

Article

A Structure-Preserving Power-Adaptive Colored-Noise Kalman Filtering Algorithm for Four-Wire Pendulum Velocity Monitoring

Dehui Meng ^{1,2} , Yongkun Chen ^{1,2,*}, Tao Yu ¹ , Huadong Li ^{1,2}, Qi Li ¹ and Zhi Wang ¹

¹ Changchun Institute of Optics, Fine Mechanics and Physics, Chinese Academy of Sciences, Changchun 130033, China; mdhui813@163.com (D.M.); yut@ciomp.ac.cn (T.Y.); lihuadong1995@163.com (H.L.); 15948157799@163.com (Q.L.); wz070611@126.com (Z.W.)

² University of Chinese Academy of Sciences, Beijing 100049, China

* Correspondence: chenrongkun@ciomp.ac.cn

Abstract

Ground testing of space-based gravitational-wave missions requires causal and low-noise horizontal-velocity estimation for vibration-isolation platforms. Fixed sensor fusion becomes suboptimal when sensor noise power varies, especially when complementary sensors have strongly frequency-dependent colored-noise spectra. This paper proposes a colored-noise power-adaptive Kalman filtering algorithm (CN-PAKF) for fusing laser interferometric displacement sensor (IFO) and broadband seismometer (SEIS) observations. The method constructs a motion-canceling synchronous difference channel and uses band-decomposed power statistics with asymmetric spectral priors to estimate sensor noise-power scales online. Scaling only the corresponding colored-noise drive covariances preserves the normalized spectral shapes of the modeled colored components and avoids unrestricted covariance adaptation; bounded physical-process and residual measurement-covariance safeguards may alter the total modeled error spectrum. In paired open-loop simulations, CN-PAKF reduced the target-band velocity RMSE by 31.3% and 70.5% relative to the IFO and SEIS observations under nominal noise conditions, while remaining within 1% of causal colored-noise baseline filters. When the IFO noise power increased 30-fold, CN-PAKF reduced the stage-wise target-band RMSE by 36.8% compared with a fixed colored-noise Kalman filter. These results indicate that CN-PAKF provides a causal, structure-preserving, and performance-tested fusion algorithm for synchronized sensors with time-varying noise power and approximately fixed colored-noise spectral shapes.

Keywords: adaptive Kalman filtering; colored-noise modeling; adaptive sensor fusion; online covariance adaptation; algorithm performance testing



Academic Editor: Wanquan Liu

Received: 31 August 2026

Revised: 22 September 2026

Accepted: 28 September 2026

Published: 30 September 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article

distributed under the terms and

conditions of the [Creative Commons](https://creativecommons.org/licenses/by/4.0/)

[Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

Space-based gravitational-wave missions rely on precision inertial sensors and optical telescopes whose intrinsic performance is difficult to reproduce and evaluate on the ground [1–4]. Low-disturbance platforms and pendulum-based apparatuses help separate instrument behavior from environmental motion [5–7]. The four-wire pendulum provides one such platform [8]. Its residual horizontal motion must be monitored across a broad low-frequency band. Sky-hook damping and related feedback strategies also require a causal estimate of horizontal absolute velocity [9–11]. Estimation noise therefore limits monitoring sensitivity and may re-enter a future feedback loop.

The target is a causal, low-noise velocity estimate over 0.01–1 Hz. The laser interferometric displacement sensor (IFO) has lower velocity-equivalent noise at the low-frequency end. By contrast, the broadband seismometer (SEIS) is more informative at middle and higher frequencies [10–13]. Neither channel is uniformly preferable across the target band. Their complementary spectra require a fusion rule that preserves frequency-dependent confidence instead of assigning fixed broadband weights. This complementarity motivates a fusion algorithm that uses a motion-canceling synchronous difference statistic to remove shared platform motion and spectral-prior-guided band-power statistics to attribute noise-power changes to individual sensors.

Active and hybrid vibration-isolation systems commonly combine inertial and relative sensors to extend usable bandwidth and mitigate the limitations of individual instruments [9–11,14–16]. Classical complementary fusion derives frequency-dependent weights from sensor noise spectra, and optimized formulations can approach a minimum-variance combination when these spectra are known [10,11]. These methods establish the value of spectral complementarity, but their weights generally remain fixed at nominal operating conditions. Sensor-condition monitoring can indicate degradation [17], and observer-based sensor-fusion architectures can improve the reliability of noisy measurements through model-based virtual sensors [18]. However, a separate condition indicator does not provide a causal, frequency-structured fusion rule.

State observers and Kalman filters provide a causal framework for vibration estimation and can directly incorporate mechanical dynamics [9,19–21]. A white measurement-noise approximation becomes inadequate when sensor spectra vary strongly with frequency. Colored measurement noise can instead be generated by shaping filters and represented through additional states, allowing structured sensor noise to be distinguished from the physical state [22,23]. Adaptive colored-noise filters further show that such models can accommodate changing statistics [24]. However, augmented states alone do not determine how frequency-dependent confidence should change when one sensor degrades.

Nominal spectra become insufficient when operating conditions change. Alignment, local disturbances, and sensor operating states can alter overall noise power while leaving the principal spectral shape approximately unchanged. Innovation matching, Sage–Husa-type updates, maximum-likelihood estimation, and related schemes infer uncertain process or measurement covariances from online estimation statistics [9,25–27]. Robust gates can prevent unusually large innovations from immediately contaminating these updates [28]. Related adaptive-filtering algorithms also emphasize the trade-off between online update selectivity, convergence, and computational cost [29]. Such approaches are useful when covariance structure is unknown, but unrestricted adaptation may absorb platform motion or model mismatch into the sensor-noise model and alter spectral relationships fitted offline. A constrained formulation is therefore needed to identify a degraded sensor while retaining the offline spectral information. The present problem is restricted to sensor-level power variation under an approximately fixed colored-noise shape.

This problem is addressed by the colored-noise power-adaptive Kalman filtering algorithm (CN-PAKF). Offline spectral fits determine a causal state-space realization of each sensor’s nominal colored-noise shape. Online, synchronous IFO and SEIS observations are differenced to cancel their shared platform velocity. The squared low- and high-frequency powers of this difference provide motion-independent statistics, and fixed asymmetric spectral priors attribute the observed power changes to the IFO and SEIS channels. These statistics recover the two sensor noise-power scales. Only the corresponding colored-noise drive covariances are scaled, which changes their power without moving their poles or altering their normalized spectral shapes. Bounded updates of the physical process

covariance and small residual measurement covariance serve only as implementation safeguards. They are not treated as independent sources of performance improvement.

Three contributions are made. First, two-sensor fusion is formulated for a time-varying-power colored-noise model with approximately fixed colored-noise spectral shapes. The method uses a motion-canceling difference statistic to reject shared platform motion and a spectral-prior-guided band-power inversion to distinguish sensor noise-power changes. The resulting spectral priors are embedded in a causal augmented-state estimator. Second, a motion-canceling, band-decomposed difference channel estimates sensor noise-power scales. Neither the unknown reference velocity nor Kalman innovations are used in these scale estimates. Third, CN-PAKF is evaluated in 30 paired open-loop simulations with independent sensor-noise realizations driven by the same measured ground-motion test record. Nominal target-band performance remains within 1% of the causal colored-noise baselines. The time-varying evaluation quantifies stage-wise target-band improvement under IFO power changes. Sensor noise-power-scale tracking is also examined under SEIS and dual-channel power changes. The scope is limited to synchronized observations, a nominal start-up interval, and approximately fixed colored-noise spectral shapes. Hardware timing and closed-loop control remain outside the present study.

2. Problem Formulation

2.1. Platform Dynamics and Synchronous Observations

Horizontal motion within the target band is represented by a base-driven second-order system [30]. The state comprises displacement $q(t)$ and absolute velocity $v(t)$ for one horizontal isolation mode. The base motion is unmeasured because only two sensor observations are available to the estimator. The sampling interval is 0.01 s.

$$\mathbf{x}_c(t) = \begin{bmatrix} q(t) \\ v(t) \end{bmatrix}. \quad (1)$$

The reference plant is governed by

$$\dot{\mathbf{x}}_c(t) = \mathbf{A}_c \mathbf{x}_c(t) + \mathbf{G}_c u_b(t), \quad (2)$$

where

$$\mathbf{A}_c = \begin{bmatrix} 0 & 1 \\ -\omega_n^2 & -2\zeta\omega_n \end{bmatrix}, \quad \mathbf{G}_c = \begin{bmatrix} 0 \\ 1 \end{bmatrix}, \quad (3)$$

$$u_b = \omega_n^2 q_b + 2\zeta\omega_n v_b.$$

The corresponding base-to-table displacement transmissibility is

$$H(j\omega) = \frac{\omega_n^2 + j2\zeta\omega_n\omega}{\omega_n^2 - \omega^2 + j2\zeta\omega_n\omega}. \quad (4)$$

Equation (4) describes the horizontal response of the four-wire pendulum to base motion. It predicts resonant amplification near the natural frequency and increasing isolation above resonance.

