

Article

Dynamic Ranging-Error Compensation and Consistent Cooperative Localization in TDMA UWB UAV Swarms

Zheng Xie ¹, Chenxin Tu ¹, Yiding Zhan ¹, Gang Liu ^{1,*}, Xiaowei Cui ^{1,2} and Mingquan Lu ^{1,2}

¹ Department of Electronic Engineering, Tsinghua University, Beijing 100084, China; xiez24@mails.tsinghua.edu.cn (Z.X.); tcx22@mails.tsinghua.edu.cn (C.T.); zhanyd23@mails.tsinghua.edu.cn (Y.Z.); cxw2005@tsinghua.edu.cn (X.C.); lumq@tsinghua.edu.cn (M.L.)

² State Key Laboratory of Space Network and Communications, Tsinghua University, Beijing 100084, China

* Correspondence: liu_gang@tsinghua.edu.cn

Highlights

What are the main findings?

- The proposed SCI+MPLS co-design achieves cooperative position RMSE of 0.55 m (2-anchor) and 0.96 m (anchor-free) in 16-node UAV networks at the cross-point ($v = 25$ m/s, $T_{\text{round}} = 64$ ms), with the MPLS front-end independently validated (Experiment 1) to hold distance RMSE below 0.15 m across the polling-period envelope. Median ANEES stays within the strict $[0.5, 2]$ consistency band at every tested point across the full velocity (10–40 m/s) and polling-period (16–160 ms) sweep.
- A five-method ablation along front-end (MPLS vs. SDS-TWR) and back-end (SCI vs. EKF/EKF+Inflate/CI) axes isolates each module's contribution: front-end ranging quality dominates position accuracy, while back-end correlation handling dominates covariance consistency; the two effects are empirically decoupled within the tested operating envelope.

What are the implications of the main findings?

- Reliable uncertainty bounds support trustworthy downstream decision-making (formation control, collision-avoidance margin computation) in GNSS-denied UAV swarm operations, where conventional EKF-based methods report overconfident covariance unsuitable for risk-aware planning.
- The calibrated per-link variance interface keeps the front-end and back-end independently replaceable, so the co-design—demonstrated here on agile UAV swarms—transfers to other UWB-based cooperative-localization platforms, where ranging quality and correlation handling are coupled.



Academic Editor: Emmanouel T. Michailidis

Received: 9 July 2026

Revised: 29 August 2026

Accepted: 6 September 2026

Published: 9 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 Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

Abstract

Unmanned aerial vehicle (UAV) swarms increasingly operate in GNSS-denied environments, where cooperative localization provides relative positioning by fusing onboard odometry with inter-agent range measurements, for which ultra-wideband (UWB) two-way ranging is a common infrastructure-free choice. Beyond accuracy, safe swarm autonomy needs a trustworthy measure of positioning uncertainty, since collision-avoidance and formation-keeping decisions derive their safety margins from the reported covariance. Under sustained agile flight, both are hard to achieve at once: motion within each time-division multiple access (TDMA) polling round induces a ranging bias well above the UWB noise floor, and reusing shared information across the network drives the reported covariance below the true error. To address these two problems jointly rather than in isolation, we propose a modular architecture coupling an online maximum-likelihood

polynomial least-squares (MPLS) ranging front-end with a fading split covariance intersection (SCI) cooperative back-end through a per-link variance interface: online MPLS compensates the motion-induced bias and reports a calibrated, time-varying variance that fading SCI consumes as measurement noise while its continuous-time fading factor bounds the reused-information covariance. Monte Carlo simulation over anchored and anchor-free 16-node swarms shows the two effects to be empirically decoupled, with front-end ranging quality governing positioning accuracy and back-end correlation handling governing estimator consistency. The method attains sub-meter positioning accuracy in both settings and, without per-scenario tuning, keeps consistency—quantified by the average normalized estimation error squared (ANEES)—within a trusted band; comparably accurate extended Kalman filter and covariance-intersection baselines fall outside it, becoming overconfident and overconservative, respectively. It thus delivers the trustworthy uncertainty that safety-critical swarm decisions require in GNSS-denied flight.

Keywords: ultra-wideband; TDMA ranging; cooperative localization; UAV swarm; GNSS-denied navigation; dynamic ranging error; estimation consistency

1. Introduction

From disaster response to infrastructure inspection, unmanned aerial vehicle (UAV) swarms are moving into environments where global navigation satellite systems (GNSSs) cannot be relied upon, such as building interiors, areas beneath forest or orchard canopy during precision-agriculture and crop-protection missions, and obstructed or contested airspace [1]. Without a shared global reference, the swarm determines its members' positions relative to one another from onboard motion and inter-agent measurements, a capability known as cooperative localization (CL). Safe autonomy, however, needs more than an accurate relative-position estimate: collision-avoidance, formation-keeping, and chance-constrained planning decisions draw their safety margins from the reported position covariance. Those margins are sound only when the covariance neither understates nor overstates the true error—a property captured by the average normalized estimation error squared (ANEES) [2].

Securing accuracy and trustworthy uncertainty together is what makes the UAV setting demanding. Size, weight, and power (SWaP) limits [3] commonly leave each aircraft with a single non-coherent ultra-wideband (UWB) transceiver and no external time reference, so inter-agent distance is obtained by two-way ranging (TWR) and fused with onboard odometry at autopilot rates [4]. Under the sustained, agile flight of a real swarm, however, this standard TWR pipeline fails on two coupled fronts—a front-end ranging error that erodes accuracy, and a back-end inconsistency that renders the reported covariance untrustworthy.

The front-end failure scales with the swarm itself. Because a time-division multiple-access (TDMA) polling round lengthens with the number of UAVs, the unmodeled distance a fast-moving aircraft covers within one round—the *dynamic ranging error*—grows with both flight speed and swarm size [5]. In highly dynamic swarms, this motion-induced bias can exceed the UWB timestamp-noise floor [6] by an order of magnitude, dominating the ranging error budget.

The back-end failure emerges once these ranges circulate through the swarm. Reusing a pairwise range introduces statistical correlations along multi-hop measurement chains [7]: ignoring them yields overconfident covariance ($\text{ANEES} \gg 1$), whereas assuming maximal correlation produces the opposite, overconservative failure ($\text{ANEES} \ll 1$). The two failures

are mechanistically coupled: motion-induced bias systematically violates the unbiased-range assumption underlying consistent cooperative filters. The ranging front-end and the cooperative back-end are therefore best co-designed rather than developed in isolation.

This paper proposes a modular architecture in which two co-designed modules are connected by a feedforward variance interface that propagates calibrated per-link uncertainty between layers. The *online MPLS ranging front-end* extends the batch maximum-likelihood polynomial least-squares (MPLS) model [5], representing the time-varying propagation delay as a polynomial within overlapping sliding windows and emitting, for each measurement, a bias-corrected distance estimate together with a calibrated per-link variance. The *fading SCI cooperative back-end* extends the Split Covariance Intersection (SCI) filter [8]. It decomposes each agent's covariance into an independent and a dependent component, applies covariance intersection (CI) only to the dependent one, takes the front-end variance as its measurement-noise input, and damps the dependent component with a continuous-time fading factor to prevent unbounded accumulation. Figure 1 previews the pipeline.




Figure 1. Proposed SCI+MPLS framework. The physical scenario (top) supplies two-way-ranging timestamps and onboard odometry to two co-designed modules separated by a per-link variance interface (centre): an online MPLS ranging front-end (left) that emits a calibrated per-link variance $\hat{\sigma}_{ij}^2(t_k)$, and a fading SCI cooperative back-end (right) that consumes it as measurement noise and damps the reused-information covariance $\mathbf{P}_D$ by a continuous-time fading factor $e^{-T_{\text{round}}/\tau}$ (dashed loop). Module internals are detailed in the text.

The main contributions are as follows, each elaborating one module or their joint validation:

1. **Dynamic ranging-error compensation for TDMA UWB UAV swarms.** We recast the batch MPLS ranging framework as an *online MPLS* front-end: a Legendre basis keeps the per-window condition number bounded, an overlapping sliding window enables continuous-time tracking, a sign-consistency gate rejects transient clock-fusion failures, and a two-step clock-then-distance estimation fuses clocks across the network by inverse-variance weighting. Together, these components suppress the motion-induced ranging bias and produce a per-link variance calibrated to the realized error—a trustworthy uncertainty signal that the cooperative back-end requires, but that conventional static-model TWR front-ends do not provide.
2. **Consistent CL under sustained high-speed motion.** We extend the standard SCI formulation into a *fading SCI* back-end: it takes in the front-end's calibrated time-varying variance as its independent-component noise, damps the dependent covariance component with a continuous-time fading factor, and selects the mixing weight by determinant minimization through a deterministic two-stage grid search. The fading factor keeps the reused-information covariance bounded, so the estimator stays

consistent under sustained motion—where a standard extended Kalman filter (EKF) turns overconfident and CI overconservative.

3. **Modular architecture and ablation study.** The online MPLS front-end and fading SCI back-end are connected by a feedforward per-link variance interface, forming a fully integrated co-designed CL pipeline. A five-method ablation confirms that the proposed SCI+MPLS configuration achieves balanced estimator consistency in both anchored and anchor-free scenarios without per-scenario parameter tuning; alternative configurations—substituting symmetric double-sided TWR (SDS-TWR) for MPLS, or EKF/EKF+Inflate/CI for SCI—exhibit characteristic overconfidence or over-conservatism.

2. Related Works

The proposed framework tackles two interrelated challenges that arise when CL is deployed on a fast-moving UAV swarm: dynamic ranging error in time-division TWR and estimation inconsistency under multi-hop information reuse. This section reviews each in turn and highlights the gaps that motivate the present work.

2.1. TWR Error Correction

TWR extracts inter-node distance from round-trip time-of-flight in UWB positioning [4], bounded by oscillator-driven clock error [5] and motion-induced distance change during the exchange interval. Single-link protocols (SDS-TWR [9], AltDS-TWR [6,10]) compensate for clock drift through symmetric timestamp pairs; hardware calibration [11] targets only the clock channel. Concurrent and swarm protocols [12] raise the ranging rate but inherit the same static-distance assumption. Shared across all these baselines, this assumption emerges as the dominant error source at high relative velocity, motivating motion-aware compensation.

Existing approaches to motion-induced TWR error can be grouped into four broad categories, distinguished by *where* they intervene (radio hardware, learned mapping, estimator state, or signal model) and *which assumption* they relax. *Channel-level compensation* estimates Doppler-driven carrier-frequency and phase offsets at the radio layer [13] but requires phase-coherent hardware, which is unavailable on standard non-coherent IEEE 802.15.4z chips. *Learning-based bias correction* [14] trains neural networks on pose-dependent ranging error but requires an instrumented training environment matching the deployment velocity domain. *Closed-form joint estimation* [15,16] extends the state vector to include position, velocity, and clock parameters under a constant-velocity assumption, limiting applicability under substantial acceleration. *Polynomial signal-model* approaches [5] represent the delay as a polynomial in time and jointly estimate clock and distance parameters via MPLS, with estimator variance approaching the Cramér–Rao lower bound. Subsequent extensions [17] establish this approach’s generality across kinematic orders. These formulations, however, remain within a batch paradigm whose monomial (Vandermonde) basis becomes ill-conditioned as polynomial order or observation count grows; the same work [17] identifies continuous online operation across successive windows as an open problem.

We adopt the polynomial signal-model method because it runs on standard non-coherent IEEE 802.15.4z timestamps and natively yields the per-link variance for the fusion stage. Section 3.2 details how the proposed module extends it for continuous online operation. Tightly coupled UWB/IMU [18,19] and data-driven channel impulse response (CIR) feature methods [20,21], along with UWB-based UAV motion sensing [22], address adjacent problems, complementary to the in-exchange dynamic bias treated here.

