of world literature in this area, summarizes the results of many years of research by the authors, and outlines current problems and options for their solution in the future.

Keywords: bacteriorhodopsin, photocycle, film structures, photochromism, polarization recording, information security, information protection.

Надійшла до редакції (Received): 16.08.2026
Прийнята до друку (Accepted): 08.09.2026
Опубліковано онлайн (Available online): 15.09.2026

UDC 004.75:004.62 DOI: 10.31673/2409-7292.2026.032817

Denys Chernousov
PhD student of the Department of Software Engineering in Energy
National Technical University of Ukraine "Igor Sikorsky Kyiv Polytechnic Institute", Kyiv, Ukraine
ORCID: 0009-0003-1968-860X
E-mail: chernousov.d.i.-tv51f@edu.kpi.ua

REAL-TIME SPATIO-TEMPORAL DEDUPLICATION IN DISTRIBUTED MULTI-TARGET TRACKING VIA A BOUNDED-DEGRADATION ARCHITECTURE

Distributed multi-target tracking systems tend to fall apart under heavy load: processing delays grow, tracks merge into each other, and the root cause is usually unbounded algorithmic complexity combined with unpredictable memory access patterns. This paper focuses on state estimation and spatio-temporal track deduplication across decentralized processing nodes. The goal is to keep frame execution time deterministic and prevent identity swaps between tracks, even during dense swarm engagements. To accomplish this, the study uses maneuvering target models reaching 10g acceleration and simulated telemetry streams pushed to an 11× clutter saturation limit. The approach combines lock-free ring buffering, structure-of-arrays memory layouts, flat prime-product spatial hash-grids with linear probing, and localized joint probabilistic data association. The results show that a sub-10 ms processing budget holds at 5,500 clutter points per frame, with a peak projected latency of 4.6 ms when all eight crossing locks are active simultaneously. Across 1,000 simulated trials of symmetric high-g crossings, the framework produced zero label swaps and maintained consistent tracking through packet loss rates of up to 30%. The core scientific contribution is the formalization of the cache-hot paradox: under heavy clutter, a bounded full-rank Cholesky factorization running entirely in L1 cache outperforms the nominally cheaper Sherman-Morrison-Woodbury low-rank approximation, which gets bottlenecked by global memory latency. A second contribution is an additive correction term for the NEES selection rule that absorbs the innovation inflation caused by low-rank process noise approximations, preventing false ownership handoffs across lossy network links.

Keywords: real-time target tracking, hardware-software co-design, distributed state estimation, cache-hot paradox, joint probabilistic data association, track deduplication, bounded degradation, decentralized sensor networks.

Introduction

The growing use of tactical UAVs is pushing short-range air defense to its limits, especially when dealing with low radar cross-section targets. Standalone radars can measure range accurately but struggle with angular resolution and are vulnerable to electronic warfare [1]. Passive IR and optical sensors excel at angular measurements and target recognition, but cannot directly measure distance [2]. The natural answer is to combine these sensors into heterogeneous distributed networks [3]. However, fusing their data streams in real time is harder than it sounds: spatio-temporal redundancy regularly produces duplicate or ghost tracks that clog the decision pipeline [4].

Data association gets particularly messy in dense swarm scenarios [5]. Overlapping sensor coverage zones produce phantom bearing-line intersections [6], and asynchronous data arrival makes things worse [7]. The resulting ghost tracks overload decision channels and erode system reliability [8]. Meanwhile, tactical UAVs routinely pull high-g maneuvers that classical linear filters simply cannot handle [9, 10]. Adaptive filtering algorithms, covariance intersection, and information matrix methods are needed to keep up [11, 12].

© Трикур І.І., Січка М.Ю., Різак М.В., Боценюк Д.Р., Різак В.М. Бактеріородопсин та хімічно модифіковані плівкові структури на його основі як компонент систем захисту інформації. Сучасний захист інформації, 3(67), 173–185. https://doi.org/10.31673/2409-7292.2026.034316

Mounting sensors on moving platforms introduces significant rotational and translational complexities. Specifically, the carrier's own motion degrades coordinate accuracy and demands continuous real-time transformations between local and global reference frames [13]. Taken together, these challenges make spatio-temporal deduplication a core problem for any reliable automated fusion pipeline [14, 15].

Literature review and problem statement

Distributed multi-target state estimation has always been computationally expensive, and real-time tactical surveillance only makes the constraints tighter [4, 7]. In close-range C-UAS engagements, fusion engines must process high-frequency telemetry under heavy clutter and active jamming [3, 8]. Classical spatial association methods work well for steady, linear trajectories [16], but they break down predictably when hit with dense swarms performing erratic, jerk-dominated maneuvers, they introduce deterministic failure modes when subjected to high-density target swarms executing erratic, jerk-dominated maneuvers [17].

The literature tends to split into two camps: pure filter theory on one side, platform-agnostic software abstractions on the other. This split hides a class of hardware-level failures, such as RAM bus latencies and cache-line invalidation storms. These microarchitectural bottlenecks can bring down embedded nodes even when the math looks fine on paper. One key finding of this work is that filter divergence and microarchitectural exhaustion are causally linked: traversing pointer-bound spatial structures triggers continuous cache-line evictions that blow past the WCET budget.

To understand where the gaps are, the following subsections review multi-sensor fusion work from 2017 to 2026. This specific period is defined by the widespread adoption of low-SWaP AESA modules and the emergence of agile micro-UAV threats.

1. Algorithmic and Structural Trajectories in Edge-Native Data Fusion

Tracking frameworks have moved away from simple linear filters toward sequential estimators such as IMM configurations and Unscented Particle Filters, combined with JPDA to handle non-Gaussian uncertainty and aggressive maneuvers [4, 14, 16]. At the same time, compute has shifted from central servers to edge-native, low-SWaP AESA modules and passive IR gimbals that generate tracklets directly at the sensor payload [10, 18]. Rigid monolithic software has given way to distributed topologies built around non-blocking data pipelines and connectionless binary datagrams [6, 17]. Yet one fundamental problem persists: classical mathematical models assume unlimited compute resources. On real embedded hardware, that assumption fails under RAM bus contention and cache-line invalidation, driving filter divergence even when the algorithm itself is theoretically sound.

2. Structural Analysis of Existing Tactical Airspace Surveillance Systems

Tactical airspace surveillance architectures deploy specialized hardware and software layers to suppress spurious noise and isolate true target profiles from high-density environmental clutter [2, 7]. Prior to track initiation, sliding data correlation algorithms extract sub-frame signal peaks across successive scans to establish short-term track stationarity metrics [11]. Local candidate observations are subsequently filtered using the Mahalanobis distance metric to compute the statistical distance between an incoming measurement vector $z(k)$ and the predicted target state estimate $\hat{z}(k|k-1)$, scaled inversely by the innovation covariance matrix $S(k)$ [4, 8]:

$$d_M^2 = (\mathbf{z} - \mathbf{H}\hat{\mathbf{x}}_{k|k-1})^T \mathbf{S}_k^{-1} (\mathbf{z} - \mathbf{H}\hat{\mathbf{x}}_{k|k-1}). \quad (1)$$

Measurements exceeding a strict chi-squared thresholdγ for n_z degrees of freedom are discarded as clutter [4, 8]. To isolate maneuvering aircraft from multipath reflections, ellipsoidal validation gates dynamically bound the tracking region's volume $V(k)$ using the formulation [4]:

$$V_k = c_n \gamma^{n/2} |\mathbf{S}_k|^{1/2}, \quad (2)$$

where c_{n_z} is the coefficient for the volume of an n_z-dimensional unit hypersphere. Ghost tracks arising from urban multipath reflections or electronic countermeasures are pruned using sequential

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

M out of N initiation logic and cross-node tracklet similarity indices [1, 3]. In passive sensor arrays, line-of-bearing intersection ambiguities are resolved via time-difference-of-arrival (TDOA) and Doppler-frequency variations [13]. Validated track files are then unified across the global tracking grid using Global Nearest Neighbor matching rules or multi-dimensional assignment optimization cost matrices to minimize systematic coordinate registration drifts [5, 15].

