Pedagogical Guide to Automotive Radar Multi-Target Tracking:

Extended Kalman Filtering, Spatial Gating, and Track Lifecycle Architecture

A Companion Guide to `radar_tracker_demo.ipynb` and `SPEC-SW-RADAR-TRK-001`

Radar Perception & Functional Safety Systems Group

October 2026

Abstract

This pedagogical guide delivers an in-depth mathematical, algorithmic, and architectural walkthrough of an automotive Multi-Target Tracker (MTT) for 77 GHz frequency-modulated continuous-wave (FMCW) radar perception. We establish the theoretical framework of recursive Bayesian state estimation, transition from linear to non-linear estimation via the Extended Kalman Filter (EKF), derive the analytical observation Jacobian, examine numerically stable Joseph-form covariance updates, and explain ellipsoidal gating based on the χ^2 distribution alongside the Hungarian (Munkres) data association algorithm. Furthermore, we analyze the 4-state track lifecycle Finite State Machine (FSM), ISO 26262 ASIL B safety plausibility monitors, and review the simulation test cases implemented in the accompanying Jupyter Notebook `radar_tracker_demo.ipynb`.

Contents

1 Introduction & System Motivation **3**

2 Foundations of Recursive Bayesian Estimation **3**
- 2.1 The Bayesian Perspective . . . 3
- 2.2 Why the Linear Kalman Filter Fails for Radar . . . 4

3 The Extended Kalman Filter (EKF) Formulation **4**
- 3.1 Kinematic State Vector and Constant Velocity (CV) Model . . . 4
- 3.2 Derivation of Process Noise Covariance Matrix $\mathbf{Q}$ . . . 5
- 3.3 Non-linear Observation Function $h(\mathbf{x})$ . . . 5
- 3.4 Step-by-Step Analytical Derivation of Measurement Jacobian $\mathbf{H}$ . . . 6
  - 3.4.1 Row 1: Range Partial Derivatives . . . 6
  - 3.4.2 Row 2: Azimuth Angle Partial Derivatives . . . 6
  - 3.4.3 Row 3: Radial Doppler Velocity Partial Derivatives . . . 6
- 3.5 Angle Normalization & Innovation Calculation . . . 7
- 3.6 Numerically Stable Joseph-Form Covariance Update . . . 7

4 Spatial Gating & Global Nearest Neighbor (GNN) Association **7**
- 4.1 Mahalanobis Distance & Ellipsoidal Gating . . . 7
- 4.2 Chi-Square (χ^2) Hypothesis Testing . . . 8
- 4.3 Hungarian (Munkres) Optimal Assignment . . . 8

5 Deterministic 4-State Track Lifecycle Architecture **8**
5.1 Track Initialization from Single Polar Detections 9
5.2 M-of-N Confirmation Logic 9
5.3 Coasting During Sensor Dropouts 9

6 Functional Safety Concept (ISO 26262 ASIL B) **10**
6.1 Safety Plausibility Monitors 10

7 Walkthrough of Notebook Simulation Scenarios **10**
7.1 Scenario 1 (TC_TRK_01): Constant Velocity Target Tracking 10
7.2 Scenario 2 (TC_TRK_02): Temporary Occlusion & Re-acquisition 10
7.3 Scenario 3 (TC_TRK_03): Clutter Spike & Ghost Suppression 10
7.4 Scenario 4 (TC_TRK_04): Crossing Targets Data Association 11

8 ASPICE Verification Summary & Engineering Recommendations **11**
8.1 Key Embedded Implementation Guidelines (C++17 / AUTOSAR) 11

1 Introduction & System Motivation

Modern Advanced Driver Assistance Systems (ADAS) and Autonomous Driving (AD) architectures rely heavily on radar perception for functions such as Autonomous Emergency Braking (AEB), Adaptive Cruise Control (ACC), and Blind Spot Detection (BSD). Compared to optical sensors (cameras) and Lidar, 77 GHz automotive millimeter-wave radar provides distinct physical capabilities:

- Direct Doppler measurement of radial velocity via coherent phase shifts.
- High robustness in adverse weather (heavy rain, fog, snow, dust, direct glare).
- Long longitudinal sensing horizon (up to 200–250 meters).