2.2. CL Techniques

CL fuses inter-agent ranges with onboard motion to jointly estimate agent positions [23,24], and recent UAV-swarm systems combine UWB ranging with visual-inertial sensing [25–27]. The central difficulty in such fusion is correlation: when agent i updates on a range to j and j later updates on a range back to i , the reused measurement is no longer independent of j 's prior [7], so treating updates as independent corrupts the reported covariance. Whether that covariance faithfully bounds the true error—the filter's *consistency*—is measured by the ANEES, ideally around unity [2,28,29]. This section surveys the recursive Bayesian filters that tackle this correlation online at autopilot (kilohertz) rates—the class of filters to which the proposed back-end belongs. It excludes the batch factor-graph smoothers [30] and one-shot anchor self-localization [31], which address complementary problems. These filters differ primarily in how they handle correlation, which positions each on a spectrum ranging from overconfident to overconservative.

At the overconfident end of this spectrum, the distributed EKF [32] treats measurements as independent and yields $\text{ANEES} \gg 1$ [28]. Observability-preserving variants [29] correct linearization but retain the independence assumption, while explicit correlation-tracking [7] restores consistency only at the cost of history bookkeeping. Covariance inflation [33,34] rescales process noise by $\alpha > 1$ toward unity, but the required α is scenario-dependent. At the overconservative extreme, CI [35,36] assumes maximal correlation and yields $\text{ANEES} \ll 1$. Selective CI [37] recovers some lost information through measurement-provenance bookkeeping, again at a storage and communication cost.

Between these extremes, the Split Covariance Intersection Filter (SCIF) decomposes each agent's covariance into an independent component (uncorrelated sources) and a dependent component (shared or reused measurements), applying CI only to the latter through a mixing weight that minimizes a scalar functional of the updated covariance. SCI originates with Julier and Uhlmann [36], was instantiated for vehicle localization by Li et al. [38], and extended to range-only cooperative UAV localization with an explicit independent/dependent noise decomposition by Pierre et al. [8]. We adopt this formulation as the closest published instantiation to our range-only setting and extend it to the highly dynamic regime; Section 3.3 details the extensions. Contemporary SCIF-family peers [39–41] extend in orthogonal directions: bounded-error interval propagation, fault-tolerant Gaussian belief propagation, and robust decentralized EKF with explicit correlation tracking. Each, however, presumes a bias-mitigated ranging input rather than addressing the joint front-end/back-end design pursued here, making a direct head-to-head comparison outside the scope of their intended use.

Two gaps motivate this work. First, the survey of TWR error correction above reveals that motion-aware ranging and calibrated per-link variance have largely been addressed in isolation, leaving their joint treatment unresolved. Second, the recursive filters surveyed above (EKF, CI, and SCIF variants) consume constant-statistics measurement noise rather than the time-varying variance that motion-aware ranging produces. The architecture proposed in Section 3 bridges these two gaps through a feedforward variance interface.

3. System Model and Method Framework

3.1. Problem Formulation and System Overview

Consider a UAV swarm modeled as a wireless network of N nodes, comprising N_a anchor nodes with known fixed positions and $N_m = N - N_a$ mobile agents with unknown, time-varying positions. Each node carries a TWR transceiver and an onboard velocity sensor, and a pair (i, j) can exchange TWR measurements only when $d_{ij}(t) < r_{\text{comm}}$, inducing a time-varying neighbor set $\mathcal{N}_i(t)$. Throughout, the indices i and j range over network nodes, anchors, or mobile agents alike; where a cooperative update is described,

i denotes the agent performing the estimate and j its ranging neighbor, which may itself be an anchor. Each agent i aims to estimate its own position $\hat{\mathbf{x}}_i \in \mathbb{R}^2$ and the associated error covariance $\mathbf{P}_i$ in a decentralized manner, fusing TWR ranges from neighbors with onboard velocity while maintaining consistency under correlated information from multi-hop measurement exchanges. The state is modeled in the horizontal plane $\mathbb{R}^2$, appropriate for the co-altitude swarm operation considered here, in which the inter-agent geometry is dominated by horizontal motion; extension to full three-dimensional localization is a direct state-vector augmentation.

The N nodes share the channel through a contention-free TDMA protocol with round-robin scheduling, completing one polling round every $T_{\text{round}} = N \cdot T_s$ seconds; T_{round} scales linearly with N , so larger networks incur proportionally greater dynamic ranging error.

TWR exchanges rely on timestamps from independent local clocks, each modeled by an affine deviation from global time t with clock skew α_i and offset β_i [5]:

$$T_i(t) = \alpha_i t + \beta_i, \quad (1)$$

and the timestamp observation recorded by node j when node i broadcasts is

$$T_j^{\text{rx}} = \alpha_j \left(\frac{T_i^{\text{tx}} - \beta_i}{\alpha_i} + \frac{d_{ij}(t)}{c} \right) + \beta_j + \varepsilon_{ij}, \quad (2)$$

where c is the speed of light, $d_{ij}(t) = \|\mathbf{x}_i(t) - \mathbf{x}_j(t)\|$, and the timestamp noise $\varepsilon_{ij} \sim \mathcal{N}(0, \sigma_T^2)$ with standard deviation σ_T . This observation couples the clock parameters of both nodes with the time-varying propagation delay $d_{ij}(t)/c$; a static delay model incurs systematic bias proportional to the relative radial velocity and is dominant in highly dynamic scenarios. The ranging model assumes line-of-sight (LoS) propagation; non-line-of-sight (NLOS) multipath is left to a separate preprocessing stage, outside the scope of this work.

In addition to ranging, each agent obtains an onboard velocity estimate $\mathbf{u}_i$ with measurement-noise covariance $\mathbf{Q}_u$ (the odometry uncertainty). Propagating it over one prediction step of duration Δt by dead reckoning yields the per-step position process-noise covariance, quantifying how fast position uncertainty grows between updates:

$$\mathbf{Q}_x = \Delta t^2 \mathbf{Q}_u. \quad (3)$$

The proposed method comprises two co-designed modules connected by a feedforward variance interface (Figure 1). The *online MPLS ranging compensation module* processes raw TWR timestamps and outputs, for each pair (i, j) , distance polynomial coefficients $\hat{\gamma}_{ij}$ and covariance $\Sigma_{\hat{\gamma}}^{(ij)}$, which the interface layer evaluates at each update time t_k as

$$\hat{d}_{ij}(t_k) = c \cdot \mathbf{v}(\tilde{\tau}_k)^\top \hat{\gamma}_{ij}, \quad \hat{\sigma}_{ij}^2(t_k) = c^2 \cdot \mathbf{v}(\tilde{\tau}_k)^\top \Sigma_{\hat{\gamma}}^{(ij)} \mathbf{v}(\tilde{\tau}_k), \quad (4)$$

subject to three quality checks that screen for extrapolation, non-physical distances, and degenerate variance. The *fading SCI cooperative estimation module* consumes $(\hat{d}_{ij}, \hat{\sigma}_{ij}^2)$, fuses with onboard velocity, and produces consistent position estimates under multi-hop inter-agent correlations.

3.2. Online MPLS Ranging Compensation Module

The TWR observation (2) couples clock parameters with time-varying propagation delays, preventing direct range extraction. The batch MPLS framework of Ref. [5] fits a polynomial delay model within a temporal window and jointly estimates clock and dis-

tance parameters by least squares. Our *online MPLS* front-end extends it for the continuous, autopilot-rate tracking required by a UAV swarm in sustained flight, through four additions: a Legendre polynomial basis for numerical stability, an overlapping sliding-window scheme, a sign-consistency rejection gate, and a two-step clock-then-distance estimator with network-level inverse-variance fusion. The resulting per-link variance calibration, which forms the interface to the SCI module, is presented in Section 3.2.3.

Figure 2 previews the temporal relationship among the four processing layers—TDMA slot allocation, polling rounds, MPLS sliding windows, and SCI update sampling—that the rest of this section formalizes.




Figure 2. Temporal structure of the SCI+MPLS pipeline (**top to bottom**): the physical two-way-ranging exchange between nodes i and j ; TDMA slots of duration T_s ; polling rounds $T_{\text{round}} = N \cdot T_s$; overlapping MPLS sliding windows $T_w = 6 T_{\text{round}}$ advancing one round per step; and SCI back-end updates, one per round with prediction-only propagation between them.

3.2.1. Polynomial Delay Model and Linear System

Within an estimation window of duration T_w centered at time t_c , the time-varying propagation delay between nodes i and j is modeled as a polynomial expansion in normalized time:

$$\tau_{ij}(t) = \sum_{l=0}^{L-1} \gamma_{ij,l} P_l\left(\frac{t - t_c}{T_{\text{half}}}\right), \quad (5)$$

where P_l denotes the l -th Legendre polynomial, $T_{\text{half}} = T_w/2$ is the window half-width, and $\gamma_{ij} = [\gamma_{ij,0}, \dots, \gamma_{ij,L-1}]^T$ collects the delay polynomial coefficients. The zeroth-order coefficient $\gamma_{ij,0}$ represents the mean propagation delay (proportional to mean distance via d_{ij}/c), while $\gamma_{ij,1}$ and $\gamma_{ij,2}$ capture (up to a basis-dependent scaling) radial velocity and acceleration effects, respectively.

For the k -th TWR exchange with timestamps $T_{ij}[k]$, $T_{ji}[k]$ and direction flag $\delta_k \in \{+1, -1\}$, substituting Equations (5) and (1) into (2) yields a linear-in-parameters system [5]; stacking K exchanges within the window yields

$$\mathbf{A}_{ij} \boldsymbol{\theta}_{ij} = \mathbf{b}_{ij} + \mathbf{n}_{ij}, \quad (6)$$

where the unknown vector $\boldsymbol{\theta}_{ij} \in \mathbb{R}^{L+2}$ stacks node j 's two clock-calibration parameters (skew α_j and offset β_j) and the L delay-polynomial coefficients γ_{ij} ; $\mathbf{A}_{ij}$ is the timestamp design matrix; $\mathbf{b}_{ij}$ enforces the reference-clock anchor ($\alpha_1 = 1, \beta_1 = 0$); and $\mathbf{n}_{ij} \sim \mathcal{N}(\mathbf{0}, 2\sigma_T^2 \mathbf{I}_K)$ is the stacked timestamp noise over the K exchanges. The weighted least-squares (WLS) solution is

$$\hat{\boldsymbol{\theta}}_{ij} = (\mathbf{A}_{ij}^T \mathbf{W} \mathbf{A}_{ij})^{-1} \mathbf{A}_{ij}^T \mathbf{W} \mathbf{b}_{ij}, \quad \Sigma_{\hat{\boldsymbol{\theta}}} = \hat{\sigma}^2 (\mathbf{A}_{ij}^T \mathbf{A}_{ij})^{-1}, \quad (7)$$

with $\hat{\sigma}^2$ the residual variance over the retained (inlier) rows and the weights iteratively reweighted (IRLS) via Tukey's bisquare [42]; iteration terminates at $\|w^{(k+1)} - w^{(k)}\|_{\infty} < 10^{-3}$ or five iterations. The weights typically settle within 3–4 iterations, so the cap rarely binds

and keeps the per-link solve within the real-time budget, while the 10^{-3} tolerance already lies below the weight change induced by residual timestamp noise. The point estimate uses the IRLS weights, while the covariance is scaled by the Gram matrix of the inlier rows, consistent with Equation (11); any residual variance underestimation this introduces is bounded by the measured calibration ratio $\rho_{\text{cov}} \leq 2.83$.