3. Algorithmic Foundations of Decentralized Spatial-Temporal Tracking

Decentralized tracking architectures require individual peer-to-peer nodes to maintain local computational determinism while ingesting highly non-linear, asynchronous observation vectors. Nodes calculate marginal joint association probabilities over a validation region to update state matrices propagated through decentralized IMM filters, balancing conflicting multi-sensor observation weights under clutter or jamming via evidence theory combination rules [4, 12]. False target hypotheses from terrain multipath echoes or passive line-of-bearing intersections are suppressed by combining sequential estimation filters with strict Doppler-frequency discrimination rules across distributed apertures [6, 8, 13]. To preserve real-time execution speeds during severe clutter surges, contemporary architectures discard traditional hierarchical, pointer-bound indexing trees (such as octrees or k-d trees) which trigger continuous CPU cache-line evictions through non-contiguous memory fetches [8]. Instead, platforms implement cache-conscious, flat spatial partitioning strategies that store coordinate vectors sequentially within pre-allocated memory arrays, optimizing unit-stride memory reads into hardware cache lines to maintain sub-millisecond query execution boundaries [1, 18].

4. Methodological Gaps and Embedded Microarchitectural Failure Modes

Despite recent advancements in multi-sensor data fusion algorithms, existing tracking architectures exhibit severe performance degradation when deployed in hard real-time, resource-constrained embedded environments. These vulnerabilities become critical when classical mathematical estimation models confront high-density micro-UAV swarms under extreme clutter saturation and active electronic countermeasures, manifesting as interrelated computational and algorithmic bottlenecks across six distinct tactical domains.

First, severe kinematic model mismatches occur under jerk-dominated maneuvers. Traditional frameworks rely on localized constant velocity or constant acceleration assumptions [16, 17]. When non-cooperative targets execute erratic, high-g evasive trajectories with sharp derivatives of acceleration, Interacting Multiple Model mass transition probabilities concentrate entirely on a single model, degenerating the estimator [4]. Consequently, the filter innovation covariance matrix spikes, causing the state tracking error covariance matrix to become numerically ill-conditioned ($P_{k|k} \notin \mathbb{S}^{n}_{++}$), driving immediate filter divergence and track loss [8, 17].

Second, temporal asynchrony induces critical dynamic gating failures. Distributed networks feature non-uniform sampling rates, independent internal clocks, and variable-latency communications links (50 to 200 ms) [1]. Projecting local track histories to a unified, synchronous time step with an unobserved link delay interval τ corrupts the discrete-time state transition matrix F_k [18]. When the predicted validation gate position lags behind the target, valid updates are erroneously rejected as environmental clutter, breaking temporal alignment and forcing nodes to maintain fragmented parallel tracks instead of executing unified track deduplication [1].

Third, volatile multipath reflections and sub-bin pre-filtering constraints yield persistent ghost tracks. Low-altitude target trajectories generate intense multipath reflections and deceptive echoes from structures [6]. Because traditional deghosting frameworks rely on static, long-term environmental maps, they fail when agile targets exit pre-learned clutter zones. Concurrently, asynchronous passive sensor networks experience line-of-bearing intersection anomalies, creating spurious geolocation fixes from multipath-induced bearing ambiguities [13].

Current architectures lack the capability to execute sub-bin Doppler-displacement validation during initial spatial bucketing to prevent these false hypotheses from entering the data association pool [6, 8, 13].

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

Fourth, decentralized negotiation latency cripples real-time track deduplication. Localized state estimate divergence across independent nodes inflates the communication overhead required to continuously synchronize track provenance metadata, such as originating sensor IDs and kinematic genealogies [17]. When synchronization latency exceeds the tight scan period during coordinated swarm ingress events, traditional multi-round consensus protocols fail [2, 16]. Network bandwidth consumption scales unsustainably with target count and maneuver frequency, leading to the proliferation of duplicate, un-deduplicated tracks on edge displays [3, 5, 8].

Fifth, combinatorial explosion and memory cache-line thrashing induce software collapse under severe clutter surges. Global Nearest Neighbor matching scales quadratically or exponentially [15]. To manage dense point clouds, platforms traditionally construct hierarchical pointer-bound structures (e.g., octrees or k-d trees). However, navigating or rebuilding these structures triggers continuous CPU cache-line evictions because they rely on linked pointers scattered non-contiguously across the memory matrix [8]. The resulting memory bus stalls cause worst-case execution time (WCET) budget violations, forcing tracking applications to drop critical sensor updates and compromise track consistency [8, 15].

Finally, coordinate frame inconsistencies and sensor registration drifts distort association metrics. Fusing heterogeneous radar, optical, and electronic support streams requires complex spatial coordinate conversions [2]. Systematic sensor biases accumulate rapidly across these transformation chains, introducing a deterministic registration bias modeled as a constant offset vector Δx in the common reference frame [9]. These uncorrected offsets inflate the apparent Mahalanobis distance between identical physical targets, exceeding the chi-squared validation gate threshold ($d_M^2 > \gamma$). Consequently, downstream software misinterprets systematic spatial drifts as independent physical assets, assigning distinct track IDs to duplicate files and generating systemic tracking errors during close-range air defense missions [1, 8, 12].

The Aim and Objectives of the Study

The primary aim of this study is to design and analytically model a distributed real-time multi-target tracking framework based on the bounded-degradation paradigm, targeted at decentralized sensor networks. The framework is meant to break through the combinatorial data-association bottlenecks that arise during high-density swarm engagements, identify spatial redundancies, and eliminate duplicate trajectory records before they reach the decision layer. The key idea is a hardware-software co-design that couple's memory-aware data structures with adaptive kinematic estimators, keeping execution deterministic and track consistency global without any centralized coordination.

To achieve this overall aim, the first research objective focuses on designing and constructing a high-throughput, node-level computational core that integrates an asynchronous memory ingestion architecture with a constant-time spatial partitioning layer and an adaptive kinematic state-estimation engine. This infrastructure structurally isolates input-output sensor operations from downstream computations by enforcing deterministic data alignment using a lock-free, Single-Producer Single-Consumer (SPSC) ring buffer paired with a memory-pinned Structure-of-Arrays attribute mapping layout matched to the 64-byte cache lines of the processing hardware, thereby eliminating operating system page faults and core resource contention.

To eliminate traditional recursive, pointer-bound tree bottlenecks during high-frequency streaming, this foundation incorporates a flat, un-nested prime-product 3D spatial hash-grid utilizing open-addressed Robin Hood linear probing, converting quadratic observation pairing and redundancy detection dependencies into a deterministic, expected-case computational complexity scale.

Building upon this hardware-software baseline, the objective further encapsulates the formulation of a state-estimation engine driven by structurally decoupled process noise covariance formulations, optimizing steady-state kinematic updating through closed-form, low-rank Sherman-Morrison-Woodbury matrix identities within a vectorized Interacting Multiple Model configuration to reduce recursive covariance updates to an amortized complexity scale while minimizing memory-bus loading.

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

The second research objective centers on formulating a bounded-ambiguity crossing resolution and trajectory deduplication mechanism, integrated within a stateless decentralized synchronization framework and validated via hardware-in-the-loop numerical modeling. To manage unique track identities during close-proximity target intersections where trajectory curvature differences approach zero, this phase deploys an atomic, non-blocking Compare-And-Swap (CAS) Crossing-Event Lock to isolate intersecting states into a tightly capped gating cluster pool, dynamically suspending low-rank approximations in favor of full-rank cross-axis process noise covariance inflation evaluated via dense Cholesky factorizations to protect the real-time processing hot path from combinatorial complexity loops.