Here, ω_n and ζ denote the natural angular frequency and damping ratio. The base displacement and velocity are $q_b(t)$ and $v_b(t)$, respectively. Table 1 lists the apparatus parameters. The reference trajectory is generated using $u_b(t)$ from the measured ground record. This forcing is unavailable to the online filter as a known input. Instead, process noise represents the forcing, parameter mismatch, and model error. The matrix $\mathbf{G}_c$ indicates that the forcing enters through velocity.

Table 1. Core platform and CN-PAKF parameters.

Category	Parameter	Value
Pendulum platform	Platform mass m , pendulum length l_p	167.70 kg, 1.49 m
Pendulum platform	Natural frequency f_n , damping ratio ζ	0.40817 Hz, 1.2088×10^{-3}
Sampling	f_s	100 Hz
Band decomposition	f_{lo}, f_{hi}	0.2 Hz, 1.0 Hz
Power estimation	N_{lo}, N_{hi}	20,000, 2000
Start-up calibration	N_{cal}	40,000
Sensor noise-power scale	$[\alpha_{min}, \alpha_{max}]$	$[10^{-2}, 10^3]$
Auxiliary bounds	$[\beta_{min}, \beta_{max}], \gamma$	$[1/3, 3], 9$

With $\mathbf{F}_p = \exp(\mathbf{A}_c T_s)$, matrix-exponential discretization at T_s gives the online propagation model

$$\mathbf{x}_{k+1} = \mathbf{F}_p \mathbf{x}_k + \mathbf{w}_{p,k}. \quad (5)$$

Here, $\mathbf{x}_k = [q_k, v_k]^T$ denotes the discrete physical state. The process noise is $\mathbf{w}_{p,k} \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_p)$.

The baseline covariance $\mathbf{Q}_p^0$ is calibrated offline from the independent reference record described in Section 4.1. It represents unmeasured forcing, parameter mismatch, and state-propagation error. Velocity-only forcing produces an anisotropic, nearly rank-one covariance. The online scalar update in Section 3 preserves this calibrated shape.

Both the control readout and the performance metric are based on v_k . Process-noise adaptation affects only the physical covariance $\mathbf{Q}_p$.

The observation model expresses both sensors in the velocity domain. IFO displacement is converted to velocity during preprocessing, whereas the SEIS measures velocity directly [11,12]. Both observations share a common sampling grid. Their online construction requires neither a reference state nor future samples.

Let $\mathcal{S} = \{\text{IFO}, \text{SEIS}\}$, and let $i \in \mathcal{S}$ index a sensor throughout. With $\mathbf{H}_p = \begin{bmatrix} 0 & 1 \end{bmatrix}$ and $\mathbf{1}_2 = \begin{bmatrix} 1 & 1 \end{bmatrix}^T$, the two observations can be stacked as

$$\mathbf{y}_k = \begin{bmatrix} y_k^{(\text{IFO})} \\ y_k^{(\text{SEIS})} \end{bmatrix} = \mathbf{1}_2 \mathbf{H}_p \mathbf{x}_k + \boldsymbol{\varepsilon}_k. \quad (6)$$

Here, $\boldsymbol{\varepsilon}_k = \text{col}(\varepsilon_k^{(\text{IFO})}, \varepsilon_k^{(\text{SEIS})})$. Let $\mathbf{J} = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ exchange the two channels, and let $\mathbf{T} = 2^{-1/2} \begin{bmatrix} 1 & 1 \\ 1 & -1 \end{bmatrix}$ be the orthogonal symmetry transform. Then,

$$\begin{bmatrix} y_k^{(+)} \\ y_k^{(-)} \end{bmatrix} = \mathbf{T} \mathbf{y}_k = \begin{bmatrix} \sqrt{2} \mathbf{H}_p \mathbf{x}_k + \varepsilon_k^{(+)} \\ \varepsilon_k^{(-)} \end{bmatrix}. \quad (7)$$

Here, $\varepsilon_k^{(\pm)} = (\varepsilon_k^{(\text{IFO})} \pm \varepsilon_k^{(\text{SEIS})})/\sqrt{2}$. Since $\mathbf{T}\mathbf{J} = \text{diag}(1, -1)\mathbf{T}$, channel exchange leaves $y_k^{(+)}$ unchanged and reverses the sign of $y_k^{(-)}$. Thus, $y_k^{(+)}$ is the symmetric common-mode component and contains the shared platform velocity. The antisymmetric component $y_k^{(-)}$ contains only sensor error, so it cancels the unknown platform motion.

IFO differentiation is causal and uses the same first-order discrete difference as the spectral realization. Before fusion, the SEIS record is aligned to the same time stamps. These operations preserve the causal interpretation of the estimator. They also prevent temporal misalignment from introducing apparent colored noise.

Both observation errors are colored, but their spectra differ. IFO velocity noise increases with frequency, whereas SEIS velocity noise increases toward low frequencies. A white-noise model discards this contrast and may assign confidence incorrectly across the target band.

2.2. Approximately Fixed Spectral-Shape Noise Model and Estimation Objective

Let $S_{n,i}^0(f)$ denote the nominal one-sided power spectral density (PSD) of the structured colored-noise component for sensor i . With $\alpha_{i,k} > 0$, the operating regime is defined by

$$S_{n,i,k}(f) = \alpha_{i,k} S_{n,i}^0(f). \quad (8)$$

The sensor noise-power scale $\alpha_{i,k}$ changes the power of the structured component without changing its normalized spectral shape. The total measurement error also includes the small residual component defined in Equation (12). The structured IFO and SEIS noises, together with their residual components, are assumed to be mutually uncorrelated across channels. The sensor noise-power scale may remain constant for long intervals or change between operating stages. Because changes are observed through a finite-memory estimator, abrupt transitions produce a finite tracking delay. Equation (8) is a modeling assumption rather than a claim that all in-operation disturbances preserve their spectral shape.

For compact notation, the subscripts L and I denote the IFO (laser) and SEIS (inertial) channels, respectively. The start-up normalization interval is assumed to represent the nominal sensor condition, with $\alpha_{L,k} = \alpha_{I,k} = 1$ throughout the normalization interval. Otherwise, the recovered quantities should be interpreted as effective scales relative to the start-up reference rather than absolute multipliers of the datasheet priors.

The estimator alone cannot distinguish a degraded start-up interval from a changed absolute noise reference without additional information. In practice, nominality of the start-up interval can be checked against archived band-power references or an independent sensor-health assessment. If no verified nominal interval is available, the normalization should be treated explicitly as a reference-state calibration, with the recovered scales interpreted relative to that reference. Unequal initial degradation of the two channels can also bias channel attribution when the nominal mixing matrix is retained. A later verified nominal interval can be used to re-estimate b_B and rebase the scales.

Given synchronous observations through time k , the objective is to estimate $\hat{v}_k$, $\hat{\alpha}_{L,k}$, and $\hat{\alpha}_{I,k}$ causally. Neither the reference velocity nor the unmeasured base forcing is available. The estimator must preserve nominal frequency dependence while reducing the contribution of a channel whose sensor noise-power scale increases. Auxiliary covariance updates must remain within prescribed bounds. Frequency-shape changes, asynchronous observations, hardware timing, and closed-loop control performance remain outside the present formulation. Controlled spectral-shape perturbations are evaluated separately in Section 5.4 to delimit this operating assumption; they do not extend the estimator to arbitrary time-varying spectral shapes.

The two-sensor observation and approximately fixed spectral-shape colored-noise construction are summarized in Figure 1. Channels 1 and 2 represent the IFO and SEIS, respectively, and both observe the shared platform velocity v . For each channel, white noise w_i drives the fitted shaping filter $G_i(z)$. The sensor noise-power scale α_i multiplies the drive covariance, so the filter output n_i changes in power while retaining its fitted spectral shape. Adding n_i to v gives the velocity-domain observation y_i . The small white residual e_i remains in the mathematical model but is omitted from the schematic.




Figure 1. Two-sensor observation model illustrating the motion-canceling difference statistic and sensor spectral priors. Channels 1 and 2 are the laser interferometric displacement sensor (IFO) and broadband seismometer (SEIS), respectively.

3. Proposed CN-PAKF

3.1. Overview, Spectral-Prior Realization, and Augmented-State Model

CN-PAKF separates offline spectral characterization from online sensor-power adaptation. Offline, the nominal IFO and SEIS spectra determine the fixed colored-noise shaping filters and the band-mixing relation. During operation, synchronized IFO and SEIS observations are preprocessed and stacked into the velocity-domain observation vector. Figure 2 provides a system-level overview of the signal path from the four-wire pendulum measurements to the causal velocity estimate.




Figure 2. System-level overview of CN-PAKF sensor fusion for four-wire pendulum velocity monitoring.

