a contradiction. Its output is linked to the same context, residual history, and transition record as the lower-order confidence arms.

Protocol and Data-Flow Architecture

Execution stages

The pipeline is deliberately ordered. Source acquisition precedes physical normalisation; normalisation precedes peer comparison; confidence qualifies the result; and cryptographic issuance records what was actually accepted. An implementation may stop at any stage, but it must retain the reason and the input artifact for that stop.

Each stage consumes and emits a typed context rather than an unlabelled timestamp. The context carries the execution identifier, observation epoch, coordinate frame, source class, and current evidence state. Physical stages may add position, range, or propagation terms; epistemic stages may qualify applicability and source dependence; the attestation stage records only what the preceding predicates actually established.

Core/profile boundary

TTTPS core fields remain unchanged. The profile adds an out-of-band context identifier and evidence references rather than silently changing record size or semantics. This permits a core verifier to validate the record while a deep-space verifier additionally evaluates applicability and propagation.

This separation supports two verifier modes. A core verifier can validate the cryptographic temporal record without implementing deep-space navigation. A profile verifier can require the additional context and return a qualified state when that context is absent. Neither verifier should infer the other layer's result from a field that it does not understand.

Figure 1 separates the physical, epistemic, and cryptographic responsibilities. The arrows are data dependencies, not claims that every deployment must use every optional source.




Figure 1 — Layered data flow from an event observation to an evidence-bearing disposition. Optional positioning sources enter the same context interface.

The complete profile is:

```
source -> position/time -> OWLT -> TCB -> peer aggregation
      -> confidence -> PoT -> integrity -> verification ->
                                quarantine/recovery
```

The TTTPS core remains the cryptographic temporal-attestation layer. The deep-space profile supplies physical context and delayed-link policy around that core. A receipt binds the observation, ordering, and hashes of an evidence package; it does not certify the physical claim contained in it.

The operational ordering is also a safety boundary. Position and epoch normalisation precede OWLT projection; projection precedes TCB composition; peer aggregation precedes confidence qualification; and issuance follows the qualification result. A failure at an earlier stage cannot be hidden by a successful later cryptographic operation. The receipt may preserve a failure artifact, but it does not transform it into an attested success.

Profile negotiation and failure routing

The profile is negotiated from declared capability and observed link context, not inferred from a node name. A node advertises the coordinate frame, position-source classes, propagation model, peer capacity, and evidence version it can process. If negotiation fails, the record is routed to **HOLD** or **UNVERIFIABLE**; it is not downgraded to a terrestrial profile merely to obtain a **PASS**.

Evidence and receipt boundary

The evidence package is append-only at the logical level. It contains the source observation, context projection, peer inputs, confidence decisions, attestation bytes, and final disposition, each identified by a content hash. TTTPS then seals the package’s observed ordering and digest. The receipt proves package continuity and observation time; physical correctness, applicability, and claim identity remain separate predicates.

The package is suitable for delayed delivery and replay. A receiver can verify the digest and reconstruct the recorded transition even when the source is no longer reachable. If a required physical or authority artifact is unavailable, the package remains evidence of an incomplete evaluation and is labelled accordingly. This is the intended distinction between durable observation and verified physical truth.

Reference Implementation

Module responsibilities

ctx binds position, velocity, potential, epoch, and source identity. **tcb** performs integer-preserving conversion. **ephemeris** and **jpl** provide offline and authority-backed position paths. **pulsar** solves the XNAV system, while **median** and **peer_health** handle peer projection and selection. **residual_history** supplies aligned cross-round correlation, and **verify** orchestrates the final predicate and evidence result.

The module boundaries are also evidence boundaries. **ctx** is the single carrier for frame, epoch, source, and uncertainty metadata; **tcb** must not reconstruct that metadata from a serialized integer; and **verify** must not infer a successful physical context from a cryptographic signature alone. The navigation modules can therefore be tested independently while the integration path checks that their provenance is preserved end to end.

Table 6 — Reference implementation components and evidence classes.

Module	Responsibility	Primary evidence
ctx	frame, epoch, source context	unit/integration
tcb	integer-preserving projection	deterministic test
ephemeris/jpl	position model and authority	reference fixture
pulsar	XNAV solver interface	synthetic test
median	peer projection/aggregation	simulation
peer_health	roster and link state	simulation
residual_history	aligned dependence data	benchmark
verify	predicates and disposition	integration

Build and dependency boundary

The deep-space tier is imported from the shared Hydra-KV type rather than redeclared in the server. A reproducible build therefore pins both the server revision and the linked dependency revision. Feature-gated authority calls must report unavailable when the feature or network is absent.

The implementation is organised in a deep-space module with context, TCB, ephemeris, JPL authority, pulsar navigation, peer median, peer health, and verification orchestration components. The shared Hydra-KV type carries the deep-space tier, so the server does not define a second incompatible tier enum.

The shared type is a compatibility constraint: adding a deep-space profile must not create a second tier enum or silently change the core wire contract. The server records the selected profile and its context reference alongside the core attestation, allowing an older core verifier to reject or ignore unsupported profile context explicitly rather than misinterpreting it.