3.2.2. Sliding-Window Two-Step Algorithm and Sign-Consistency Gate

Per-link application of (7) yields independent per-pair estimates; multi-neighbor clock redundancy motivates a two-step procedure with intermediate clock fusion.

In Step 1, per-link clock estimates are corrected to absolute offsets ($\hat{\beta}_j^{\text{abs}} = \hat{\beta}_j^{\text{rel}} + (1 - \hat{\alpha}_j) t_c$) and fused through inverse-variance weighting:

$$\hat{\alpha}_j^{\text{fused}} = \frac{\sum_{i \in \mathcal{N}_j} w_{ij} \hat{\alpha}_{ji}}{\sum_{i \in \mathcal{N}_j} w_{ij}}, \quad \hat{\sigma}_{\alpha_j^{\text{fused}}}^2 = \frac{1}{\sum_{i \in \mathcal{N}_j} w_{ij}}, \quad (8)$$

with $w_{ij} = 1/\hat{\sigma}_{\alpha_{ji}}^2$, reducing variance by $\sim |\mathcal{N}_j|$ -fold. Cross-correlations from shared timestamps scale as $O(1/K)$ in the window-sample count K and are neglected relative to the per-link timestamp variance. Nodes beyond the reference's direct reach obtain calibrated clocks by composing the affine clock maps (1) along a connecting path (the conditional chain step in Algorithm 1), with negligible multiplicative-error accumulation over the few-hop paths of our network, given UWB-grade oscillators ($|\alpha_i - 1| \lesssim 20$ ppm).

In Step 2, corrected global-time differences relate to the propagation delay by

$$\Delta t_k = \delta_k \mathbf{v}(\tilde{\tau}_k)^\top \gamma_{ij} + n_k, \quad (9)$$

with $\mathbf{v}(\tilde{\tau}_k)$ being the Legendre basis vector. The sign-consistency ratio

$$\rho = \frac{1}{K} \sum_{k=1}^K \mathbf{1}[\delta_k \cdot \Delta t_k > 0] \quad (10)$$

validates physical plausibility: $\rho \rightarrow 1$ under correct clock parameters, and $\rho < \rho_{\text{min}} = 0.7$ triggers window rejection. The threshold sits above the 0.5 random-sign baseline to require a clear physical-plausibility majority, while remaining below unity to tolerate residual timestamp noise after clock correction. When the gate fires, the Step-1 clock estimates are retained, but the Step-2 distance update is skipped for that window; the SCI module continues prediction-only steps until the next window passes the consistency check.

The original MPLS [5] uses the monomial basis, whose Vandermonde design, built on unnormalized timestamps, is poorly scaled and ill-conditioned. The Legendre basis $\{P_l(\tau)\}_{l=0}^{L-1}$, orthogonal on $[-1, 1]$, gives a near-orthogonal—and, after the normalization introduced below, well-scaled—design whose condition number stays bounded across the repeated per-window solves. Its margin over the monomial basis widens as L grows. Timestamps are normalized as $\tilde{\tau}_k = (T_{ij}[k] - t_c)/T_{\text{half}}$. Order $L = 3$ captures up to acceleration whenever $\omega_{\text{rot}} T_w/(2\pi) \ll 1$ ($\omega_{\text{rot}} = v/R$); this ratio stays at most ≈ 0.26 across the tested envelope, peaking at the highest-curvature node ($R = 15$ m, $v = 25$ m/s, $T_{\text{round}} = 160$ ms). Extending to $L = 4$ would raise the unknown count to $L + 2 = 6$, exhausting the $K = 6$ per-window budget and leaving no residual degrees of freedom for variance estimation, whereas $L = 3$ retains one while still resolving acceleration. The order thus trades kinematic fidelity against the finite per-window sample count, with transient super-acceleration maneuvers caught instead by the sign-consistency gate introduced above.

While Ref. [5] processes timestamps in a single batch, continuous tracking requires updating estimates over successive windows. We use overlapping sliding windows of length $T_w = 6 T_{\text{round}}$ (≈ 0.4 s at $T_{\text{round}} = 64$ ms; $K = 6 > L + 2$ at every tested T_{round}), advancing by $5 T_s$ slots; overlapping query times use the window with the nearest center. For each pair and window, the module outputs $(\hat{\gamma}_{ij}, \Sigma_{\hat{\gamma}}^{(ij)}, t_c, T_{\text{half}})$, consumed by Equation (4). Skew is treated as constant within each window; smooth intra-window drift is absorbed into the polynomial delay coefficients alongside motion, so only variation of order higher than L remains unmodeled over the sub-second windows.

The complete two-step procedure is summarized in Algorithm 1.

Algorithm 1 MPLS Two-Step Sliding Window Estimation

Require: TWR timestamps $\{T_{ij}[k], \delta_k\}$; window width T_w ; step size Δ_{step} ; polynomial order L ; consistency threshold $\rho_{\min}$

Ensure: Distance coefficients $\hat{\gamma}_{ij}$ and covariance $\Sigma_{\hat{\gamma}}^{(ij)}$ for all active links

```

1: for each sliding window centered at  $t_c$  do
2:   — Step 1: Clock Estimation —
3:   for each communicating pair  $(i, j)$  do
4:     Construct  $\mathbf{A}_{ij}, \mathbf{b}_{ij}$  (Equation (6)); solve via IRLS-WLS (Equation (7))
5:     Extract  $(\hat{\alpha}_{j|i}, \hat{\beta}_{j|i})$  and convert offset to absolute
6:   end for
7:   for each non-reference node  $j$  do
8:     Fuse per-link clock estimates via inverse-variance weighting (Equation (8))
9:     if  $j$  not adjacent to reference then chain via intermediate node
10:    end if
11:  end for
12:  — Step 2: Distance Estimation —
13:  for each communicating pair  $(i, j)$  do
14:    Correct timestamps to global time using fused clock parameters
15:    Solve  $\hat{\gamma}_{ij}$  via (9)
16:    Compute sign consistency  $\rho$  (Equation (10))
17:    if  $\rho < \rho_{\min}$  then reject window; retain Step 1 estimates
18:    end if
19:  end for
20: end for
  
```

3.2.3. Variance Calibration

Mapping the distance-coefficient covariance through the Legendre basis vector at query time $\tilde{\tau}_k$ yields the per-link, time-varying analytical range variance—the calibrated uncertainty that couples the front-end to the SCI back-end:

$$\hat{\sigma}_{ij}^2(t_k) = c^2 \cdot \mathbf{v}(\tilde{\tau}_k)^\top \mathbf{Cov}(\hat{\gamma}_{ij}) \mathbf{v}(\tilde{\tau}_k), \quad \mathbf{Cov}(\hat{\gamma}_{ij}) = s^2 \cdot (\mathbf{A}_{ij}^\top \mathbf{A}_{ij})^{-1}, \quad (11)$$

where s^2 is the IRLS residual variance and $(\mathbf{A}_{ij}^\top \mathbf{A}_{ij})^{-1}$ inherits the constrained Fisher information structure of Ref. [5]. This $\hat{\sigma}_{ij}^2(t_k)$ is the variance consumed by the SCI module through Equation (14). Experiments bound the calibration ratio $\rho_{\text{cov}} = \text{MSE}(\hat{a}_{ij}) / \hat{\sigma}_{ij}^2 \leq 2.83$ across the envelope ($v \in [10, 40]$ m/s, $T_{\text{round}} \in [16, 160]$ ms; median 1.32); MPLS thus supplies a bias-corrected range whose uncertainty tracks the realized error to within a bounded factor over the tested envelope. This boundedness is structural rather than incidental: writing $\text{MSE} = \hat{\sigma}_{ij}^2 + b_{ij}^2$ with b_{ij} the residual truncation bias, the excess $\rho_{\text{cov}} - 1 = b_{ij}^2 / \hat{\sigma}_{ij}^2$ is the normalized model bias, governed by the same unmodeled kinematic ratio $\omega_{\text{rot}} T_w / (2\pi) \leq 0.26$ that bounds the retained polynomial order; the resulting $\rho_{\text{cov}} \leq 2.83$ is only a $\sqrt{2.83} \approx 1.68 \times$ discrepancy in standard deviation terms.

3.3. Fading SCI Cooperative Estimation Module

Our *fading SCI* back-end adopts the SCIF framework of Ref. [8], which decomposes each agent's error covariance into an *independent* component $\mathbf{P}_I$ and a *dependent* component $\mathbf{P}_D$. We extend it with three design choices not specified there: (i) substitution of the online MPLS calibrated time-varying ranging variance $\hat{\sigma}_{ij}^2(t_k)$ into Pierre's structural noise routing ([8] Equation (21)), in place of an *a priori* constant; (ii) a continuous-time fading factor on $\mathbf{P}_D$; (iii) a deterministic two-stage grid search for the mixing weight ω^* .

3.3.1. Covariance Decomposition and Range Observation Model

The SCIF framework [8] decomposes each agent's error covariance as

$$\mathbf{P}_i = \mathbf{P}_I[i] + \mathbf{P}_D[i], \quad (12)$$

with $\mathbf{P}_I[i]$ collecting agent-local uncertainty (velocity-process noise, hardware ranging noise) and $\mathbf{P}_D[i]$ collecting uncertainty from shared or reused information (inter-agent ranges, MPLS-derived clock parameters); the decomposition is invariant under all prediction and update steps. Initially $\mathbf{P}_I[i](0) = \mathbf{P}_i(0)$, $\mathbf{P}_D[i](0) = \mathbf{0}$; anchor nodes carry $\mathbf{P}_I[j] = \mathbf{P}_D[j] = \mathbf{0}$.

At each SCI update time t_k , the MPLS interface delivers a scalar range $z_{ij}(t_k) = \|\mathbf{x}_i - \mathbf{x}_j\| + n_{ij}$ with variance $\hat{\sigma}_{ij}^2(t_k)$; the linearized observation Jacobian is

$$\mathbf{H}_i = \frac{(\hat{\mathbf{x}}_i - \hat{\mathbf{x}}_j)^\top}{\hat{d}_{ij}}, \quad \mathbf{H}_j = -\mathbf{H}_i. \quad (13)$$

Following the structural routing of ([8] Equation (21)), the noise is split by origin. Pierre states this routing in block-matrix form, with $R_i = \text{diag}(\sigma_r^2, 0)$ and $R_d = \text{diag}(0, \mathbf{H}_j(\mathbf{P}_I[j] + \mathbf{P}_D[j])\mathbf{H}_j^\top)$; for the scalar innovation $\tilde{z}_{ij}$ used here, the off-diagonal blocks collapse, reducing it to the scalar form of Equation (14) below. The scalar range variance $\hat{\sigma}_{ij}^2(t_k)$, private to the link and independent across agents, is assigned to the independent component. The agent- j position covariance, correlated with agent i through the shared inter-agent history, is propagated into the dependent component via the linearized Jacobian:

$$R_i^{\text{eff}} = \hat{\sigma}_{ij}^2(t_k), \quad R_d^{\text{eff}} = \mathbf{H}_j(\mathbf{P}_I[j] + \mathbf{P}_D[j])\mathbf{H}_j^\top, \quad (14)$$

reducing to $R_i^{\text{eff}} = \hat{\sigma}_{ij}^2(t_k)$, $R_d^{\text{eff}} = \mathbf{0}$ when j is an anchor. The form requires no scalar tuning parameter: it is determined by MPLS's calibrated $\hat{\sigma}_{ij}^2(t_k)$ and by $(\mathbf{P}_I[j], \mathbf{P}_D[j])$.