The synchronous difference d_k suppresses the common platform-velocity component. Its low- and high-band powers are jointly interpreted using the fixed asymmetric spectral priors to estimate the sensor noise-power scales. These scales modify only the corresponding colored-noise drive covariances, preserving the normalized spectral shapes, before augmented-state Kalman filtering produces the causal velocity estimate. The physical-

process scale and small residual measurement covariance are also updated within prescribed bounds; these auxiliary safeguards do not contribute to identification of the sensor noise-power scales. The following subsections formalize the spectral-prior realization, the difference-channel scale estimation, and the causal covariance recursion.

Nominal amplitude spectral density (ASD) curves are obtained from the IFO and SEIS datasheets. Each curve is converted from its native output to velocity and then squared to obtain the one-sided velocity PSD. These spectra serve as offline priors rather than in-operation measurements. Their slopes determine the sensor that is more informative within each frequency range. Log-power fitting is used because the PSDs span several decades. This representation prevents the largest values from dominating the fit.

The nominal PSD shape is represented by a non-negative repeated-pole model:

$$S_{n,i}^0(f) = \sum_{m=0}^{M_i} c_{i,m} \Psi_{i,m}(f). \quad (9)$$

The non-negative weights $c_{i,m}$ combine displacement-domain bases with the velocity factor $(2\pi f)^2$, where

$$\Psi_{i,m}(f) = \begin{cases} (2\pi f)^2, & m = 0, \\ \frac{(2\pi f)^2}{|1 - \rho_{i,m} e^{-j2\pi f / f_s}|^{2r_{i,m}}}, & m \geq 1. \end{cases} \quad (10)$$

The $m = 0$ branch represents displacement white noise. Each $m \geq 1$ branch has order $r_{i,m}$ and radius $\rho_{i,m} = \exp(-2\pi f_{i,m} / f_s)$. Log power is fitted by non-negative least squares, and negligible branches are removed by thresholding. A first-order backward difference is applied within each retained branch. This construction avoids differentiating the combined noisy observation. At the 100 Hz sampling rate, the discrete power gain differs from $(2\pi f)^2$ by less than 4% over 0.01–10 Hz.

Each branch is expressed through the state-space realization [22,23]

$$\begin{aligned} \xi_{i,m,k+1} &= \mathbf{F}_{i,m} \xi_{i,m,k} + \mathbf{G}_{i,m} w_{i,m,k}, \\ \eta_{i,m,k} &= \mathbf{H}_{i,m} \xi_{i,m,k}. \end{aligned} \quad (11)$$

The drive satisfies $w_{i,m,k} \sim \mathcal{N}(0, q_{i,m})$, where $q_{i,m} = c_{i,m} f_s / 2$. Each branch contains $r_{i,m}$ pole states and one memory state. Stacking the branches gives $\xi_{i,k+1} = \mathbf{F}_{n,i} \xi_{i,k} + \mathbf{w}_{n,i,k}$. The corresponding output is $n_{i,k} = \mathbf{H}_{n,i} \xi_{i,k} = \sum_m \eta_{i,m,k}$. The fitted branch-drive variances are stored in the diagonal matrix $\mathbf{Q}_{n,i}^0$.

The sensor error is decomposed into fitted colored noise and a small unmodeled residual:

$$\begin{aligned} \varepsilon_k^{(i)} &= n_{i,k} + e_k^{(i)}, \\ \mathbf{e}_k &= \text{col}(e_k^{(\text{IFO})}, e_k^{(\text{SEIS})}) \sim \mathcal{N}(\mathbf{0}, \mathbf{R}). \end{aligned} \quad (12)$$

The residual $e_k^{(i)}$ represents fitting error, quantization, and other effects that are not captured by the shaping filters. The initial diagonal of $\mathbf{R}$ is 10^{-6} times the corresponding nominal sensor variance.

The resulting total measurement-error PSD is

$$S_{\varepsilon,i,k}(f) = \alpha_{i,k} S_{n,i}^0(f) + S_{e,i,k}(f). \quad (13)$$

Here, $S_{e,i,k}(f)$ is the small white-residual contribution associated with $\mathbf{R}_k$. Consequently, structure preservation applies to the modeled colored component. It does not apply to the total error spectrum when the residual term changes.

Figure 3 shows that the fitted models reproduce the principal frequency dependence used by the estimator. This comparison supports the nominal shape model but does not test whether in-operation disturbances preserve that shape.

The physical state and both colored-noise states are stacked as

$$\mathbf{z}_k = \text{col}(\mathbf{x}_k, \boldsymbol{\xi}_{\text{IFO},k}, \boldsymbol{\xi}_{\text{SEIS},k}). \quad (14)$$

The augmented dynamics are

$$\mathbf{z}_{k+1} = \mathbf{F}\mathbf{z}_k + \mathbf{w}_k. \quad (15)$$

The transition matrix and nominal covariance have the block-diagonal forms

$$\begin{aligned} \mathbf{F} &= \text{blkdiag}(\mathbf{F}_p, \mathbf{F}_{n,\text{IFO}}, \mathbf{F}_{n,\text{SEIS}}), \\ \mathbf{Q}^0 &= \text{blkdiag}(\mathbf{Q}_p^0, \mathbf{Q}_{n,\text{IFO}}^0, \mathbf{Q}_{n,\text{SEIS}}^0). \end{aligned} \quad (16)$$

The three blocks correspond to the physical state, the IFO colored-noise state, and the SEIS colored-noise state, respectively.




Figure 3. Nominal and fitted velocity-domain amplitude spectral density (ASD) curves for the IFO and SEIS. Solid and dashed lines show targets and fits; annotations report root-mean-square (RMS) fitting errors in decibels.

The two observations are described by

$$\mathbf{y}_k = \mathbf{H}\mathbf{z}_k + \mathbf{e}_k. \quad (17)$$

The first row of $\mathbf{H}$ is $[\mathbf{H}_p, \mathbf{H}_{n,\text{IFO}}, \mathbf{0}]$. Its second row is $[\mathbf{H}_p, \mathbf{0}, \mathbf{H}_{n,\text{SEIS}}]$. Thus, each observation explicitly contains the shared physical velocity and its sensor-specific colored-noise state. The disturbances satisfy $\mathbf{w}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_k)$ and $\mathbf{e}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{R}_k)$. Quantities computed from sample k determine $\mathbf{Q}_k$ and $\mathbf{R}_k$ for the next prediction. The Joseph form is used for covariance correction to preserve symmetry and positive semidefiniteness to numerical precision [31].

3.2. Motion-Canceling Difference Statistics and Sensor Noise-Power Scale Estimation

Kalman innovations contain the common physical prediction error and depend on the previously selected covariances. Estimating sensor noise-power scales from these innovations would couple the scale statistic to the estimator it modifies. The synchronous difference instead yields

$$\begin{aligned} d_k &= \sqrt{2} y_k^{(-)} = y_k^{(\text{IFO})} - y_k^{(\text{SEIS})} \\ &= n_{\text{IFO},k} - n_{\text{SEIS},k} + e_k^{(\text{IFO})} - e_k^{(\text{SEIS})}. \end{aligned} \quad (18)$$

Under the synchronous observation model, the shared platform velocity cancels exactly. The resulting difference is an open-loop statistic of the two sensor errors rather than an estimation residual. Under the cross-channel uncorrelatedness assumption, the expected band power of the difference equals the sum of the two sensor contributions. The start-up normalization includes the nominal residual power, whereas subsequent residual variations are neglected in the scale inversion because the residual covariance is tightly bounded.

Equation (7) shows that channel exchange reverses the sign of $y_k^{(-)}$ while leaving $y_k^{(+)}$ unchanged. Equation (18) uses the corresponding unnormalized antisymmetric component. Because each $B_B(z)$ is linear, channel exchange maps both d_k and $d_{B,k}$ to their negatives. Their squared values and hence each $P_{B,k}$ remain unchanged. Thus, the adaptive power statistic is exchange invariant and does not prefer a sensor by label.

SEIS and IFO noise dominate below and above their 0.26 Hz crossover, respectively. Two low-pass stages suppress IFO leakage into the low-frequency statistic. A complementary high-pass output provides the second statistic. Let $\mathcal{B} = \{\text{lo}, \text{hi}\}$, and let $B \in \mathcal{B}$ index a band. Define $G_B(z) = \kappa_B / [1 - (1 - \kappa_B)z^{-1}]$, where $\kappa_B = 1 - \exp(-2\pi f_B / f_s)$. Let $D(z)$ and $D_B(z)$ denote the transforms of d_k and $d_{B,k}$, respectively. Then,

$$\begin{aligned} B_{\text{lo}}(z) &= G_{\text{lo}}^2(z), & B_{\text{hi}}(z) &= 1 - G_{\text{hi}}(z), \\ D_B(z) &= B_B(z)D(z). \end{aligned} \quad (19)$$

The band powers are estimated recursively by

$$P_{B,k} = \left(1 - \frac{1}{N_B}\right) P_{B,k-1} + \frac{d_{B,k}^2}{N_B}. \quad (20)$$

The low band contains fewer effectively independent samples; therefore, $N_{\text{lo}} > N_{\text{hi}}$. The nominal energy contributed by sensor i to each statistic is

$$E_{B,i} = \int_0^{f_s/2} S_{n,i}^0(f) \left| B_B(e^{j2\pi f / f_s}) \right|^2 df. \quad (21)$$

The corresponding fractional contribution is

$$f_{B,i} = \frac{E_{B,i}}{\sum_{j \in \mathcal{S}} E_{B,j}}. \quad (22)$$

The row-stochastic mixing matrix is defined by $M_{B,i} = f_{B,i}$. The sensor noise-power scale vector is $\alpha_k = [\alpha_{L,k}, \alpha_{I,k}]^T$.