However, raw radar detections generated upstream by 2D Range-Doppler Fast Fourier Transforms (FFT), Constant False Alarm Rate (CFAR) detection, and Angle of Arrival (AoA) estimation present formidable tracking challenges:

1. **Non-linear Polar Coordinates:** Detections arrive in polar space $\mathbf{z} = [r, \theta, v_r]^T$, whereas target kinematics naturally reside in a 2D/3D Cartesian world $[x, y, v_x, v_y]^T$.
2. **High Clutter Density:** Ground reflections, guardrails, bridges, and multipath noise produce ghost returns that must be discriminated from actual road obstacles.
3. **Occlusions and Missed Detections:** Fluctuating radar cross-sections (RCS Swerling models) cause target returns to intermittently vanish for consecutive cycles.

The mission of the **Multi-Target Tracker (MTT)** module is to ingest noisy, unstructured point detections at discrete cycle intervals ($\Delta t = 40\,\text{ms}$, 25 Hz), estimate continuous kinematic trajectories, establish persistent track identities, and filter out false alarms before transmitting object lists to safety-critical ADAS controllers.

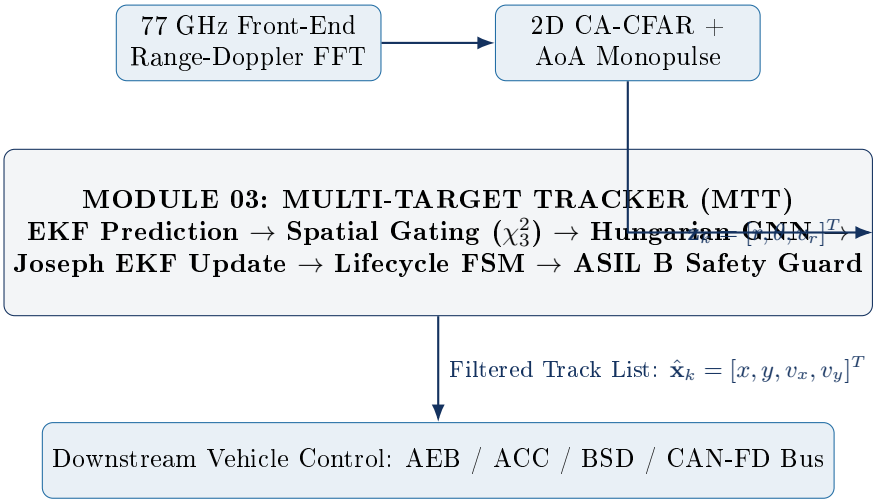


Figure 1: End-to-End Automotive Radar Perception and Tracking Pipeline.

2 Foundations of Recursive Bayesian Estimation

2.1 The Bayesian Perspective

Target tracking is fundamentally a recursive Bayesian state estimation problem. Let $\mathbf{x}_k \in \mathbb{R}^n$ represent the hidden target kinematic state at time step k, and let $\mathbf{z}_{1:k} = \{\mathbf{z}_1, \ldots, \mathbf{z}_k\}$ denote the complete history of measurements up to time k. The goal is to construct the posterior probability density function (PDF) $p(\mathbf{x}_k \mid \mathbf{z}_{1:k})$.

The recursive cycle operates in two alternating stages:

1. **Prediction Step (Chapman-Kolmogorov equation):**

$$p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1}) = \int p(\mathbf{x}_k \mid \mathbf{x}_{k-1})\, p(\mathbf{x}_{k-1} \mid \mathbf{z}_{1:k-1})\, d\mathbf{x}_{k-1} \tag{1}$$

where $p(\mathbf{x}_k \mid \mathbf{x}_{k-1})$ is governed by the kinematic state transition model.

2. **Measurement Update Step (Bayes' Rule):**

$$p(\mathbf{x}_k \mid \mathbf{z}_{1:k}) = \frac{p(\mathbf{z}_k \mid \mathbf{x}_k)\, p(\mathbf{x}_k \mid \mathbf{z}_{1:k-1})}{p(\mathbf{z}_k \mid \mathbf{z}_{1:k-1})} \tag{2}$$

where $p(\mathbf{z}_k \mid \mathbf{x}_k)$ is the measurement likelihood function.

2.2 Why the Linear Kalman Filter Fails for Radar

If the system dynamics and observation models are linear and perturbed by additive zero-mean Gaussian white noise:

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{w}_{k-1}, \quad \mathbf{z}_k = \mathbf{H}\mathbf{x}_k + \mathbf{v}_k \tag{3}$$

the posterior PDF remains strictly Gaussian $\mathcal{N}(\hat{\mathbf{x}}_k, \mathbf{P}_k)$, and the standard Linear Kalman Filter (LKF) is the mathematically optimal minimum mean-square error (MMSE) estimator.