Although Pierre's SCIF [8] assumes an *a priori* constant measurement variance, substituting the data-driven $\hat{\sigma}_{ij}^2(t_k)$ preserves consistency to first order, neglecting the $O(1/K)$ within-window correlation of the reported variance: in the calibrated regime, $\hat{\sigma}_{ij}^2(t_k)$ is dominated by state-independent timestamp noise and adds no cross-agent dependence beyond the explicit Jacobian routing. The property holds whenever $\hat{\sigma}_{ij}^2(t_k) \geq \sigma_T^2$ (the timestamp-noise floor), empirically satisfied across the envelope ($\rho_{\text{cov}} \leq 2.83$).

3.3.2. Prediction with Fading Memory

The process noise $\mathbf{Q}_x$ (Equation (3)) enters only the independent component, $\mathbf{P}_I[i]^- (k) = \mathbf{P}_I[i]^+ (k-1) + \mathbf{Q}_x$, since dead reckoning adds genuinely new, uncorrelated uncertainty.

The dependent component requires the opposite treatment. Because SCI never discounts a past correlation, $\mathbf{P}_D[i]$ accumulates monotonically as inter-agent ranges are reused over a UAV's sustained flight, and an undamped $\mathbf{P}_D[i]$ drives the filter toward the overcon-

servative regime ($\text{ANEES} \ll 1$) of plain CI [8]. We therefore attenuate only the dependent branch—adapting the classical *fading-memory* idea [43,44], also called a *forgetting factor* in recursive least squares—while leaving the independent component, which carries legitimately uncorrelated information, intact:

$$\mathbf{P}_D[i]^{-}(k) = e^{-T_{\text{round}}/\tau} \mathbf{P}_D[i]^{+}(k-1), \quad (15)$$

where $T_{\text{round}} = N \cdot T_s$ is the polling-round duration.

The decay is deliberately parametrized in continuous time rather than per update step. A fixed per-step factor $\lambda \in (0, 1)$ would impose an effective half-life $T_{\text{round}} \ln 2 / (-\ln \lambda)$ that scales with the polling interval; since $T_{\text{round}} = N \cdot T_s$ grows with network size and varies across the sweep, the correlation memory horizon—and hence the attained consistency—would then depend on N and the operating point. Tying the decay to elapsed time instead makes the half-life $t_{1/2} = \tau \ln 2$ a function of τ alone, invariant to T_{round} , so a single τ holds across the entire $v \times T_{\text{round}}$ envelope without per-scenario retuning. We adopt $\tau = 2.0$ s ($t_{1/2} \approx 1.39$ s), a horizon spanning many polling rounds so that genuine correlations persist long enough for consistency, while $\mathbf{P}_D[i]$ remains bounded. The analytical T_{round} -invariance is confirmed empirically in Section 4.5.

3.3.3. SCI Update and Weight Optimization

The complete SCI estimation cycle—prediction, ω -optimization, and split covariance update—is summarized in Algorithm 2.

Algorithm 2 SCI Cooperative Estimation Cycle

Require: State $\hat{\mathbf{x}}_i$, covariances $\mathbf{P}_I[i]^+$, $\mathbf{P}_D[i]^+$; process noise $\mathbf{Q}_x$; fading time constant τ ; polling-round duration T_{round}

Ensure: Updated state $\hat{\mathbf{x}}_i^+$ and decomposition $\mathbf{P}_I[i]^+$, $\mathbf{P}_D[i]^+$

— **Prediction** —

1: $\mathbf{P}_I[i]^{-} \leftarrow \mathbf{P}_I[i]^+ + \mathbf{Q}_x$; $\mathbf{P}_D[i]^{-} \leftarrow e^{-T_{\text{round}}/\tau} \cdot \mathbf{P}_D[i]^+$

— **Measurement Update** — (for each valid range z_{ij})

2: Compute Jacobians $\mathbf{H}_i$, $\mathbf{H}_j$ (Equation (13))

3: Route measurement noise via Equation (14)

4: Optimize ω^* via two-stage grid search (Equation (16); 16 evaluations)

5: Scale: $\mathbf{P}_1 = \mathbf{P}_I[i]^{-} + (\omega^*)^{-1} \mathbf{P}_D[i]^{-}$, $\mathbf{P}_2 = R_i^{\text{eff}} + (1 - \omega^*)^{-1} R_d^{\text{eff}}$

6: Apply Joseph-form update to obtain $\hat{\mathbf{x}}_i^+$, $\mathbf{P}_I[i]^+$, $\mathbf{P}_D[i]^+$ (Equations (17)–(20))

Following Ref. [8], scaled covariances $\mathbf{P}_1 = \mathbf{P}_I[i]^{-} + \omega^{-1} \mathbf{P}_D[i]^{-}$ and $\mathbf{P}_2 = R_i^{\text{eff}} + (1 - \omega)^{-1} R_d^{\text{eff}}$ ensure consistency for any $\omega \in (0, 1)$; the optimal weight minimizes the position-covariance determinant:

$$\omega^* = \arg \min_{\omega \in (0, 1)} \det(\mathbf{P}_i^+(\omega)). \quad (16)$$

Each range update is scalar, so the determinant objective $\det(\mathbf{P}_i^+(\omega))$ is a smooth function of the single bounded variable $\omega \in (0, 1)$, minimized by grid search rather than an interior-point solve. Two properties of this objective underpin the ± 0.025 figure. First, it is empirically *unimodal*: across 10^5 randomized tuples spanning the operational envelope—the noise routing of Equation (14) evaluated over anchored and anchor-free geometries—every case shows either a unique interior minimum or strict monotonicity, with no secondary basins. Under this unimodality, the 11-point coarse stage brackets the minimizer and the five-point fine stage refines within that bracket, so the returned $\hat{\omega}$ lies within the fine-grid resolution of ± 0.025 of ω^* . The figure is thus a grid-resolution guarantee conditional on the observed unimodality, not a closed-form worst-case bound—no closed form for ω^* exists. Second, at an interior minimum, the objective is first-order stationary and hence

locally quadratic, so a residual weight error of ± 0.025 perturbs the updated covariance, and thereby the position estimate, only to second order; the 16 evaluations accordingly suffice, avoiding interior-point overhead on control-rate UAVs. The Kalman gain and state correction follow the standard form:

$$\mathbf{K} = \mathbf{P}_1 \mathbf{H}_i^\top (\mathbf{H}_i \mathbf{P}_1 \mathbf{H}_i^\top + \mathbf{P}_2)^{-1}, \quad \hat{\mathbf{x}}_i^+ = \hat{\mathbf{x}}_i^- + \mathbf{K} \tilde{z}_{ij}, \quad \tilde{z}_{ij} = z_{ij} - \|\hat{\mathbf{x}}_i - \hat{\mathbf{x}}_j\|. \quad (17)$$

Here, $\tilde{z}_{ij}$ is the range innovation, the measured range minus its predicted value at the prior state estimates. The covariance update uses the Joseph form [45]:

$$\mathbf{P}_i^+ = (\mathbf{I} - \mathbf{K} \mathbf{H}_i) \mathbf{P}_1 (\mathbf{I} - \mathbf{K} \mathbf{H}_i)^\top + \mathbf{K} \mathbf{P}_2 \mathbf{K}^\top, \quad (18)$$

$$\mathbf{P}_1[i]^+ = (\mathbf{I} - \mathbf{K} \mathbf{H}_i) \mathbf{P}_1[i]^- (\mathbf{I} - \mathbf{K} \mathbf{H}_i)^\top + \mathbf{K} R_i^{\text{eff}} \mathbf{K}^\top, \quad (19)$$

$$\mathbf{P}_D[i]^+ = \mathbf{P}_i^+ - \mathbf{P}_1[i]^+. \quad (20)$$

Equation (18) uses the *inflated* ($\mathbf{P}_1, \mathbf{P}_2$) for consistency, while Equation (19) uses the *unscaled* ($\mathbf{P}_1[i]^-, R_i^{\text{eff}}$), admitting only genuinely independent information. Positive semi-definiteness of $\mathbf{P}_D[i]^+$ follows because the SCI inflation makes $\mathbf{P}_1 \succeq \mathbf{P}_1[i]^-$ and $\mathbf{P}_2 \geq R_i^{\text{eff}}$, so the independent update subtracted in Equation (20) is dominated in the Löwner order by the total update (both share the gain $\mathbf{K}$), leaving a valid positive-semidefinite dependent covariance. Both agents update from their prior-round ($\mathbf{P}_1, \mathbf{P}_D$) decompositions held as frozen inputs, making the bilateral update order-independent. One SCI update is performed per polling round, matching the MPLS window-output rate so that each fused range is statistically distinct rather than a reuse of the same polynomial fit.

3.4. Co-Design Analysis

This section analyzes the MPLS–SCI co-design: how ranging accuracy propagates to cooperative estimation quality, and how the five compared methods differ in their consistency guarantees. The two problems are addressed jointly because they share a single mechanism: an uncompensated motion bias violates the unbiased-measurement assumption on which every consistent cooperative filter relies. The architecture nonetheless keeps them modular, with front-end and back-end retaining independent internal designs and interacting only through the feedforward variance interface, so joint consistency is obtained without coupling their implementations.

The per-measurement variance $\hat{\sigma}_{ij}^2$ propagates through the SCI noise routing, scaled covariance, and Kalman gain (Equations (14)–(17)); consistent covariance estimates require $\hat{\sigma}_{ij}^2$ to track the true mean-squared error (MSE) $b_{ij}^2 + \sigma_{ij}^2$ within a bounded factor. MPLS is bias-corrected and residual-calibrated, giving $b_{ij} \approx 0$ and bounded $\rho_{\text{cov}} = \text{MSE}(\hat{d}_{ij}) / \hat{\sigma}_{ij}^2$ across the envelope. SDS-TWR, by contrast, leaves the motion bias $v_r \cdot T_{\text{round}}$ unmodeled ($b_{ij}^2 \gg \hat{\sigma}_{ij}^2$), so correlation handling alone cannot recover consistency.

A filter is consistent when $\mathbf{P}_i \succeq \mathbb{E}[\mathbf{e}_i \mathbf{e}_i^\top]$, quantified by the ANEES [2] (ideally unity in $\mathbb{R}^2$). The four fusion strategies span the EKF–CI conservatism spectrum (Table 1); SCI’s $\mathbf{P}_1/\mathbf{P}_D$ decomposition is designed to attain ANEES ≈ 1 when the measurement variance is faithful.

The five methods form a structured ablation along two axes: front-end (MPLS vs. SDS-TWR) with SCI back-end, and back-end (SCI, EKF, EKF+Inflate, CI) with MPLS front-end. Table 1 summarizes their theoretical properties, where *variance fidelity* denotes whether the reported measurement variance reflects the true measurement error and *correlation handling* denotes each method’s treatment of inter-agent information overlap.

Table 1. Theoretical properties of the five compared methods.

Method	Front-End	Correlation Handling	Dynamic Bias	Variance Fidelity	Predicted ANEES
SCI+MPLS	MPLS	Bounded ($\mathbf{P}_I/\mathbf{P}_D$)	Low	Accurate	≈ 1
EKF+MPLS	MPLS	Ignored	Low	Accurate	$\gg 1$
EKF+Inflate+MPLS	MPLS	Ignored (inflated $\mathbf{Q}$)	Low	Inflated	Condition-dep. [†]
CI+MPLS	MPLS	Maximal	Low	Accurate	$\ll 1$
SDS-TWR+SCI	SDS-TWR	Bounded ($\mathbf{P}_I/\mathbf{P}_D$)	Moderate	Nominal	> 1 [‡]

[†] A fixed inflation factor α can reduce ANEES toward unity at one operating point but cannot guarantee consistency across varying velocities, topologies, and communication rates; even a per-scenario oracle α matches the scalar ANEES only pointwise, as isotropic inflation cannot track the anisotropic measurement correlation and over-inflates benign conditions. [‡] SDS-TWR's unmodeled within-exchange distance change leaves a residual bias scaling with relative radial velocity.

