

Evaluation of Lagrange fractional-delay interpolation orders for time-varying delay processing in space-based gravitational-wave arm locking: Monte Carlo mean performance and maximum error over a frozen-response grid

Zhengrong Xiang^{a,b}, Jie Luo^{a*}, Jun Ke^{c**}, Haoxin Xu^a, Tianchi Zhu^a

^a School of Mechanical Engineering and Electronic Information, China University of Geosciences, 430074, Wuhan, China

^b School of Physics and Information Engineering, Guangxi Science & Technology Normal University, 546199, Laibin, China

^c National Gravitation Laboratory, MOE Key Laboratory of Fundamental Physical Quantities Measurement, and School of Physics, Huazhong University of Science and Technology, 430074, Wuhan, China

ABSTRACT

Arm locking and the associated time-delay signal-processing chains in space-based gravitational-wave detectors require the reconstruction of delayed signals at noninteger sampling instants as the interspacecraft geometry evolves. The order of fractional-delay interpolation determines not only the local reconstruction error but also the number of interpolation nodes and the arithmetic cost of the digital implementation. We develop a lightweight numerical benchmark for arm-locking-oriented time-varying delay processing and evaluate second-, fourth-, and sixth-order backward Lagrange fractional-delay interpolation in terms of mean error, maximum frozen-response grid error, residual decomposition, tail risk, and deterministic arithmetic workload. The benchmark incorporates an analytic composite source consisting of multiple low-frequency quasi-monochromatic components and a linear chirp, linearized ephemerides for three spacecraft with static Gaussian position offsets, a fixed-point light-travel-time solver, a weighted combination of three links, and in-band power spectral evaluation. For each link, the code separately computes the analytically exact delayed signal, the analytic signal evaluated using the estimated delay, and the discrete Lagrange-interpolated signal. This procedure enables a rigorous decomposition of the output into the physical light-travel-time residual, interpolation reconstruction error, and measurement noise. The formal Monte Carlo simulations yield mean band-limited relative interpolation errors of 2.9821×10^{-4} , 4.9143×10^{-6} , and 1.2333×10^{-7} for the second-, fourth-, and sixth-order interpolators, respectively. The in-band energy of the total residual including noise is approximately 8.37×10^{-4} for all three orders. Increasing the interpolation order substantially reduces the numerical reconstruction error but produces little change in the noisy total residual under the present measurement-noise-dominated configuration, while also increasing the arithmetic workload. Under the prescribed 10^{-3} error budget, the second-order interpolator is the lowest baseline order satisfying the formal Monte Carlo mean-error criterion, whereas a supplementary contiguous frozen-response scan over $q = 1-6$ identifies third-order interpolation as the lowest order satisfying the full-domain grid-error criterion. Because the third-order configuration has not yet undergone the same formal Monte Carlo, PSD, and CVaR analyses as the baseline orders, fourth-order interpolation is retained as the lowest-order conservative reference configuration supported by the complete stochastic evidence currently available.

Keywords: space-based gravitational-wave detection; arm locking; fractional-delay interpolation; Lagrange interpolation; frozen-response grid maximum error; conditional value at risk

INTRODUCTION

Space-based gravitational-wave detection integrates precision interferometry, space optics, control, signal processing, and fundamental physics. By extending interferometric baselines far beyond those achievable on the ground, space-based observatories provide access to the low-frequency gravitational-wave spectrum. The Laser Interferometer Space Antenna (LISA) employs a three-spacecraft constellation and long-baseline inter-spacecraft laser interferometry for gravitational-wave observations [1]. The full-mission analysis of LISA Pathfinder demonstrated the free-fall performance and low-noise

* Corresponding author.

E-mail addresses: zhengrongxiang0913@163.com (Z. Xiang), luojiethanks@126.com (J. Luo), junke@hust.edu.cn (J. Ke), 2356167242@qq.com (H. Xu), 19107289352@163.com (T. Zhu).

measurement technologies required for LISA and provided an extensive characterization of the residual acceleration-noise environment [2]. Complementary space-based concepts are being developed in China: the Taiji program adopts inter-spacecraft laser interferometry and drag-free control within a heliocentric constellation [3], while the TianQin project is advancing toward a three-spacecraft geocentric gravitational-wave observatory through its staged technology-development roadmap [4].

Laser frequency noise in space-based laser interferometry is many orders of magnitude larger than the target interferometric signals. Because the inter-spacecraft arm lengths are neither exactly equal nor constant, direct common-mode cancellation is insufficient. Time-delay interferometry (TDI) suppresses laser frequency noise by combining appropriately time-shifted inter-spacecraft measurements to synthesize virtual equal-arm interferometers [5,6]. For moving constellations, second-generation TDI accounts for the time dependence and noncommutativity of successive delay operators [7], while geometric TDI provides a complementary formulation in terms of synthesized optical paths [8]. These theoretical developments are increasingly accompanied by high-fidelity numerical processing frameworks. Modern LISA instrument simulations incorporate realistic orbital motion, laser locking, clock behavior, Doppler shifts, onboard processing, and updated noise models [9], while numerical investigations of TDI combinations and dynamical delay construction have quantified the effects of instrument noise, constellation dynamics, and data-processing choices on the resulting observables [10,11].

Arm locking provides a complementary means of reducing laser frequency noise before TDI postprocessing by using the long inter-spacecraft optical arms as frequency references. Sheard et al. introduced laser-frequency stabilization by locking to a LISA arm and demonstrated theoretically that useful feedback bandwidth can be achieved despite the long propagation delay [12]. Subsequent studies quantified the effects of Doppler uncertainty, spacecraft motion, clock noise, measurement noise, and controller design on arm-locking performance [13], and time-domain simulations examined dynamic-range requirements and dominant noise limitations [14]. More recently, Wang et al. explicitly investigated arm locking in conjunction with TDI and derived TDI combinations for measurements that had first undergone arm-locking suppression, thereby linking front-end laser stabilization to subsequent TDI processing [15]. Experimental work has further demonstrated a combined arm- and cavity-locking architecture using an optical cavity and a delayed interferometric reference, including suppression of cavity fluctuations and mitigation of Doppler pulling [16]. Together, these developments indicate that arm-locking performance and its interaction with downstream processing depend not only on the feedback controller but also on the treatment of delayed measurements throughout the signal-processing chain.

For spacecraft in relative motion, the light-travel time $\tau_\ell(t)$ of link ℓ varies continuously with time. In a uniformly sampled system, the required delayed epoch $t_n - \tau_\ell(t_n)$ therefore generally falls between available samples, so the corresponding signal value must be reconstructed numerically. Precise delay knowledge is itself an essential input to delay-based LISA processing. Recent work on LISA ranging sensor fusion has treated ranging ambiguities, measurement noise, and onboard delays within a unified data-processing framework, emphasizing the practical importance of accurate inter-spacecraft ranging for TDI [17].

Fractional-delay filtering provides the digital-signal-processing framework required to reconstruct signals at nonsinteger sampling locations. Laakso et al. established a broad treatment of finite-impulse-response, all-pass, and related fractional-delay structures [18], while the Farrow structure introduced a practical framework for continuously variable digital delay [19]. These foundational methods continue to motivate modern variable fractional-delay designs in which numerical approximation accuracy must be balanced against implementation complexity. For example, Zhao and Tay showed that variable fractional-delay filter structures can be designed explicitly around accuracy–complexity trade-offs, with computational complexity treated as a first-class implementation consideration rather than as an incidental consequence of filter design [20].

These issues become particularly important when fractional-delay processing is embedded in a time-varying LISA measurement chain. Page and Littenberg applied fractional-delay interpolation to TDI for orbiting LISA and showed that time-dependent spacecraft separations substantially increase the computational burden because the interpolation operation must be updated along the time series; they accordingly developed a computationally reduced fractional-delay formulation for the time-dependent problem [21]. Staab et al. subsequently analyzed laser-noise residuals arising from onboard processing and TDI within a unified framework that explicitly includes ranging and interpolation errors, and further showed that laser locking modifies the correlations entering the final TDI residuals [22]. These results establish a direct connection among laser stabilization, delay knowledge, fractional-delay reconstruction, and residual noise after digital delay processing.