The normalized band-power vector is $\mathbf{p}_k = [P_{\text{lo},k}/b_{\text{lo}}, P_{\text{hi},k}/b_{\text{hi}}]^T$.

For finite-memory band-power estimates, the scale model and its bounded inversion are

$$\mathbf{p}_k \simeq \mathbf{M} \alpha_k, \quad \hat{\alpha}_k = \text{clip}(\mathbf{M}^+ \mathbf{p}_k, \alpha_{\min}, \alpha_{\max}). \quad (23)$$

The operator $\text{clip}(\cdot, a, b)$ applies the lower and upper bounds a and b componentwise. The offline spectra and band responses determine $\mathbf{M}$. A start-up window of length N_{cal} estimates b_{B} while both sensor noise-power scales are nominally one. This normalization requires neither the reference velocity nor the estimation error. The operator $\mathbf{M}^+$ denotes the inverse for nonsingular $\mathbf{M}$ and the Moore–Penrose pseudoinverse otherwise. Hats are omitted when the recovered scales are inserted into the covariance recursion below. The IFO sensor noise-power scale primarily follows the N_{hi} memory. The SEIS sensor noise-power scale responds according to the longer N_{lo} memory.

The band powers are invariant under channel exchange, but they do not by themselves identify which sensor changed. For the row-stochastic 2×2 matrix $\mathbf{M}$, define the spectral-attribution index

$$\Delta_{\text{spec}} \equiv \det(\mathbf{M}) = f_{\text{lo},L} - f_{\text{hi},L}. \quad (24)$$

The second equality follows by expanding the determinant and using $f_{\text{B},I} = 1 - f_{\text{B},L}$. Exchanging the sensor labels permutes the columns, $\mathbf{M} \mapsto \mathbf{M}\mathbf{J}$. Since $\det(\mathbf{J}) = -1$, this operation reverses the sign of Δ_{spec} while leaving $|\Delta_{\text{spec}}|$ invariant. The two scale factors are separately identifiable when $\Delta_{\text{spec}} \neq 0$. If the normalized sensor spectra produced identical fractions in the selected bands, the columns of $\mathbf{M}$ would coincide and Δ_{spec} would vanish. CN-PAKF therefore combines motion-canceling difference statistics with spectral-prior-guided channel attribution.

For the offline spectra and bands used here, with rows ordered as the low- and high-frequency bands and columns as the IFO and SEIS channels, the realized mixing matrix was

$$\mathbf{M} = \begin{bmatrix} 0.0049376 & 0.9950624 \\ 0.9998809 & 0.0001191 \end{bmatrix}. \quad (25)$$

Its signed determinant was -0.9949433 , and its singular values were $\sigma_{\text{max}} = 1.000967$ and $\sigma_{\text{min}} = 0.993982$. Hence, $\kappa_2(\mathbf{M}) = 1.007027$ and $\|\mathbf{M}^{-1}\|_2 = 1.006054$, indicating that the spectral-attribution inversion is well conditioned for the selected spectra and band decomposition. This conclusion is specific to the present offline spectra and does not follow from $\det(\mathbf{M}) \neq 0$ alone.

3.3. Structure-Preserving Covariance Adaptation and Online Recursion

Consistent with the nominal block structure in Equation (16), the online process covariance is

$$\mathbf{Q}_k = \text{blkdiag}(\beta_k \mathbf{Q}_{\text{p}}^0, \alpha_{\text{L},k} \mathbf{Q}_{\text{n,IFO}}^0, \alpha_{\text{I},k} \mathbf{Q}_{\text{n,SEIS}}^0). \quad (26)$$

The term “structure-preserving” follows directly from the linear shaping-filter model. For fixed $\mathbf{F}_{\text{n},i}$ and $\mathbf{H}_{\text{n},i}$, the output PSD depends linearly on the drive covariance. Therefore,

$$\mathbf{Q}_{\text{n},i,k} = \alpha_{i,k} \mathbf{Q}_{\text{n},i}^0 \implies S_{\text{n},i,k}(f) = \alpha_{i,k} S_{\text{n},i}^0(f). \quad (27)$$

Scalar covariance multiplication changes the structured-noise power without altering the shaping-filter poles, spectral slopes, or normalized PSD. This invariance is an algebraic property of the estimator. The separate residual contribution in Equation (13) remains bounded but lies outside the shape-preservation statement. Accordingly, “structure-preserving” refers only to the modeled colored components in Equation (27); changes in β_k or $\mathbf{R}_k$ may alter the total predicted error spectrum and are not covered by this claim. Compliance of a physical sensor with the approximately fixed spectral-shape model remains the operating assumption in Equation (8). Unity sensor noise-power scales recover the offline colored-noise model.

Two auxiliary updates accommodate physical-model mismatch and small unmodeled white residuals. Both are bounded to prevent them from replacing the sensor-specific colored-noise model. Their aggregate contribution relative to the sensor noise-power-scale update is evaluated by the component ablation in Section 5.2. They are treated as implementation safeguards rather than independent contributions.

The physical process-noise scale β_k adjusts the physical block within a trust region. It changes disturbance power while preserving the nearly rank-one shape of $\mathbf{Q}_p^0$. A heuristic mismatch trend is

$$\tilde{\mathbf{Q}}_k = \mathbf{K}_k \mathbf{v}_k \mathbf{v}_k^T \mathbf{K}_k^T + \mathbf{P}_{k|k} - \bar{\mathbf{P}}_{k|k-1}, \quad (28)$$

where $\bar{\mathbf{P}}_{k|k-1} = \mathbf{F} \mathbf{P}_{k-1|k-1} \mathbf{F}^T$ omits process noise. Equation (28) indicates the mismatch direction but is not interpreted as a consistent covariance estimate [26,27]. Its physical-state block gives

$$\beta_k^{\text{raw}} = \frac{\text{tr}(\tilde{\mathbf{Q}}_{p,k})}{\text{tr}(\mathbf{Q}_p^0)}. \quad (29)$$

For $n_p = 2$, $\tilde{\mathbf{Q}}_{p,k} = (\tilde{\mathbf{Q}}_k)_{1:n_p, 1:n_p}$. Trace matching avoids a full-matrix update when finite-sample candidates are indefinite. The bounded update is

$$\beta_k = \text{clip}(\lambda_k \beta_{k-1} + (1 - \lambda_k) \beta_k^{\text{raw}}, \beta_{\min}, \beta_{\max}). \quad (30)$$

For $k \geq k_a$, let $N_k = \min(k - k_a + 1, L)$ and $\lambda_k = (N_k - 1)/N_k$, where k_a is the first eligible auxiliary-update sample and $L = N_{\text{hi}}$. Thus, the first eligible update has $\lambda_k = 0$. The averaging therefore transitions from cumulative memory to the fixed factor $1 - 1/L$. The interval is $[1/3, 3]$ [24]. Rejected candidates retain the previous value.

The residual measurement covariance $\mathbf{R}_k$ is reserved for a small diagonal white residual after the colored components are represented by augmented states. Innovation matching gives [25,27]

$$r_{i,k}^{\text{cand}} = \lambda_k r_{i,k-1} + (1 - \lambda_k) (\mathbf{C}_{R,k})_{ii}, \quad (31)$$

where $\hat{\mathbf{C}}_{v,k} = \lambda_k \hat{\mathbf{C}}_{v,k-1} + (1 - \lambda_k) \mathbf{v}_k \mathbf{v}_k^T$ and $\mathbf{C}_{R,k} = \hat{\mathbf{C}}_{v,k} - \mathbf{H} \mathbf{P}_{k|k-1} \mathbf{H}^T$. The candidate is clipped as

$$r_{i,k}^{\text{clip}} = \text{clip}(r_{i,k}^{\text{cand}}, r_i^{\min}, r_i^{\max}). \quad (32)$$

The initial, lower, and upper levels are 10^{-6} , 10^{-8} , and 10^{-5} of the corresponding nominal sensor variance. Only the diagonal terms are updated. A normalized per-sensor innovation gate is defined as [28]

$$\tau_{i,k} = \frac{v_{i,k}^2}{(\mathbf{S}_k)_{ii}}. \quad (33)$$

The gate uses the 3σ threshold $\gamma = 9$. The residual entry $r_{i,k}$ is updated only when $\tau_{i,k} \leq \gamma$. The factor β_k is updated only when both channels pass. Failed gates retain their previous values. The open-loop sensor noise-power scales $\alpha_{i,k}$ do not use this innovation gate.