Within this comparison set, SCI+MPLS satisfies both requirements: (1) MPLS provides low-bias estimates with calibrated time-varying $\hat{\sigma}_{ij}^2(t_k)$; (2) SCI bounds inter-agent correlations via the $\mathbf{P}_I/\mathbf{P}_D$ decomposition. The other four combinations relax at least one requirement, with the corresponding degradation in accuracy, consistency, or both. Section 4 quantifies these predictions head-to-head, reporting every method's RMSE and ANEES under identical Monte-Carlo conditions ($N_{MC} = 100$ trials) across all three experiments.

Computational complexity. MPLS solves a $(L+2) \times (L+2)$ WLS problem per link per window; SCI performs a 2×2 covariance update with a scalar ω -grid search per neighbor. The per-node cost is $O(|\mathcal{N}_i|)$ per step, matching standard EKF scaling and avoiding centralized $O(N^3)$ or exact-tracking decentralized $O(N^2)$ alternatives. Measured in MATLAB R2025a (The MathWorks, Inc., Natick, MA, USA) at the cross-point ($N = 16$, average degree ≈ 4): MPLS $\approx 5 \mu\text{s}$ per link (≈ 0.16 ms per full sliding window); SCI $\approx 54 \mu\text{s}$ per single update (≈ 0.22 ms per node per polling round, summed over neighbors). The combined per-node cost stays orders of magnitude below the polling-round duration, indicating substantial headroom for onboard real-time execution on the embedded flight-control hardware of a UAV.

4. Experiments and Results

4.1. Simulation Setup

This subsection specifies the platform, common parameters, compared methods, and evaluation metrics shared by the three experiments, which are ordered by increasing scenario complexity (Table 2). Experiment 1 (E1) isolates the ranging front-end on a two-node counter-rotating pair, comparing MPLS against SDS-TWR and a zero-order-hold baseline. Experiment 2 (E2) embeds the front-end in a 16-node anchored swarm (two anchors and fourteen mobile agents) and evaluates the five CL methods. Experiment 3 (E3) removes the anchors, repeating the five-method comparison in a fully mobile 16-node formation to probe anchor-free consistency.

4.1.1. Setup, Parameters, and Compared Methods

Platform and parameters. All experiments run in MATLAB R2025a (Windows 11; Intel Core Ultra 7 255H CPU; 32 GB RAM), using Monte Carlo with $N_{MC} = 100$ trials, each with independently drawn clock parameters, initial trajectory phases, and noise realizations. Common parameters are listed in Table 3; design rationale appears in Sections 3.2 and 3.3. Each agent is initialized at its true position perturbed by $\sigma_{\text{init}} = 1$ m per axis, where $\mathbf{P}_D[i](0) = \mathbf{0}$. LoS propagation is assumed; the wider ± 40 ppm automotive clock-skew band is left to deployment validation. The $\sigma_T = 100$ ps timestamp noise reflects temperature-compensated crystal oscillator (TCXO) and Decawave-class timing; the corresponding

range-equivalent noise floor $\sigma_{Tc} \approx 0.03$ m scales linearly with σ_T . For reproducibility and paired statistical testing, all compared methods are reset to a common per-trial seed before fusion, sharing identical trajectory, clock, and noise realizations within each trial.

Odometry noise model. The velocity sensor follows a $\sqrt{\text{distance}}$ random walk: its position-error standard deviation grows with the square root of the distance traveled, governed by cross-/along-track coefficients (κ_{ct}, κ_{at}) of unit $\sqrt{\text{m}}$ (Table 3), equivalent to $\sim 0.5\%$ drift over 225 m, at the favorable end of the commodity MEMS-IMU band typical of small UAV autopilots [18,24]. Each prediction step over path increment $\Delta s = \|\mathbf{u}_i\| \Delta t$ injects zero-mean Gaussian noise with along-/cross-track standard deviations $\kappa_{at}\sqrt{\Delta s}$ and $\kappa_{ct}\sqrt{\Delta s}$, plus an isotropic $\sigma_{\text{odom}} \Delta t$ floor.

Table 2. Overview of the three experiments.

	Experiment 1	Experiment 2	Experiment 3
Nodes	2	16 (2 anchors + 14 mobile)	16 (all mobile)
Methods	3 ranging methods	5-way ablation	5-way + DR baseline
v range (m/s)	10–40 (step 5)	10–40 (step 5)	10–40 (step 5)
T_{round} range (ms)	16–160	16–160	16–160
v^{ref}	25 m/s	25 m/s	25 m/s
$T_{\text{round}}^{\text{ref}}$	64 ms	64 ms	64 ms
Focus	Ranging accuracy, variance calibration	Position accuracy, estimation consistency	Absolute accuracy, outperforming DR

Table 3. Common simulation parameters shared across the three experiments.

Category	Parameter	Symbol	Value
Protocol	Node count	N	2 (E1)/16 (E2, E3)
	Slot duration	T_s	T_{round}/N
	Simulation duration	T_{sim}	20 s
	Comm. range	r_{comm}	200 m (E2, E3)
Clock	Skew range	α_i	1 ± 20 ppm
	Offset range	β_i	± 1 ms
Noise	Timestamp std. dev.	σ_T	100 ps
	Velocity std. dev.	σ_{odom}	0.1 m/s per axis
	Initial position std. dev.	σ_{init}	1 m per axis
Odometry	Along-track drift coeff.	κ_{at}	$0.04 \sqrt{\text{m}}$
	Cross-track drift coeff.	κ_{ct}	$0.075 \sqrt{\text{m}}$
MPLS	Polynomial order	L	3
	Window duration	T_w	6 rounds (≈ 0.4 s at $T_{\text{round}} = 64$ ms)
	Window step	Δ_{step}	5 slots ($5 T_s$)
SCI	Fading time constant	τ	2.0 s ($t_{1/2} \approx 1.39$ s)
	Update interval	—	one update per polling round T_{round}
Robustification	MPLS Tukey bisquare const.	c_T	4.685 (95% Gaussian efficiency)
	MPLS quality gates	—	$ \tilde{\tau} \leq 1.5, \hat{d} > 0.01$ m, $\hat{\sigma}^2 < 10^{10}$ m ²
Sensitivity grid	Fading τ	—	7 points spanning $[0.4, 40]$ s
Monte Carlo	Trial count	N_{MC}	100

Compared methods. The proposed online MPLS front-end and fading SCI back-end are denoted MPLS and SCI in the ablation labels for brevity. The five methods of Table 1 are evaluated, with EKF+Inflate using a fixed order-of-magnitude process-noise inflation

($\alpha = 10$), as multiplicative inflation lacks a canonical value [33,34]. SDS-TWR [9] cancels first-order clock drift via a symmetric three-message exchange:

$$\hat{d}_{ij}^{\text{SDS}} = \frac{c}{4} [(T_{\text{round},1} - T_{\text{reply},1}) + (T_{\text{round},2} - T_{\text{reply},2})], \quad (21)$$

where $T_{\text{round},k}$ and $T_{\text{reply},k}$ are the round-trip and reply intervals; this estimator assumes static distance within each exchange and reports only a fixed nominal variance $\sigma_T^2 c^2$ (AltDS-TWR [10] shares this limitation). Zero-order hold (ZOH, included only in E1) assigns the most recent TWR distance with the same nominal variance; E3 additionally includes a dead-reckoning (DR) baseline (odometry-only, no range updates).

Evaluation metrics. Metrics are computed over $[0.4 \text{ s}, T_{\text{sim}}]$ (K_{eval} samples, one per TDMA slot; excluding the MPLS window-initialization transient). The per-trial position RMSE averages over N_m agents and the evaluation interval (and is applied to the scalar distance error in E1):

$$\text{RMSE}^{(m)} = \sqrt{\frac{1}{N_m K_{\text{eval}}} \sum_{i,k} \|\mathbf{e}_i^{(m)}(t_k)\|^2}. \quad (22)$$

Cross-trial $\text{RMSE}^{(m)}$ over the N_{MC} trials is reported as mean $\pm$ std (Tables 4 and 5); as its per-trial distribution is near-symmetric (unlike the right-skewed ANEES below), the mean is representative and is the value quoted in the abstract and highlights. ANEES [2] measures filter consistency (ideal = 1 in $\mathbb{R}^2$):

$$\text{ANEES}^{(m)} = \frac{1}{n_x N_m K_{\text{eval}}} \sum_{i,k} \epsilon_i^{(m)}(t_k), \quad \epsilon_i^{(m)} = \mathbf{e}_i^{(m)\top} [\mathbf{P}_i^{(m)}]^{-1} \mathbf{e}_i^{(m)}. \quad (23)$$

Unless noted, reported ANEES is the median of $\text{ANEES}^{(m)}$ across the N_{MC} trials (with a 95% bootstrap confidence interval from 1000 resamples and interquartile range (IQR), as the per-trial distribution is right-skewed (Table 5). NEES 95% coverage reports the fraction satisfying $\epsilon_i \leq \chi_{n_x, 0.95}^2$ (ideal $\approx 95\%$). Consistency is judged by ANEES and NEES 95% coverage jointly, as neither alone is sufficient. All reported p -values derive from the $N_{\text{MC}} = 100$ independent per-trial scalars; the per-trial NEES coverage and bootstrap intervals are descriptive diagnostics. Calibration ratio $\rho_{\text{cov}} = \text{MSE}(\hat{d}_{ij}) / \hat{\sigma}_{ij}^2$ (E1 only) verifies MPLS variance calibration; ideal $\rho_{\text{cov}} = 1$; larger values indicate the reported variance underestimates realized MSE.

4.1.2. Sweep Design and Operational Envelope

Each experiment adopts a two-phase sweep decoupling velocity and frequency sensitivity: Phase 1 fixes $T_{\text{round}} = T_{\text{round}}^{\text{ref}}$ and sweeps $v \in \{5, 10, \dots, 40\}$ m/s; Phase 2 fixes $v = v^{\text{ref}} = 25$ m/s and sweeps T_{round} over a graded sequence spanning 16–160 ms. The lowest sweep point $v = 5$ m/s is simulated for completeness but excluded from reported sweep statistics, as initialization transients dominate at that operating point; the reported envelope is therefore $v \in [10, 40]$ m/s. The cross-point ($v^{\text{ref}}, T_{\text{round}}^{\text{ref}}$) is shared between the two phases. Table 2 summarizes the three experiments.

The velocity sweep $v \in [10, 40]$ m/s spans the range from multirotor cruise to the high-speed agile regime of autonomous racing UAVs [46]. At the cross-point $v = 25$ m/s, the cluster-circular orbits ($R = 15$ m) reach a sustained $\sim 4.3 g$ centripetal load ($v^2/R = 25^2/15 \approx 41.7 \text{ m/s}^2$)—within the maneuvering envelope demonstrated by agile quadrotors, yet deliberately chosen to stress the polynomial-smoothness assumption of the sliding window.

4.2. Experiment 1: MPLS Ranging Validation

This experiment evaluates MPLS ranging accuracy and variance calibration in isolation from the SCI module. Two UAVs perform counter-rotating circular motion ($R_1 = 25$ m, $R_2 = 15$ m, matching the cluster-circular orbits of E2/E3) in an anti-diagonal layout (Figure 3), generating maximum relative radial velocity ($\approx 2v$). Since $N = 2$, raw and ZOH strategies are equivalent (labeled ZOH hereafter). Phase 1 fixes $T_{\text{round}} = 64$ ms and sweeps velocity; Phase 2 fixes $v = 25$ m/s and sweeps $T_{\text{round}} \in \{16, 32, 64, 96, 128, 160\}$ ms.