The temporal behavior of the interpolation kernel introduces an additional constraint. Staab et al. showed that, when a time-varying fractional delay passes through an integer-sample boundary, insufficient smoothness of the interpolation kernel can generate transient discontinuities and glitch-like spectral features; their analysis specifically identified limitations of conventional Lagrange interpolation for such transitions and motivated smoother interpolation-kernel designs [23]. Interpolation order can also have system-level consequences beyond arithmetic workload. Burke et al. recently demonstrated that missing samples can propagate into substantially longer gaps in second-generation TDI variables and that the additional gap duration increases approximately linearly with the order of the fractional-delay filter [24]. Although integer-sample boundary crossings and data-gap propagation are not modeled in the benchmark developed here, these results reinforce the broader conclusion that increasing interpolation order can incur practical processing costs beyond the reduction of numerical interpolation error alone.

Interpolation order should therefore not be selected from a single mean-error metric. A low-order interpolator requires fewer samples and less arithmetic work but may produce larger reconstruction errors near the upper end of the frequency range or within unfavorable fractional-delay intervals. Increasing the order can substantially reduce numerical interpolation error, but the associated increase in node count, coefficient evaluation, sample access, and arithmetic workload does not necessarily produce a corresponding improvement in the end-to-end residual. This distinction becomes particularly important when other stochastic error sources dominate the overall measurement. Moreover, evaluating only a sparse, prespecified subset of integer orders can lead to an incorrect inference regarding the lowest order satisfying a prescribed error requirement. An interpolation-order assessment should therefore consider typical stochastic performance, high-error regions of the frequency–fractional-delay domain, upper-tail behavior, implementation workload, the relative magnitude of other error sources, and the completeness of the candidate design space.

This work does not attempt to model a complete closed-loop arm-locking system. Instead, we construct a lightweight numerical benchmark motivated by the time-varying digital delay-processing problem encountered in arm-locking-oriented applications. The benchmark retains the elements required to isolate interpolation-order effects: an analytically defined continuous-time source, time-varying light-travel times, spacecraft-position perturbations, fractional-delay reconstruction, multi-link combination, signal-level residual decomposition, in-band spectral evaluation, and statistical uncertainty analysis. In contrast to studies aimed at reproducing a complete TDI observable or closed-loop arm-locking architecture, the present benchmark uses an analytical delayed signal as the reference and isolates the numerical reconstruction problem from a full mission-level control and measurement model.

The main contributions of this work are as follows:

1. A numerical benchmark combining a continuous-time analytical reference with discrete fractional-delay reconstruction is established, enabling the physical light-travel-time residual, interpolation reconstruction error, and measurement noise to be separated at the signal level.
2. Monte Carlo mean performance is evaluated for the baseline interpolation orders $q=2, 4$, and 6 using common random numbers across both interpolation orders and spacecraft-position uncertainty scenarios. Statistical uncertainty is quantified using a paired trial-ID block bootstrap that preserves the joint dependence structure across the σ strata.
3. The stochastic mean-performance analysis is complemented by an independent frequency–fractional-delay frozen-response grid scan. A supplementary contiguous scan over $(q=1)–6$ is used to assess candidate-set completeness and to determine whether the original sparse baseline set omits a lower-order configuration satisfying the prescribed grid-error budget.
4. Mean interpolation accuracy, upper-tail conditional value at risk (CVaR), maximum frozen-response grid error, node count, and an implementation-derived deterministic arithmetic-work proxy are considered together to characterize the trade-off between numerical accuracy, implementation simplicity, system-noise dominance, and the limitations of the currently available stochastic evidence.

TIME-VARYING DELAY BENCHMARK FOR ARM-LOCKING-ORIENTED PROCESSING

2.1. Link signals and residual decomposition

Let $h(t)$ denote the analytically defined continuous-time source. The reference and estimated light-travel times for link ℓ are denoted by

$$\tau_{\ell}^{\text{true}}(t) \quad \text{and} \quad \tau_{\ell}^{\text{est}}(t)$$

respectively.

At sampling time t_n , the analytical true-delay signal is defined as

$$x_\ell^{\text{true}}[n] = h(t_n - \tau_\ell^{\text{true}}(t_n)) \quad (1)$$

and the analytical estimated-delay signal is

$$x_\ell^{\text{est}}[n] = h(t_n - \tau_\ell^{\text{est}}(t_n)) \quad (2)$$

Applying a q -th-order backward Lagrange fractional-delay interpolator to the uniformly sampled sequence $h[n] = h(t_n)$ gives

$$\hat{x}_{\ell,q}[n] = \mathcal{D}_q[h; \tau_\ell^{\text{est}}(t_n)] \quad (3)$$

The physical light-travel-time residual for link ℓ is defined by

$$d_\ell[n] = x_\ell^{\text{true}}[n] - x_\ell^{\text{est}}[n] \quad (4)$$

and the interpolation reconstruction error relative to the analytical estimated-delay signal is

$$e_{\ell,q}^{\text{interp}}[n] = \hat{x}_{\ell,q}[n] - x_\ell^{\text{est}}[n] \quad (5)$$

With measurement noise $\eta_\ell[n]$, the total residual on the link is

$$r_{\ell,q}[n] = x_\ell^{\text{true}}[n] + \eta_\ell[n] - \hat{x}_{\ell,q}[n] \quad (6)$$

It follows that

$$r_{\ell,q}[n] = d_\ell[n] - e_{\ell,q}^{\text{interp}}[n] + \eta_\ell[n] \quad (7)$$

The three links are combined according to

$$r_q[n] = \sum_{\ell=1}^3 w_\ell r_{\ell,q}[n] \quad (8)$$

where

$$\mathcal{L} = \{(1,2), (2,3), (3,1)\}, \quad \mathbf{w} = [1, -1, 1]^T \quad (9)$$

Here, the ordered pair (i, j) denotes a signal emitted by spacecraft j and received by spacecraft i . We further define

$$d[n] = \sum_{\ell=1}^3 w_\ell d_\ell[n] \quad (10)$$

$$e_q[n] = \sum_{\ell=1}^3 w_\ell e_{\ell,q}^{\text{interp}}[n] \quad (11)$$

$$\eta[n] = \sum_{\ell=1}^3 w_\ell \eta_\ell[n] \quad (12)$$

and the combined analytical reference signal

$$x^{\text{est}}[n] = \sum_{\ell=1}^3 w_\ell x_\ell^{\text{est}}[n] \quad (13)$$

The combined total residual therefore satisfies

$$r_q[n] = d[n] - e_q[n] + \eta[n] \quad (14)$$

Equation (8) is used solely to construct a numerical metric for delayed-signal reconstruction. It should not be interpreted as a complete TDI observable or as the output of a closed-loop arm-locking sensor.

In the formal numerical experiment, the maximum relative reconstruction error associated with Eq. (14) was

$$2.452 \times 10^{-15},$$

showing that the decomposition identity is satisfied to within double-precision floating-point roundoff.

2.2. Low-frequency composite analytical source

The benchmark input consists of NGB slowly varying quasi-monochromatic components and one linear chirp:

$$h(t) = \sum_{k=1}^{N_{\text{GB}}} h_k(t) + h_{\text{chirp}}(t) \quad (15)$$

The k -th quasi-monochromatic component is

$$h_k(t) = A_k \cos \left[2\pi \left(f_{k,0}t + \frac{1}{2} \dot{f}_k t^2 \right) + \phi_k \right] \quad (16)$$

The number of components is drawn uniformly from the integer set $\{5, \dots, 20\}$. Their amplitudes and initial phases satisfy

$$A_k \sim \mathcal{U}(0.2, 1.0), \quad \phi_k \sim \mathcal{U}(0, 2\pi) \quad (17)$$

For each quasi-monochromatic component, the initial and final frequencies are independently drawn from

$$\left[5.0 \times 10^{-4}, 8.0 \times 10^{-3} \right] \text{ Hz}$$

The corresponding linear frequency-sweep rate is then determined from the 6 h record length so that the instantaneous frequency remains within this interval over the complete record.