Source-level review confirmed the following properties: the tier selector widens its tolerance from the fixed microsecond/millisecond-scale tiers to a propagation-dependent envelope; the hysteresis transition is tested; the ephemeris and XNAV calculations are real algorithms rather than empty success stubs; the peer-health path uses its roster, trend, and relay-code components; and the TCB conversion preserves the absolute integer time value.

The implementation record freezes the source commit, linked dependency revision, Rust toolchain, and **Cargo.lock**; these identifiers are part of the implementation claim rather than post-hoc publication metadata.

The build boundary is tested in layers. A locked dependency build establishes that the source and shared types

resolve consistently; library tests exercise the pure transformations; and integration tests exercise the composed evidence path. A pass in one layer is not substituted for a pass in another, and an unavailable feature-gated authority path is recorded as unavailable.

Operational boundary

The implementation separates the pure transformation path from authority and transport adapters. Deterministic unit tests run without network access; authority-backed inputs are represented by frozen fixtures when reproducible evaluation is required. A live adapter that cannot provide its source, revision, or measurement epoch returns an explicit unavailable result rather than an internally generated substitute.

Reproducible command surface

An artifact record names the repository revision, dependency lock, compiler, feature set, test command, input hashes, random seeds, and output hashes. The same record can therefore distinguish a source-level test, a synthetic navigation experiment, and an externally supplied observation. The command surface is part of the result: a successful binary without a reconstructable input and configuration is not a reproducible protocol claim.

The recorded full-pipeline and attack examples exercise the production functions in deterministic simulations. They do not establish live interplanetary deployment. In particular, the following remain explicit evidence boundaries: no live deep-space mesh was connected, no flight X-ray detector was operated, and peer independence is an assumption to be recorded and assessed rather than inferred from a peer count.

Experimental Evaluation

Measurement protocol

Each experiment freezes its input, configuration, seed, source revision, and expected evidence class before execution. Unit tests measure local functions; synthetic tests use constructed truth; simulations measure behavior under a declared model; and live black-box tests measure only the exposed endpoint. The denominator is never enlarged by silently treating an unavailable hardware or flight observation as a failure of the algorithm.

Evaluation groups

The nine groups are reported independently so that a passing TCB conversion does not mask a missing ephemeris authority, and an attack simulation does not become a live deployment claim. The end-to-end journey is a composition test, not evidence that a spacecraft followed the simulated route.

The evaluation is divided into nine independently reproducible groups: TCB conversion; OWLT propagation; ephemeris/reference-vector comparison; Byzantine

median under $f < N/2$; low-peer and partition cases; attack simulations; VN residual-history benchmarks; D3 calibration; and an end-to-end relay journey. Each result is tagged as IMPLEMENTED, SIMULATED, SYNTHETIC, LIVE BLACK-BOX, or HARDWARE-MEASURED. The current evidence uses the first four tags in different components and contains no hardware-measured result.

The deterministic examples call the production functions rather than a separate mock implementation. Their purpose is to establish code-path reproducibility and failure behavior. They do not establish interplanetary flight performance.

Table 7 — Evaluation result classes and current interpretation.

Group	Measured outcome	Evidence class
TCB/OWLT	deterministic conversion and scale checks	tested
Ephemeris	frozen reference-vector residuals	offline validated
XNAV	synthetic position recovery	synthetic
Median	honest-interval behavior under model	simulated
Partition	holdover and linked recovery	simulated
VN/D3	correlation/calibration behavior	synthetic
Relay path	context and disposition continuity	implemented/tested

The table reports what each experiment establishes rather than assigning one overall success label. Numeric values belong to the frozen result artifacts; where no hardware or live-link denominator exists, the corresponding claim is not made. This prevents a strong local test from being presented as evidence for an unmeasured physical deployment.

Reproducible evaluation and denominators

Each result is reported with its population and denominator. Retrieval coverage counts eligible source claims, whereas kernel-dispatch reduction counts candidates that reached the kernel boundary. An unavailable authority or detector is not silently included as an algorithmic failure. The ledger therefore reports NOT_MEASURED, HOLD, and UNVERIFIABLE separately from negative outcomes.

TCB conversion tests

Evidence tag: implemented/tested. The conversion tests cover epoch ordering, frame metadata, sign convention, integer preservation, and repeated application of a bounded correction. The principal invariant is that the large integer epoch is unchanged by floating-point conversion; only the explicitly rounded correction may alter the result. The test output includes the input epoch, correction terms, model revision, and output hash.

OWLT propagation tests

Evidence tag: **implemented/synthetic.** Frozen distances produce the illustrative one-way and round-trip scales in Table 2. Tests separately vary distance, processing delay, queue delay, and observation uncertainty. This distinguishes a propagation calculation from a measured radio exchange and checks that OWLT is not inserted into the local Error Bound.

Ephemeris and JPL validation

Evidence tag: **implemented/reference fixture.** The analytic ephemeris is compared at a recorded epoch with frozen reference vectors. An authority adapter is tested both when its response is available and when the authority path is unavailable. The latter must produce an explicit unavailable state; the experiment does not turn an offline approximation into a live JPL claim.

Byzantine median simulations

Evidence tag: **simulated.** The simulation varies N , the number of faulty peers f , honest residual spread, delayed observations, and arbitrary faulty values. It measures whether the aggregate remains within the honest interval under $f < N/2$ and records the failure boundary when that assumption is violated. It does not measure physical ranging error or prove peer independence.