Figure 3. Experiment 1 scenario: Two UAVs in counter-rotating circular motion with distinct radii ($R_1 = 25$ m, $R_2 = 15$ m) and anti-diagonal centers at (215, 285) and (285, 215) m.

Figure 4a,b shows the distance RMSE across both sweep phases, with the MPLS self-reported $\hat{\sigma}_{ij}(t_k)$ envelope (Equation (11)) overlaid. MPLS rises gently from 0.051 m to 0.063 m across the velocity sweep, tracking $\hat{\sigma}_{ij}$ within $\rho_{\text{cov}} \leq 2.83$, while ZOH and SDS-TWR degrade steadily. At $v = 40$ m/s, MPLS outperforms ZOH by $\sim 10\times$ and SDS-TWR by $\sim 21\times$. In Phase 2, as T_{round} increases from 16 ms to 160 ms, MPLS grows moderately from 0.051 m to 0.144 m, paralleling the rising $\hat{\sigma}_{ij}$ envelope while preserving a $\sim 3\text{--}10\times$ advantage over ZOH that widens with T_{round} . SDS-TWR outperforms ZOH at moderate dynamic stress but reverses at high stress due to its doubled temporal baseline ($2T_{\text{round}}$). Figure 5 shows the cross-point time series, where MPLS residuals stay near the ranging noise floor while ZOH and SDS-TWR trace quasi-periodic envelopes peaking at maximum relative radial velocity.

The calibration ratio ρ_{cov} (Figure 4c,d) confirms variance fidelity. MPLS maintains ρ_{cov} within $[1.27, 2.83]$ across all tested conditions (median 1.32, mean 1.60), approaching the upper bound only at the high-velocity/long- T_{round} corner, with the $\pm 1\sigma$ bands clustering near the median rather than the worst case. MPLS's self-reported uncertainty thus tracks the realized error to within a bounded factor, satisfying the SCI input contract (Equation (14)). ZOH and SDS-TWR report only the fixed nominal variance $\sigma_T^2 c^2$ and miss the motion-induced bias.




Figure 4. Experiment 1 distance RMSE (a,b) and calibration ratio ρ_{cov} (c,d): (a,c) sweep v at $T_{\text{round}} = 64$ ms, (b,d) sweep T_{round} at $v = 25$ m/s. $N_{\text{MC}} = 100$, $\pm 1\sigma$ bands; dotted line in (a,b) is the MPLS $\hat{\sigma}_{ij}$ envelope.




Figure 5. Experiment 1: Distance error time series at the cross-point ($v = 25$ m/s, $T_{\text{round}} = 64$ ms), with per-method RMSE annotated.

4.3. Experiment 2: Anchored CL

This experiment evaluates end-to-end CL (ranging + fusion) in a 16-node network with two anchors at (0, 0) and (500, 500) m. The 14 mobile UAVs execute heterogeneous trajectories in a 500×500 m workspace—a mid-scale patrol and search-and-rescue footprint—with $r_{\text{comm}} = 200$ m, within the typical UWB ranging envelope at the 6.8 Mbps data-rate class [4]. The link set churns frequently.

The trajectory mix (Figure 6) emulates concurrent UAV missions: six high-curvature circular orbits in two clusters α/β ($R = 15$ m) for coordinated loiter, two elliptical orbits, two Lissajous figure-eights, and four non-smooth patrol nodes—three rectangular perimeter-surveillance circuits with sharp 90° corners and one Boustrophedon lawnmower-survey U-turn. The patrol nodes deliberately violate the polynomial-smoothness assump-

tion, stress-testing the MPLS sliding window. All five methods of Table 1 are compared under the two-phase sweep. Full trajectory parameters (radii, frequencies, phases) are available from the corresponding author on request.




Figure 6. Experiment 2 scenario: Two anchors at (0,0)/(500,500) m and 14 mobile UAVs on the heterogeneous trajectory mix. Gray: $r_{\text{comm}} = 200$ m links at mid-time.

The ranging strategy dominates position accuracy (Figure 7): SDS-TWR+SCI produces RMSE of 31.7–58.2 m due to unmodeled distance dynamics, while the four MPLS-based methods stay in 0.36–1.30 m—a 50–100-fold improvement. The four MPLS RMSEs cluster at every operating point because they share identical calibrated ranges; the fusion choice manifests in covariance quality, not point estimates. In Phase 2, SCI+MPLS RMSE grows from 0.52 m at $T_{\text{round}} = 16$ ms to 0.66 m at $T_{\text{round}} = 160$ ms, reflecting per-update odometry drift $\sim \kappa_{\text{ct}} \sqrt{v T_{\text{round}}}$ rather than the MPLS residual.

The ANEES results (Figure 7c,d) confirm the conservatism spectrum of Table 1: median ANEES for EKF+MPLS spans 11–37, EKF+Inflate+MPLS 1.7–13.8, CI+MPLS stays below the strict band (0.17–0.53), and SCI+MPLS lies entirely within the strict [0.5, 2] band (0.59–0.91 across both sweeps). SDS-TWR+SCI reaches $\sim 9 \times 10^3$ – 2.6×10^4 despite SCI fusion, indicating that accurate ranging is a prerequisite for consistency. The χ^2 NEES 95% coverage shows the same ranking (cross-point values in Table 5). SCI+MPLS and CI+MPLS attain comparable coverage, but median ANEES separates them: SCI sits in the lower half of the strict band, mildly conservative, while CI falls below the band, uniformly overconservative. The non-smooth patrol nodes did not destabilize SCI+MPLS: violations are confined to offending links, and the variance interface lets SCI down-weight them automatically. Per-condition RMSE in Table 4.




Figure 7. Experiment 2 position RMSE (a,b; m) and ANEES (c,d; log scale) under velocity (a,c) and T_{round} (b,d) sweeps. IQR bands: SCI+MPLS and CI+MPLS only. Shaded [0.5, 2]: Consistent band (above overconfident, below overconservative).

Table 4. Experiment 2 position RMSE (mean $\pm$ std, m; $N_{\text{MC}} = 100$) at the sweep endpoints and the cross-point.

Method	Phase 1 ($T_{\text{round}} = 64$ ms), v in m/s		
	10	25	40
SCI+MPLS	0.491 ± 0.128	0.548 ± 0.113	0.611 ± 0.145
EKF+MPLS	0.375 ± 0.168	0.435 ± 0.202	0.490 ± 0.181
EKF+Inflate+MPLS	0.391 ± 0.202	0.458 ± 0.234	0.518 ± 0.210
CI+MPLS	1.006 ± 0.142	1.163 ± 0.143	1.286 ± 0.171
SDS-TWR+SCI	31.72 ± 10.90	51.21 ± 9.15	58.23 ± 9.22
Method	Phase 2 ($v = 25$ m/s), T_{round} in ms		
	16	64	160
SCI+MPLS	0.517 ± 0.112	0.548 ± 0.113	0.660 ± 0.158
EKF+MPLS	0.357 ± 0.148	0.435 ± 0.202	0.493 ± 0.162
EKF+Inflate+MPLS	0.448 ± 0.200	0.458 ± 0.234	0.510 ± 0.174
CI+MPLS	0.856 ± 0.115	1.163 ± 0.143	1.304 ± 0.141
SDS-TWR+SCI	50.53 ± 11.33	51.21 ± 9.15	49.48 ± 10.49

4.4. Experiment 3: Anchor-Free CL

This experiment converts Experiment 2's two anchors into mobile UAVs executing small-radius circular orbits ($R_{\text{anc}} = 25$ m at 30% target velocity, same family as the cluster-circular agents); the other 14 UAVs retain the E2 heterogeneous trajectory mix. All 16 UAVs are mobile, no absolute reference is available, and metrics are evaluated in absolute coordinates. The global frame is pinned only by the known initial positions ($\sigma_{\text{init}} = 1$ m), so the network is anchor-free in steady state and the absolute-coordinate RMSE is meaningful only relative to that initial pin; the rotation/translation gauge freedom is not otherwise resolved. Six methods (Table 2), including a dead-reckoning (DR) baseline, are compared under the same two-phase sweep.

MPLS-based methods lead in position accuracy (Figure 8): SCI+MPLS reaches 0.81–1.14 m across the velocity sweep and 0.91–1.03 m across $T_{\text{round}} = 16$ –160 ms, below the DR baseline (1.56–2.07 m and 1.83–1.86 m, respectively). E3 reveals a phenomenon

absent in the anchored setting: SDS-TWR+SCI produces RMSE of 5.3–13.0 m, exceeding the DR baseline at every operating point. The uncompensated systematic bias degrades the odometry-only trajectory, indicating that under highly dynamic anchor-free conditions ranging compensation is needed for the pipeline to outperform pure prediction. The four MPLS RMSEs cluster within 0.79–1.47 m at the cross-point (SCI 0.963, EKF 0.788, EKF+Inflate 0.855, CI 1.469; mean $\pm$ SD in Table 5), again decoupling the fusion choice from point accuracy.




Figure 8. Experiment 3: Position RMSE (a,b; m) and ANEES (c,d; log scale) under velocity (a,c) and T_{round} (b,d) sweeps; DR baseline dashed. Band and IQR conventions as in Figure 7.

The same conservatism ordering persists in the anchor-free setting (Figure 8c,d): SCI+MPLS median ANEES spans 1.16–1.74 across both sweeps (cross-point 1.38). This level is elevated relative to E2 owing to the absent absolute position reference, yet the median falls inside the strict $[0.5, 2]$ band at every tested operating point—a point-wise empirical observation rather than a simultaneous-band guarantee. The $\sqrt{\text{ANEES}}$ underestimation factor $\leq 1.32\times$ is substantially smaller than EKF+MPLS ($\sqrt{60.1} \approx 7.8\times$) or EKF+Inflate+MPLS ($\sqrt{12.2} \approx 3.5\times$). The DR baseline achieves median ANEES ≈ 0.95 since pure prediction incurs no measurement-step double-counting, but with $\sim 2\times$ worse RMSE. Cross-point descriptive statistics are reported alongside E2 in Table 5; Section 4.5 synthesizes findings across all three experiments.

4.5. Summary and Discussion

The three experiments progressively validate the proposed architecture along two independent ablation axes: front-end quality (MPLS vs. SDS-TWR) governs position RMSE, while back-end correlation handling (SCI vs. EKF/CI) governs ANEES. Within the five-method comparison, SCI+MPLS avoids both failure modes—neither the overconfidence of EKF/EKF+Inflate ($\text{ANEES} \gg 1$) nor the over-conservatism of CI ($\text{ANEES} \ll 1$)—across all tested conditions, consistent with the theoretical predictions of Table 1.

Table 5. Cross-point ($v = 25$ m/s, $T_{\text{round}} = 64$ ms, $N_{\text{MC}} = 100$) descriptive statistics: RMSE mean $\pm$ SD, ANEES median [95% CI] and IQR, NEES95 coverage.