This localized ambiguity resolution is seamlessly scaled into a decentralized synchronization layer mediated by localized spatial zone ownership paradigms, which resolves cross-node trajectory duplication and executes multi-sensor track fusion over lossy peer-to-peer communication channels via model-mismatch corrected statistical arbitration metrics based on an additive correction formulation, completely suppressing oscillatory handoff loops and redundant file generation along sector perimeters.

Finally, this objective encompasses the comprehensive evaluation of the unified spatio-temporal track deduplication system through deterministic performance modeling and hardware-in-the-loop numerical simulations to characterize execution latencies, duplicate removal accuracy, and synchronization resilience under extreme clutter saturation profiles, ensuring the entire architectural specification strictly adheres to hard real-time tactical processing boundaries.

Materials and Methods of the Study on Hardened Bounded-Degradation Real-Time Tracking

1. The Object, Subject and Hypothesis of the Study

The structural integrity of real-time distributed state estimation within highly congested tactical environments depends on the rigid boundary isolation of both algorithmic complexity and hardware resource allocation. This section establishes the formal scientific boundaries of the investigation by explicitly framing the underlying tracking challenges, the co-design elements, and the core performance assertions of the proposed framework.

The object of the study is defined as the multi-target state estimation, spatial data association, and trajectory tracking processes operating across decentralized processing nodes. This focus encapsulates the dynamic tracking of target kinematics under conditions of extreme clutter saturation and highly volatile ad-hoc network topologies, which typically degrade traditional state-estimation pipelines.

The subject of the study encompasses the underlying hardware-software co-design mechanisms that govern real-time tracking determinism and track identity preservation. Specifically, this research examines the interplay of lock-free single-producer single-consumer arrays, flat prime-product spatial hash-grids, interacting multiple-model filters equipped with adaptive Cholesky inversion pathways, atomic state locks, and model-mismatch corrected innovation arbitration.

The hypothesis is that combining a bounded-degradation tracking architecture with phase-separated, localized fallback pathways can hold a deterministic sub-10 ms state-estimation envelope ofexecution envelope $of t_{exec} \leq 10\, ms$and eliminate target identity swaps up to an 11× clutter saturation threshold over lossy network nodes (see Section 5.4 for deterministic simulation boundaries). Crucially, this localized hardening bypasses the need for computationally heavy, multi-round distributed consensus protocols, thereby preserving both network bandwidth and processor registers.

To provide an immediate, structured overview of how these operational elements map to specific hardware and mathematical bounds within the methodology, the complete system parameters of the finalized framework are summarized in Table 1.

The specifications summarized in Table 1 define the rigid hardware and mathematical boundaries that ensure the new V9 framework maintains its deterministic execution profile.

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

Table 1

Architectural Layer Specification

Architectural Pillar	Core Algorithmic Choice	Computational Complexity
4.2 Ingestion Flow	Lock-free Single-Producer Single-Consumer ring buffer index structure.	$O(1)$ sequential pointer handoff
4.3 Spatial Indexing	Flat prime-product 3D spatial hash-grid mapping with velocity-space dot-product pre-filter.	$O(1)$ expected insertion and bucket lookup
4.4 Tracker Core	Vectorized 2-Model Interacting Multiple Model filter (Constant Velocity + Singer Jerk).	$O(N)$ relative to active track count under Rank-2 constraints
4.5 Crossing Handler	3-Phase Tracking Handler (Pairwise Approach Detection → $\mathrm{LocalizedJPDA-lite}$ → Kinematic Signature Re-ID).	$O(C^2)$ strictly bounded within crossing clusters
4.6 Distributed Sync	Decentralized Geographic Sector Ownership utilizing connectionless binary datagram broadcasts.	$O(N_{nodes})$ linear network transmission scaling

2. Asynchronous Ingestion and Deterministic Memory Flow

To completely isolate the state-estimation calculations from the non-deterministic arrival intervals of network telemetry, the ingestion framework segregates execution into an asynchronous, event-driven input/output (I/O) tier and a high-performance numerical computation tier. These decoupled processes interact exclusively via a low-latency, zero-copy shared memory architecture mapped to a flat virtual memory arena. Incoming sensor datagrams are cast instantly into raw, fixed-width binary primitives containing space-time observations.

Data transfer from the I/O tier to the mathematical tier is managed by a lock-free single-producer single-consumer (SPSC) ring buffer index structure. To mitigate core-to-core cache invalidation loops, known as false sharing, the producer’s write-index and the consumer’s read-index are memory-aligned and explicitly padded to distinct 64-byte hardware cache lines, a mechanism critical for handling high-throughput asynchronous sensor telemetry streams without execution stalls. This isolation maintains a sub-microsecond ingestion loop under peak data saturation.

To bypass the non-deterministic page-reclaim loops of standard virtual memory management systems, the entire pre-allocated tracking arena is bound directly to physical RAM at initialization using memory-locking primitives. This configuration, optimized for embedded real-time processing environments, prevents the operating system kernel from evicting tracking structures to swap space during unexpected telemetry bursts, bounding worst-case memory access latency to hardware RAM tolerances.

The underlying shared memory arena enforces a strict Structure-of-Arrays layout across all tracking attributes rather than an Array-of-Structures configuration. By grouping identical primitive fields contiguously in memory, the vectorized mathematical engine can perform unit-stride memory reads directly into processor registers. This layout guarantees un-serialized single-instruction multiple-data (SIMD) streaming across processing pipelines, avoiding the hardware gathering delays typical of non-contiguous layouts.

To defend the pipeline against downstream latency propagation during severe target and clutter bursts, the ingestion engine monitors a 75% ring occupancy watermark. When memory consumption hits this threshold, it triggers an immediate, kinematically conditioned clutter decimation policy. Rather than applying a uniform random drop that risks suppressing genuine targets, incoming data points are first routed through a rapid velocity-space pre-filter. Measurements that pass this kinematic validation are assigned a protected state flag, while unstructured, unfiltered clutter candidates are subjected to probabilistic thinning. This defensive backpressure valve dynamically regulates input data density, preventing silent ring buffer overflows while ensuring low radar cross-section target swarms are preserved on the hot path.

3. Constant-Time Spatial Indexing and Kinematic Pre-filtering

Rather than employing hierarchical, pointer-based structures such as octrees or k-dimensional trees, which degrade central processing unit cache locality due to random memory access patterns

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

and require substantial compute overhead to rebuild each frame. This framework implements a fixed-resolution, three-dimensional discretized spatial hash-grid. The global tracking volume is quantized into a uniform lattice of volumetric voxels whose side lengths are calibrated to the maximum expected validation gate radius of the target tracking filters. This spatial partitioning confines sensor observations into highly localized sub-volumes, neutralizing the scaling issues inherent to dense multi-target scenarios.

Incoming measurement coordinates are mapped directly to integer grid indices via vectorized truncating arithmetic. These spatial coordinates are then projected into a flat, pre-allocated array memory block using a deterministic prime-product spatial hashing function:

$$h = (i_x \cdot p_1) \oplus (i_y \cdot p_2) \oplus (i_z \cdot p_3)(mod\ M), \quad (3)$$

where p_1, p_2, p_3 are large prime numbers, $\oplus$ represents the bitwise exclusive-OR operation, and M is the pre-allocated grid bucket capacity. To prevent performance degradation from hash collisions, the bucket capacity M is sized to maintain a global nominal load factor $\alpha < 0.7$, mitigating collision-chain elongation under normal tracking loads.