Figure 4 details the causal recursion. The current observation pair drives the augmented-state filter and, in parallel, the difference-based band-power path. The two band-filter outputs form the recursive power statistics that are mapped through $\mathbf{M}^+$ to $\hat{\boldsymbol{\alpha}}_k$. These estimates enter $\mathbf{Q}_k$ after the state update and are used for the prediction from k to $k + 1$, avoiding an algebraic loop. The bounded updates of β_k and $\mathbf{R}_k$ remain auxiliary safeguards within the same loop.




Figure 4. Detailed CN-PAKF online recursion.

Algorithm 1 gives the corresponding pseudocode.

Algorithm 1 CN-PAKF online recursion

Require: $F, H, Q_p^0, Q_{n,i}^0, M$, band filters, memories, and bounds

- 1: Initialize $\hat{\mathbf{z}}_{0|0}, \mathbf{P}_{0|0}, \alpha_{i,0} = 1, \beta_0 = 1$, and $\mathbf{R}_0$
 - 2: **for** each synchronous observation pair $\mathbf{y}_k$ **do**
 - 3: Predict $\hat{\mathbf{z}}_{k|k-1}$ and $\mathbf{P}_{k|k-1}$ using $\mathbf{Q}_{k-1}$
 - 4: Form ν_k and $\mathbf{S}_k$, then correct the augmented state and covariance in Joseph form
 - 5: Compute d_k , its low- and high-band outputs, and the recursive powers $P_{B,k}$
 - 6: Recover and clip $\alpha_{L,k}$ and $\alpha_{I,k}$ using Equation (23)
 - 7: Apply the innovation gates, then update or retain β_k and the diagonal entries of $\mathbf{R}_k$
 - 8: Assemble $\mathbf{Q}_k$ from Equation (26) for the next prediction
 - 9: Output $\hat{v}_k$ from the corrected physical-state block
 - 10: **end for**
-

4. Evaluation Methodology

The evaluation addresses nominal sensor complementarity, target-band adaptation to changing sensor power, and the cost of finite-memory tracking. All the comparisons use matched observations and fixed parameters.

4.1. Reference Data, Noise Scenarios, and Fixed Parameters

The available setup does not provide an experimental sensor with sufficiently precise ground truth across the complete target band. Evaluation therefore uses a numerical reference trajectory while retaining measured environmental excitation. Measured ground motion drives the four-wire pendulum model to generate this trajectory. The ground excitation and reference trajectory are used only to construct and score the offline evaluation; neither is available to the online estimator.

Figure 5 characterizes this reference construction in the frequency domain. The simulated table response exhibits resonant amplification near $f_n = 0.408$ Hz and increasing isolation above resonance. The data-derived transmissibility follows the theoretical frequency dependence, providing a consistency check that the numerical reference trajectory obeys the adopted four-wire pendulum dynamics.

An independent ground-motion record determines the offline quantities, including $\mathbf{Q}_p^0$. For the one-step reference residual,

$$\begin{aligned} \mathbf{r}_{p,k} &= \mathbf{x}_{k+1}^{\text{ref}} - \mathbf{F}_p \mathbf{x}_k^{\text{ref}}, \\ \mathbf{Q}_p^0 &= \frac{1}{N_Q - 1} \sum_{k=1}^{N_Q} (\mathbf{r}_{p,k} - \bar{\mathbf{r}}_p)(\mathbf{r}_{p,k} - \bar{\mathbf{r}}_p)^T + \epsilon \mathbf{I}, \end{aligned} \quad (34)$$

where $\epsilon > 0$ ensures positive definiteness. This calibration represents unmeasured forcing and model error as state-propagation error. The test record is excluded from spectral fitting and error-based tuning.

Start-up power normalization uses only current sensor observations and assumes that the initial interval represents nominal operation. Each observation adds an independent realization of its target colored-noise spectrum to the same reference velocity. Ground excitation, mechanics, spectra, calibration, and algorithm parameters remain fixed across 30 realizations. Within each realization, every method receives the same observation pair.

A non-overlapping ground-motion record from another period at the same site provided an additional qualitative check without retuning the spectral priors, filter structure, or parameters.




Figure 5. Frequency-domain characterization of the numerical reference trajectory: (a) acceleration ASDs; (b) base-to-table transmissibility. The dash-dotted line marks the natural frequency.

The nominal scenario keeps both sensor power multipliers equal to one. The first time-varying scenario comprises four 900 s stages. The SEIS multiplier remains $m_{\text{SEIS}} = 1$,

while the IFO multiplier m_{IFO} follows 1, 10, 30, and 1. The dual-channel scenario uses IFO/SEIS multiplier pairs of (1, 1), (10, 1), (1, 10), and (5, 5). For either sensor,

$$n_{i,k}^{\text{tv}} = \sqrt{m_i(k)} n_{i,k}. \quad (35)$$

The multiplier m_i changes the noise power, whereas the square root in Equation (35) scales the noise amplitude. These nominal and staged power-change scenarios satisfy the approximately fixed spectral-shape assumption by construction. They evaluate the estimator under this assumption rather than testing the assumption itself.

Table 1 lists the shared parameters, and the offline fits provide the remaining coefficients. Memory lengths, bounds, and crossover bands are fixed before the methods are compared.

An additional controlled sensitivity test examined departures from the fixed-shape and synchronous-observation assumptions without retuning the offline spectral priors or **M**. Both sensor noise powers were changed by factors of 4 and 9 after calibration. The test imposed opposing PSD-slope changes of $\delta = \pm 0.1$ and ± 0.2 about 1 Hz, SEIS delays of 0.1, 0.5, and 1.0 ms, and a joint case with $\delta = 0.2$ and a 1.0 ms delay. The 0.01–1 Hz error was evaluated over the post-transition plateau. All cases used the same reference motion and one paired noise realization; they are therefore reported as a sensitivity study rather than an inferential robustness assessment.

The parameters in Table 1 define a prespecified reference configuration. At 100 Hz, its memory lengths correspond to ideal single-stage EMA t_{90} values of 460.5 s and 46.0 s. The 400 s nominal-power calibration follows a 50 s start-up period, so sensor-scale adaptation begins near 450 s; auxiliary updates are eligible from 50 s. The sensor-scale bounds cover all injected gains, while the β and **R** bounds constrain the auxiliary updates. A one-factor scan varied band boundaries, memories, calibration length, and bounds using matched observations and a fixed offline model. Varying N_{hi} also changes the auxiliary averaging factor.

4.2. Comparison Methods, Error Metrics, and Statistical Reporting

The comparison includes the raw IFO and SEIS velocity observations, which define the single-sensor performance limits, and several fusion methods. The Kalman filter (KF) uses a white-noise state-space model with fixed covariances. The Sage–Husa Kalman filter (SH-KF) applies conventional adaptive covariance estimation to the same white-noise formulation. The minimum-variance complementary filter (MVCF) is included as a non-causal, fixed-spectrum reference rather than an online competitor.

Three causal baselines use colored-noise-augmented states. The colored-noise Kalman filter (CN-KF) retains fixed covariances. The colored-noise Sage–Husa Kalman filter (CN-SHKF) uses Sage–Husa adaptation. The colored-noise maximum-likelihood-estimation Kalman filter (CN-MLEKF) uses maximum-likelihood adaptation. CN-PAKF is the proposed sensor-power-adaptive colored-noise method under approximately fixed spectral shapes. All causal state-space methods use identical platform dynamics, sampling grids, and observation pairs. All colored-noise methods also share the same offline spectral priors.

The time-varying IFO schedule also includes a component ablation. CN-PAKF- α adapts $\alpha_{i,k}$ while fixing $\beta_k = 1$ and $\mathbf{R}_k = \mathbf{R}_0$. CN-PAKF- (β, R) fixes $\alpha_{i,k} = 1$ while adapting β_k and $\mathbf{R}_k$. Full CN-PAKF adapts all three quantities. Oracle-CN-KF uses the injected multipliers $m_i(k)$ with fixed auxiliary covariances and therefore serves as a non-implementable reference. All variants use the same four-stage schedule and offline spectral priors.