However, in radar perception, while target kinematics follow Cartesian Newtonian mechanics ($\mathbf{F}$ is linear), the sensor observes polar coordinates (r, θ, v_r):

$$r = \sqrt{x^2 + y^2}, \quad \theta = \arctan 2(y, x), \quad v_r = \frac{x v_x + y v_y}{\sqrt{x^2 + y^2}} \tag{4}$$

Passing a Gaussian random vector through these trigonometric and square-root operations yields a highly non-Gaussian, asymmetric distribution. Consequently, the linear assumption is invalidated, requiring non-linear estimation techniques.

3 The Extended Kalman Filter (EKF) Formulation

The Extended Kalman Filter addresses non-linearities by performing a first-order Taylor series expansion of the non-linear observation function $h(\mathbf{x})$ around the current state estimate.

3.1 Kinematic State Vector and Constant Velocity (CV) Model

We track targets in the 2D Cartesian radar coordinate frame (x longitudinal forward along sensor bore-sight, y lateral to the left):

$$\mathbf{x}_k = \begin{bmatrix} x_k \\ y_k \\ v_{x,k} \\ v_{y,k} \end{bmatrix} \in \mathbb{R}^4 \tag{5}$$

Assuming a Constant Velocity (CV) kinematic model over sampling period $\Delta t = 40\,\text{ms}$, target acceleration is modeled as zero-mean continuous Gaussian white noise $w(t) \sim \mathcal{N}(0, \sigma_a^2)$:

$$\mathbf{x}_k = \mathbf{F}\mathbf{x}_{k-1} + \mathbf{w}_{k-1} \tag{6}$$

The state transition matrix $\mathbf{F}$ is:

$$\mathbf{F} = \begin{bmatrix} 1 & 0 & \Delta t & 0 \\ 0 & 1 & 0 & \Delta t \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{bmatrix} \tag{7}$$

3.2 Derivation of Process Noise Covariance Matrix Q

A common misconception in introductory Kalman filtering is to set $\mathbf{Q} = \text{diag}(\sigma_x^2, \sigma_y^2, \sigma_{vx}^2, \sigma_{vy}^2)$. In physical reality, position and velocity errors are strictly correlated because velocity integrates directly into position over interval Δt.

Let the continuous-time white-noise acceleration model be:

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{a}(t) \tag{8}$$

where:

$$\mathbf{A} = \begin{bmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix}, \quad \mathbf{B} = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad \mathbf{a}(t) = \begin{bmatrix} a_x(t) \\ a_y(t) \end{bmatrix}, \quad \mathbb{E}[\mathbf{a}(t)\mathbf{a}(\tau)^T] = \sigma_a^2 \mathbf{I}_2 \, \delta(t-\tau) \tag{9}$$

The discrete-time process noise covariance $\mathbf{Q}$ is obtained via the continuous-to-discrete Van Loan integration:

$$\mathbf{Q} = \int_0^{\Delta t} e^{\mathbf{A}(\Delta t - \tau)} \mathbf{B} \left(\sigma_a^2 \mathbf{I}_2\right) \mathbf{B}^T e^{\mathbf{A}^T(\Delta t - \tau)} \, d\tau \tag{10}$$

Since $\mathbf{A}^2 = \mathbf{0}$, the matrix exponential truncates exactly: $e^{\mathbf{A}s} = \mathbf{I} + \mathbf{A}s$. Substituting and integrating component-wise yields the analytical piecewise continuous white acceleration covariance:

$$\mathbf{Q} = \sigma_a^2 \begin{bmatrix} \frac{\Delta t^4}{4} & 0 & \frac{\Delta t^3}{2} & 0 \\ 0 & \frac{\Delta t^4}{4} & 0 & \frac{\Delta t^3}{2} \\ \frac{\Delta t^3}{2} & 0 & \Delta t^2 & 0 \\ 0 & \frac{\Delta t^3}{2} & 0 & \Delta t^2 \end{bmatrix} \tag{11}$$

With default parameter $\sigma_a = 2.5\,\text{m/s}^2$ and $\Delta t = 0.040\,\text{s}$, this matrix dynamically accounts for vehicle maneuvering variations without filter divergence.