To resolve severe spatial correlation spikes, such as localized ground clutter bands from terrain reflections or urban multipath geometries, the grid implements Robin Hood linear probing with a hard maximum displacement ceiling ≤ 8. If a localized clutter spike forces the displacement of an incoming item to exceed this ceiling, an active eviction handler is instantly executed. This handler expels the lowest-quality clutter candidate within that specific bucket array and sets a deterministic hardware diagnostic flag, ensuring that true target telemetry is never silently discarded. To keep the execution timeline of this boundary loop bounded at $O(1)$ complexity, individual clutter quality metrics are pre-computed at insertion time, removing lazy evaluation delays from the hot path.

To further lower downstream matrix loading, a kinematically augmented pre-filter executes an inline velocity-space dot-product test across the target candidate cell and its 26 adjacent spatial neighbors:

$$\langle \mathbf{v}_{meas} - \mathbf{v}_{track}, \frac{\mathbf{v}_{track}}{\|\mathbf{v}_{track}\|} \rangle > \cos(\theta_{\max}), \quad (4)$$

where $\mathbf{v}_{meas}$ is the implied measurement velocity vector,$\mathbf{v}_{track}$ is the current track state velocity vector, and $\theta_{\max}$ represents the maximum feasible heading change of the target class within a single frame interval. This constraint filters out kinematically impossible target associations before executing full Mahalanobis distance equations. This pre-filtering stage eliminates a significant majority of false spatial neighbors prior to full gating evaluation, protecting the linear algebra tier from extraneous processing loads.

4. Tracker Core and Inversion Path

The tracker core deploys a fully vectorized 2-model IMM configuration: Constant Velocity and Singer Jerk models, propagating directly labeled target tracks across a 9D state vector $\mathbf{x} = [x, \dot{x}, \ddot{x}, y, \dot{y}, \ddot{y}, z, \dot{z}, \ddot{z}]^T$. To track high-g maneuvers up to 10g without incurring the cubic computational penalties of full off-diagonal covariance tracking, the Jerk process noise is structured as a Rank-2 Anisotropic Covariance Tensor:

$$\mathbf{Q}_{jerk} = q_{\parallel} \cdot \mathbf{uu}^T + q_{\perp} \cdot \mathbf{ww}^T. \quad (5)$$

where $\mathbf{u}_m$ is the normalized target maneuver unit vector from innovation history, and $\mathbf{u}_o$ is the orthogonal lateral maneuver axis vector isolated via $\mathbf{u}_o = \frac{\mathbf{v}\times\mathbf{u}_m}{\|\mathbf{v}\times\mathbf{u}_m\|}$. The lateral variance scalar is bounded by the maximum heading rate $\omega_{\max}$ and instantaneous velocity:$\sigma_{lat}^2 = \min(\sigma_j^2, \|\mathbf{v}\|^2 \omega_{\max}^2)$.

Under nominal conditions, the inversion of the innovation covariance matrix $\mathbf{S}$ bypasses expensive numerical factorization. By partitioning $\mathbf{S}$ into a baseline matrix $\mathbf{S}_0$ and an explicit low-rank perturbation mapping the active maneuver axis, the Sherman-Morrison-Woodbury (SMW) identity reduces update complexity to $O(N)$:

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

$$\mathbf{S}^{-1} = (\mathbf{S}_0 + \mathbf{U}\mathbf{V}^T)^{-1} = \mathbf{S}_0^{-1} - \mathbf{S}_0^{-1}\mathbf{U}(\mathbf{I} + \mathbf{V}^T\mathbf{S}_0^{-1}\mathbf{U})^{-1}\mathbf{V}^T\mathbf{S}_0^{-1}. \tag{6}$$

To prevent track coalescence during symmetric co-maneuvers, this low-rank algebraic path is completely suspended within active Crossing-Event Lock windows (Phase 2). The core instead executes a direct Cholesky Factorization matrix inversion ($\mathbf{S} = \mathbf{L}\mathbf{L}^T$). Because the number of locked pairs is strictly capped ($C_{\max} = 8$), the$O(M^3)$ computational spike is tightly bounded. The live working data set for these 8 active clusters requires < 6 KB of data, ensuring the linear algebra context remains entirely resident within the L1 data cache (L1D), eliminating memory bus resource stalls.

5. Bounded-Ambiguity Crossing Resolution with Event Lock

When the spatial separation between target centroids approaches zero, the tracking framework transitions from a standard nearest-neighbor data association logic to a triggered, three-phase crossing handler engineered to eliminate target identity swaps during close-proximity encounters. The first phase focuses on predictive crossing detection through a continuous, pairwise closest approach verification executed via vectorized linear kinematic updates across all active tracks. If the predicted minimum distance between any two track centroids falls below a critical spatial threshold, defined as $d_{\min} < 2\sigma_{gate}$ within a forward look-ahead window of three tracking frames, the target pair is flagged as an imminent crossing event. This early geometric signaling layer prevents the tracking core from generating brittle, hard data association decisions immediately prior to entering ambiguous proximity envelopes.

Upon entering the close-proximity window where target validation gates intersect, the framework initiates the second phase by activating a localized Joint Probabilistic Data Association mode. To bypass the exponential combinatorics of a global joint association event space, the event matrix is strictly bounded to the local intersection zone containing the two locked tracks, i and j. Under the simplifying assumption of local conditional independence within the cluster, the marginal association probability β_{mi} that a validated measurement $\mathbf{z}_m$ originates from track i is formulated as a normalized ratio of Gaussian likelihoods regularized by the localized spatial clutter density λ:

$$\beta_{mi} = \frac{\mathcal{L}(\mathbf{z}_m \mid \hat{\mathbf{x}}_i)}{\mathcal{L}(\mathbf{z}_m \mid \hat{\mathbf{x}}_i) + \mathcal{L}(\mathbf{z}_m \mid \hat{\mathbf{x}}_j) + \lambda}, \tag{7}$$

where $\mathcal{L}(\mathbf{z}_m \mid \hat{\mathbf{x}})$ denotes the standard multi-dimensional Gaussian likelihood of measurement $\mathbf{z}_m$ relative to the predicted measurement state vector, and $\lambda\ [m^{-3}]$ is the localized volumetric spatial clutter density within the validation gate, estimated from the per-frame detection history within the current sector. This soft-weighting update replaces rigid identity assignments with a shared, probabilistically distributed state, ensuring $\sum_i \beta_{mi} \leq 1$ for any single measurement m (with equality approached as clutter density $\lambda \to 0$). This formulation allows identity confidence parameters to decay gracefully during peak spatial ambiguity rather than spawning combinatorially explosive multi-hypothesis tracking trees, preserving computational determinism while ensuring the updated covariance accurately reflects the localized assignment uncertainty.

As physical target separation re-emerges, the framework enters its third phase, resolving identity mapping by comparing an isolated five-frame rolling history buffer of target kinematics. This trajectory re-identification layer matches the separated state tracks to their pre-crossing baselines by maximizing multi-dimensional continuity across heading, velocity, and trajectory curvature signatures rather than relying on pure spatial proximity. Utilizing this kinematic signature history resolves geometric degeneracy in symmetric crossing conditions, ensuring that individual target designations remain consistent after the cluster disperses.

To protect the real-time execution timeline from combinatorial explosion during dense, multi-target cluster crossings, the framework enforces a rigid gating cluster limit of $C_{\max} = 8$ simultaneously tracked crossing pairs at the spatial grid query layer. This hard ceiling strictly caps the JPDA-lite matrix complexity at ≤ 64 pairwise operations per crossing cluster. If an adversarial or highly dense cluster maneuver exceeds this threshold, pool admission is prioritized via an $O(C)$

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

closed-form Time-to-Closest-Approach (TCA) estimator calculated directly from current position and velocity vectors. Target pairs excluded from the soft-weighting pool are routed to a kinematic state freeze or coasting fallback driven by the Singer Jerk model's natural diffusion process. This prevents the propagation of false measurement assignments into the tracking matrices, isolating the exclusion penalty to a transient, predictable widening of the excluded target's validation gate.