The linear chirp is

$$h_{\text{chirp}}(t) = A_c \cos \left[2\pi \left(f_{c,0}t + \frac{1}{2} \dot{f}_c t^2 \right) + \phi_c \right] \quad (18)$$

with

$$A_c = 0.8, \quad f_{c,0} = 1.2 \times 10^{-3} \text{ Hz}, \quad \dot{f}_c = 1.0 \times 10^{-7} \text{ Hz s}^{-1} \quad (19)$$

Over the 6 h record, its instantaneous frequency increases from

$$1.2 \times 10^{-3} \text{ Hz}$$

to

$$3.36 \times 10^{-3} \text{ Hz}$$

The source model is used only to provide a controlled spectral composition for the numerical benchmark. Its amplitude is not calibrated to physical strain, optical-frequency, or phase units.

2.3. Linearized ephemeris and position uncertainty

The position of spacecraft i is modeled as

$$\mathbf{r}_i(t) = \mathbf{r}_{i,0} + \mathbf{v}_i t + \delta \mathbf{r} \quad (20)$$

where $\mathbf{r}_{i,0}$ and $\mathbf{v}_i$ are the initial position and constant velocity, respectively. The reference ephemeris is defined by

$$\delta \mathbf{r}_i = \mathbf{0}$$

For each Monte Carlo trial, a standardized three-dimensional Gaussian position-offset vector is first generated and then scaled by the scenario parameter σ :

$$\delta \mathbf{r}_i(\sigma) = \sigma \mathbf{z}_i, \quad \mathbf{z}_i \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_3) \quad (21)$$

The position-uncertainty scenarios are

$$\sigma \in \{1, 10, 100, 1000\} \text{ m} \quad (22)$$

where σ is the standard deviation of the position error along each Cartesian axis.

Table 1. Linearized spacecraft states

Spacecraft	Initial position (x, y, z) [m]	Velocity (v _x , v _y , v _z) [m/s]
SC1	(0, 0, 0)	(0, 0, 0)
SC2	(2.5×10^9 , 1.0×10^8 , 0)	(120, -80, 0)
SC3	(-1.2×10^9 , 2.1×10^9 , 5.0×10^7)	(-90, 140, 0)

For a given trial ID, all interpolation orders use the same analytical source realization, link-noise realization, and spacecraft-position perturbation. Across different σ scenarios, the standardized perturbation direction is also shared, with only its magnitude being rescaled. This common-random-number design reduces unnecessary stochastic variation in comparisons across both interpolation order and spacecraft-position uncertainty.

2.4. Fixed-point light-travel-time solution

For a signal emitted by spacecraft j and received by spacecraft i at time t , the light-travel time satisfies

$$\tau_{ij}(t) = \frac{\|\mathbf{r}_i(t) - \mathbf{r}_j(t - \tau_{ij}(t))\|}{c} \quad (23)$$

where

$$c = 299792458 \text{ m s}^{-1}$$

The initial estimate is the simultaneous geometric separation divided by the speed of light:

$$\tau_{ij}^{(0)}(t) = \frac{\|\mathbf{r}_i(t) - \mathbf{r}_j(t)\|}{c} \quad (24)$$

The fixed-point iteration is

$$\tau_{ij}^{(r+1)}(t) = \frac{\|\mathbf{r}_i(t) - \mathbf{r}_j(t - \tau_{ij}^{(r)}(t))\|}{c} \quad (25)$$

The iteration terminates when

$$|\tau_{ij}^{(r+1)}(t) - \tau_{ij}^{(r)}(t)| < 10^{-12} \text{ s} \quad (26)$$

with a maximum of 50 iterations. Light-travel times for the reference ephemeris are precomputed and cached before the Monte Carlo evaluation. Those associated with the perturbed ephemeris are recomputed within each trial.

2.5. Backward Lagrange fractional-delay interpolation

The sampling frequency, sampling interval, and record duration are

$$f_s = 0.2 \text{ Hz}, \quad \Delta t = 5 \text{ s}, \quad T = 6 \text{ h} \quad (27)$$

The light-travel time expressed in units of the sampling interval is

$$s_n = \frac{\tau(t_n)}{\Delta t} \quad (28)$$

Let

$$k_n = \lfloor s_n \rfloor, \quad \lambda_n = s_n - k_n, \quad 0 \leq \lambda_n < 1 \quad (29)$$

so that

$$t_n - \tau(t_n) = (n - k_n - \lambda_n)\Delta t \quad (30)$$

A q -th-order backward Lagrange interpolator uses the $q + 1$ samples

$$x[n - k_n], x[n - k_n - 1], \dots, x[n - k_n - q]$$

and produces

$$\mathcal{D}_q x[n] = \sum_{m=0}^q x[n - k_n - m] L_m(\lambda_n) \quad (31)$$

where

$$L_m(\lambda) = \prod_{\substack{r=0 \\ r \neq m}}^q \frac{\lambda - r}{m - r} \quad (32)$$

The formal Monte Carlo baseline set is

$$\mathcal{Q}_{\text{MC}} = \{2, 4, 6\} \quad (33)$$

corresponding to 3, 5, and 7 interpolation nodes, respectively. This set was the prespecified baseline design space of the original numerical experiment; it is not an even-order restriction imposed by the Lagrange formulation itself. To avoid a biased inference about the lowest qualifying order from a sparse candidate set, the subsequent deterministic frozen-response analysis expands the order range to $q = 1 - 6$.

Zero padding is applied at the input boundary. To prevent the padding from affecting the performance metrics, 20 samples are removed from each end of the record. With 4320 total samples, the effective evaluation length is therefore

$$N_{\text{eval}} = 4320 - 2 \times 20 = 4280 \quad (34)$$

PERFORMANCE METRICS AND STATISTICAL METHODS

3.1. In-band residual energy

One-sided power spectral densities (PSDs) are estimated for the combined physical light-travel-time residual $d[n]$, interpolation reconstruction error $e_q[n]$, noisy total residual $r_q[n]$, combined analytical reference signal $x^{\text{est}}[n]$, and combined measurement noise $\eta[n]$.

For any evaluated component $a[n]$, the in-band energy is defined as

$$J_a = \int_{f_{\min}}^{f_{\max}} S_a(f) df \quad (35)$$

with

$$f_{\min} = 5.0 \times 10^{-4} \text{ Hz}, \quad f_{\max} = 1.2 \times 10^{-2} \text{ Hz} \quad (36)$$

The principal quantities reported are

$$J_{\text{total}}, \quad J_{\text{delay}}, \quad J_{\text{interp}}, \quad J_{\text{reference}}$$

This frequency interval was prescribed as the numerical benchmark evaluation band before the interpolation-order selection was performed. Its lower bound coincides with the minimum designed source frequency, while its upper bound extends beyond the 8.0×10^{-3} Hz maximum designed instantaneous frequency of the quasi-monochromatic components. The additional high-frequency range is included to probe the increase in interpolation error toward higher frequencies. The interval is therefore a design choice of the present numerical benchmark and should not be interpreted as the control bandwidth or a mission-level requirement of an operational arm-locking system.

The coarse stage uses a 1024-point Hann window with 512-point overlap, whereas the fine stage uses a 4096-point Hann window with a configured overlap of 2048 points. Before PSD estimation, the sample mean over the effective evaluation interval is removed from every evaluated sequence, and the same spectral-estimation procedure is then applied. Because the effective record contains only 4280 samples, only one complete 4096-point segment is available in the fine stage. The

resulting estimator is therefore statistically closer to a single-segment Hann-windowed periodogram than to a conventional multi-segment Welch average.

3.2. Band-limited relative interpolation error

Using the combined analytical estimated-delay signal $x^{\text{est}}[n]$ as the reference, the band-limited relative interpolation error is defined as

$$E_{\text{interp}} = \sqrt{\frac{J_{\text{interp}}}{J_{\text{reference}}}} \quad (37)$$

This quantity measures the band-limited root-mean-square interpolation reconstruction error relative to the continuous-time analytical delay reference and serves as the primary numerical-accuracy metric.