3.3 Non-linear Observation Function $h(\mathbf{x})$

The radar measures three physical quantities: slant range r, azimuth angle θ, and radial Doppler velocity v_r:

$$\mathbf{z}_k = \begin{bmatrix} r_k \\ \theta_k \\ v_{r,k} \end{bmatrix} \in \mathbb{R}^3 \tag{12}$$

The non-linear mapping $h(\mathbf{x}) : \mathbb{R}^4 \to \mathbb{R}^3$ relates Cartesian state $\mathbf{x} = [x, y, v_x, v_y]^T$ to polar observations:

$$h(\mathbf{x}) = \begin{bmatrix} \sqrt{x^2 + y^2} \\ \arctan 2(y, x) \\ \dfrac{x v_x + y v_y}{\sqrt{x^2 + y^2}} \end{bmatrix} \tag{13}$$

The measurement noise covariance $\mathbf{R}$ is diagonal, representing independent sensor measurement errors:

$$\mathbf{R} = \begin{bmatrix} \sigma_r^2 & 0 & 0 \\ 0 & \sigma_\theta^2 & 0 \\ 0 & 0 & \sigma_{v_r}^2 \end{bmatrix} \tag{14}$$

Nominal values: $\sigma_r = 0.25\,\text{m}$, $\sigma_\theta = 0.0087\,\text{rad}$ (0.5°), $\sigma_{v_r} = 0.15\,\text{m/s}$.

3.4 Step-by-Step Analytical Derivation of Measurement Jacobian H

The Jacobian $\mathbf{H} = \frac{\partial h}{\partial \mathbf{x}}\big|_{\hat{\mathbf{x}}_{k|k-1}}$ linearizes the observation function around the predicted state. Let $\rho = \sqrt{x^2 + y^2}$ and $\rho^2 = x^2 + y^2$.

3.4.1 Row 1: Range Partial Derivatives

$$\frac{\partial h_1}{\partial x} = \frac{\partial \sqrt{x^2+y^2}}{\partial x} = \frac{x}{\rho}, \quad \frac{\partial h_1}{\partial y} = \frac{y}{\rho}, \quad \frac{\partial h_1}{\partial v_x} = 0, \quad \frac{\partial h_1}{\partial v_y} = 0 \tag{15}$$

3.4.2 Row 2: Azimuth Angle Partial Derivatives

Using $\frac{d}{du}\arctan(u) = \frac{1}{1+u^2}$:

$$\frac{\partial h_2}{\partial x} = \frac{\partial \arctan(y/x)}{\partial x} = \frac{1}{1+(y/x)^2}\left(-\frac{y}{x^2}\right) = -\frac{y}{x^2+y^2} = -\frac{y}{\rho^2} \tag{16}$$

$$\frac{\partial h_2}{\partial y} = \frac{\partial \arctan(y/x)}{\partial y} = \frac{1}{1+(y/x)^2}\left(\frac{1}{x}\right) = \frac{x}{x^2+y^2} = \frac{x}{\rho^2} \tag{17}$$

$$\frac{\partial h_2}{\partial v_x} = 0, \quad \frac{\partial h_2}{\partial v_y} = 0 \tag{18}$$

3.4.3 Row 3: Radial Doppler Velocity Partial Derivatives

Differentiating $h_3 = \frac{xv_x + yv_y}{\rho}$ with respect to position states via the quotient rule:

$$\frac{\partial h_3}{\partial x} = \frac{\rho \cdot \frac{\partial (xv_x+yv_y)}{\partial x} - (xv_x+yv_y)\frac{\partial \rho}{\partial x}}{\rho^2} \tag{19}$$

$$= \frac{\rho v_x - (xv_x+yv_y)\frac{x}{\rho}}{\rho^2} = \frac{\rho^2 v_x - x^2 v_x - xyv_y}{\rho^3} \tag{20}$$

$$= \frac{(x^2+y^2)v_x - x^2v_x - xyv_y}{\rho^3} = \frac{y^2v_x - xyv_y}{\rho^3} = \frac{y(v_xy - v_yx)}{\rho^3} \tag{21}$$