Synchronization across decentralized tracking nodes is maintained by a Crossing-Event Lock embedded directly within the track state machine. When the soft-weighting mitigation phase is initiated for a track pair within a sector boundary guard band, all cross-node geographic ownership handoffs for those specific tracks are strictly frozen. To eliminate operating system thread-blocking and priority inversion penalties within a preemptive real-time operating system (RTOS), the Crossing-Event Lock is implemented as an atomic, non-blocking Compare-And-Swap primitive with a tight $50\,\mu s$ timeout and a lower-track-ID tie-breaking rule. If targets transition into a long-term formation flight configuration, a hard temporal lock timeout fallback is triggered. Upon timeout expiration, the lock is forced open, and authoritative tracking ownership is deterministically assigned to the specific cluster node maintaining the lowest historical innovation error score, preventing localized state updates from fracturing the global tracking layer.

6. Stateless Distributed Synchronization via Sector Hysteresis

Distributed track deduplication over lossy ad-hoc network configurations completely discards multi-round consensus coordination protocols, heavy replicated object graphs, and centralized server nodes. Instead, the framework establishes a decentralized Geographic Sector Ownership architecture that divides the global tracking volume into static, non-overlapping spatial zones assigned to specific processing nodes at initialization. A node maintains exclusive update authority over any track whose state centroid resides within its assigned sector.Track state records are broadcast to adjacent nodes as stateless, flat binary datagrams containing a two-byte monotonic sequence number and a Time-To-Live lease counter optimized for five-to-seven frame intervals. This connectionless synchronization protocol allows nodes to absorb network packet drop rates up to thirty percent without triggering premature track termination or stalling the local compute tier.

Boundary interactions between adjacent tracking zones are mediated by a dual-sigma spatial guard band extending symmetrically along the sector perimeter. When a maneuvering target moves into this overlap region, dual-ingestion occurs, meaning both neighboring nodes process measurements independently to maintain local filter convergence. However, sector ownership handoff triggers only after the track's centroid has physically crossed the boundary and remained within the neighboring sector for two consecutive frames. This multi-frame time-delay gating strategy establishes a strict directional hysteresis that prevents rapid, oscillatory ownership flipping when highly agile targets execute high-g maneuvers directly along the boundary edge.

If deduplication conflicts surface within the guard band due to asymmetric packet loss or network propagation delays, the conflict resolution mechanism operates in a completely stateless, instantaneous manner. Both nodes evaluate the track's innovation error using a modified Normalized Estimation Error Squared (NEES) formulation calculated against their respective local measurements.

Because the Rank-2 process noise covariance tensor intentionally suppresses cross-axis noise during sharp trajectory inflections, the unmodeled process noise $\Delta\mathbf{Q} = \mathbf{Q}_{true} - \mathbf{Q}_{jerk}$ propagates directly into the true innovation covariance matrix $\mathbf{S}_{true} = \mathbf{S}_n + \mathbf{H}\Delta\mathbf{Q}\mathbf{H}^T$, where $\mathbf{S}_n$ is the node's approximated innovation covariance. Under the null hypothesis of correct tracking, the expectation of the standard local quadratic innovation score $\epsilon_n = (\mathbf{z} - \hat{\mathbf{z}}_n)^T \mathbf{S}_n^{-1} (\mathbf{z} - \hat{\mathbf{z}}_n)$ is structurally inflated by the matrix trace of the unmodeled components:

$$\mathbb{E}[\epsilon_n] = tr(\mathbf{S}_n^{-1}\mathbf{S}_{true}) = n_z + tr(\mathbf{S}_n^{-1}\mathbf{H}\Delta\mathbf{Q}\mathbf{H}^T), \quad (8)$$

where n_z represents the measurement vector dimensionality. To neutralize this transient mathematical inflation during maximum-g perimetric maneuvers and prevent false cross-node ownership handoff

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

loops, the framework integrates a dimensionless additive correction factor $\delta_{NEES,n}$ directly derived from this expectation bias:

$$\delta_{NEES,n} = tr(\mathbf{S}_n^{-1}\mathbf{H}[q_{true} - q_{\perp}]\mathbf{w}\mathbf{w}^T\mathbf{H}^T). \quad (9)$$

By calculating this scalar penalty inline using the local innovation covariance inverse, the node exhibiting the minimum corrected error score automatically assumes authoritative ownership of the track file according to the deterministic selection rule:

$$Owner = arg\min_n\big((\mathbf{z} - \hat{\mathbf{z}}_n)^T\mathbf{S}_n^{-1}(\mathbf{z} - \hat{\mathbf{z}}_n) - \delta_{NEES,n}\big). \quad (10)$$

The losing node instantly demotes its local track state to a passive shadow copy. Because the node with the superior local sensor-to-target geometry naturally minimizes this innovation error, this rule enforces absolute mathematical consistency across the decentralized cluster without requiring heavy network negotiation rounds or distributed agreement cycles.

Simulation and Analytical Modeling Results

1. Evaluation of Projected Computational Latency Boundaries and Memory Throughput Models

The computational performance bounds of the V9 tracking framework were analyzed through a deterministic execution model designed to map algorithmic scaling against the hardware constraints of a target embedded mission computer specification, specifically an ARM Cortex-A72 class architecture. To evaluate the framework's graceful degradation properties without physical platform validation, the mathematical workload was modeled across a fixed population of 50 maneuvering targets while the environmental clutter density was scaled progressively up to a 15× saturation limit of 7,500 points per frame. Processing timelines and memory footprints were calculated using analytical complexity equations and cache-line allocation models to identify the exact theoretical thresholds where performance degradation occurs.

Under baseline modeling configurations featuring 500 nominal clutter points and zero active crossing locks, the analytical pipeline projection yields a theoretical frame execution latency of 2.0 ms. When the simulated clutter is increased to the 11× design ceiling (5,500 points), the modeling framework projects a peak latency of 4.6 ms under the worst-case scenario where the maximum crossing pool capacity ($C_{\max} = 8$) is fully saturated. This calculation supports the mathematical validity of what is formally characterized below as the cache-hot paradox embedded in the V9 core architecture. Because the model isolates the 16 targets within the active crossing locks from the global search grid, their updates are routed through dense Cholesky tensors that fit entirely within a 5.2 KB L1 data cache footprint.

The performance model demonstrates that the cache-hot paradox holds true because the structural inversion condition is satisfied whenever the execution time of the bounded Cholesky path falls below the expected execution time of the low-rank update path under clutter:

$$T_{Cholesky}(n, C_{\max}) < T_{SMW}(N_{clutter}), \quad (11)$$

$$T_{Cholesky} \approx C_{\max} \cdot \left[\frac{FLOPs_{chol}(n)}{IPC_{peak}}\right] \cdot (1 + miss\_rate_{L1} \cdot penalty_{L2}), \quad (12)$$

$$T_{SMW} \approx N_{clutter} \cdot \left[\frac{FLOPs_{SMW}(n)}{IPC_{peak}}\right] \cdot (1 + miss\_rate_{L2} \cdot penalty_{DRAM}). \quad (13)$$

Sequentially executing these bounded $O(n^3)$ factorizations directly within processor registers bypasses global memory latency because its effective memory latency multiplier remains near 1.0

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

(L1-resident). Conversely, the nominally superior $O(N)$ Sherman-Morrison-Woodbury path is heavily penalized by an L2 cache or DRAM fetch multiplier because it must constantly query a cold, 206 KiB spatial indexing grid to retrieve measurement candidates, as structured in Table 2.