The prescribed error budget is

$$\mathcal{E}_{\text{budget}} = 10^{-3} \quad (38)$$

3.3. Light-travel-time error diagnostic

For link ℓ , the light-travel-time error is

$$e_{\tau,\ell}[n] = \tau_{\ell}^{\text{est}}(t_n) - \tau_{\ell}^{\text{true}}(t_n) \quad (39)$$

The mean three-link RMS diagnostic is

$$J_{\tau} = \frac{1}{3} \sum_{\ell=1}^3 \left[\frac{1}{N_{\text{eval}}} \sum_n e_{\tau,\ell}^2[n] \right]^{1/2} \quad (40)$$

Because the light-travel-time solution is independent of the interpolation order q , J_{τ} is used only to verify the response to spacecraft-position perturbations and the consistency of the numerical implementation. It is not used to select the interpolation order.

3.4. Monte Carlo design

The coarse stage evaluates all

$$3 \times 4 = 12$$

baseline (q, σ) combinations using 20 trials per combination. This stage is used for workflow auditing and preliminary trend assessment.

The formal fine stage evaluates $q = 2, 4$, and 6 for all four spacecraft-position uncertainty scenarios using

$$M_{\text{fine}} = 200 \quad (41)$$

trials per (q, σ) combination. Each baseline order therefore contributes

$$4 \times 200 = 800$$

scenario-trial observations. The coarse and fine stages do not constitute a Pareto-elimination sequence; all three baseline orders proceed to the formal fine-stage evaluation. Because the two stages use different trial counts and different PSD settings, their results should not be interpreted as a numerical-grid convergence sequence.

3.5. Paired trial-ID block bootstrap across σ strata

The same trial ID corresponds to the same analytical source realization, link-noise realization, and standardized spacecraft-position perturbation across both interpolation orders and σ scenarios. To preserve this common-random-number structure, the statistical analysis uses trial ID as the joint bootstrap resampling unit.

For each bootstrap replicate, one set of indices is drawn with replacement from the $M_{\text{fine}} = 200$ trial IDs. The same resampled indices are then applied simultaneously to all interpolation orders and all σ scenarios. Sample means are first computed separately within each of the four σ strata and are then equally averaged:

$$\bar{J}_a(q) = \frac{1}{4} \sum_{\sigma} \bar{J}_a(q, \sigma) \quad (42)$$

This procedure preserves the dependence induced by common random numbers both across interpolation orders and across spacecraft-position uncertainty scenarios.

The number of bootstrap replicates is

$$B = 2000 \quad (43)$$

and 95% confidence intervals are obtained using the percentile method.

For the formal Monte Carlo baseline set, the mean-error selection rule is

$$q_{\text{mean}} = \min_{q \in \mathcal{Q}_{\text{MC}}} \{q : E_{\text{interp}, q}^{\text{upper}} \leq 10^{-3}\} \quad (44)$$

where $E_{\text{interp}, q}^{\text{upper}}$ denotes the upper bound of the corresponding 95% confidence interval.

3.6. Upper-tail conditional value at risk

Each of the four σ scenarios contains 200 samples. Equal-weight pooling of the four scenarios therefore gives 800 observations for each baseline order.

Let the ordered metric values be

$$x_{(1)} \leq \dots \leq x_{(800)}$$

The empirical 95% value at risk (VaR) is defined as

$$\text{VaR}_{95} = x_{(\lceil 0.95N \rceil)} \quad (45)$$

For a metric in which larger values indicate poorer performance, the upper-tail conditional value at risk (CVaR) is

$$\text{CVaR}_{95} = \text{mean} \{x_i : x_i \geq \text{VaR}_{95}\} \quad (46)$$

The CVaR_{95} is evaluated for both the in-band energy of the noisy total residual and that of the interpolation reconstruction error.

3.7. Deterministic arithmetic work

Long-duration Monte Carlo wall-clock measurements can be affected by just-in-time compilation, operating-system load, cache state, and execution order. Measured execution time is therefore not used as a formal order-selection metric.

Instead, a deterministic arithmetic-work proxy is derived from the loop structure of the current scalar Lagrange implementation:

$$W(q) = 4q^2 + 7q + 3 \quad (47)$$

One scalar addition, subtraction, multiplication, or division is counted as one arithmetic work unit. The count includes computation of the fractional part, generation of the Lagrange weights, and the final weighted summation. Operations associated with floor, integer indexing, conditional branches, memory access, array allocation, cache behavior, and processor vectorization are excluded.

For the formal Monte Carlo baseline orders,

$$W(2) = 33, \quad W(4) = 95, \quad W(6) = 189 \quad (48)$$

Normalized to $q = 2$,

$$W_{\text{norm}} = \{1, 2.879, 5.727\} \quad (49)$$

For comparison, the third-order interpolator uses four nodes and has

$$W(3) = 60$$

which is

$$\frac{60}{95} \approx 0.632$$

of the arithmetic-work proxy of the fourth-order implementation. This metric characterizes only the growth in scalar arithmetic work with interpolation order in the present implementation. It is not equivalent to execution time, power consumption, storage requirements, or memory bandwidth on a target processor.

3.8. Frequency–fractional-delay frozen-response grid scan

To complement the Monte Carlo assessment of typical performance under random source realizations, we also examine the complex frequency response of the Lagrange interpolator for fixed fractional delays.

The normalized angular frequency is

$$\omega = 2\pi f / f_s$$

For a fixed fractional delay λ , the complex response of a q -th-order interpolator is

$$G_q(\omega, \lambda) = \sum_{m=0}^q L_m(\lambda) e^{-j\omega m} \quad (50)$$

while the ideal fractional-delay response is

$$G_{\text{ideal}}(\omega, \lambda) = e^{-j\omega\lambda} \quad (51)$$

Because the ideal response has unit magnitude, the complex-response error is defined as

$$\varepsilon_q(f, \lambda) = |G_q(\omega, \lambda) - G_{\text{ideal}}(\omega, \lambda)| \quad (52)$$

The scan uses 801 logarithmically spaced frequency points and 2001 uniformly spaced fractional-delay points. The frequency range is

$$5.0 \times 10^{-4} \leq f \leq 1.2 \times 10^{-2} \text{ Hz}$$

and the full fractional-delay domain is

$$0 \leq \lambda < 1$$

An extended domain of interest,

$$\lambda \in [0.60, 0.82]$$

is also considered to provide a margin around the fractional-delay range realized in the formal simulation,

$$\lambda \in [0.6139, 0.8103]$$

The extended domain is used only to assess the interpolation-error margin around the realized physical range; it is not interpreted as a mission-specified fractional-delay interval.

Equations (50)–(52) describe the **frozen frequency response** for a fixed λ . The resulting maximum is therefore the maximum error over a finite, discrete frequency–fractional-delay grid. It is neither a rigorous analytical upper bound over the continuous (f, λ) domain nor a characterization of coefficient modulation, sidebands, or integer-sample boundary transitions caused by a time-varying $\lambda(t)$.

The original Monte Carlo design considered the three baseline orders $\{2, 4, 6\}$. To determine whether this set omitted a feasible intermediate order, the frozen-response scan was extended to

$$Q_{\text{scan}} = \{1, 2, 3, 4, 5, 6\} \quad (53)$$

The grid-based selection rule is

$$q_{\text{grid}} = \min_{q \in Q_{\text{scan}}} \left\{ q : \max_{(f, \lambda) \in \mathcal{G}} \varepsilon_q(f, \lambda) \leq 10^{-3} \right\} \quad (54)$$

where $\mathcal{G}$ denotes the prescribed discrete scan grid.