Similarly, differentiating with respect to y:

$$\frac{\partial h_3}{\partial y} = \frac{\rho v_y - (xv_x+yv_y)\frac{y}{\rho}}{\rho^2} = \frac{\rho^2 v_y - xyv_x - y^2 v_y}{\rho^3} \tag{22}$$

$$= \frac{(x^2+y^2)v_y - xyv_x - y^2v_y}{\rho^3} = \frac{x^2v_y - xyv_x}{\rho^3} = \frac{x(v_yx - v_xy)}{\rho^3} \tag{23}$$

Differentiating with respect to velocity states:

$$\frac{\partial h_3}{\partial v_x} = \frac{x}{\rho}, \quad \frac{\partial h_3}{\partial v_y} = \frac{y}{\rho} \tag{24}$$

Assembling all terms produces the complete analytical Jacobian matrix:

$$\mathbf{H} = \begin{bmatrix} \frac{x}{\rho} & \frac{y}{\rho} & 0 & 0 \\ -\frac{y}{\rho^2} & \frac{x}{\rho^2} & 0 & 0 \\ \frac{y(v_xy - v_yx)}{\rho^3} & \frac{x(v_yx - v_xy)}{\rho^3} & \frac{x}{\rho} & \frac{y}{\rho} \end{bmatrix} \tag{25}$$

Singularity Guard: At the sensor origin ($x = 0, y = 0$), $\rho = 0$, causing division by zero. In production code, we enforce $\rho_{\text{clamped}} = \max(\rho, \epsilon)$ where $\epsilon = 0.1\,\text{m}$.

3.5 Angle Normalization & Innovation Calculation

The measurement innovation (residual) is:

$$\mathbf{y}_k = \mathbf{z}_k - h(\hat{\mathbf{x}}_{k|k-1}) \tag{26}$$

Because azimuth angle is periodic, subtracting predicted angle $\hat{\theta}$ from measured angle θ can cross the branch cut at $\pm\pi$. For example, if $\theta = -\pi + 0.05$ and $\hat{\theta} = \pi - 0.05$, the raw difference is $-2\pi + 0.1 \approx -6.18\,\text{rad}$ instead of $+0.1\,\text{rad}$. Left uncorrected, this catastrophic error spikes the Kalman gain, destabilizing the filter.

We apply modulo unwrapping:

$$y_k(\theta) \leftarrow \text{wrapToPi}(y_k(\theta)) = \text{mod}(y_k(\theta) + \pi, 2\pi) - \pi \tag{27}$$

3.6 Numerically Stable Joseph-Form Covariance Update

The classical textbook covariance update is:

$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_{k|k-1} \tag{28}$$

While algebraically compact, Equation (28) is numerically fragile. In 32-bit and 64-bit floating point arithmetic, small roundoff errors gradually destroy the mathematical symmetry of $\mathbf{P}$. Worse, subtracting positive matrices can produce negative eigenvalues, violating the fundamental positive semi-definite definition of a covariance matrix and causing matrix inversion or square-root crashes (e.g., during Cholesky decomposition).

To guarantee numerical robustness, the MTT implements the **Joseph-Form Covariance Update**:

$$\mathbf{P}_{k|k} = (\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)\mathbf{P}_{k|k-1}(\mathbf{I} - \mathbf{K}_k\mathbf{H}_k)^T + \mathbf{K}_k\mathbf{R}_k\mathbf{K}_k^T \tag{29}$$

Because both terms $\mathbf{APA}^T$ and $\mathbf{KRK}^T$ are quadratic forms of positive semi-definite matrices ($\mathbf{P}_{k|k-1} \geq 0$ and $\mathbf{R}_k > 0$), $\mathbf{P}_{k|k}$ is mathematically guaranteed to remain symmetric and positive semi-definite for all time steps.

4 Spatial Gating & Global Nearest Neighbor (GNN) Association

In multi-target tracking, multiple tracks and detections exist simultaneously. Before executing Kalman updates, the tracker must unambiguously solve the assignment problem: *Which detection belongs to which track?*