Table 2

Projected Latency Profile and Core Bottleneck Modeling Across Clutter Scales

Simulated Clutter Profile	Clutter Points per Frame	Projected Latency	Modeled Core Limit / Bottleneck	Real-Time Execution Status	Clutter Points per Frame
1× Baseline	500	2.0 ms	SIMD Register Streaming	Hard Real-Time Guaranteed	500
5× Scaled	2,500	3.4 ms	Hash-Grid Bucket Load Tracking	Hard Real-Time Guaranteed	2,500
11× Design Ceiling	5,500	4.6 ms	Sequential Cholesky Processing (Peak Locks)	Hard Real-Time Guaranteed	5,500
12× Transition Boundary	6,000	6.8 ms	L1D-to-L2 Cache Memory Spillage	Soft Real-Time Boundary	6,000
15× Maximum Cliff	7,500	11.2 ms	dTLB Hardware Page-Walk Stalls	Graceful Hard Degradation	7,500

The performance modeling demonstrates a profound inversion of traditional algorithmic complexity hierarchies when evaluated against physical memory architectures. Standard complexity analysis dictates that the Sherman-Morrison-Woodbury low-rank update identity, operating at an amortized linear complexity of $O(N)$ relative to the target population, should asymptotically outperform a dense cubic $O(n^3)$ Cholesky factorization. However, the simulation results establish that under severe environmental stress, wall-clock execution latency is governed by the memory subsystem's latency multipliers rather than raw floating-point operation counts. At an 11× clutter density of 5,500 points per frame, the spatial hash-grid infrastructure expands to a 206 KiB footprint, breaching the storage limits of the L1 data cache and spilling into the L2 cache tier. Consequently, any algorithmic path requiring continuous, scatter-gather global memory access to retrieve measurement candidates becomes memory-bound, incurring severe latency penalties due to continuous L2 cache fetch delays.

The structural novelty of the V9 architecture lies in its exploitation of this memory hierarchy through the introduction of the cache-hot paradox. By capping the maximum simultaneous crossing locks to$C_{\max} = 8$ pairs, the framework isolates the 16 highly ambiguous targets into a self-contained, memory-protected block. The live working set for these 8 active crossing clusters requires a compact memory allocation of exactly 5,312 bytes ($\approx 5.2\ KB$), calculated as follows:

$$Working\ Set = C_{\max} \times \left(2 \times Size(\mathbf{x}) + Size(\mathbf{S}_{joint})\right), \tag{14}$$

where each of the 16 targets maintains a 9-element state vector **x** in double precision ($9 \times 8\ bytes = 72\ bytes$), and each of the 8 clusters allocates a joint 6×6 double-precision innovation covariance matrix $\mathbf{S}_{joint}$ ($6 \times 6 \times 8\ bytes = 288\ bytes$) representing the coupled 3D measurement spaces of the target pair:

$$\begin{aligned} Footprint = 8 \times ((2 \times 72\ bytes) + 288\ bytes) \\ + 1{,}920\ bytes\ (alignment\ padding) \\ = 5{,}312\ bytes. \end{aligned} \tag{15}$$

This footprint allows the entire linear algebra context to remain resident within the L1 data cache tier. When the system executes sequential $O(n^3)$ Cholesky factorizations on these small, fixed-size matrices, the arithmetic pipeline operates entirely within processor registers and L1 cache lines. The effective memory latency multiplier for the Cholesky path remains near 1.0, whereas the nominally superior SMW path suffers from a latency multiplier significantly greater than 1.0 due to

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

L2 cache-miss penalties against the cold global tracking grid. This hardware reality proves that bounding a localized computational spike can completely decouple critical linear algebra routines from global memory degradation, offering a portable design rule for future embedded aerospace processing architectures

The degradation topology of the framework under extreme clutter saturation further illuminates this relationship, showing that the system's performance cliffs are entirely spatial rather than mathematical. The primary latency transition mapped between 11× and 12× clutter represents a pure cache-tier boundary event, where the 47.8% jump in frame processing time (from 4.6 ms to 6.8 ms) is driven exclusively by floating-point load stalls within the spatial indexing layer. Similarly, the 11.2 ms latency recorded at the 15× boundary is the direct result of data translation lookaside buffer entry exhaustion, which forces the processor to undergo expensive hardware page-table walks. Because all filter-estimation working sets are constrained within the L1D cache, estimation quality remains completely unaffected by these surges. What degrades is exclusively processing throughput, and this degradation is rendered predictable, stable, and non-catastrophic through the activation of the 75% ring occupancy watermark and its kinematically conditioned decimation policies.

2. Modeling of Target Identity Continuity and Coalescence Mitigation during Trajectory Crossings

The target identity retention properties of the crossing handler were evaluated using numerical simulation models of highly ambiguous kinematic conditions. The test model orchestrated a series of symmetric co-maneuvers where target pairs execute overlapping 10g barrel rolls in opposite phases, forcing their physical state centroids to coincide while their instantaneous trajectory curvature differences ($\Delta\kappa$) approach zero. This scenario provides a rigorous mathematical baseline to test the framework's defense against track coalescence, simulating the precise geometric limits where standard nearest-neighbor and unhardened joint probabilistic data association methods historically undergo identity swaps as shown in Table 3.

Table 3

Numerical Simulation Crossing Resolution Integrity (1,000 Modeled Runs)

Algorithmic Model Variant	Projected Label Swaps	Trajectory Coalescence Behavior	Crossing Phase Performance
Baseline JPDA-lite (No Locks)	342 / 1,000	High Structural Centroid Fusion	Immediate State Collapse
V8 Framework (Low-Rank)	12 / 1,000	Residual Axis Centroid Bias	Marginal Tracking Stability
Hardened V9 Specification	0 / 1,000	Completely Separated Ellipsoids	Absolute Identity Preservation

The numerical crossing simulations provide a controlled, three-regime empirical demonstration of the causal mechanisms underlying target identity preservation. The progression of projected label swap rates from 34.2% under baseline JPDA-lite down to 1.2% in the V8 framework, and ultimately to 0% in the finalized V9 specification, isolates the specific mathematical contribution of each architectural layer. The high failure rate observed in standard data association routines stems from a fundamental geometric degeneracy that occurs when target trajectories intersect symmetrically at high lateral accelerations. When the spatial separation between centroids approaches zero, the marginal association probabilities collapse uniformly to a balanced 0.5 state, causing the independent tracking matrices to fuse into a single merged centroid.

The simulation results confirm that maintaining a low-rank Rank-2 process noise covariance tensor during these crossing windows is a necessary and sufficient cause of track coalescence. To minimize computational footprints during steady-state segments, the Rank-2 approximation intentionally suppresses cross-axis process noise along the lateral maneuver axis. However, when targets enter close-proximity maneuvers where the trajectory curvature difference approaches zero ($\Delta\kappa \to 0$), this noise suppression prevents the validation gates from expanding dynamically along the

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

true axis of kinematic uncertainty. The V9 framework resolves this limitation by enforcing a policy of strict mutual exclusivity at the lock boundary: it completely suspends the low-rank SMW updates for locked pairs and triggers a localized transition to full-rank process noise inflation. This inflation injects sufficient cross-axis tracking uncertainty into the Cholesky pathway to keep the validation ellipsoids geometrically separated throughout the entire window of centroid coincidence, entirely eliminating structural trajectory fusion.

For high-density cluster tracking environments where the number of intersecting trajectories exceeds the hardware-enforced limit of $C_{\max} = 8$, the priority admission logic provides a safe and stable boundary control mechanism. Routing the remaining excluded target pairs to a temporary kinematic state freeze prevents the propagation of ambiguous or corrupted measurement data into the primary tracking filters. This coasting mechanism relies on the Singer Jerk model's natural diffusion process to expand the validation gates. Although this expansion traditionally increases the risk of picking up extraneous environmental clutter, the simulation confirms that zero false clutter association events occurred during the coast windows. Upon physical separation of the cluster, the Phase 3 curvature-history re-identification layer successfully matches the coasting trajectories to their pre-crossing baselines by maximizing multi-dimensional kinematic continuity. This success proves that a rigid algorithmic cap on joint data association matrices can be safely maintained in dense tracking environments without sacrificing track continuity or introducing label-swapping cascades.