Exp.	Method	RMSE (m)		ANEES		NEES95 (%)
		Mean	SD	Median [95% Conf. Int.]	IQR	
E2	SCI+MPLS	0.548	0.113	0.68 [0.64, 0.73]	0.28	93.8
E2	EKF+MPLS	0.435	0.202	20.2 [17.6, 26.5]	27.8	20.7
E2	EKF+Inflate+MPLS	0.458	0.234	4.02 [3.09, 4.84]	5.07	61.6
E2	CI+MPLS	1.163	0.143	0.404 [0.388, 0.423]	0.099	93.1
E2	SDS-TWR+SCI	51.21	9.15	2.0×10^4 [1.7×10^4 , 2.1×10^4]	9.5×10^3	4.5
E3	SCI+MPLS	0.963	0.183	1.38 [1.23, 1.61]	0.85	87.7
E3	EKF+MPLS	0.788	0.418	60.1 [49.3, 74.6]	71.8	7.7
E3	EKF+Inflate+MPLS	0.855	0.556	12.2 [10.9, 14.4]	11.4	34.8
E3	CI+MPLS	1.469	0.181	0.492 [0.456, 0.518]	0.147	93.5
E3	SDS-TWR+SCI	10.30	1.79	201 [174, 218]	120	22.6
E3	DR	1.856	0.230	0.953 [0.927, 1.015]	0.303	94.4

Although the four MPLS-based methods cluster in RMSE, their median ANEES spans up to two orders of magnitude—widest in E3 (Table 5). From E2 to E3, the SCI+MPLS median rises from 0.68 to 1.38 but stays within the strict [0.5, 2] band, reflecting the loss of absolute reference. The tabulated spread and coverage tell the same story: SCI+MPLS keeps a tight ANEES interquartile range (0.28 in E2, 0.85 in E3) and near-nominal 95% NEES coverage (93.8% and 87.7%), whereas the EKF-family entries exhibit order-of-magnitude interquartile ranges and NEES coverage falling to 8–62%. ANEES is summarized by its median rather than mean throughout, as the EKF-family distribution is heavily right-skewed.

4.5.1. Statistical and Practical Significance

Among MPLS-based methods, the principal differentiator is covariance quality rather than point accuracy. At the E2 cross-point, EKF+MPLS's median ANEES of 20.2 corresponds to a reported σ underestimating the true error by $\sqrt{20.2} \approx 4.5\times$, while SCI+MPLS (median ANEES 0.68) overestimates the true error by only $\sqrt{1/0.68} \approx 1.21\times$. This provides a calibrated, typical-case bound for downstream decision-making (e.g., formation safety margins).

Pairwise Wilcoxon signed-rank tests on the $N = 100$ per-trial scalars, with Holm–Bonferroni correction across the comparison family, confirm three patterns across both E2 and E3; the verdicts are insensitive to the exact family size, as all adjusted p remain far below 10^{-3} . The consistency comparisons are one-sided: the conservatism spectrum of Table 1 fixes the expected deviation *direction* a priori (SCI nearer unity than its over/under-confident competitors), whereas the RMSE comparisons, lacking such a directional prior, are two-sided. (i) On consistency ($|\log(\text{ANEES})|$), SCI+MPLS is significantly better than every EKF-family and CI competitor (all $p_{\text{Holm}} < 10^{-3}$, at the $N = 100$ exact-test resolution floor; effect size $r \in [0.48, 0.87]$ [47]); CI's E3 median ANEES of 0.49 sits at the lower band edge. (ii) SCI is likewise significantly better on RMSE than EKF and EKF+Inflate, and outperforms CI+MPLS by 34–53%. (iii) In E3, DR matches SCI+MPLS on consistency but with $\sim 2.0\times$ worse RMSE. The trade-off: SCI+MPLS accepts ~ 13 –26% RMSE cost versus EKF variants to gain ~ 0.4 –1.6 orders in $|\log_{10}(\text{ANEES})|$ (The RMSE cost is computed on means and the consistency gain on medians, reflecting the differing skew of the two metrics), with consistency effect size materially exceeding RMSE effect size.

4.5.2. Robustness and Deployment Considerations

Fading time constant. Sweeping τ over the seven-point grid of Table 3 ($[0.4, 40]$ s) traces a monotone consistency–accuracy trade-off at the E3 cross-point: a short memory (small τ) leaves the estimate overconfident (median ANEES 5.86 at $\tau = 0.4$ s), a long memory (large τ) turns it overconservative (median ANEES 0.56 at $\tau = 40$ s), and RMSE varies by only $\sim 36\%$ across the span. The chosen $\tau = 2.0$ s sits in the centering region (median ANEES 1.38). A complementary two-configuration comparison across the polling-period range shows that the continuous-time parametrization of Equation (15) keeps the median ANEES within a $\leq 1.6\times$ range factor, versus up to $5.4\times$ for a fixed-per-step decay, confirming the T_{round} -invariance claim of Section 3.3.2.

Packet loss. Figure 9 sweeps i.i.d. packet loss $p_{\text{loss}} \in [0, 20]\%$ at the cross-point ($N_{\text{MC}} = 100$), reported as median + IQR. Among the four MPLS-based methods, SCI+MPLS's median ANEES remains within one order of magnitude of unity across the entire sweep (0.68, 0.88, 1.31, 1.82, 3.30 at $p_{\text{loss}} = 0, 5, 10, 15, 20\%$). The strict $[0.5, 2]$ band is maintained up to $p_{\text{loss}} = 15\%$ and exited only at $p_{\text{loss}} = 20\%$, as missing rounds accumulate prediction-only drift; this band figure is for i.i.d. loss, and bursty loss—concentrating missed rounds—would degrade consistency at a lower mean rate. The swept $[0, 20]\%$ span brackets the loss levels representative of dense multi-UAV operation; higher sustained loss reflects a link-budget shortfall, better resolved at the MAC/PHY layer than by the estimator. Degradation stays graceful rather than catastrophic on both axes: at the fusion stage, it is the monotone consistency drift above, exiting the band near 20% without abrupt divergence; at the ranging stage, a window that violates the polynomial-smoothness assumption is degraded per-link by the sign-consistency gate rather than destabilizing the estimate. Other fusion methods stay at their structural conservatism levels across $p_{\text{loss}} = 0\text{--}20\%$ (EKF+MPLS rising from ≈ 20 to 118, EKF+Inflate 4 to 37, CI+MPLS 0.40 to 0.44). SDS-TWR+SCI remains motion-bias-dominated, omitted from the figure for axis readability. Per-trial inspection shows the MC-mean SCI+MPLS ANEES at higher p_{loss} is dominated by an outlier minority (per-trial maxima $10^2\text{--}10^4$); the IQR remains clustered around the median.




Figure 9. Robustness to i.i.d. packet loss at the cross-point ($v = 25$ m/s, $T_{\text{round}} = 64$ ms, $N_{\text{MC}} = 100$): (a) position RMSE (median, IQR); (b) median ANEES (log scale; dashed: ideal).

Simulation-to-real gap. Deployment on physical UAVs introduces effects orthogonal to motion-induced bias: NLOS multipath from airframes and surrounding structures [20,21], antenna patterns that vary with vehicle attitude, bursty (rather than i.i.d.) packet loss under contested spectrum, and minute-scale crystal-oscillator wander. These act on the same

two axes (variance fidelity and correlation handling) and can be staged as preprocessing modules without architectural change.

Section 5 concludes the paper and discusses future research directions.

5. Conclusions

This paper proposed a modular architecture for dynamic ranging error compensation and consistent CL in TDMA UWB UAV swarm networks, evaluated in Monte Carlo simulation. The online MPLS ranging front-end, extending Ref. [5], delivers a bias-corrected range together with a calibrated time-varying analytical variance $\hat{\sigma}_{ij}^2(t_k)$ that serves as the per-link variance interface to the fading SCI cooperative back-end. The latter, extending Ref. [8] with a continuous-time fading factor on $\mathbf{P}_D$, consumes $\hat{\sigma}_{ij}^2(t_k)$ as its measurement-noise input. A five-method ablation across both anchored and anchor-free 16-node scenarios confirmed that front-end quality governs RMSE while back-end correlation handling governs ANEES. SCI+MPLS keeps median ANEES within the strict $[0.5, 2]$ band at every tested operating point across the envelope ($v \in [10, 40]$ m/s, $T_{\text{round}} \in [16, 160]$ ms), closing the ANEES gap against EKF variants by 0.4–1.6 orders of magnitude at 13–26% additional RMSE cost. Uncompensated SDS-TWR degrades below dead reckoning under anchor-free conditions, indicating that dynamic ranging compensation is needed for CL to outperform pure prediction at high relative motion. Together, these results show that the architecture performs what a GNSS-denied swarm needs. Without per-scenario tuning and within a UAV's onboard compute budget, it delivers, in both anchored and anchor-free flight, not only sub-meter positioning but also a covariance that downstream collision-avoidance and formation keeping can trust—one that the overconfident EKF and overconservative CI baselines evaluated here fail to provide.

Several directions remain. Flight validation on a UWB-equipped UAV testbed would confirm the framework's real-time feasibility and robustness under realistic NLOS, antenna-pattern, bursty-packet-loss, and crystal-wander conditions. Within the algorithmic framework, three-dimensional localization is a direct extension (chiefly state-vector augmentation), and adaptive polynomial-order selection could absorb impulsive maneuvers that exceed the smoothness regime of $L = 3$. Because the polling round $T_{\text{round}} = N \cdot T_s$ grows with network size, the T_{round} sweep already probes the per-round-duration degradation that larger swarms induce at fixed slot length. Channel saturation and neighbor-set growth at substantially larger N remain for future work.

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

Funding: This research was funded by the National Key R&D Program of China, grant number 2021YFA0716603.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data presented in this study are available on request from the corresponding author. The data are not publicly available due to the data being part of an ongoing research project.

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

Abbreviations

The following abbreviations are used in this manuscript:

UAV	Unmanned Aerial Vehicle
UWB	Ultra-Wideband
GNSS	Global Navigation Satellite System
IMU	Inertial Measurement Unit
TWR	Two-Way Ranging
SDS-TWR	Symmetric Double-Sided Two-Way Ranging
TDMA	Time-Division Multiple Access
ZOH	Zero-Order Hold
DR	Dead Reckoning
LoS	Line-of-Sight
NLOS	Non-Line-of-Sight
MPLS	Maximum-Likelihood Polynomial Least Squares
SCI	Split Covariance Intersection
SCIF	Split Covariance Intersection Filter
EKF	Extended Kalman Filter
CI	Covariance Intersection
RMSE	Root Mean Square Error
NEES	Normalized Estimation Error Squared
ANEES	Average NEES
IQR	Interquartile Range