RESULTS

4.1. Simulation configuration

Table 2. Configuration of the formal Monte Carlo simulation

Parameter	Value
Sampling frequency	0.2 Hz
Sampling interval	5 s
Record duration	6 h
Total number of samples	4320
Edge trimming	20 samples at each end
Effective number of samples	4280
Per-axis position uncertainty σ	1, 10, 100, 1000 m
Monte Carlo baseline orders	2, 4, 6
Coarse-stage trials	20 per (q, σ)
Fine-stage trials	200 per (q, σ)
Bootstrap replicates	2000
Evaluation band	5×10^{-4} – 1.2×10^{-2} Hz
Measurement-noise standard deviation	0.05
Fine-stage PSD window/configured overlap	4096/2048
Common random numbers	enabled across q and σ
Parallel execution	disabled

4.2. Numerical consistency checks

Table 3. Numerical validation results for the formal run

Validation item	Result
Analytical-source zero-delay relative error	0
Maximum relative difference in physical light-travel-time residual across orders	0
Maximum residual-decomposition error	2.452×10^{-15}
Maximum relative PSD-integration error for the total residual	2.748×10^{-10}
Maximum relative PSD-integration error for the interpolation error	1.323×10^{-9}
Monotonic decrease in interpolation error across baseline orders	PASS
Overall critical validation	PASS

These checks show that the continuous-time analytical reference, residual decomposition, common-random-number design, PSD integration, and numerical responses of the three formal Monte Carlo baseline orders satisfy the prescribed internal consistency requirements.

4.3. Monte Carlo mean interpolation accuracy




Figure 1. Mean band-limited relative interpolation error for the Lagrange baseline orders, with 95% confidence intervals from the paired trial-ID block bootstrap across σ strata.

Figure 1 presents the mean band-limited relative interpolation errors for the second-, fourth-, and sixth-order Lagrange interpolators together with their 95% confidence intervals. The horizontal dashed line denotes the 10^{-3} mean-error budget.

Table 4. Fine-stage mean interpolation results

q	Nodes	Mean E_{interp} and 95% CI	Mean J_{interp} and 95% CI
2	3	2.982097×10^{-4} ([2.868742, 3.102497] $\times 10^{-4}$)	2.985670×10^{-7} ([2.696462, 3.299430] $\times 10^{-7}$)
4	5	4.914274×10^{-6} ([4.633056, 5.214315] $\times 10^{-6}$)	9.000038×10^{-11} ([7.837217, 10.31898] $\times 10^{-11}$)
6	7	1.233260×10^{-7} ([1.137834, 1.329412] $\times 10^{-7}$)	6.413247×10^{-14} ([5.314408, 7.578646] $\times 10^{-14}$)

Increasing the order from 2 to 4 reduces the mean relative interpolation error by

$$\frac{2.982097 \times 10^{-4}}{4.914274 \times 10^{-6}} \approx 60.7$$

whereas increasing it from 4 to 6 provides a further reduction of approximately 39.8.

The in-band interpolation reconstruction-error energy decreases by approximately

$$4.66 \times 10^6$$

from $q = 2$ to $q = 6$, demonstrating the strong reduction in numerical reconstruction error relative to the continuous-time analytical reference as the interpolation order increases.

The upper bounds of the 95% confidence intervals for all three formal Monte Carlo baseline orders remain below 10^{-3} . Thus, within the baseline set and under the mean-error criterion,

$$q_{\text{mean}} = 2 \quad (55)$$

The corresponding threshold sensitivity is summarized below.

Mean-error budget	Lowest qualifying order in the baseline set
10^{-2}	2
10^{-3}	2
10^{-4}	4
10^{-5}	4

The Monte Carlo conclusions in this section apply only to $q = 2, 4$, and 6 , for which the formal stochastic simulations were performed. They cannot be directly extrapolated to $q = 3$ or $q = 5$, which have not yet undergone the same Monte Carlo evaluation.

4.4. Residual decomposition and response to spacecraft-position uncertainty




Figure 2. Fine-stage residual decomposition and deterministic arithmetic-work proxy.

Figure 2 shows the noisy total residual, physical light-travel-time residual, interpolation reconstruction error, and deterministic arithmetic-work proxy for the formal Monte Carlo baseline orders. The curves in the first three panels represent the means of the 200 trials for each σ scenario. The lower-right panel shows the normalized deterministic arithmetic work of the current scalar implementation.

The equal-weight means over the four σ scenarios for the in-band noisy total residual energy are

$$q = 2: 8.3737 \times 10^{-4}; q = 4: 8.3741 \times 10^{-4}; q = 6: 8.374 \times 10^{-4}$$

Using the corrected joint trial-ID bootstrap, the corresponding 95% confidence intervals are

$$q = 2: [8.236, 8.517] \times 10^{-4}$$

$$q = 4: [8.234, 8.516] \times 10^{-4}$$

and

$$q = 6: [8.234, 8.516] \times 10^{-4}$$

The relative difference between the $q = 2$ and $q = 4$ point estimates is approximately 0.0048%, while that between $q = 4$ and $q = 6$ is approximately 0.0012%. Under the jointly paired bootstrap, the 95% confidence interval for the difference in mean total residual between $q = 2$ and $q = 4$ includes zero. The present Monte Carlo data therefore do not show an identifiable reduction in the noisy total residual as the interpolation order is increased.

Figure 2 further shows that J_{delay} increases with the spacecraft-position uncertainty σ , whereas J_{interp} is governed primarily by interpolation order and is only weakly sensitive to σ . The total residual J_{total} is comparatively insensitive to both interpolation order and σ . The sixth-order implementation requires approximately twice the deterministic arithmetic work of the fourth-order implementation, while producing essentially no change in the noisy total residual under the present noise conditions.

The measurement-noise parameters provide an independent order-of-magnitude check of this noise-dominated behavior. If the measurement noises on the three links are mutually independent, each with standard deviation 0.05, and the link-combination weights are $[1, -1, 1]$, the variance of the combined noise is

$$\sigma_\eta^2 = (1^2 + (-1)^2 + 1^2)(0.05)^2 = 7.5 \times 10^{-3}$$

For ideal discrete white noise sampled at $f_s = 0.2$ Hz, the corresponding one-sided PSD is approximately

$$S_\eta \approx \frac{2\sigma_\eta^2}{f_s} = 0.075 \text{ a.u.}^2/\text{Hz}$$

Integration over the evaluation band gives

$$J_{\eta, \text{ideal}} \approx 0.075(0.012 - 0.0005) = 8.625 \times 10^{-4}$$

This estimate is close to the numerically obtained total-residual energy of approximately 8.37×10^{-4} , supporting the interpretation that the present benchmark is dominated by measurement noise.

4.5. Spectral results




Figure 3. Equal-weight average power spectral densities of the noisy total residual and interpolation reconstruction error over the four spacecraft-position uncertainty scenarios.

Figure 3 presents the noisy-total-residual PSD and interpolation-reconstruction-error PSD after equal weighting of the four spacecraft-position uncertainty scenarios. The blue vertical dashed lines indicate the boundaries of the formal evaluation band. The three baseline noisy-total-residual spectra are nearly indistinguishable, consistent with the in-band integrated results in Section 4.4 and indicating that changing the interpolation order has little effect on the present noise-dominated output.

In contrast, within the principal source band, the interpolation reconstruction-error PSDs satisfy

$$q = 2 > q = 4 > q = 6$$

with the curves separated by several orders of magnitude. Above approximately 8×10^{-3} Hz, the interpolation-error PSD decreases markedly, consistent with the maximum designed instantaneous frequency of the quasi-monochromatic source components.

The spectral results therefore reinforce the conclusion that the principal benefit of increasing the interpolation order is a reduction in the numerical interpolation error itself, rather than a reduction in the noisy total residual under the present noise configuration.

4.6. Upper-tail risk




Figure 4. Upper-tail empirical exceedance distributions of the noisy total residual and interpolation reconstruction error.

Figure 4 shows the empirical exceedance probabilities of the noisy total residual and interpolation reconstruction error for the equal-weight mixture of the four spacecraft-position uncertainty scenarios.