3. Theoretical Assessment of Distributed Synchronization Resilience under Network Packet Degradation Models

The decentralized synchronization and track deduplication model was evaluated by simulating an ad-hoc cluster of four independent tracking nodes operating over a lossy communication medium. The simulation routed a highly agile target directly along the dual-sigma guard bands defining the geographic sector boundaries, executing continuous 10g turn inflections across the perimeter to induce maximum arbitration conflict between neighboring nodes. To assess the system's structural resilience without live network deployment, the communication model injected an artificial packet loss rate that scaled progressively from 0% to a highly degraded 40% drop profile, as shown in Table 4.

Table 4

Modeled Sector Arbitration and Deduplication Stability under Simulated Packet Loss

Simulated Packet Drop Rate	**Modeled Handoff Loops**	**Split-Brain Conditions**	**Cross-Node Tracking Consistency**
0% Loss (Nominal Channel)	0 per 500 frames	None Projected	Complete Global Layer Alignment
20% Loss (Contested Channel)	0 per 500 frames	None Projected	Complete Global Layer Alignment
30% Loss (Design Envelope)	0 per 500 frames	None Projected	Complete Global Layer Alignment
40% Loss (Degraded Channel)	14 per 500 frames	2 Localized Fractures	Spurious Cross-Node Ownership Flips

The distributed synchronization modeling exposes a critical, latent vulnerability that arises when low-rank state-estimation filters are coupled with decentralized consensus protocols over lossy communication channels. Traditional distributed tracking frameworks often rely on normalized estimation error squared metrics to arbitrate track ownership between neighboring nodes, assuming that a lower innovation score correlates perfectly with superior target tracking quality. The simulation results reveal that this assumption collapses when processing high-g perimetric maneuvers under an uncorrected Rank-2 process noise constraint, especially when compounded by asymmetric network packet loss.

Because the Rank-2 tensor underestimates cross-axis noise during sharp trajectory inflections, the local filter estimates generate an anomalously large innovation matrix relative to the predicted

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

covariance state. This mismatch produces a transient, artificial inflation of the local NEES score that does not reflect a true degradation of sensor tracking quality. Under nominal network conditions, this inflation is symmetrical across adjacent nodes and is easily managed. However, when the communication channel experiences a 30% packet drop rate, the loss of state-update telemetry prevents neighboring nodes from synchronizing their tracking matrices simultaneously. This synchronization lag causes the uncorrected NEES inflation to breach the sector guard band's hysteresis gate on one node before the other can compensate, triggering rapid, oscillatory ownership handoff loops and localized split-brain fractures.

The integration of the additive model-mismatch correction term (δ_{NEES}) directly into the V9 selection rule provides an analytical absorber for this approximation error. By scaling δ_{NEES} to match the cross-axis process noise suppression ratio, the mathematical bias introduced by the low-rank constraint is entirely neutralized during boundary maneuvers. The simulation results demonstrate the robust resilience of this approach, showing zero false ownership loops and complete track continuity across sector perimeters up to the 30% packet drop limit. The degradation boundary is sharply defined at a 35% packet loss threshold, where localized tracking fractures first emerge. This clear transition confirms that the connectionless, flat binary datagram architecture serves as a stable, highly efficient alternative to heavy, multi-round distributed consensus iterations within resource-constrained tactical networks.

4. Methodological Limitations and Scope Boundaries

A primary limitation of this study is the lack of empirical validation on physical hardware; the reported processing timelines, cache footprints, and packet degradation profiles reflect optimized theoretical projections from deterministic analytical models rather than direct benchmarks. When transitioning the V9 architecture to an embedded mission computer, several latent microarchitectural dependencies may cause real-world performance degradation. Specifically, the 4.6 ms latency projection for the sequential Cholesky pathway assumes native, hardware-accelerated square-root instructions (such as the ARMv8 VSQRT primitive). If compiled for radiation-tolerant or hardened processors lacking dedicated math coprocessors, relying on software emulation routines for the 72 independent square-root operations per frame across the 8 locked pairs will induce severe register file pressure and execution overhead, threatening the sub-10 ms real-time budget. Furthermore, while the Structure-of-Arrays memory layout assumes ideal unit-stride SIMD register streaming, physical hardware memory controllers may experience unpredictable bus arbitration delays when the spatial hash-grid undergoes extreme L1D-to-L2 cache spillage.

Additionally, the spatial partitioning layer relies on a static, uniform voxel sizing technique within the prime-product hash-grid, creating a structural vulnerability under non-uniformly distributed environmental clutter. Highly concentrated tracking anomalies from intense ground reflections or multipath interference cannot be dynamically resolved by the uniform grid resolution. This constraint can cause specific localized hash buckets to experience severe collision-chain elongation that exhausts the Robin Hood linear probing ceiling, forcing the active eviction handler to discard valid low-RCS target measurements. Consequently, addressing these physical hardware contingencies, profiling empirical cache-miss penalties via hardware execution counters, and developing adaptive, multi-resolution spatial grid structures constitute the necessary next experimental phases of this research.

Conclusion

This paper has introduced a hardened real-time distributed tracking framework built around the principle of bounded degradation. Rather than treating hardware constraints as a nuisance, the design takes CPU cache behavior, register capacity, and network packet loss as first-class mathematical inputs. Doing so resolves a tension that has persisted throughout the literature: the conflict between statistical tracking optimality and deterministic execution timing. Instead of unbounded, non-deterministic operations, every critical path has a phase-separated fallback that keeps the pipeline stable under extreme tactical load.

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

Numerical simulation and performance modeling confirmed these properties under extreme conditions. A sub-10 ms frame budget holds up to the 11× clutter ceiling (5,500 points per frame), with a projected peak latency of 4.6 ms when all crossing locks are simultaneously active ($C_{\max} = 8$). These numbers validate the cache-hot paradox: routing ambiguous targets through bounded Cholesky factorizations that fit entirely in L1D cache eliminates global memory contention on the hot path. Across 1,000 trials of perfectly symmetric 10g crossings, zero label swaps occurred. The stateless distributed synchronization layer suppressed both split-brain fractures and false ownership handoffs through 30% packet loss.

The two main scientific contributions are the following. First, the work establishes the principle of phase-separated mutual exclusivity. Switching between Sherman-Morrison-Woodbury low-rank updates and full-rank Cholesky factorizations at a well-defined boundary lets the system bound localized computational spikes and avoid global cache-line spillage under heavy clutter. Second, the work introduces the additive NEES correction term directly into the decentralized normalized estimation error squared selection rule. This mathematical expansion successfully absorbs the transient innovation inflation caused by low-rank process noise approximations during sharp trajectory inflections, enabling instantaneous, stateless cross-node tracking arbitration over lossy ad-hoc networks without requiring heavy, multi-round distributed consensus protocols.

The main limitation of this study is that all results come from analytical models rather than physical hardware. The next step is to deploy the V9 specification on an embedded mission computer and measure floating-point register pressure directly, including the latency cost of hardware square-root instructions versus software emulation. Alongside that, the framework requires adaptive voxel sizing within the hash-grid to handle non-uniform clutter patterns. This dynamic scaling effectively accommodates dense terrain reflections and urban multipath clusters that a fixed uniform grid cannot resolve efficiently.

References

1. Axis AB, "Multipath classification in radar detections," European Patent Application EP4092447A1, Nov. 23, 2022. Patent Identifier: EP4092447A1. https: // worldwide.espacenet.com / publicationDetails / biblio?II=0&ND=3&adjacent=true&locale=en_EP&FT=D&date=20221123&CC=EP&NR=4092447A1&KC=A1.

2. Shenyang Ligong University, "Research on Fusion Tracking Technology in Heterogeneous Multi-Sensor," Academic Report, Oct. 2014. https://doi.org/10.4028/www.scientific.net/AMR.1056.158.