Low-peer and partition simulations

Evidence tag: **simulated.** Peer loss, link partition, stale observations, quorum recovery, and bounded holdover are exercised as explicit state transitions. The measurements include time in holdover, peer count, last accepted epoch, and recovery linkage. A successful recovery creates a new observation transition; it does not rewrite the partition interval.

Attack simulations

Evidence tag: **simulated.** Forgery, replay, delay manipulation, Sybil inflation, collusion, common-source failure, stale ephemeris, compromised anchor, key compromise, and partition abuse are tested independently. Each scenario identifies the control under test and the attacks it cannot detect; for example, a receipt preserves history but cannot repair a compromised source, and a median cannot remove correlated honest-looking error.

VN residual-history benchmark

Evidence tag: **synthetic/simulated.** Aligned residual histories are constructed with independent, correlated, and partially missing source observations. The benchmark reports the marginal entropy arm separately from the density-matrix correlation arm, including the alignment window and source-dependency classification. Missing alignment is measured as insufficient evidence rather than zero correlation.

D3 calibration

Evidence tag: **synthetic.** The absolute-value attempt and signed third-order calibration are separate generations. The recorded absolute candidate reached $R^2 \approx 0.133$ and 0.060 , while the signed diagnostic reached quadratic $R^2 = 0.994027$ and 0.975254 on its isolated design, with bootstrap lower bounds 0.991293 and 0.917285 . These results quantify the fixture’s calibration only; they do not establish universal physical collusion detection.

End-to-end relay journey

The relay experiment composes the modules in operational order: a source event is timestamped, projected through a declared link, aggregated, qualified, and encoded as an attestation. The journey checks that each stage preserves the context identifier and that a later verifier can reconstruct the decision. The final routing stage records PASS, HOLD, UNVERIFIABLE, or quarantine before any recovery transition. It is a software integration experiment with synthetic or frozen inputs, not a measurement of an interplanetary radio link.

Security Analysis

Security versus applicability

Cryptographic authenticity answers who signed an observation. Applicability answers whether the observation can be compared in the declared physical context. A valid signature on a stale or wrong-frame position therefore does not imply a valid time correction.

The security objective is layered: preserve authenticity and ordering, prevent unauthorised sources from silently entering the peer set, expose physical-context failures, and avoid promotion when applicability evidence is absent. The objective is not to make an authenticated source truthful, remove light-time delay, or infer peer independence from roster size.

Collusion and common source

The median protects against a bounded number of arbitrary values; it does not detect every correlated honest-looking source. VN and D3 are diagnostic layers for this residual risk. When their evidence is insufficient, the correct outcome is a preserved HOLD or escalation, not an automatic rejection.

The threat model covers forged source values, replay, delay manipulation, Sybil identity inflation, collusion, common-source correlation, compromised anchors, stale ephemerides, key compromise, and partition/recovery abuse. Signatures and hash links protect record history but do not make a compromised source correct. OWLT and range checks constrain delay claims but do not remove common physical error. Median robustness requires $f < N/2$ and an adequate peer set. Authority failure must produce HOLD or UNVERIFIABLE evidence rather than a fabricated positive.

Threat classes and controls

Table 8 — Threat classes and corresponding control boundaries.

Threat	Control	Residual limitation
Forged source	signature and identity	compromised key
Replay	epoch/nonce and lineage	valid old observation
Delay attack	OWLT/range checks	asymmetric unknown delay
Sybil	authenticated roster	shared operator/source
Collusion	median and dependence diagnostics	common-source error
Stale ephemeris	age and authority gate	unavailable authority
Partition	holdover and recovery	prolonged isolation
Key compromise	rotation and provenance	pre-rotation history

Replay and delay manipulation

Replay protection binds an observation to its issuance epoch, source identity, and execution lineage. A retransmission may be retained as transport evidence, but it cannot silently count as a fresh peer sample. Delay checks compare observed timing with the declared propagation context; they do not pretend to recover unknown asymmetric delay. Insufficient timing evidence therefore remains **HOLD** or **UNVERIFIABLE**.

Identity, keys, and anchors

Roster authentication limits unauthorised membership changes, while source independence remains a separate provenance assertion. Key rotation creates a new identity epoch and preserves pre-rotation history; it does not retroactively validate observations signed before the relevant trust state. An authenticated but compromised anchor can still produce consistent values, so anchor dependency is recorded and exposed to confidence and applicability.

Security claims and non-claims

The claim is evidence-preserving containment under the declared threat model. The paper does not claim physical truth, detector authenticity, Byzantine consensus, or global attack detection without the required assumptions. A receipt establishes that a result was recorded in a given order; it cannot establish that its frame, ephemeris, or source model was correct.

Failure containment

Unexpected source transitions, inconsistent frames, replayed epochs, and quorum loss are routed to a bounded state with preserved evidence. New evidence may create a linked reassessment, but cannot erase the prior observation or relabel it retroactively.

Error Bounds and Failure Semantics

Bound composition

The declared bound is assembled from local clock drift, position error, range-rate error, ephemeris residual, correction-model residual, and hop accumulation. OWLT is retained as a propagation quantity and is not counted twice as local clock error. The bound is a contract for comparison, not a claim that the physical variables themselves were measured without error.