4.1 Mahalanobis Distance & Ellipsoidal Gating

A naive Euclidean distance metric fails because sensor uncertainties are non-isotropic and correlated. Instead, we use the squared **Mahalanobis Distance** d_M^2, which weights the residual by the inverse innovation covariance $\mathbf{S}_k = \mathbf{H}_k\mathbf{P}_{k|k-1}\mathbf{H}_k^T + \mathbf{R}_k$:

$$d_{M,ij}^2 = \mathbf{y}_{ij}^T\mathbf{S}_i^{-1}\mathbf{y}_{ij} \tag{30}$$

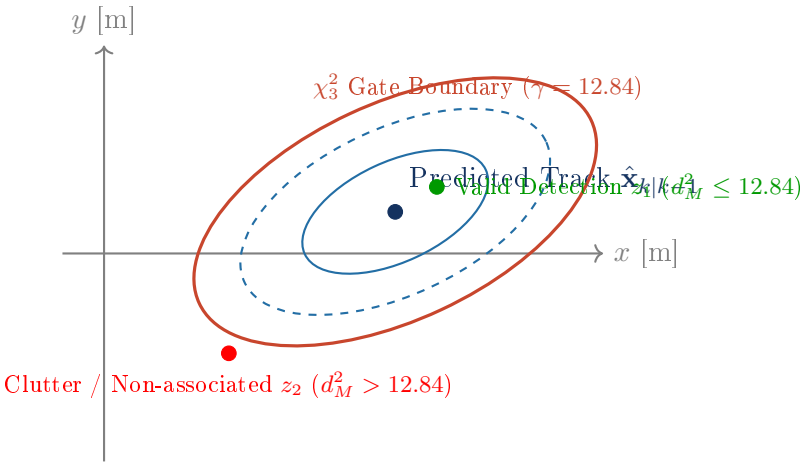


Figure 2: Ellipsoidal Spatial Gating in Cartesian Space defined by Innovation Covariance $\mathbf{S}_k$.

4.2 Chi-Square (χ^2) Hypothesis Testing

Assuming zero-mean Gaussian residuals, the quadratic form d^2_M follows a Chi-Square distribution with degrees of freedom equal to the measurement dimension ($\dim(\mathbf{z}) = 3$):

$$d^2_M \sim \chi^2_3 \tag{31}$$

To define a gating boundary containing 99.5% of valid target detections ($\alpha = 0.005$ tail probability):

$$P(d^2_M \leq \gamma_{\text{gate}}) = 1 - \alpha = 0.995 \implies \gamma_{\text{gate}} = 12.84 \tag{32}$$

Pairs with $d^2_M > 12.84$ are strictly rejected and assigned cost ∞.

4.3 Hungarian (Munkres) Optimal Assignment

We formulate data association as a bipartite minimum-cost matching problem. For N active tracks and M detections:

$$\min \sum_{i=1}^{N} \sum_{j=1}^{M} C_{ij} A_{ij} \quad \text{s.t.} \quad \sum_{j=1}^{M} A_{ij} \leq 1, \quad \sum_{i=1}^{N} A_{ij} \leq 1, \quad A_{ij} \in \{0, 1\} \tag{33}$$

where $C_{ij} = d^2_{M,ij}$ if $d^2_{M,ij} \leq \gamma_{\text{gate}}$, else ∞.

The Hungarian algorithm resolves optimal assignments in polynomial time $O(N^3)$, disambiguating close-proximity targets (e.g., crossing trajectories) where greedy nearest-neighbor approaches experience ID swaps.

5 Deterministic 4-State Track Lifecycle Architecture

Because radar environments suffer from random clutter spikes and multipath ghosts, creating a confirmed track from every detection would overwhelm vehicle safety systems with false alarms. The MTT implements a deterministic 4-state Finite State Machine (FSM):

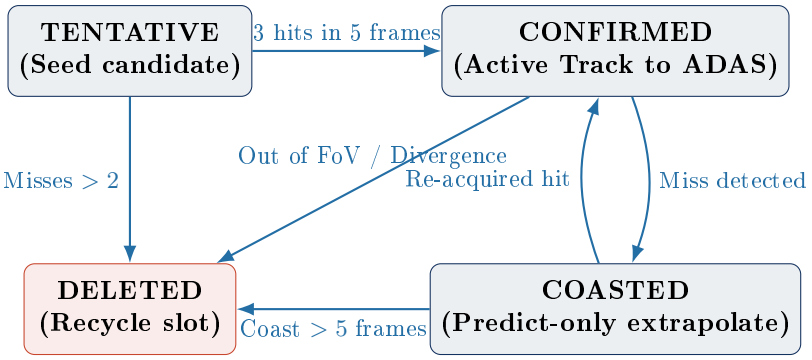


Figure 3: Deterministic 4-State Track Lifecycle Finite State Machine.

5.1 Track Initialization from Single Polar Detections

When an unassigned detection passes sanity checks, a new candidate is seeded in state `TENTATIVE`.