Table 5. Upper-tail CVaR results

q	$\text{CVaR}_{95}(\mathbf{J}_{\text{total}})$	$\text{CVaR}_{95}(\mathbf{J}_{\text{interp}})$
2	1.035273×10^{-3}	8.815559×10^{-7}
4	1.036145×10^{-3}	3.444009×10^{-10}
6	1.036138×10^{-3}	3.207119×10^{-13}

The upper-tail point estimates of the noisy total residual are very similar for all three baseline orders. The total-residual CVaR differs by approximately 0.084% between $q = 2$ and $q = 4$, and by less than 0.001% between $q = 4$ and $q = 6$. The present results therefore provide no evidence that any of the baseline orders has a clear advantage in the upper-tail risk of the noisy total residual.

By contrast, the upper-tail interpolation reconstruction error decreases substantially with increasing order. The CVaR is reduced by approximately 2560 times from $q = 2$ to $q = 4$, and by a further factor of approximately 1074 from $q = 4$ to $q = 6$. Thus, increasing the interpolation order reduces not only the mean interpolation error but also the high-error tail of its empirical distribution.

4.7. Frequency–fractional-delay frozen-response grid scan

To determine whether the original $\{2, 4, 6\}$ candidate set omitted a feasible intermediate integer order, a supplementary frozen-response scan was performed for $q = 1 - 6$ using the same frequency and fractional-delay grids as the original analysis.

Table 6. Maximum complex-response error over the frozen-response grids

q	Nodes	Full fractional-delay domain	Extended domain [0.60, 0.82]	Realized physical domain [0.6139, 0.8103]
1	2	1.77127×10^{-2}	1.70007×10^{-2}	1.67880×10^{-2}
2	3	3.40971×10^{-3}	2.97697×10^{-3}	2.91045×10^{-3}
3	4	8.29731×10^{-4}	6.69392×10^{-4}	6.50630×10^{-4}
4	5	2.25705×10^{-4}	1.70511×10^{-4}	1.65051×10^{-4}
5	6	6.55686×10^{-5}	4.68339×10^{-5}	4.51902×10^{-5}
6	7	1.98945×10^{-5}	1.35317×10^{-5}	1.30230×10^{-5}

For every order considered, the maximum grid error occurs at the upper boundary of the evaluation band,

$$f = 0.012 \text{ Hz}$$

As the interpolation order increases, the fractional delay associated with the full-domain maximum shifts toward smaller λ . For example, the maxima for $q = 2$ and $q = 3$ occur at approximately $\lambda = 0.4228$ and $\lambda = 0.3823$, respectively.

Under the 10^{-3} error budget,

$$E_{2,\max} = 3.4097 \times 10^{-3} > 10^{-3}$$

whereas

$$E_{3,\max} = 8.2973 \times 10^{-4} < 10^{-3}$$

Hence,

$$q_{\text{grid}} = 3 \tag{56}$$

is the lowest order within the scanned range $q = 1 - 6$ that satisfies the prescribed full-domain grid-error budget.

This result revises the earlier inference obtained from the restricted set $\{2, 4, 6\}$. Although the fourth-order interpolator does satisfy the budget, its apparent status as the lowest qualifying order resulted from the omission of the third-order candidate rather than from any even-order restriction inherent to Lagrange interpolation.

The sensitivity to the grid-error budget after expanding the order set is

Grid-error budget	Lowest qualifying order among $q = 1-6$
10^{-2}	2
10^{-3}	3
10^{-4}	5
10^{-5}	none

The discrete choice of candidate orders can therefore directly affect the inferred lowest acceptable order. Candidate-set completeness should be checked explicitly whenever minimum node count or minimum order is used as part of the design criterion.

4.8. Order evaluation and interpretation of the current configuration

The Monte Carlo mean-performance analysis and the frozen-response grid scan characterize complementary aspects of interpolation performance.

First, among the formally simulated baseline orders $q = 2, 4$, and 6 , the upper bounds of the 95% confidence intervals for the mean interpolation error are all below 10^{-3} . Under the random-source mean-error criterion alone, the lowest baseline order is therefore

$$q_{\text{mean}} = 2$$

Second, the full frequency–fractional-delay grid scan shows that the second-order interpolator does not satisfy the 10^{-3} budget, whereas the third-order interpolator does. When the integer-order design space is treated as contiguous over the scanned range,

$$q_{grid} = 3$$

rather than 4.

Third, the available Monte Carlo data show that increasing the order from 2 to 4 or 6 substantially reduces the interpolation error but does not produce an identifiable improvement in the noisy total residual. This behavior is consistent with the measurement-noise-dominated error budget of the present benchmark.

Fourth, the third-order interpolator uses four nodes and has a deterministic arithmetic-work proxy of $W(3) = 60$, which is lower than $W(4) = 95$ for the fourth-order interpolator. The third-order candidate therefore merits further evaluation from both frozen-response accuracy and implementation-workload perspectives.

However, the formal 20/200 Monte Carlo analysis, PSD evaluation, and CVaR analysis in this study were performed only for $q = 2, 4$, and 6 . Without an equivalent statistical evaluation of $q = 3$, the available data cannot establish that the third-order configuration satisfies the same set of stochastic performance criteria already evaluated for the fourth-order configuration.

Accordingly, this work no longer interprets

$$q = 4$$

as the lowest acceptable integer order or as a uniquely optimal configuration. The evidence instead supports the following distinction: $q = 3$ is the lowest qualifying order identified by the frozen-response grid criterion, while $q = 4$ is the lowest-order conservative reference configuration supported by the complete Monte Carlo, PSD, and CVaR evidence currently available. Before an engineering interpolation order is finally fixed, $q = 3$ should be subjected to the same common-random-number Monte Carlo and statistical analysis used for $q = 2, 4$, and 6 .

DISCUSSION

5.1. Distinction from a complete arm-locking system

The present study addresses a numerical benchmark motivated by time-varying delay processing in arm locking rather than a complete spaceborne closed-loop arm-locking model. The benchmark does not include a feedback controller, free-running laser-frequency noise, Doppler pulling, actuators, a phasemeter, clock-noise calibration, attitude dynamics, or a mission-level high-fidelity orbit model.

Consequently, J_{total} should not be interpreted as the final noise budget or sensitivity of an operational spaceborne arm-locking loop. The conclusion supported directly by the present analysis is narrower: for the specified digital delay-processing chain, increasing the Lagrange interpolation order can substantially reduce the numerical interpolation error relative to a continuous-time analytical reference, but such an improvement does not necessarily produce an observable reduction in the end-to-end residual when measurement noise dominates.

5.2. Analytical reference and error decomposition

The delayed reference is evaluated directly from the analytically defined continuous-time source,

$$h[t - \tau(t)]$$

rather than being generated by one of the candidate interpolators. Consequently, J_{interp} represents the in-band energy of the error between the discrete Lagrange reconstruction and the continuous-time analytical delayed signal.

The physical light-travel-time residual, interpolation reconstruction error, and measurement noise are constructed separately at the signal level and related exactly through Eq. (14). This decomposition prevents ephemeris-induced light-travel-time errors from being merged with numerical interpolation errors into a single residual metric and makes the independent effect of interpolation order directly observable.

5.3. Mean and grid-based criteria

The Monte Carlo analysis characterizes typical statistical performance for the prescribed random source model, spacecraft-position perturbation model, and finite set of trial realizations. A mean relative interpolation error below 10^{-3} in these random trials does not imply that the same order satisfies the 10^{-3} constraint at every frequency and fractional delay.

The frozen-response grid scan complements the Monte Carlo analysis but has two important limitations.

First, the scan covers only a finite discrete grid. Its maximum value is therefore not a rigorous upper bound over the continuous parameter domain. Adaptive gridding, interval analysis, or analytical error bounds could be used in future work to provide a stronger mathematical guarantee.

Second, Eq. (50) describes the frozen frequency response at a fixed λ , whereas $\lambda(t)$ varies in an actual time-varying delay system. The scan therefore does not characterize spectral modulation caused by time-varying coefficients and cannot replace a full time-varying operator analysis.

Throughout this work, the term *grid maximum error* therefore refers specifically to the maximum complex-response error over the prescribed discrete frozen-response grid.