For a route of k hops, a conservative decomposition is

$$E_{\text{total}} \leq E_{\text{clock}} + E_{\text{position}} + E_{\text{range-rate}} + E_{\text{ephemeris}} + E_{\text{model}} + \sum_{h=1}^k E_{h,\text{residual}}. \quad (24)$$

The terms are retained by origin so that a shared ephemeris or anchor is not mistaken for independent random noise. Where only a deterministic bound is available, the result is a bound and not a confidence interval. Where a term cannot be estimated, the record cannot silently report a smaller total.

Table 9 — Error sources and their failure implication.

Source	Recorded quantity	Missing state
Clock	drift and local bound	HOLD
Position	frame, epoch, residual	HOLD/UNVERIFIABLE
Range-rate	derivative/Doppler residual	HOLD
Ephemeris	model revision and error	UNVERIFIABLE
Peer set	count and independence	HOLD
Hop route	per-hop residuals	UNVERIFIABLE

Figure 2 shows the intended failure semantics. In particular, **HOLD** and **UNVERIFIABLE** are terminal observations for the current execution but remain eligible for a linked recovery transition; neither is silently changed to **PASS** or **REJECT** by a retry.




Figure 2 — Disposition state machine. Reassessment creates a linked transition; it does not rewrite the historical observation.

A verifier should distinguish at least four outcomes: accepted under the stated evidence, rejected by a proved incompatibility, held because required evidence is unavailable, and unverifiable because an authority path failed. The profile must not collapse the latter two into a false negative or a false positive.

Each deep-space decision should bind the local input, context identifier, position source, peer set, OWLT estimate, TCB correction, confidence result, transition, and final receipt. A receipt seals observation and ordering of the evidence package; it does not prove that a physical claim is true. This is the same distinction as:

$$\text{formal validity} \neq \text{claim identity}. \quad (25)$$

Failure-state interpretation

PASS means that the declared predicate and its prerequisites were measured. REJECT means that a specified incompatibility or policy violation was established. HOLD means that evidence is insufficient for a strong decision, while UNVERIFIABLE identifies a missing or unavailable authority path. A retry creates a new transition linked to the original state; it does not rewrite history.

Transition rules

The state machine is monotone with respect to historical evidence. A PASS transition requires all mandatory predicates and admissible evidence; a proved incompatibility may produce REJECT; missing or ambiguous evidence produces HOLD; and an unavailable authoritative source produces UNVERIFIABLE. A subsequent observation can create a new transition from either non-success state, but it cannot change the state of the earlier execution.

The distinction between REJECT and UNVERIFIABLE is operationally significant. The former is evidence of an established failed predicate, while the latter records that the predicate could not be evaluated under the available authority. Likewise, a propagation delay that is large but correctly modelled is not itself a rejection; it becomes a rejection only when the declared policy or bound is violated.

Implementation Status and Limitations

This paper reports an implementation and verification profile. It does not claim that the system has achieved a particular absolute time accuracy in interplanetary space. The $1/\sqrt{N}$ intuition for aggregation is not a guarantee against correlated error. The Byzantine condition fails when the faulty population reaches the stated bound. XNAV requires an actual detector; an abstraction and synthetic solver test are not detector telemetry. A live JPL authority call can be unavailable. These limitations are part of the protocol evidence, not footnotes to be removed from a deployment claim.

The resulting architecture is deliberately layered:

1. physical context: OWLT, position, ephemeris, and relativistic terms;
2. peer aggregation: projection, ranging constraints, and median;
3. confidence: marginal spread and pairwise correlation;

4. higher-order shadow: D3 residual and escalation;
5. cryptographic record: unchanged Proof-of-Time issuance and receipt.

Table 10 — Implementation status and evidence boundary.

Item	Status
TCB / OWLT JPL / ephemeris	implemented and unit-tested implemented; offline reference validated
XNAV	algorithmic abstraction; synthetic solver test
Ranging / median	implemented; simulated under $f < N/2$
Peer independence	assumption; not proved by count
Live deep-space mesh	not claimed
Flight hardware D3	not measured isolated calibration; informative shadow

What is implemented and what is not claimed

The implementation status is deliberately finer than a binary deployment label. TCB and OWLT paths are executable and tested against deterministic fixtures. Ephemeris and JPL paths are validated against frozen reference vectors or authority responses when available. XNAV is an algorithmic abstraction with a synthetic solver test. Peer resilience is simulated under its stated fault assumptions. No row in this table should be read as evidence of a flight-qualified interplanetary service.

Open engineering boundaries

The remaining boundaries are operational rather than hidden correctness claims: detector calibration, live ephemeris authority, peer independence, radio-link asymmetry, and spacecraft hardware. Each boundary has an explicit state and evidence class so that a future deployment can replace it with a measurement without changing the paper’s semantics.

Evidence maturity levels

The status table uses a deliberately conservative vocabulary. *Implemented* means that the code path exists; *tested* means that a deterministic test exercised it; *offline validated* means comparison with a frozen reference artifact; *simulated* means behavior under an explicit model; and *hardware-measured* means an observation from the relevant physical instrument or vehicle. These labels are not interchangeable. In particular, an implemented XNAV solver is not an X-ray detector, and a Byzantine simulation is not a deployed peer network.

