

Time-Based Target Tracking from TOA and TDOA Measurements:

A Unified Model, Consistent Initialisation, and Efficiency of the Extended Kalman Filter

Ömer Kahyaoglu¹

Department of Mathematical Engineering, İstanbul Teknik Üniversitesi

Abstract. We study the tracking of a manoeuvring emitter from the arrival times of its signal at a network of sensors, in time-of-arrival (TOA) and time-difference-of-arrival (TDOA) form. Both modes are cast as one linear observation of the per-sensor timestamps, $z = \mathbf{D}\mathbf{r}(\mathbf{p})/c + \mathbf{D}\mathbf{w}$. This yields the correlated TDOA covariance directly, and it gives an exact characterisation of the information TDOA forgoes: one rank-one term along the resultant of the line-of-sight vectors. We pair the model with a least-squares initialiser that resolves the ambiguous three-sensor and collinear geometries explicitly, an extended Kalman filter (EKF) over four exactly discretised motion models, and two extensions: NIS-gated adaptive process noise and map-aware measurement noise. We prove that the steady-state error of an n -fold-integrator tracker scales as $\sigma_t^{(2n-1)/(2n)}$ in the timing noise. In a controlled Monte Carlo study with run-level bootstrap intervals and common random numbers, the EKF is consistent and attains the posterior Cramér–Rao bound (RMSE/bound 0.967–1.004). The measured exponents match theory: 0.749 for position against $3/4$, and 0.239 for velocity against $1/4$. TDOA ties TOA inside a surrounding array and loses up to $2.06\times$ outside it. Adaptive process noise reduces the error on an abruptly manoeuvring target 1.6–2.6-fold. A filter that models the propagation environment remains consistent where a blind one is over-confident by two orders of magnitude. All results are reproducible from the accompanying open software.

Index terms. Target tracking, time of arrival, time difference of arrival, extended Kalman filter, Cramér–Rao bound, filter consistency, geometric dilution of precision, manoeuvring targets.

1 Introduction

Locating and following a radio, acoustic or optical emitter from the times at which its signal reaches a set of receivers is a classical problem of statistical signal processing. It underlies satellite navigation [1], cellular and ultra-wideband positioning [2–4], wireless sensor networks [5] and passive radar [6]. With propagation speed $c = 299\,792\,458$ m/s [7], each nanosecond of timing error is 0.3 m of range. The question is how such errors, the geometry of the receivers, the motion of the target and the propagation environment combine into tracking accuracy, and whether a practical recursive estimator realises the accuracy the data allow.

1.1 Related work

Localisation from arrival times. Static positioning from TOA and TDOA has a long history. Foy’s Taylor-series method [8] and Torrieri’s statistical analysis [9] established iterative least squares and geometric dilution of precision (GDOP) as the basic tools. Closed-form range-difference solutions followed [10–12]. Chan and Ho’s estimator [11] attains the Cramér–Rao bound at moderate noise and remains the standard initialiser. Fundamental limits are treated in [13, 14]. Surveys cover the wireless setting [2–5] and the non-line-of-sight (NLOS) problem [15].

Tracking. Recursive estimation goes back to Kalman [16]. The extended Kalman filter, its numerically robust Joseph-form update [17, 18] and the target motion models in common use are covered in [19, 20]

¹Department of Mathematical Engineering, Faculty of Science and Letters, İstanbul Teknik Üniversitesi, 34469 Maslak, İstanbul, Türkiye. E-mail: kahyaoglu20@itu.edu.tr. Term project of MTH 407 (CRN 23577), May 2026; revised October 2026.

and surveyed in [21]. These include white-noise acceleration and jerk, the Singer model [22] and the coordinated turn. Manoeuvres are usually handled by multiple-model filters [23] or by adapting the noise statistics [19, 24]. For stronger non-linearities, unscented and particle filters are used [25, 26].