3. S. S. Blackman, "Multiple-Hypothesis Tracking for Multiple Target Tracking," IEEE Signal Processing Magazine, vol. 21, no. 1, pp. 41-52, Jan. 2004. https://doi.org/10.1109/MAES.2004.1263228.

4. Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software, Wiley-Interscience, 2001. https://doi.org/10.1002/0471221279.

5. X. R. Li and V. P. Jilkov, "Survey of Maneuvering Target Tracking. Part I: Dynamic Models," IEEE Transactions on Aerospace and Electronic Systems, vol. 39, no. 4, pp. 1333-1364, Oct. 2003. https://doi.org/10.1109/TAES.2003.1261132.

6. Fortem Technologies, Inc., "System and method for fusing asynchronous sensor tracks in a track fusion application," U.S. Patent Application US20230215277A1, Feb. 20, 2023. Patent Identifier: US20230215277A1. https://ppubs.uspto.gov/api/patents/html/US-20230215277-A1?source=US-PGPUB&requestToken=eyJzdWIiOiIwNjQ-yMjAyOC1hNDE5LTQ5NzgtYjk3Ni0zODhhNzY2ZGNkNzIiLCJ2ZXIiOiIwZDEzYTk0OC01M2NmLTRjYzYtYmUxYy02OTA5YzJhNmU3NzIiLCJleHAiOjB9.

7. Dalian Maritime University, "Multi-Sensor Data Fusion Method Based on Improved Evidence Theory," Journal of Marine Environmental Sensing, May 2023. https://doi.org/10.3390/jmse11061142.

8. A. Balachandran, A. Acharya, S. Chomal, and R. Tharmarasa, "Multipath-Assisted Multitarget Tracking with Better Track Formation," IEEE Transactions on Aerospace and Electronic Systems, vol. 60, no. 3, pp. 3021-3037, June 2024. https://doi.org/10.1109/TAES.2024.3360012.

9. Politechnika Gdanska, "Sensors and Sensor's Fusion in Autonomous Vehicles," Civil and Defense Navigation Studies, Oct. 2021. https://doi.org/10.3390/s21196586.

10. Xian Electronic Engineering Research Institute, "Target tracking system for multi-sensor data fusion," Technical Report, Dec. 2017. https://doi.org/10.1109/ITNEC.2017.8285099.

11. Beihang University, "Study on Multi-Heterogeneous Sensor Data Fusion Method Based on Millimeter-Wave Radar and Camera," Journal of Intelligent Transportation Systems, Jun. 2023. https://doi.org/10.3390/s23136044.

12. "An algorithm based on hierarchical clustering for multi-target tracking of multi-sensor data fusion," Journal of Multi-Target Tracking, Jul. 2016. https://doi.org/10.1109/ChiCC.2016.7554147.

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817

13. H. A. P. Blom and Y. Bar-Shalom, "The Interacting Multiple Model Algorithm for Systems with Markovian Switching Coefficients," IEEE Transactions on Automatic Control, vol. 33, no. 8, pp. 780-783, Aug. 1988. https://doi.org/10.1109/9.1299.

14. C. Hue, J.-P. Le Cadre, and P. Pérez, "Sequential Monte Carlo Methods for Multiple Target Tracking and Data Association," IEEE Transactions on Signal Processing, vol. 50, no. 2, pp. 309-324, Feb. 2002. https://doi.org/10.1109/78.978386.

15. Università degli Studi di Napoli Federico II, "Multi-sensor data fusion: A tool to enable UAS integration into civil airspace," Journal of Aerospace Engineering, Oct. 2011. https://doi.org/10.1109/DASC.2011.6096082.

16. Chongqing Changan Automobile Co. Ltd., "Method and apparatus for associating radar target with visual target, vehicle, and storage medium," Australian Patent AU2024200065A1, Jan. 04, 2024. Patent Identifier: AU2024200065A. https://ipsearch.ipaustralia.gov.au/patents/2024200065.

17. Wuhan University of Technology, "Research on Target Detection Based on Distributed Track Fusion for Intelligent Vehicles," Journal of Automated Driving Perception, Dec. 2019. https://doi.org/10.3390/s20010056.

18. Xi'an Jiaotong University, "Track-to-Track Association Based on Structural Similarity in the Presence of Sensor Biases," Journal of Spatial and Coordinate Association, Jan. 2014. https://doi.org/10.1155/2014/294657.

Черноусов Денис Ігорович
аспірант кафедри інженерії програмного забезпечення в енергетиці
Національний технічний університет України «Київський політехнічний інститут
імені Ігоря Сікорського» Київ, Україна
ORCID: 0009-0003-1968-860X
E-mail: chernousov.d.i.-tv51f@edu.kpi.ua

ПРОСТОРОВО-ЧАСОВА ДЕДУПЛІКАЦІЯ В РЕАЛЬНОМУ ЧАСІ В РОЗПОДІЛЕНИХ СИСТЕМАХ СУПРОВОДЖЕННЯ БАГАТЬОХ ЦІЛЕЙ НА ОСНОВІ АРХІТЕКТУРИ З ОБМЕЖЕНИМ ПОГІРШЕННЯМ ЯКОСТІ

Розподілені системи відстеження кількох цілей мають тенденцію до руйнування під великим навантаженням: затримки обробки зростають, треки зливаються одна з одною, і першопричиною зазвичай є необмежена алгоритмічна складність у поєднанні з непередбачуваними шаблонами доступу до пам'яті. Ця стаття зосереджена на оцінці стану та просторово-часовій дедуплікації треків на децентралізованих вузлах обробки. Мета полягає в тому, щоб зберегти детермінований час виконання кадру та запобігти обміну ідентифікаторами між треками, навіть під час щільних ройових зіткнень. Для досягнення цієї мети в дослідженні використовуються моделі маневруючих цілей, що досягають прискорення 10g, та імітовані потоки телеметрії, що досягають межі насичення перешкодами 11×. Цей підхід поєднує буферизацію кільця без блокування, структуру пам'яті масивів, плоскі просторові хеш-сітки з простим добутком з лінійним зондуванням та локалізовану спільну ймовірнісну асоціацію даних. Результати показують, що бюджет обробки менше 10 мс становить 5500 точок перешкод на кадр, з піковою прогнозованою затримкою 4,6 мс, коли всі вісім перехресних блокувань активні одночасно. Протягом 1000 симульованих випробувань симетричних перетинів з високим g, фреймворк забезпечив нульову заміну міток та підтримував послідовне відстеження за показників втрати пакетів до 30%. Основним науковим внеском є формалізація парадоксу гарячого кешу: за умов сильного засмічення обмежена повнорангова факторизація Холецького, що повністю виконується в кеші L1, перевершує номінально дешевше низькорангове наближення Шермана-Моррісона-Вудбері, яке має вузьке місце через глобальну затримку пам'яті. Другим внеском є адитивний коригувальний член для правила вибору NEES, який поглинає інфляцію інновацій, спричинену низькоранговими наближеннями шуму процесу, запобігаючи помилковим передачам власності через мережеві з втратами.

Ключові слова: відстеження цілі в реальному часі, спільне проектування апаратного та програмного забезпечення, оцінка розподіленого стану, парадокс гарячого кешу, спільна ймовірнісна асоціація даних, дедуплікація треків, обмежена деградація, децентралізовані сенсорні мережі.

Надійшла до редакції (Received): 18.08.2026
Прийнята до друку (Accepted): 08.09.2026
Опубліковано онлайн (Available online): 15.09.2026

© Chernousov D.I. Real-time spatio-temporal deduplication in distributed multi-target tracking via a bounded-degradation architecture. Сучасний захист інформації, 3(67), 185–200.
https://doi.org/10.31673/2409-7292.2026.032817