Deployment interpretation

The profile is ready to serve as a reproducible software research artifact when its source, toolchain, inputs, outputs, and evidence class are frozen. Deployment in an interplanetary mission would additionally require calibrated detectors, authoritative ephemeris operations, authenticated key lifecycle, link-level time-transfer characteriza-

tion, independent peer governance, and flight qualification. None of those requirements is discharged by the software tests reported here.

Limitations as protocol states

Limitations are represented as states that affect admission. Missing detector observations, stale authority data, insufficient peers, or unknown common source dependence can preserve a record while preventing a strong PASS. This is intentional: the system’s safe behavior under uncertainty is part of the implementation claim. Future measurements may replace an explicit unknown with evidence, but may not be inferred from the existence of an interface.

Reproducibility and Artifacts

Artifact identity

An artifact reference includes a content hash, source revision, execution identifier, correlation identifier, timestamp, and evidence class. A receipt binds these references in order. Missing metadata yields NOT_MEASURED rather than a guessed value. This rule applies equally to positive results and to quarantine or recovery results.

Every reported result should retain the source commit, toolchain, dependency lock file, configuration and ruleset hashes, simulation seed, input/output hashes, exact test command, and receipt for the evidence package. The TeX source and generated PDF are retained together. Historical test counts are not substituted for a current run without an explicit evidence timestamp.

The minimum identity tuple is

$$(h_{\text{artifact}}, h_{\text{source}}, h_{\text{input}}, h_{\text{config}}, \text{execution_id}, \text{correlation_id}). \quad (26)$$

The tuple prevents a result from one source revision being silently paired with an input or configuration from another. Timestamps provide ordering metadata; hashes provide content identity. Neither is treated as a substitute for the other.

Artifact bundle layout

A reproducibility bundle contains a machine-readable manifest, source and dependency identifiers, frozen input fixtures, test outputs, simulation seeds, generated records, and verification receipts. Files are addressed by content hash and the manifest records whether each item is implemented, synthetic, simulated, live black-box, or hardware-measured. A missing item is retained as an explicit unknown rather than replaced by a link to a mutable latest artifact.

K-Lean or kernel certificates are retained as references to the formal verification artifact, not as a blanket certificate for the surrounding physical or semantic claim. Similarly, a TTTPS receipt binds the evidence package’s digest and ordering but does not certify the truth of an ephemeris,

Table 11 — Minimum reproducibility-bundle contents.

Class	Required contents
Source	commit, dirty-state record, TeX and code hashes
Environment	Rust/compiler, platform, feature flags, <code>Cargo.lock</code>
Inputs	fixture identifiers, source revisions, input hashes, seeds
Execution	exact commands, execution/correlation IDs, logs
Verification	K-Lean certificates, test results, disposition states
Integrity	manifest, artifact hashes, TTTPS receipt
Publication	PDF, TeX, bibliography, version and known unknowns

position, or claim correspondence.

Replay procedure

Replay first verifies the manifest and hashes, then reconstructs the declared toolchain and invokes the recorded commands in a fresh process. It compares semantic state and evidence references, not only presentation text. A different environment may produce a controlled UNVERIFIABLE result; it must not be reported as a matching run merely because a PDF was produced.

The freeze procedure is ordered: capture source and environment, freeze the input and configuration manifest, execute the recorded commands, hash the outputs, generate the receipt, and finally re-read every hash from the stored bundle. If the source revision changes between index creation and checksum generation, the bundle is invalidated and regenerated rather than annotated as equivalent. A replay is successful only when the semantic state, evidence references, and declared disposition agree under the same frozen identities.

Relationship to TTTPS and Confidence-00

Citation and extension policy

The companion cites the core TTTPS specification for cryptographic semantics, Confidence-00 for confidence vocabulary, and P2/D3 for higher-order dependence. It does not duplicate their normative text. Any later TTTPS draft should link this document as an Informative Reference for the non-terrestrial profile, while the profile must remain usable as a separate research artifact.

TTTPS core is cryptographic temporal attestation. The Deep-space profile is propagation, positioning, peer resilience, and delayed-link failure policy. Confidence-00 is the oracle-confidence and AdaptiveSwitch layer. P2/D3 is the higher-order dependence theory and diagnostic layer. K-Lean is the formal contract and certificate layer. The layers can be cited together without turning confidence into integrity or D3 into a mandatory wire-level field.

The dependency direction is intentionally one-way:

TTTPS core $\leftarrow$ Deep-space context ,
Confidence/P2-D3 $\rightarrow$ qualification evidence .

The profile consumes the core attestation primitive and produces additional context and qualification evidence. It does not make the core depend on an optional navigation source, and a confidence diagnostic cannot alter the cryptographic meaning of an already-issued record.

Table 12 — Boundary between companion layers.