Performance is quantified from the velocity error $\tilde{v}_k = \hat{v}_k - v_k$ and its one-sided PSD $S_{\tilde{v}}(f)$. The primary target band is $[f_1, f_2] = [0.01, 1]$ Hz, and the corresponding root-mean-square error (RMSE) is

$$\text{RMSE}_{[f_1, f_2]} = \left(\int_{f_1}^{f_2} S_{\tilde{v}}(f) df \right)^{1/2}. \quad (36)$$

All percentage reductions and time-varying performance comparisons use the 0.01–1 Hz target-band RMSE in Equation (36). For the IFO and dual-channel power schedules, Equation (36) is applied separately to each 900 s stage. Each stage therefore includes the post-change adjustment interval.

Transient behavior at each step is additionally represented by the causal 20 s trailing-window time-domain statistic, with $N_w = 20f_s = 2000$:

$$\text{RMSE}_{20s}(k) = \left(\frac{1}{N_w} \sum_{j=k-N_w+1}^k \tilde{v}_j^2 \right)^{1/2}. \quad (37)$$

This local diagnostic is used only to describe the time-domain transition; it does not replace the stage-wise target-band RMSE.

The reported statistics are the mean, sample standard deviation (s.d.), and Student- t 95% confidence interval (CI), $\bar{E} \pm t_{0.975, 29} s_E / \sqrt{30}$. Method differences and percentage reductions are calculated from paired within-realization values.

Conditioning sensitivity was evaluated using every stored normalized band-power vector after alpha calibration (315,001 samples). With fixed $\mathbf{M}$ and calibration powers, the unbounded scale estimate obeys the exact local relation $\delta\alpha = \mathbf{M}^{-1}\delta\mathbf{p}$. Independent uniformly distributed bounded relative perturbations were applied to the low- and high-band powers, with three draws per sample and error level, and the resulting deviations from the unperturbed unbounded scale estimates were summarized as pooled RMS values. This analysis quantifies propagation of power-statistic perturbations for the saved $\mathbf{M}$; it does not estimate uncertainty in the offline spectra, calibration powers, or $\mathbf{M}$ itself.

5. Results and Discussion

The results compare nominal and time-varying performance, isolate the contributions of adaptive components, and examine tracking delays and sensitivity to model assumptions.

5.1. Nominal Sensor Complementarity and Adaptation Cost

Figure 6 and Table 2 compare nominal error spectra and integrated errors across 30 paired realizations.

CN-PAKF reduced the target-band RMSE by 31.3% relative to IFO and 70.5% relative to SEIS; the paired 95% CI for the IFO reduction was 30.8–31.8%. All colored-noise methods outperformed the white-noise KF and SH-KF, indicating the importance of spectral modeling. CN-PAKF differed from the three causal colored-noise baselines by less than 1% in mean RMSE, although this does not establish statistical equivalence. MVCF achieved the lowest error as a non-causal, fixed-spectrum reference.

For $n = 39$ states and $m = 2$ observations, dense CN-PAKF requires $\mathcal{O}(n^3)$ computation per sample and $\mathcal{O}(n^2)$ streaming memory, with 115.0 KiB of working arrays, excluding histories, interpreter overhead, and temporary allocations. One run over 360,001 samples averaged 327.715 μs per sample (3.3% of the 10 ms interval), including adaptation and history writes (Python 3.12.4/NumPy 1.26.4, Windows 10 Intel x64, four MKL threads; no separate warm-up), without establishing worst-case latency.

Table 2. Velocity root-mean-square error (RMSE) under time-invariant noise across 30 paired realizations. Values are mean $\pm$ sample standard deviation (s.d.) in 10^{-10} m s $^{-1}$; brackets give the Student- t 95% confidence interval (CI) of the target-band mean.

Method	Target-Band RMSE	Full-Record RMSE	95% CI of the Band Mean
IFO	1.508 $\pm$ 0.017	265.080 $\pm$ 0.361	[1.502, 1.515]
SEIS	3.516 $\pm$ 0.171	4.497 $\pm$ 0.229	[3.452, 3.580]
KF	11.096 $\pm$ 0.268	18.875 $\pm$ 0.314	[10.996, 11.196]
SH-KF	3.350 $\pm$ 0.173	19.496 $\pm$ 0.530	[3.285, 3.414]
MVCF	0.590 $\pm$ 0.005	2.036 $\pm$ 0.003	[0.588, 0.592]
CN-KF	1.044 $\pm$ 0.010	2.258 $\pm$ 0.005	[1.040, 1.048]
CN-SHKF	1.043 $\pm$ 0.010	2.242 $\pm$ 0.005	[1.039, 1.047]
CN-MLEKF	1.034 $\pm$ 0.010	2.238 $\pm$ 0.005	[1.030, 1.038]
CN-PAKF	1.036 $\pm$ 0.017	2.241 $\pm$ 0.007	[1.029, 1.042]




Figure 6. Velocity-error amplitude spectral densities under time-invariant noise (30 independent sensor-noise realizations). Curves and shading show cross-realization means and pointwise 95% CIs.

5.2. Performance Under Time-Varying Sensor Power

Table 3 and Figure 7 compare stage-wise target-band errors under the four-stage IFO power schedule across 30 paired realizations.

At IFO power multipliers of 10 and 30, CN-PAKF reduced the target-band RMSE relative to CN-KF by 17.0% and 36.8%, respectively. The paired 95% CIs were 16.5–17.6% and 36.0–37.5%, with lower errors in all 30 realizations. CN-PAKF also achieved lower mean errors than CN-SHKF and CN-MLEKF in both stages.

Table 3. Stage-wise 0.01–1 Hz velocity-error RMSE under time-varying IFO noise. Values are mean $\pm$ sample s.d. across 30 paired realizations in 10^{-10} m s $^{-1}$.

Method	Stage 1	Stage 2	Stage 3	Stage 4
CN-KF	1.0275 $\pm$ 0.0257	1.8672 $\pm$ 0.0471	2.9495 $\pm$ 0.0602	1.0320 $\pm$ 0.0214
CN-SHKF	1.0271 $\pm$ 0.0261	1.8412 $\pm$ 0.0462	2.8973 $\pm$ 0.0587	1.0317 $\pm$ 0.0217
CN-MLEKF	1.0274 $\pm$ 0.0257	1.7092 $\pm$ 0.0445	2.2655 $\pm$ 0.0483	1.0360 $\pm$ 0.0212
CN-PAKF	1.0269 $\pm$ 0.0274	1.5494 $\pm$ 0.0447	1.8647 $\pm$ 0.0700	1.0397 $\pm$ 0.0252

The component ablation evaluates which adaptive quantities produce the target-band improvement under the same IFO power schedule. Table 4 reports the configurations and 0.01–1 Hz RMSEs, while Figure 8 compares the corresponding error spectra.




Figure 7. Stage-wise 0.01–1 Hz velocity-error RMSE under time-varying IFO noise (30 independent sensor-noise realizations). Boxes show interquartile ranges, medians, whiskers extending 1.5 interquartile ranges, and outliers.

Table 4. Component ablation under the four-stage time-varying IFO-noise schedule. Descriptive 0.01–1 Hz RMSE values are in $10^{-10} \text{ m s}^{-1}$. Gain denotes improvement relative to CN-KF.

Method	α_i	β	R	RMSE	Gain (dB)
CN-KF	1	1	$\mathbf{R}_0$	1.9122	0.000
CN-PAKF- α	Adaptive	1	$\mathbf{R}_0$	1.4293	2.528
CN-PAKF- (β, R)	1	Adaptive	Adaptive	1.9100	0.010
Full CN-PAKF	Adaptive	Adaptive	Adaptive	1.4293	2.528
Oracle-CN-KF	$m_i(k)$	1	$\mathbf{R}_0$	1.4180	2.598

Adapting only α matched full CN-PAKF to the reported precision, yielding a 25.3% RMSE reduction (2.528 dB) relative to CN-KF. Adapting only (β, R) yielded 0.010 dB, identifying sensor-scale adaptation as the main source of improvement in this schedule. Oracle-CN-KF provided a slightly lower multiplier-informed reference (Table 4).




Figure 8. Velocity-error amplitude spectral densities for the time-varying IFO-noise ablation. CN-PAKF- α , CN-PAKF- (β, R) , and Oracle-CN-KF adapt only sensor noise-power scales, adapt only auxiliary quantities, and use injected multipliers, respectively.

Across 355001 eligible samples, the IFO, SEIS, and joint innovation-gate pass rates were 99.9963%, 100.0000%, and 99.9963%, respectively. The actual parameter-change rates were 81.5068% for $\mathbf{Q}_p$, 0.1839% for R_L , and 0.0783% for R_I , reflecting smoothing and clipping after gate acceptance. The realized β spanned $[1/3, 3]$ (median 1.006), spending 12.88% and 5.62% of eligible samples at its lower and upper bounds. Corresponding lower/upper-bound occupancies were 33.11%/66.75% for R_L and 64.67%/35.33% for R_I , indicating frequent saturation of these small residual terms.

Figure 9 tests whether the estimated scales assigned IFO-only power changes to the correct sensor.