5.4. Benefit of higher order under noise-dominated conditions

From $q = 2$ to $q = 6$, the interpolation reconstruction-error energy decreases by more than six orders of magnitude, whereas the noisy total residual remains nearly unchanged. Together with the measurement-noise estimate and the PSD results, this behavior supports the relation

$$J_{noise} \gg J_{interp}$$

under the present configuration.

Once the interpolation error falls below the dominant stochastic noise contribution, increasing the interpolation order primarily provides additional numerical margin rather than improving the final output. Order selection should therefore not follow a simple “smaller numerical error is always better” rule, but should account for the other error sources and implementation resources of the overall delay-processing chain.

5.5. Interpretation of the complexity metric

Wall-clock time from long MATLAB simulations is not used as a formal complexity measure because it may be affected by just-in-time compilation, system load, memory state, cache behavior, and execution order.

The quantity $W(q)$ is instead a deterministic arithmetic-work proxy derived from the current scalar Lagrange implementation. Scalar additions, subtractions, multiplications, and divisions are each counted as one work unit. Indexing, data movement, cache behavior, SIMD vectorization, and processor-pipeline effects are excluded. The metric therefore cannot be extrapolated directly to execution time or power consumption on a spaceborne processor.

A mission-level implementation would require additional measurements or estimates of multiply-accumulate counts, coefficient storage, sample access, memory bandwidth, processing latency, throughput, and power consumption on a target processor or representative embedded platform.

5.6. Statistical interpretation

The common-random-number design causes the same trial ID to share the random inputs across both interpolation orders and spacecraft-position uncertainty scenarios. The bootstrap therefore uses trial ID as the joint resampling unit and preserves dependence across both order and σ .

Correcting the bootstrap pairing does not change the point estimates, but it does change the widths of the confidence intervals. For the interpolation reconstruction error, the three formal baseline orders remain separated by several orders of magnitude and the effect of interpolation order is clear. For the noisy total residual, the differences between order-dependent point estimates remain much smaller than the trial-to-trial stochastic variability.

Overlap of marginal confidence intervals should not be interpreted as a formal demonstration of statistical equivalence. If engineering equivalence between two interpolation orders is required for a system-level metric, an equivalence margin should be defined a priori and an appropriate equivalence or noninferiority analysis should be performed.

The empirical CVaR₉₅ is also determined primarily by approximately 5% of the samples, corresponding to about 40 observations for each formal baseline order. CVaR should therefore be treated as supporting evidence alongside the mean and spectral results rather than as the sole basis for order selection.

5.7. Integer-sample boundaries and the applicability of the present benchmark

Previous work has shown that, when a time-varying delay crosses an integer-sample boundary, the temporal continuity and smoothness of the interpolation kernel can affect delayed-signal reconstruction and can introduce additional transients or spectral features [20].

The fractional delays realized in the formal simulation lie within

$$\lambda \in [0.6139, 0.8103]$$

The simulation therefore does not include a transition in which the fractional part approaches unity and wraps to values near zero across an integer-sample boundary. The present results consequently characterize interpolation accuracy within the realized physical domain and the frozen response at fixed fractional delays. They do not validate the kernel smoothness of the different interpolation orders at integer-delay transitions and cannot be used to establish suppression of such glitch-like effects.

Future studies of more general arm-locking and TDI delay trajectories should therefore include explicit crossings of integer-sample boundaries and assess the resulting interpolation behavior in both the time and frequency domains.

5.8. Candidate-set completeness and the third-order configuration

The original formal Monte Carlo analysis used $q = 2, 4$, and 6 as baseline orders. This choice provides a useful view of the order-dependent reduction in interpolation error but does not guarantee that the lowest qualifying member of this sparse set is also the lowest feasible integer order.

The extended frozen-response scan gives a full-domain maximum error of

$$8.2973 \times 10^{-4}$$

for $q = 3$, which satisfies the 10^{-3} budget. Its deterministic arithmetic-work proxy,

$$W(3) = 60$$

is also appreciably lower than the 95 work units associated with $q = 4$. The third-order configuration is therefore a lower-complexity candidate that cannot be excluded from the design space.

At the same time, the present study does not yet contain formal random-source mean-error, PSD, or CVaR results for $q = 3$. The available evidence is therefore insufficient to establish that the third-order configuration is superior to the fourth-order configuration under the complete set of statistical metrics considered here. More generally, these results show that candidate-set design is itself part of the interpolation-order optimization problem: before a minimum-order conclusion is drawn from a finite set, the candidate set should be checked to ensure that it contains the potentially lowest feasible integer order.

5.9. Reproducibility

The formal Monte Carlo analysis used the following fixed configuration:

- base random seed: 20260401;
- common random numbers enabled across interpolation orders and σ ;
- 20 coarse-stage trials per formal baseline (q, σ) combination;
- 200 fine-stage trials per formal baseline (q, σ) combination;
- 2000 bootstrap replicates;
- joint trial-ID blocks as the bootstrap resampling units;
- 4280 effective samples;
- parallel execution disabled;
- formal Monte Carlo baseline orders $q = 2, 4$, and 6 .

The frozen-response analysis used 801 frequency points and 2001 fractional-delay points, and the candidate-set completeness check covered $q = 1-6$.

The code, configuration files, and aggregated results supporting the conclusions of this study are planned for release through a public data repository after acceptance. During peer review, they may be made available from the corresponding author upon reasonable request. No unverified repository link is provided before the archive has been formally established.

CONCLUSION

This study developed a lightweight numerical benchmark for time-varying digital delay processing motivated by space-based gravitational-wave arm locking. The benchmark comprises an analytically defined continuous-time composite source, a three-spacecraft linearized ephemeris, static position uncertainty, a fixed-point light-travel-time solution, backward Lagrange fractional-delay interpolation, weighted multi-link combination, residual decomposition, and spectral statistics.

Rather than evaluating only the noisy total residual, the proposed benchmark uses the continuous-time analytical delayed signal as a reference and explicitly separates the physical light-travel-time residual, interpolation reconstruction error, and measurement noise. The formal Monte Carlo analysis evaluated the baseline orders $q = 2, 4$, and 6 under four spacecraft-position uncertainty scenarios using 20 coarse-stage trials and 200 fine-stage trials per (q, σ) combination. Statistical uncertainty was evaluated using a paired trial-ID block bootstrap that preserves the dependence induced by common random numbers across both interpolation order and σ .

The mean band-limited relative interpolation errors for the second-, fourth-, and sixth-order baseline configurations were

$$2.9821 \times 10^{-4}, 4.9143 \times 10^{-6}, 1.2333 \times 10^{-7}$$

respectively. Both the mean and the upper-tail interpolation reconstruction errors decreased sharply with increasing order. In contrast, the in-band noisy total residual remained close to

$$8.37 \times 10^{-4}$$

for all three baseline orders, and their corresponding PSDs and upper-tail distributions were nearly identical. An independent estimate based on the measurement-noise level supports the conclusion that the current benchmark is strongly measurement-noise dominated.

Under the 10^{-3} error budget, the second-order interpolator satisfies the mean-error criterion within the formal Monte Carlo baseline set. The frozen-response grid scan, however, gives a full-domain maximum error of

$$3.4097 \times 10^{-3}$$

for $q = 2$, exceeding the prescribed grid budget. After expanding the original candidate set, the corresponding maximum error for $q = 3$ is

$$8.2973 \times 10^{-4}$$

which satisfies the criterion. Thus, within the scanned range $q = 1-6$,

$$q_{\text{grid}} = 3$$

is the lowest order satisfying the 10^{-3} frozen-response grid budget. The earlier conclusion that $q = 4$ was the lowest qualifying order arose from restricting the comparison to $q = 2, 4$, and 6 and should not be generalized to the complete integer-order design space.