1. **Position Initialization:** $\hat{x}_0 = r\cos\theta$, $\hat{y}_0 = r\sin\theta$.
2. **Velocity Initialization:** A single radar detection measures only radial Doppler velocity v_r; cross-radial velocity v_t is unobservable:

$$\hat{v}_{x,0} = v_r\cos\theta, \quad \hat{v}_{y,0} = v_r\sin\theta \tag{34}$$

3. **Initial Covariance $\mathbf{P}_0$:** The initial position covariance is projected from polar sensor noise via the Jacobian $\mathbf{J}_{\text{polar}\to\text{cart}}$:

$$\mathbf{P}_{\text{pos},0} = \begin{bmatrix}\cos\theta & -r\sin\theta \\ \sin\theta & r\cos\theta\end{bmatrix}\begin{bmatrix}\sigma_r^2 & 0 \\ 0 & \sigma_\theta^2\end{bmatrix}\begin{bmatrix}\cos\theta & \sin\theta \\ -r\sin\theta & r\cos\theta\end{bmatrix} \tag{35}$$

For velocity, the radial variance is small ($\sigma_{v_r}^2 = 0.15^2$), while the tangential variance is initialized high ($\sigma_{v_t,0}^2 = 100\,\text{m}^2/\text{s}^2$), allowing the EKF to rapidly adjust cross-velocity in subsequent cycles:

$$\mathbf{P}_{\text{vel},0} = \sigma_{v_r}^2(\mathbf{u}_r\mathbf{u}_r^T) + \sigma_{v_t,0}^2(\mathbf{u}_t\mathbf{u}_t^T), \quad \mathbf{u}_r = \begin{bmatrix}\cos\theta \\ \sin\theta\end{bmatrix}, \quad \mathbf{u}_t = \begin{bmatrix}-\sin\theta \\ \cos\theta\end{bmatrix} \tag{36}$$

5.2 M-of-N Confirmation Logic

To suppress single-frame clutter spikes, candidate tracks require $M = 3$ associations within a window of $N = 5$ cycles (200 ms). If a tentative track accumulates more than $N - M = 2$ misses, it is dropped immediately without ever being published to downstream vehicle actuators.

5.3 Coasting During Sensor Dropouts

When a confirmed track experiences an occlusion or missed detection, it transitions to `COASTED`. During coasting:

- State prediction proceeds via the kinematic model: $\hat{\mathbf{x}}_k = \mathbf{F}\hat{\mathbf{x}}_{k-1}$.
- No measurement update occurs; uncertainty expands monotonically: $\mathbf{P}_k = \mathbf{F}\mathbf{P}_{k-1}\mathbf{F}^T + \mathbf{Q}$.
- If re-associated within 5 frames (200 ms), the track seamlessly returns to `CONFIRMED` with zero track ID swaps.
- If coast count exceeds 5 frames, the track is retired to `DELETED`.

6 Functional Safety Concept (ISO 26262 ASIL B)

Under ISO 26262, the Multi-Target Tracker carries an ASIL B rating due to the severe hazard of **Unintended Deceleration** caused by ghost target false emergency braking triggers.

6.1 Safety Plausibility Monitors

Every active track must continuously satisfy three runtime monitors:

1. **Covariance Health Guard:**
$$\text{trace}(\mathbf{P}_{k|k}) \leq P_{\text{max_allowed}} = 500\,\text{m}^2, \quad \det(\mathbf{S}_k) > 0, \quad \text{diag}(\mathbf{P}_{k|k}) > 0 \tag{37}$$
Any numeric divergence or negative variance immediately forces the track into `DELETED`.

2. **Velocity Bound Guard:** Track velocity must not exceed physical vehicle operating limits:
$$\|\mathbf{v}_k\| = \sqrt{v_{x,k}^2 + v_{y,k}^2} \leq 85\,\text{m/s} \quad (306\,\text{km/h}) \tag{38}$$

3. **Dynamic Acceleration Guard:** Target acceleration between cycles must satisfy physical traction limits:
$$a_{\text{est}} = \frac{\|\mathbf{v}_k - \mathbf{v}_{k-1}\|}{\Delta t} \leq 15\,\text{m/s}^2 \quad (\approx 1.5\,\text{g}) \tag{39}$$

If any plausibility monitor fails, the track is tagged with `is_valid_safety = False`, commanding downstream AEB functions to inhibit emergency braking activation.