Figure 9. Online sensor noise-power scales under time-varying IFO noise (30 independent sensor-noise realizations). Solid curves and shading show cross-realization means and pointwise 95% CIs; dashed curves show true multipliers.

The IFO scale followed the injected sequence 1, 10, 30, and 1, while the SEIS scale remained near nominal. Smoothed transitions and delayed recovery reflected finite-memory tracking. Figure 10 quantifies the first post-step transient using the causal 20 s trailing-window statistic in Equation (37).




Figure 10. Causal 20 s trailing-window velocity-error RMSE after the tenfold IFO-noise power step; the dotted line marks $t = 900$ s.

After the first IFO power increase, the CN-KF local error rose toward $3 \times 10^{-10} \text{ m s}^{-1}$, whereas CN-PAKF remained lower after a finite adjustment. This local diagnostic comple-

ments the stage-wise target-band comparison in Table 3. Following restoration of nominal IFO power, CN-PAKF incurred a 0.76% Stage-4 RMSE increase relative to CN-KF (paired 95% CI, 0.04–1.47%), consistent with delayed recovery.

Figure 11 tests scale attribution under IFO-only, SEIS-only, and simultaneous power changes.




Figure 11. Online sensor noise-power scales in the dual-channel test (30 independent sensor-noise realizations). Solid curves and shading show cross-realization means and pointwise 95% CIs; dashed curves show true multipliers.

The IFO scale responded faster than the SEIS scale, consistent with their different memory lengths; both approached 5 in the final stage. Table 5 reports the 0.01–1 Hz RMSE in $10^{-10} \text{ m s}^{-1}$ for the three elevated-noise stages of one paired realization; gain is relative to CN-KF. CN-PAKF reduced the RMSE by 14.61% and 17.53% during isolated IFO and SEIS increases, but by only 0.75% when both powers increased fivefold. This smaller benefit may reflect the unchanged ratio between sensor noise powers under proportional increases, which offers less scope for relative reweighting. During the nominal stage, performance was nearly matched, with a slight RMSE increase for CN-PAKF. These descriptive results do not establish cross-realization uncertainty.

Table 5. Target-band velocity-error RMSE during the elevated-noise stages of the dual-channel test.

$(m_{\text{IFO}}, m_{\text{SEIS}})$	CN-KF	CN-PAKF	Gain
(10, 1)	1.865	1.593	14.61%
(1, 10)	2.731	2.252	17.53%
(5, 5)	2.303	2.285	0.75%

5.3. Conditioning of Spectral Attribution

For the realized matrix, $\kappa_2(\mathbf{M}) = 1.007$, and independent band-power errors bounded by 1% produced pooled RMS scale deviations of 0.5777% (IFO) and 0.6278% (SEIS). These results pooled 315,001 post-calibration samples with three perturbations per sample and agreed with exact linear propagation (Figure 12). Along a synthetic row-similarity path, reducing ρ from 1 to 0.1 and 0.01 increased κ_2 to 10.05 and 100.51. Scale deviations under the same perturbation rose to approximately 4.1% and 41%, showing that nonzero determinant alone does not ensure robust attribution. This stress path changes matrix conditioning, not the fitted physical sensor spectra.




Figure 12. Conditioning and sensitivity of the spectral-attribution matrix.

5.4. Sensitivity to Nonideal Spectral Shape and Timing

Figure 13 summarizes the controlled nonideal tests. Under the specified spectral-shape perturbations, CN-PAKF retained a 4.27–5.85% primary-band improvement relative to the matched fixed CN-KF. Relative to the power-only case, its primary-band RMS increased by 24.48% for $\delta = -0.2$ and decreased by 20.06% for $\delta = 0.2$. The latter cross-scenario changes do not represent model-mismatch loss alone, because preserving total 0.01–10 Hz noise power redistributes noise within the 0.01–1 Hz evaluation band.




Figure 13. Controlled sensitivity to spectral-shape mismatch and inter-sensor timing offset.

Timing mismatch was more consequential: relative to the power-only case, the CN-PAKF primary-band RMS increased by 126.56%, 919.23%, and 1873.07% for SEIS delays of 0.1, 0.5, and 1.0 ms, respectively. The 1.0 ms delay combined with $\delta = 0.2$ produced a 1928.76% increase. Nevertheless, CN-PAKF retained RMSE reductions of 4.96%, 5.77%, and 8.55% relative to CN-KF under the same respective delays, and 5.96% in the joint case. Thus, timing mismatch severely degraded absolute accuracy, although CN-PAKF retained an advantage over the matched fixed filter in these tests. The delay introduces motion leakage into the synchronous difference channel, making accurate inter-sensor synchronization necessary for the intended operating conditions. Only the 100 Hz sampling rate was tested. Changing it requires re-discretizing the models and band filters and adjusting the memory and calibration sample counts to preserve their physical durations. Cross-channel noise correlation was not tested. A nonzero cross-spectrum adds $-2 \operatorname{Re} S_{LI}(f)$ to the sum of the two noise PSDs in the difference channel, violating the additive-power attribution model and potentially biasing the recovered scales.

5.5. Parameter Sensitivity and Same-Site Temporal Check

The one-factor scan used the IFO/SEIS sequence (1,1), (10,1), (30,9), and (1,1), evaluating the target-band RMSE from the first step onward. Across $N_{l_0} = 10,000$ –40,000, the measured SEIS $1 \rightarrow 9$ response delay increased from 353.35 to 481.69 s (default, 398.99 s). Delay was recorded when a 10 s trailing mean sustained the 90% threshold for 5 s, distinguishing it from ideal EMA t_{90} . Doubling N_{l_0} increased the RMSE by 0.96%; varying the low-band boundary over 0.1–0.4 Hz changed it by at most 0.38%. High-band boundaries of 0.5–2 Hz had negligible effects, and tested $\beta/\mathbf{R}$ bounds changed the RMSE by less than 0.002%.

A 200 s calibration increased the scale-tracking error; extending calibration to 800 s also increased the scale error while reducing the target-band RMSE by only 0.442%. Scale error denotes RMS relative error pooled over both channels from the first step onward. Setting $\alpha_{\max} = 10$ excluded the injected gain of 30 and caused clipping, despite an apparent reduction in scale error. These single-realization results characterize local sensitivity around the prespecified reference configuration and do not establish its optimality.

The additional same-site time segment retained numerical stability and qualitatively similar performance relative to CN-KF without retuning. This qualitative check does not establish generalization across sites, operating conditions, or sensor states.

6. Conclusions

CN-PAKF combines synchronous motion-canceling differences with spectral attribution to adapt sensor noise-power scales in causal platform-velocity estimation. Updating colored-noise drive covariances preserves their normalized spectral shapes; the bounded β and $\mathbf{R}$ safeguards can still alter the total predicted error spectrum.

Across 30 paired simulations, the nominal target-band RMSE was 31.3% and 70.5% lower than IFO and SEIS, respectively, and within 1% of causal colored-noise baselines. At 30-fold IFO noise power, CN-PAKF reduced the stage-wise RMSE by 36.8% relative to CN-KF. Ablation attributed the improvement primarily to sensor-scale adaptation, with a small recovery cost from finite memory.

The evidence remains limited to open-loop simulations driven primarily by one measured record, supplemented by a qualitative same-site temporal check. The method assumes synchronized observations, distinguishable and approximately fixed noise spectra, uncorrelated sensor noises, and nominal start-up conditions. CN-PAKF retained an advantage over matched CN-KF under the tested spectral perturbations and timing offsets, although timing mismatch severely degraded absolute accuracy. Future work should evalu-

ate multiple sites and operating conditions, hardware timing and closed-loop performance, and adaptive memory or change detection to reduce recovery delays.

Author Contributions: Conceptualization, D.M. and Y.C.; methodology, D.M. and Y.C.; software, D.M.; validation, D.M., Y.C., T.Y., H.L., Q.L., and Z.W.; formal analysis, D.M. and Y.C.; investigation, D.M., Y.C., T.Y., H.L., Q.L., and Z.W.; resources, Y.C., H.L., and Q.L.; data curation, D.M.; writing—original draft preparation, D.M. and Y.C.; writing—review and editing, D.M., Y.C., T.Y., H.L., Q.L., and Z.W.; visualization, D.M.; supervision, Y.C.; project administration, Y.C. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the National Key Research and Development Program of China, grant number 2021YFC2202404.

Data Availability Statement: The data supporting this study are available from the corresponding author upon reasonable request. Raw ground-motion records are not publicly available due to lab restrictions, but simulation settings and derived data can be provided.

Conflicts of Interest: The authors declare no conflicts of interest.