The fourth-order interpolator further reduces the full-domain grid maximum to

$$2.2570 \times 10^{-4}$$

and is supported by the complete Monte Carlo, PSD, and CVaR analyses performed in this study. It can therefore be retained as a conservative reference configuration under the currently available evidence. The third-order interpolator, however, uses only four nodes and has a deterministic arithmetic-work proxy of 60 units, compared with 95 units for $q = 4$, while already satisfying the frozen-response grid criterion. It should therefore be prioritized for the next formal Monte Carlo evaluation before the engineering interpolation order is finally fixed.

The present evidence supports a two-level conclusion: **$q = 3$ is the lowest qualifying order identified by the frozen-response grid criterion, whereas $q = 4$ is the conservative reference configuration supported by the complete stochastic evidence currently available.** More broadly, the results show that interpolation-order selection for space-based time-varying delay processing should jointly consider mean interpolation error, high-error regions of the frequency–fractional-delay domain, the system noise level, implementation workload, and candidate-set completeness.

Future work will first extend the formal common-random-number Monte Carlo, PSD, and upper-tail analyses to $q = 3$. Additional studies will consider high-fidelity mission orbits, time-correlated orbit-determination errors, nonstationary measurement noise, integer-sample boundary crossings, closed-loop arm-locking models, and representative embedded processing platforms to assess the applicability of the present order-selection conclusions under more complete mission conditions.

ACKNOWLEDGEMENTS

This work was supported by the National Key R&D Program of China (Grant No. 2022YFC2204602); the National Natural Science Foundation of China (Grant Nos. 12505072 and 12150012); the Joint Funding Program of the China Postdoctoral Science Foundation and Hubei Province (Special Funding) (Grant No. 2025T005HB); and the China Postdoctoral Science Foundation (Grant Nos. 2024M760994 and GZB20250771).

AUTHOR CONTRIBUTIONS

Zhengrong Xiang: Conceptualization, methodology, software, formal analysis, investigation, visualization, and writing—original draft. Jie Luo: Methodology, validation, statistical analysis, and writing—review and editing. Jun Ke: Software, numerical validation, data analysis, and writing—review and editing. Haoxin Xu: Methodology, supervision, interpretation of results, and writing—review and editing. Tianchi Zhu: Conceptualization, supervision, project administration, funding acquisition, and writing—review and editing. All authors reviewed and approved the final manuscript.

REFERENCES

- [1] Colpi, M., et al., “LISA Definition Study Report,” arXiv:2402.07571 (2024).
- [2] Armano, M., et al. (LISA Pathfinder Collaboration), “In-depth analysis of LISA Pathfinder performance results: Time evolution, noise projection, physical models, and implications for LISA,” *Physical Review D* 110, 042004 (2024), doi: 10.1103/PhysRevD.110.042004.
- [3] Luo, Z., Hu, W., and Jin, G., “The Taiji program: A concise overview,” *Progress of Theoretical and Experimental Physics* 2021(5), 05A108 (2021), doi: 10.1093/ptep/ptaa083.
- [4] Luo, J., et al., “Progress of the TianQin project,” *Classical and Quantum Gravity* 42(17), 173001 (2025), doi: 10.1088/1361-6382/adda8a.
- [5] Tinto, M., Estabrook, F. B., and Armstrong, J. W., “Time-delay interferometry for LISA,” *Physical Review D* 65, 082003 (2002), doi: 10.1103/PhysRevD.65.082003.
- [6] Tinto, M. and Dhurandhar, S. V., “Time-delay interferometry,” *Living Reviews in Relativity* 24, 1 (2021), doi: 10.1007/s41114-020-00029-6.
- [7] Tinto, M., Dhurandhar, S. V., and Malakar, D., “Second-generation time-delay interferometry,” *Physical Review D* 107, 082001 (2023), doi: 10.1103/PhysRevD.107.082001.
- [8] Vallisneri, M., “Geometric time delay interferometry,” *Physical Review D* 72, 042003 (2005); Erratum, *Physical Review D* 76, 109903 (2007), doi: 10.1103/PhysRevD.72.042003.
- [9] Bayle, J.-B. and Hartwig, O., “Unified model for the LISA measurements and instrument simulations,” *Physical Review D* 107, 083019 (2023), doi: 10.1103/PhysRevD.107.083019.
- [10] Hartwig, O. and Muratore, M., “Characterization of time delay interferometry combinations for the LISA instrument noise,” *Physical Review D* 105, 062006 (2022), doi: 10.1103/PhysRevD.105.062006.
- [11] Wang, G., Ni, W.-T., Han, W.-B., and Qiao, C.-F., “Algorithm for time-delay interferometry numerical simulation and sensitivity investigation,” *Physical Review D* 103, 122006 (2021), doi: 10.1103/PhysRevD.103.122006.

- [12] Sheard, B. S., Gray, M. B., McClelland, D. E., and Shaddock, D. A., “Laser frequency stabilization by locking to a LISA arm,” *Physics Letters A* 320(1), 9–21 (2003), doi: 10.1016/j.physleta.2003.10.076.
- [13] McKenzie, K., Spero, R. E., and Shaddock, D. A., “Performance of arm locking in LISA,” *Physical Review D* 80, 102003 (2009), doi: 10.1103/PhysRevD.80.102003.
- [14] Thorpe, J. I., Maghami, P., and Livas, J., “Time domain simulations of arm locking in LISA,” *Physical Review D* 83, 122002 (2011), doi: 10.1103/PhysRevD.83.122002.
- [15] Wang, P.-P., Qian, W.-L., Wu, H.-Z., Tan, Y.-J., and Shao, C.-G., “Arm locking in conjunction with time-delay interferometry,” *Physical Review D* 106, 104042 (2022), doi: 10.1103/PhysRevD.106.104042.
- [16] Valliyakalayil, J. T., Wade, A., Rabeling, D., Zhang, J., Shaddock, D., and McKenzie, K., “Experimental demonstration of the combined arm- and cavity-locking system for LISA,” *Physical Review D* 110, 062004 (2024), doi: 10.1103/PhysRevD.110.062004.
- [17] Reinhardt, J. N., Staab, M., Yamamoto, K., Bayle, J.-B., Hees, A., Hartwig, O., Wiesner, K., Shah, S., and Heinzel, G., “Ranging sensor fusion in LISA data processing: Treatment of ambiguities, noise, and onboard delays in LISA ranging observables,” *Physical Review D* 109, 022004 (2024), doi: 10.1103/PhysRevD.109.022004.
- [18] Laakso, T. I., Välimäki, V., Karjalainen, M., and Laine, U. K., “Splitting the unit delay—Tools for fractional delay filter design,” *IEEE Signal Processing Magazine* 13(1), 30–60 (1996).
- [19] Farrow, C. W., “A continuously variable digital delay element,” in *Proceedings of the 1988 IEEE International Symposium on Circuits and Systems*, Vol. 3, 2641–2645 (1988).
- [20] Zhao, R. and Tay, D. B., “A complex exponential structure for low-complexity variable fractional delay FIR filters,” *Circuits, Systems, and Signal Processing* 42, 1105–1141 (2023), doi: 10.1007/s00034-022-02169-2.
- [21] Page, J. and Littenberg, T. B., “Bayesian time delay interferometry for orbiting LISA: Accounting for the time dependence of spacecraft separations,” *Physical Review D* 108, 044065 (2023), doi: 10.1103/PhysRevD.108.044065.
- [22] Staab, M., Lilley, M., Bayle, J.-B., and Hartwig, O., “Laser noise residuals in LISA from on-board processing and time-delay interferometry,” *Physical Review D* 109, 043040 (2024), doi: 10.1103/PhysRevD.109.043040.
- [23] Staab, M., Bayle, J.-B., Hartwig, O., Hees, A., Lilley, M., Woan, G., and Wolf, P., “Optimal design of interpolation methods for time-delay interferometry,” *Classical and Quantum Gravity* 42(11), 115002 (2025), doi: 10.1088/1361-6382/add706.
- [24] Burke, O., Muratore, M., and Woan, G., “Impact of missing data on the construction of LISA time-delay-interferometry Michelson variables,” *Physical Review Applied* 25, 034041 (2026), doi: 10.1103/xpvf-syrw.