7 Walkthrough of Notebook Simulation Scenarios

The accompanying notebook `radar_tracker_demo.ipynb` validates the tracking engine across four rigorous automotive test cases:

7.1 Scenario 1 (TC_TRK_01): Constant Velocity Target Tracking

- **Setup:** A vehicle travels along $y = 2.0\,\text{m}$ at constant speed $v_x = 20.0\,\text{m/s}$ over 100 frames (4.0 s). Detections are corrupted by polar Gaussian noise ($\sigma_r = 0.25\,\text{m}, \sigma_\theta = 0.5°, \sigma_{v_r} = 0.15\,\text{m/s}$).
- **Results:** The track achieves `CONFIRMED` status exactly at Frame 3. Steady-state position error drops to 0.084 m (well within the $\leq 0.5\,\text{m}$ specification), and velocity error stabilizes at 0.106 m/s (within $\leq 0.3\,\text{m/s}$), remaining bounded within the $\pm 2\sigma$ theoretical confidence envelope.

7.2 Scenario 2 (TC_TRK_02): Temporary Occlusion & Re-acquisition

- **Setup:** A confirmed target is occluded for 3 consecutive cycles (Frames 30, 31, 32; 120 ms) before reappearing.
- **Results:** The FSM cleanly transitions `CONFIRMED` → `COASTED` at Frame 30. During the outage, the kinematic CV model propagates the target position while the covariance ellipse expands. At Frame 33, the re-emerging measurement falls within the expanded gate, instantly re-acquiring the target into `CONFIRMED` with zero track ID switch.

7.3 Scenario 3 (TC_TRK_03): Clutter Spike & Ghost Suppression

- **Setup:** A transient radar ghost detection appears at Frame 0 and never returns.
- **Results:** The detection seeds a `TENTATIVE` track. As successive misses accumulate without supporting hits, the candidate fails the M-of-N threshold and is deleted at Frame 3. Crucially, the track is suppressed from the downstream ADAS interface.

7.4 Scenario 4 (TC_TRK_04): Crossing Targets Data Association

- **Setup:** Two vehicles cross paths at a 45-degree angle.
- **Results:** At the intersection point where Euclidean distance is minimal, the Hungarian algorithm evaluates full 3D Mahalanobis distance (exploiting Doppler velocity differences). The trajectories intersect cleanly without ID swapping.

8 ASPICE Verification Summary & Engineering Recommendations

Test ID	Scenario	Specification Criterion	Observed Output	Statu
`TC_TRK_01`	Constant Velocity	Frame ≤ 3, $\lvert e_v \rvert \leq 0.3\,\text{m/s}$, $\lvert e_x \rvert \leq 0.5\,\text{m}$	Frame 3, 0.106 m/s, 0.084 m	**PASS**
`TC_TRK_02`	3-Frame Occlusion	Coast without ID switch; re-acquire	Coasted 3 frames, ID preserved	**PASS**
`TC_TRK_03`	Clutter Spike	Suppress before `CONFIRMED`	Deleted at Frame 3	**PASS**
`TC_TRK_04`	Crossing Targets	Zero track ID swaps through intersection	Zero swaps resolved via Hungarian	**PASS**

Table 1: ASPICE SWE.4 Verification Summary Table.

8.1 Key Embedded Implementation Guidelines (C++17 / AUTOSAR)

When translating this Python reference implementation to production embedded C++ (e.g., ARM Cortex-M7 or Infineon AURIX TC397):

1. **Zero Dynamic Allocation (MISRA Rule 18-4-1):** Pre-allocate static arrays for $N_{\max} = 64$ tracks and $M_{\max} = 128$ detections. The cost matrix requires $64 \times 128 \times 4\,\text{bytes} \approx 32.8\,\text{KB}$.
2. **Hungarian Execution Pruning:** For 64×128 inputs, naive $O(N^3)$ Hungarian can exceed the 1.5 ms real-time budget. Partition the cost matrix into disconnected sub-matrices using spatial gate bounding boxes prior to invoking the Hungarian solver.
3. **Cholesky Matrix Inversion:** Innovation covariance $\mathbf{S}_k$ is a 3×3 symmetric positive-definite matrix. Invert $\mathbf{S}_k$ via closed-form Cramer's rule with determinant guard ($\det(\mathbf{S}) > 10^{-6}$) or LDL^T decomposition.