Layer	Adds	Does claim	not
TTTPS core	signed record	temporal correctness	cor-
Deep-space	OWLT, TCB, position, peers	live flight operation	
Confidence-00	spread, correlation, switching	integrity proof	
P2/D3	higher-order diagnostic	automatic rejection	re-
K-Lean	formal contract/certificate	prose equivalence	equiva-

Normative and informative separation

The core protocol retains authority over wire-level temporal attestation. This paper is an informative profile and research implementation: it specifies how physical context and confidence evidence can be attached to a core record, but does not silently redefine the core record’s wire semantics. A future standards draft may normatively reference the profile after its evidence boundaries and interoperability requirements have been separately reviewed.

For a future IETF submission, the normative core fields and processing rules must remain traceable to the applicable TTTPS draft. Deep-space-specific algorithms, source adapters, and experimental evidence should be identified as profile material until their wire semantics and interoperability requirements are independently stable. This paper can therefore serve as a citable research companion without implying that publication alone changes the status of a standards-track document.

Evidence Ledger and Context Schema

The evidence boundary is made explicit in the artifact ledger. IMPLEMENTED means that a source path exists; TESTED means that a deterministic test directly exercised it; SYNTHETIC means that inputs have known constructed answers; SIMULATED means that a multi-node or attack model was used; LIVE BLACK-BOX means a bounded service endpoint was exercised; and HARDWARE-MEASURED means physical detector or spacecraft telemetry. The present paper reports no hardware-measured result.

The context associated with a record contains a context identifier, session, epoch, coordinate frame, position and velocity source, potential source, OWLT, peer roster, authority status, correction terms, confidence state, and de-

clared uncertainty. A transition stores the pre-state, observation, post-state, and reason for PASS, HOLD, UNVERIFIABLE, or REJECT. A canonical hash over this context is bound to the unchanged cryptographic record, so the result can be replayed without treating a receipt as a physical truth proof.

Reproduction Test Matrix

Reproduction proceeds by pinning the source revision, dependency lock file, compiler, feature flags, simulation seeds, and reference-vector hashes. The unit layer then exercises context, TCB, ephemeris, pulsar, median, peer-health, residual-history, and verification paths. The integration layer runs distance, peer-fault, partition, attack, VN, D3, and relay-journey examples. Finally, logs and outputs are hashed and classified using the evidence ledger.

Table 13 — Failure conditions and required state.

Condition	Required state
OWLT unavailable	HOLD or UNVERIFIABLE
JPL authority unavailable	UNVERIFIABLE
Peer set below quorum	HOLD or holdover
$f \geq N/2$	robustness claim rejected
Range-rate inconsistent	HOLD or REJECT per bound
Stale ephemeris	HOLD until refreshed
D3 residual alone	shadow escalation
Replay or duplicate	deterministic replay handling

Conclusion

Deep-space temporal ordering is feasible as a profile when delay, position, relativity, peer faults, and evidence status are represented explicitly. The key design decision is not to enlarge the wire record until it carries every physical fact. Instead, the protocol keeps its compact cryptographic core and binds the physical verification context out of band. This permits a node to extend a TCB chain across long-delay links while preserving honest HOLD and unverifiable states.

The resulting claim is deliberately bounded. The profile makes light-time delay explicit, propagates coordinate time under a declared model, combines position and peer evidence without treating count as independence, and keeps confidence, applicability, provenance, and integrity as distinct predicates. Its higher-order D3 layer remains a diagnostic rather than an automatic conformance or rejection rule.

The most important operational property is the treatment of uncertainty. A kernel-valid or cryptographically intact record is not thereby identical to a source claim, physically applicable, or flight-measured. When required evidence is unavailable, the system preserves HOLD or UNVERIFIABLE; when a specified incompatibility is established, it can produce REJECT. Each reassessment is linked

to the original transition and leaves the historical observation intact.

The implementation evidence supports the algorithms and simulated pipeline described here. It does not substitute for flight data, a live interplanetary mesh, or an independent physical review. The companion therefore provides a reproducible foundation for TIPTOP, DTN, and later aerospace evaluation rather than a claim that those evaluations have already occurred.

Future work is empirical and versioned: replace synthetic positioning inputs with calibrated observations, characterize asymmetric radio links, validate authority-backed ephemerides, establish peer-governance and independence procedures, and qualify the implementation on flight hardware. These additions can strengthen the evidence class without changing the central separation:

formal validity $\neq$ semantic identity $\neq$ physical applicability.
(27)

The profile’s final contribution is methodological as well as technical. It offers a reproducible way to state what was computed, under which physical assumptions, from which sources, with what uncertainty, and why a record was accepted, held, or rejected. That distinction is essential for long-delay systems in which a result may arrive long after the event it describes.

Reference Equations and Units

All timestamps in the propagation path are represented as integer nanoseconds unless a quantity is explicitly marked as a dimensionless correction. The following unit ledger prevents a common class of errors in which kilometres, metres, seconds, and nanoseconds are mixed inside a correction term.

Table 14 — Units and provenance of the principal quantities.

Quantity	Unit	Meaning
τ	ns	local proper-time epoch
d	m	estimated one-way range
c	m/s	speed of light
Φ	m ² /s ²	gravitational potential
v	m/s	local or relative velocity
$t_{\text{owl}t}$	s	one-way light time
T	ns	declared local tolerance
ρ	m	measured range
$d\rho/dt$	m/s	range-rate observable

The dimensionless correction is $\Phi/c^2 + v^2/(2c^2)$. The implementation adds the rounded correction to an integer epoch, rather than converting the epoch itself to floating point. A verifier should reject a context whose coordinate frame, epoch, or unit declaration is absent instead of guessing it.