Abbreviations

The following abbreviations are used in this manuscript:

ASD	Amplitude spectral density
CI	Confidence interval
CN-KF	Colored-noise Kalman filter
CN-MLEKF	Colored-noise maximum-likelihood-estimation Kalman filter
CN-PAKF	Colored-noise power-adaptive Kalman filtering algorithm
CN-SHKF	Colored-noise Sage–Husa Kalman filter
dB	Decibel
IFO	Laser interferometric displacement sensor
KF	Kalman filter
MVCF	Minimum-variance complementary filter
PSD	Power spectral density
RMS	Root mean square
RMSE	Root-mean-square error
s.d.	Sample standard deviation
SEIS	Broadband seismometer
SH-KF	Sage–Husa Kalman filter

References

1. Wang, S.; Liu, D.; Zhan, X.; Dong, P.; Shen, J.; Wang, J.; Gao, R.; Guo, W.; Xu, P.; Qi, K.; et al. Core payload of the space gravitational wave observatory: Inertial sensor and its critical technologies. *Sensors* **2024**, *24*, 7685. [\[CrossRef\]](#)
2. Luo, Z.; Wang, Y.; Wu, Y.; Hu, W.; Jin, G. The Taiji program: A concise overview. *Prog. Theor. Exp. Phys.* **2021**, *2021*, 05A108. [\[CrossRef\]](#)
3. Li, J.C.; Lin, H.A.; Luo, J.X.; Wu, Y.X.; Wang, Z. Optical design of space gravitational wave detection telescope. *Chin. Opt.* **2022**, *15*, 761–769. [\[CrossRef\]](#)
4. Yu, M.; Li, J.C.; Lin, H.A.; Huang, Y.Z.; Luo, J.X.; Wu, Y.X.; Wang, Z. Design of optical system for low-sensitivity space gravitational wave telescope. *Chin. Opt.* **2023**, *16*, 1384–1393. [\[CrossRef\]](#)
5. Wang, S.; Wang, Z.; Liu, D.; Dong, P.; Min, J.; Luo, Z.; Qi, K.; Lei, J. Torsion pendulum apparatus for ground testing of space inertial sensor. *Sensors* **2024**, *24*, 7816. [\[CrossRef\]](#)
6. Razzano, M.; Spada, F.; Balestri, G.; Basti, A.; Bellizzi, L.; De Santi, F.; Fidecaro, F.; Fiori, A.; Frasconi, F.; Gennai, A.; et al. Seismic isolation systems for next-generation gravitational wave detectors. *Nucl. Instrum. Methods Phys. Res. Sect. A* **2024**, *1067*, 169675. [\[CrossRef\]](#)
7. Akutsu, T.; Ando, M.; Arai, K.; Arai, Y.; Araki, S.; Araya, A.; Aritomi, N.; Asada, H.; Aso, Y.; Bae, S.; et al. Vibration isolation systems for the beam splitter and signal recycling mirrors of the KAGRA gravitational wave detector. *Class. Quantum Gravity* **2021**, *38*, 065011. [\[CrossRef\]](#)

8. Yang, B.X.; Liu, L.; Wu, S.C.; Li, H.Y.; Zhou, Z.B. A low frequency horizontal active vibration isolation bench for testing the performance of high-precision space inertial sensors. *Class. Quantum Gravity* **2021**, *38*, 175006. [\[CrossRef\]](#)
9. Pu, H.; Song, X.; Tian, Y.; Wang, M.; Sun, Y.; Peng, Y.; Luo, J.; Ding, J. Sensor fusion for active vibration isolation system with double noise: An adaptive Kalman filter approach. *IEEE Trans. Ind. Electron.* **2025**, *72*, 898–907. [\[CrossRef\]](#)
10. Collette, C.; Matichard, F. Sensor fusion methods for high performance active vibration isolation systems. *J. Sound Vib.* **2015**, *342*, 1–21. [\[CrossRef\]](#)
11. Tsang, T.T.L.; Li, T.G.F.; Dehaeze, T.; Collette, C. Optimal sensor fusion method for active vibration isolation systems in ground-based gravitational-wave detectors. *Class. Quantum Gravity* **2022**, *39*, 185007. [\[CrossRef\]](#)
12. Collette, C.; Janssens, S.; Fernandez-Carmona, P.; Artoos, K.; Guinchard, M.; Hauviller, C.; Preumont, A. Review: Inertial sensors for low-frequency seismic vibration measurement. *Bull. Seismol. Soc. Am.* **2012**, *102*, 1289–1300. [\[CrossRef\]](#)
13. Jin, M.; Zhao, Y.; Han, M.; Lei, J.; Cui, J. Active vibration isolation method based on multi-sensor fusion absolute displacement feedback. *J. Vib. Eng. Technol.* **2025**, *13*, 496. [\[CrossRef\]](#)
14. Xu, A.; He, S.; Huo, Q.; Wang, L.; Li, A.; Li, D.; Chen, Y.; Xu, Z. Active vibration isolation platform with distributed feed-forward feedback hybrid control. *J. Vib. Eng. Technol.* **2025**, *13*, 172. [\[CrossRef\]](#)
15. Hu, J.; Liu, R.; Xie, Z.; Zhang, D.; Zhu, G. Research on an active vibration isolation system with hybrid control strategy for the guidance system of TBM. *J. Vib. Eng. Technol.* **2024**, *12*, 773–782. [\[CrossRef\]](#)
16. Xing, X.; Chen, Z.; Feng, Z. A variable stiffness and damping control strategy for improving vibration isolation performances in low-frequency excitation. *J. Vib. Eng. Technol.* **2023**, *11*, 1595–1608. [\[CrossRef\]](#)
17. Chen, H. Sensor Condition Monitoring Incorporating Fault-tolerant Control. *J. Appl. Sci. Eng.* **2024**, *27*, 2137–2145. [\[CrossRef\]](#)
18. Begovich, O.; Lizárraga, A.; Ramírez-Treviño, A. Sensor fusion architecture for fault diagnosis with a predefined-time observer. *Algorithms* **2024**, *17*, 270. [\[CrossRef\]](#)
19. Beker, M.G.; Bertolini, A.; van den Brand, J.F.J.; Bulten, H.J.; Hennes, E.; Rabeling, D.S. State observers and Kalman filtering for high performance vibration isolation systems. *Rev. Sci. Instrum.* **2014**, *85*, 034501. [\[CrossRef\]](#)
20. Louédec, M.; Jaulin, L. Interval extended Kalman filter—Application to underwater localization and control. *Algorithms* **2021**, *14*, 142. [\[CrossRef\]](#)
21. Chien, T.H.; Chen, Y.C. An on-line tracker for a stochastic chaotic system using observer/Kalman filter identification combined with digital redesign method. *Algorithms* **2017**, *10*, 25. [\[CrossRef\]](#)
22. Bryson, A.E.; Johansen, D.E. Linear filtering for time-varying systems using measurements containing colored noise. *IEEE Trans. Autom. Control* **1965**, *10*, 4–10. [\[CrossRef\]](#)
23. Chang, G. On Kalman filter for linear system with colored measurement noise. *J. Geod.* **2014**, *88*, 1163–1170. [\[CrossRef\]](#)
24. Chen, C.; Lin, X.; Li, W.; Cheng, L.; Wang, H.; Zhang, Q.; Wang, Z. Adaptive coloured noise multirate Kalman filter and its application in coseismic deformations. *Geophys. J. Int.* **2023**, *234*, 1236–1253. [\[CrossRef\]](#)
25. Mehra, R.K. On the identification of variances and adaptive Kalman filtering. *IEEE Trans. Autom. Control* **1970**, *15*, 175–184. [\[CrossRef\]](#)
26. Vettori, S.; Di Lorenzo, E.; Peeters, B.; Luczak, M.M.; Chatzi, E. An adaptive-noise augmented Kalman filter approach for input-state estimation in structural dynamics. *Mech. Syst. Signal Process.* **2023**, *184*, 109654. [\[CrossRef\]](#)
27. Takayama, Y.; Urakubo, T.; Tamaki, H. Adaptive choice of process noise covariance in Kalman filter using measurement matrices. *IEEE Trans. Control Syst. Technol.* **2024**, *32*, 934–944. [\[CrossRef\]](#)
28. Chang, G. Kalman filter with both adaptivity and robustness. *J. Process Control* **2014**, *24*, 81–87. [\[CrossRef\]](#)
29. Mendonça, M.O.K.; Ferreira, J.O.; Tsinos, C.G.; Diniz, P.S.R.; Ferreira, T.N. On fast converging data-selective adaptive filtering. *Algorithms* **2018**, *12*, 4. [\[CrossRef\]](#)
30. Mei, B.; Ma, C.; Bai, Y.; Hu, M.; Liu, L.; Qu, S.; Wu, S.; Yu, J.; Zhou, Z. Device for evaluating the performance of high-precision accelerometers as motion sensors for an active vibration isolation system. *IEEE Trans. Instrum. Meas.* **2025**, *74*, 9510209. [\[CrossRef\]](#)
31. Zanetti, R.; DeMars, K.J. Joseph formulation of unscented and quadrature filters with application to consider states. *J. Guid. Control Dyn.* **2013**, *36*, 1860–1864. [\[CrossRef\]](#)

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.