References

- Shakhathreh, H.; Sawalmeh, A.H.; Al-Fuqaha, A.I.; Dou, Z.; Almaita, E.K.; Khalil, I.M.; Othman, N.S.; Khreishah, A.; Guizani, M. Unmanned Aerial Vehicles (UAVs): A Survey on Civil Applications and Key Research Challenges. *IEEE Access* **2019**, *7*, 48572–48634. [\[CrossRef\]](#)
- Bar-Shalom, Y.; Li, X.R.; Kirubarajan, T. *Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software*; John Wiley & Sons: Hoboken, NJ, USA, 2001. [\[CrossRef\]](#)
- Citroni, R.; Di Paolo, F.; Livreri, P. A Novel Energy Harvester for Powering Small UAVs: Performance Analysis, Model Validation and Flight Results. *Sensors* **2019**, *19*, 1771. [\[CrossRef\]](#) [\[PubMed\]](#)
- Gezici, S.; Tian, Z.; Giannakis, G.B.; Kobayashi, H.; Molisch, A.F.; Poor, H.V.; Sahinoglu, Z. Localization via Ultra-Wideband Radios: A Look at Positioning Aspects for Future Sensor Networks. *IEEE Signal Process. Mag.* **2005**, *22*, 70–84. [\[CrossRef\]](#)
- Rajan, R.T.; van der Veen, A.J. Joint Ranging and Synchronization for an Anchorless Network of Mobile Nodes. *IEEE Trans. Signal Process.* **2015**, *63*, 1925–1940. [\[CrossRef\]](#)
- Lian Sang, C.; Adams, M.; Hörmann, T.; Hesse, M.; Pörmann, M.; Rückert, U. Numerical and Experimental Evaluation of Error Estimation for Two-Way Ranging Methods. *Sensors* **2019**, *19*, 616. [\[CrossRef\]](#) [\[PubMed\]](#)
- Kia, S.S.; Rounds, S.; Martínez, S. Cooperative Localization for Mobile Agents: A Recursive Decentralized Algorithm Based on Kalman-Filter Decoupling. *IEEE Control Syst. Mag.* **2016**, *36*, 86–101. [\[CrossRef\]](#)
- Pierre, C.; Chapuis, R.; Aufrère, R.; Laneurit, J.; Debain, C. Range-Only Based Cooperative Localization for Mobile Robots. In Proceedings of the 21st International Conference on Information Fusion (FUSION), Cambridge, UK, 10–13 July 2018; pp. 1933–1939. [\[CrossRef\]](#)
- IEEE Std 802.15.4z-2020; IEEE Standard for Low-Rate Wireless Networks—Amendment 1: Enhanced Ultra Wideband (UWB) Physical Layers (PHYs) and Associated Ranging Techniques. IEEE: New York, NY, USA, 2020. [\[CrossRef\]](#)
- Neiryneck, D.; Luk, E.; McLaughlin, M. An Alternative Double-Sided Two-Way Ranging Method. In Proceedings of the 13th Workshop on Positioning, Navigation and Communications (WPNC), Bremen, Germany, 19–20 October 2016; pp. 1–4. [\[CrossRef\]](#)
- Sidorenko, J.; Schatz, V.; Scherer-Negenborn, N.; Arens, M.; Hugentobler, U. Decawave UWB Clock Drift Correction and Power Self-Calibration. *Sensors* **2019**, *19*, 2942. [\[CrossRef\]](#) [\[PubMed\]](#)
- Corbalán, P.; Picco, G.P. Ultra-wideband Concurrent Ranging. *ACM Trans. Sens. Netw.* **2020**, *16*, 41. [\[CrossRef\]](#)
- Rustamaji; Kliwati, S.; Widada, W. Two-way ranging with carrier frequency and phase offset compensation based on software defined radio. *Measurement* **2024**, *232*, 114659. [\[CrossRef\]](#)
- Zhao, W.; Panerati, J.; Schoellig, A.P. Learning-based Bias Correction for Time Difference of Arrival Ultra-wideband Localization of Resource-constrained Mobile Robots. *IEEE Robot. Autom. Lett.* **2021**, *6*, 3639–3646. [\[CrossRef\]](#)

15. Zhao, S.; Guo, N.; Zhang, X.P.; Cui, X.; Lu, M. Closed-Form Two-Way TOA Localization and Synchronization for User Devices with Motion and Clock Drift. *IEEE Signal Process. Lett.* **2022**, *29*, 100–104. [\[CrossRef\]](#)
16. Guo, N.; Zhao, S.; Zhang, X.P.; Yao, Z.; Cui, X.; Lu, M. New Closed-Form Joint Localization and Synchronization Using Sequential One-Way TOAs. *IEEE Trans. Signal Process.* **2022**, *70*, 2078–2092. [\[CrossRef\]](#)
17. Mishra, A.; Rajan, R.T. Estimation of Relative Kinematic Parameters of an Anchorless Network. *IEEE Trans. Signal Inf. Process. Netw.* **2025**, *11*, 831–844. [\[CrossRef\]](#)
18. Feng, D.; Wang, C.; He, C.; Zhuang, Y.; Xia, X.G. Kalman-Filter-Based Integration of IMU and UWB for High-Accuracy Indoor Positioning and Navigation. *IEEE Internet Things J.* **2020**, *7*, 3133–3146. [\[CrossRef\]](#)
19. Goudar, A.; Zhao, W.; Schoellig, A.P. Range-Visual-Inertial Sensor Fusion for Micro Aerial Vehicle Localization and Navigation. *IEEE Robot. Autom. Lett.* **2024**, *9*, 683–690. [\[CrossRef\]](#)
20. Angarano, S.; Mazzia, V.; Salvetti, F.; Fantin, G.; Chiaberge, M. Robust Ultra-Wideband Range Error Mitigation with Deep Learning at the Edge. *Eng. Appl. Artif. Intell.* **2021**, *102*, 104278. [\[CrossRef\]](#)
21. Coppens, D.; Van Herbruggen, B.; Shahid, A.; De Poorter, E. Removing the Need for Ground Truth UWB Data Collection: Self-Supervised Ranging Error Correction Using Deep Reinforcement Learning. *IEEE Trans. Mach. Learn. Commun. Netw.* **2024**, *2*, 1615–1627. [\[CrossRef\]](#)
22. Digulescu, A.; Despina-Stoian, C.; Popescu, F.; Stanescu, D.; Nastasiu, D.; Sburlan, D. UWB Sensing for UAV and Human Comparative Movement Characterization. *Sensors* **2023**, *23*, 1956. [\[CrossRef\]](#) [\[PubMed\]](#)
23. Wymeersch, H.; Lien, J.; Win, M.Z. Cooperative Localization in Wireless Networks. *Proc. IEEE* **2009**, *97*, 427–450. [\[CrossRef\]](#)
24. Guo, K.; Li, X.; Xie, L. Ultra-Wideband and Odometry-Based Cooperative Relative Localization with Application to Multi-UAV Formation Control. *IEEE Trans. Cybern.* **2020**, *50*, 2590–2603. [\[CrossRef\]](#) [\[PubMed\]](#)
25. Xu, H.; Zhang, Y.; Zhou, B.; Wang, L.; Yao, X.; Meng, G.; Shen, S. Omni-Swarm: A Decentralized Omnidirectional Visual-Inertial-UWB State Estimation System for Aerial Swarms. *IEEE Trans. Robot.* **2022**, *38*, 3374–3394. [\[CrossRef\]](#)
26. Luo, X.; Du, X.; Yue, S.; Lv, Y.; Zhang, L.; He, X.; Wu, W.; Mao, J. DTVMR-Swarm: A Distributed and Tightly Integrated Visual-Inertial-UWB-Magnetic System for Anchor Free Swarm Cooperative Localization. *Drones* **2026**, *10*, 49. [\[CrossRef\]](#)
27. Xu, H.; Zhao, G.; Bo, Y. UAV-UGV Collaborative Localization in GNSS-Denied Large-Scale Environments: An Anchor-Free VIO-UWB Fusion with Adaptive Weighting and Outlier Suppression. *Drones* **2026**, *10*, 175. [\[CrossRef\]](#)
28. Huang, G.P.; Trawny, N.; Mourikis, A.I.; Roumeliotis, S.I. Observability-Based Consistent EKF Estimators for Multi-Robot Cooperative Localization. *Auton. Robot.* **2011**, *30*, 99–122. [\[CrossRef\]](#)
29. Hao, N.; He, F.; Tian, C.; Yao, Y.; Xia, W. KD-EKF: A Consistent Cooperative Localization Estimator Based on Kalman Decomposition. In *2023 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*; IEEE: New York, NY, USA, 2023; pp. 11064–11070. [\[CrossRef\]](#)
30. Kaess, M.; Johannsson, H.; Roberts, R.; Ila, V.; Leonard, J.J.; Dellaert, F. iSAM2: Incremental Smoothing and Mapping Using the Bayes Tree. *Int. J. Robot. Res.* **2012**, *31*, 216–235. [\[CrossRef\]](#)
31. Corbalán, P.; Picco, G.P.; Coors, M.; Jain, V. Self-Localization of Ultra-Wideband Anchors: From Theory to Practice. *IEEE Access* **2023**, *11*, 29711–29725. [\[CrossRef\]](#)
32. Luft, L.; Schubert, T.; Roumeliotis, S.I.; Burgard, W. Recursive Decentralized Localization for Multi-Robot Systems with Asynchronous Pairwise Communication. *Int. J. Robot. Res.* **2018**, *37*, 1152–1167. [\[CrossRef\]](#)
33. Masinjala, R.; Payeur, P. Consistent Multirobot Localization Using Heuristically Tuned Extended Kalman Filter. In *2017 IEEE International Symposium on Robotics and Intelligent Sensors (IRIS)*; IEEE: New York, NY, USA, 2017; pp. 297–303. [\[CrossRef\]](#)
34. Anderson, J.L.; Anderson, S.L. A Monte Carlo Implementation of the Nonlinear Filtering Problem to Produce Ensemble Assimilations and Forecasts. *Mon. Weather Rev.* **1999**, *127*, 2741–2758. [\[CrossRef\]](#)
35. Julier, S.J.; Uhlmann, J.K. A Non-divergent Estimation Algorithm in the Presence of Unknown Correlations. In *Proceedings of the American Control Conference (ACC)*, Albuquerque, NM, USA, 4–6 June 1997; pp. 2369–2373. [\[CrossRef\]](#)
36. Julier, S.J.; Uhlmann, J.K. General Decentralized Data Fusion with Covariance Intersection. In *Handbook of Multisensor Data Fusion: Theory and Practice*, 2nd ed.; Liggins, M.E., Hall, D.L., Llinas, J., Eds.; CRC Press: Boca Raton, FL, USA, 2009; Chapter 12, pp. 319–344.
37. Bahr, A.; Walter, M.R.; Leonard, J.J. Consistent Cooperative Localization. In *2009 IEEE International Conference on Robotics and Automation (ICRA)*; IEEE: New York, NY, USA, 2009; pp. 3415–3422. [\[CrossRef\]](#)
38. Li, H.; Nashashibi, F.; Yang, M. Split Covariance Intersection Filter: Theory and Its Application to Vehicle Localization. *IEEE Trans. Intell. Transp. Syst.* **2013**, *14*, 1860–1871. [\[CrossRef\]](#)
39. Shan, X.; Cabani, A.; Chafouk, H. Interval Split Covariance Intersection Filter: Theory and Its Application to Cooperative Localization in a Multi-Sensor Multi-Vehicle System. *Sensors* **2024**, *24*, 3124. [\[CrossRef\]](#) [\[PubMed\]](#)
40. Xiong, J.; Xiong, Z.; Zhuang, Y.; Cheong, J.W.; Dempster, A.G. Fault-Tolerant Cooperative Positioning Based on Hybrid Robust Gaussian Belief Propagation. *IEEE Trans. Intell. Transp. Syst.* **2023**, *24*, 6425–6431. [\[CrossRef\]](#)

41. Yan, J.; Zhu, F.; Huang, Y.; Zhang, Y. Robust Decentralized Cooperative Localization for Multirobot System Against Measurement Outliers. *IEEE Internet Things J.* **2024**, *11*, 17934–17947. [[CrossRef](#)]
42. Beaton, A.E.; Tukey, J.W. The Fitting of Power Series, Meaning Polynomials, Illustrated on Band-Spectroscopic Data. *Technometrics* **1974**, *16*, 147–185. [[CrossRef](#)]
43. Jazwinski, A.H. Stochastic Processes and Filtering Theory. In *Mathematics in Science and Engineering*; Academic Press: Cambridge, MA, USA, 1970; Volume 64.
44. Simon, D. *Optimal State Estimation: Kalman, H_∞ , and Nonlinear Approaches*; John Wiley & Sons: Hoboken, NJ, USA, 2006.
45. Bucy, R.S.; Joseph, P.D. *Filtering for Stochastic Processes with Applications to Guidance*; Interscience Publishers: New York, NY, USA, 1968.
46. Kaufmann, E.; Bauersfeld, L.; Loquercio, A.; Müller, M.; Koltun, V.; Scaramuzza, D. Champion-Level Drone Racing Using Deep Reinforcement Learning. *Nature* **2023**, *620*, 982–987. [[CrossRef](#)] [[PubMed](#)]
47. Cohen, J. *Statistical Power Analysis for the Behavioral Sciences*, 2nd ed.; Lawrence Erlbaum Associates: Hillsdale, NJ, USA, 1988.

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.