The dimensional checks are mechanical. d/c has units of seconds; multiplying the dimensionless relativistic fac-

tor by an elapsed time has units of seconds; and only after that operation may the correction be converted to nanoseconds and rounded. Range-rate is retained in metres per second and is not confused with a time derivative expressed in nanoseconds. These checks are part of the reproduction procedure because a numerically plausible value can still be dimensionally invalid.

The unit ledger also fixes the conversion convention for distance inputs. A source expressed in kilometres is converted to metres before division by $c = 299\,792\,458$ m/s, and astronomical-unit examples are expanded before the calculation is recorded. The output stores both the normalized quantity and the original source unit so that a later audit can distinguish a conversion error from a physical-model discrepancy.

Reference Distance Calculations

For $c = 299792.458$ km/s, the recorded reference calculations are:

$$t_{\text{Earth-Moon}} \approx 1.2822 \text{ s}, \quad (28)$$

$$t_{\text{Earth-Mars,min}} \approx 182.126 \text{ s}, \quad (29)$$

$$t_{\text{Earth-Mars,max}} \approx 1337.592 \text{ s}, \quad (30)$$

$$t_{100\text{AU}} \approx 49900.488 \text{ s}. \quad (31)$$

The corresponding two-way envelopes are twice these values plus the declared base tolerance. These equations are sanity checks for the tier selector; they are not observations of a particular spacecraft at a particular epoch.

The Earth–Mars values are limiting geometric examples rather than a fixed mission distance. Actual range varies with the orbital configuration and must be supplied with an epoch and frame by the position source. Likewise, the Earth–Moon value is a nominal range-scale check, not a replacement for a measured instantaneous range. The 100 AU value is a convenient linear scale for testing propagation-dependent configuration.

The two-way expression includes only the propagation envelope and the declared base tolerance. Processing, queueing, transponder turnaround, atmospheric effects, and measurement uncertainty are separate terms in an instrumented exchange and are not silently included in these figures. A deployment must therefore recompute the complete bound from its evidence rather than reuse an illustrative number as a timing guarantee.

For reproducibility, each calculation is generated from an explicit distance, the exact value of c , and a stated rounding rule. The resulting values are stored with the input distance and unit, source revision, and calculation hash. This makes the appendix a dimensional sanity-check vector set rather than an unqualified navigation table.

Positioning and Navigation Test Vectors

The XNAV test constructs time-of-arrival values from four catalogue pulsar directions and a known position, then solves the linearised four-unknown system with partial pivoting. The expected result is the original synthetic position within 10^{-3} m in the noise-free fixture. A future hardware experiment must add detector noise, pulse-profile uncertainty, clock uncertainty, and an instrument calibration identifier to this vector; none is silently included in the present result.

The vector is defined by the pulsar catalogue identifiers and directions, the synthetic spacecraft position, the assumed emission or arrival epochs, and the noise model. For pulsar direction $\mathbf{n}_i$, a simplified arrival constraint has the form

$$\Delta t_i = \Delta t_0 + \frac{\mathbf{n}_i \cdot \Delta \mathbf{r}}{c} + \epsilon_i, \quad (32)$$

where Δt_0 is the clock offset and ϵ_i is the declared measurement residual. The catalogue revision and coordinate frame are part of the vector identity; changing either creates a new test generation.

The noise-free result isolates solver algebra and conditioning. It does not estimate the error budget of a detector, pulsar timing model, clock, or spacecraft ephemeris. A noisy vector must report photon statistics, pulse-profile uncertainty, clock noise, catalogue error, observation duration, and geometry conditioning separately rather than folding them into the solver’s position residual.

The ephemeris test uses a recorded J2000 comparison vector. Its residual is approximately 2,006 km for Earth, 1.15 million km for Jupiter, and 3.85 million km for Saturn in the documented approximate model. These residuals are not hidden by a live lookup and are not presented as spacecraft navigation accuracy. They define the model’s known reference boundary.

The ephemeris vector records the body, epoch, frame, element-set revision, reference state, analytic state, and residual norm. The residuals are a diagnostic of the documented approximate model at that reference, not a universal bound over an orbit. A future authority-backed comparison must add the service response identity, query time, interpolation method, and validity interval before the result can be interpreted as an operational navigation measurement.

Peer and Attack Matrix

The matrix separates a control that exists from a threat that has been tested in a deployed environment. For example, an identity roster can be implemented without proving that a physical constellation supplies independent clocks.

For each row, an attack vector is frozen before execution and the result records the injected condition, affected

Table 15 — Threats, control, and evidence class.

Threat	Control	Evidence
Forged record	sig/hash	tested
Replay	nonce/history	tested
Delay	OWLT	simulated
attack	gate	
Sybil	roster	partial
Collusion	median/VN/D3	synthetic
Stale	authority/age	policy
ephemeris		
Partition	holdover	simulated
Key	rotation	unmeasured
com- pro- mise		

artifact, expected control, observed state, and evidence hash. TESTED means that the deterministic implementation path was exercised; SIMULATED means that a declared network or adversary model was used; SYNTHETIC means that the constructed data have known properties; PARTIAL means that only a bounded aspect was measured; and UNMEASURED means that the control must not be presented as an empirical result.

The forged-record case changes signed bytes or their identity binding and checks that verification fails without deleting the original artifact. The replay case reuses a valid observation at a stale epoch and checks deterministic duplicate disposition. The delay case varies propagation and queue terms within and beyond the declared bound. Sybil and collusion cases separately vary authenticated membership and shared-source dependence, because identity authentication does not establish independence. The stale-ephemeris and partition cases exercise age/authority and holdover transitions. Key compromise remains an explicit unmeasured operational boundary until a key lifecycle and revocation experiment is run.

The matrix is not a claim that every threat is detected. Its purpose is to show which state is preserved when a control cannot decide. An attack that cannot be distinguished from a valid common-source failure should produce uncertainty or escalation evidence, not an unsupported hard rejection.

Artifact and Receipt Schema

The reproducibility bundle contains the following append-only records:

1. input and source provenance, including source hash and acquisition time;
2. coordinate context, epoch, frame, position source, and authority state;
3. peer observations, range/range-rate values, and aggregation inputs;
4. correction terms, confidence outputs, and D3 shadow

outputs;

5. the generated Proof-of-Time record and verification result;
6. transition evidence, final disposition, and failure reason; and
7. the receipt, source revision, toolchain, commands, and output hashes.

The receipt is a tamper-evident observation of this bundle. It proves that the bundle existed in a stated order at a stated observation point; it does not prove that the ephemeris, detector, or peer assertions were physically true.

The canonical record has four linked identities:

<code>execution_id</code>	<code>run identity</code>
<code>correlation_id</code>	<code>end-to-end lineage</code>
<code>artifact_id/hash</code>	<code>content identity</code>
<code>transition_id</code>	<code>state-change identity</code>

Every derived artifact references its parent artifacts and declares its evidence class. A record with an absent parent, mismatched hash, or unrelated correlation identifier cannot satisfy a strong verification predicate.

The lifecycle is append-only: acquisition creates the source record, normalisation creates the context record, each processing stage creates a derived artifact, and a disposition creates a transition. Reprocessing a quarantined item appends a new transition linked to the original; it does not rewrite the original failure. Duplicate input is identified by the declared idempotency key and receives a deterministic disposition.

Before accepting a bundle, a verifier checks schema version, required fields, parent references, content hashes, source and configuration revisions, and the ordering represented by the receipt. It then evaluates the declared predicate. A complete and untampered bundle can still yield `HOLD` or `UNVERIFIABLE`; structural integrity of evidence is not semantic or physical correctness.

The receipt is consequently the last integrity layer, not a truth oracle. It records the digest, ordering, observation point, and receipt version. A later verifier must retain the distinction between `receipt_valid`, `evidence_complete`, and `claim_applicable`; these fields are not collapsed into a single `VERIFIED` label.

Reproduction Limitations

The reference implementation can be rebuilt and its deterministic calculations replayed. Reproducing the software result does not reproduce a physical link. The remaining external measurements are precisely those that require a flight or laboratory setup: an operating X-ray detector, live long-delay telemetry, independent peer clocks, and an authority service available at the observation epoch. The

paper retains these as explicit unknowns so that later measurements can be added without rewriting the earlier result.

Reproduction is therefore divided into three scopes. *Software reproduction* rebuilds the pinned code and replays deterministic calculations. *Model reproduction* reruns synthetic navigation, peer, attack, and D3 experiments under the recorded assumptions. *Physical reproduction* requires the instruments, clocks, link, authority service, and environmental conditions represented by the model. Success in either of the first two scopes must not be reported as success in the third.

Some results also depend on external service availability, catalogue revision, platform arithmetic, or feature-gated network access. The bundle records these dependencies and the applicable fallback state. If a dependency cannot be reconstructed, the correct result is a bounded `UNVERIFIABLE` or `NOT_MEASURED`, not a substituted latest response.

The limitations do not invalidate the deterministic tests; they bound their interpretation. A future physical campaign should add calibrated X-ray observations, live ephemeris responses, independent clock hardware, measured link asymmetry, environmental metadata, and uncertainty estimates. Those artifacts can be attached as a new evidence generation while preserving this paper's original software and model-generation records.

Accordingly, this appendix is part of the result rather than a disclaimer outside it. It identifies which claims are reproducible now, which require a specific authority or instrument, and which remain open until an independently observable physical experiment is available.

References

- [1] H. Jorgen, "Proof-of-Time: Byzantine-Resilient Temporal Ordering in Untrusted Networks," SSRN abstract 6467220, 2026.
- [2] Helm Protocol, "TTTPS," draft-helmprotocol-tttps-06, 2026.
- [3] Helm Protocol, "TTTPS Deep-Space Extension," draft-helmprotocol-tttps-deepspace-00, 2026.
- [4] E. M. Standish, "JPL Planetary and Lunar Ephemerides," JPL, approximate planetary elements.
- [5] Manchester et al., "The Australia Telescope National Facility Pulsar Catalogue," ATNF catalogue.
- [6] V. Cerf et al., "Delay-Tolerant Networking Architecture," RFC 4838, 2007.