Performance assessment. Filter consistency is assessed with the normalised estimation error and innovation statistics (NEES, NIS) and their Monte Carlo chi-square bands [19, Sec. 5.4]. Efficiency is assessed against the posterior Cramér–Rao bound (PCRLB) [26, 27]. Steady-state accuracy of fixed-gain trackers is classically characterised by the tracking index [28].

Gap. Each ingredient above is well established. They are seldom examined together under one statistically controlled protocol, however, and a number of statements repeated in applied work are imprecise for a *tracker*. Two examples are that TDOA performs like TOA “because the extra TOA measurement only informs the emission time”, and that tracking error grows linearly with the timing noise. This study revisits the problem with that aim: to derive what can be derived, and to test the rest with quantified uncertainty.

1.2 Research questions

- RQ1** Is the EKF statistically consistent in this problem, and how close does it come to the posterior Cramér–Rao bound?
- RQ2** When, and by how much, does TDOA lose accuracy relative to TOA?
- RQ3** How does the tracking error scale with timing noise and with the number of sensors, and can the scaling be predicted?
- RQ4** How strongly does sensor geometry influence accuracy, relative to algorithmic choices?
- RQ5** How sensitive is performance to the motion model, and can abrupt manoeuvres be handled without a multiple-model filter?
- RQ6** What does knowledge of the propagation environment contribute, in accuracy and in consistency?

1.3 Contributions

- (i) A unified observation model in which TOA and TDOA differ by one matrix $\mathbf{D}$, and an exact expression for the Fisher information TDOA forgoes, with its consequences (Proposition 1 and Corollary 1).
- (ii) A scaling law for the steady-state error of n -fold-integrator trackers, $\sigma_t^{(2n-1)/(2n)}$ for position and $\sigma_t^{(2n-3)/(2n)}$ for velocity, proved by a similarity transformation of the Riccati equation (Proposition 3).
- (iii) A robust initialiser. It enforces the range constraint in the three-sensor Chan–Ho problem (Lemma 1), resolves mirror ambiguities by a prior, and anchors the window estimate at the epoch where the filter starts.
- (iv) A Monte Carlo evaluation answering **RQ1–RQ6**, with run-level bootstrap confidence intervals and paired comparisons under common random numbers.
- (v) Open, tested software that regenerates every figure and number of this paper from one command.

Section 2 states the problem and its assumptions. Sections 3 and 4 develop the measurement and motion models. Section 5 presents the estimator. Section 6 gives the performance theory, Section 7 the experimental methodology, and Section 8 the results. Sections 9 and 10 discuss and conclude.

2 Problem formulation

Vectors are bold lowercase and matrices bold uppercase. $\mathbf{1}$ is the all-ones vector and $\otimes$ the Kronecker product. S sensors at known positions $\mathbf{s}_i \in \mathbb{R}^2$ observe a target with position $\mathbf{p}_k$ and velocity $\mathbf{v}_k$ at epochs $t_k = k\Delta$. For a target at $\mathbf{p}$, $r_i(\mathbf{p}) = \|\mathbf{p} - \mathbf{s}_i\|$ is the range to sensor i and $\mathbf{u}_i = (\mathbf{p} - \mathbf{s}_i)/r_i$ the unit line-of-sight vector. $\mathbf{U}$ stacks the rows $\mathbf{u}_i^\top$. The state $\mathbf{x}_k \in \mathbb{R}^n$ contains position and velocity and, depending on the motion model, acceleration or turn rate.

Assumptions.

- A1** Planar geometry. The sensor positions are known and the sensor clocks are mutually synchronised.
- A2** Line-of-sight propagation at speed c . The arrival time at sensor i carries independent zero-mean Gaussian jitter of standard deviation σ_i . Obstructions inflate σ_i but introduce no bias.
- A3** TOA: the emission instant is known (cooperative emitter). TDOA: it is unknown.
- A4** The filter's motion model is a Markov process with additive Gaussian process noise. The true trajectory may or may not follow it.
- A5** A single target with no missed detections or clutter, observed synchronously every Δ seconds.

Objective. Recursively compute an estimate $\hat{\mathbf{x}}_{k|k}$ and covariance $\mathbf{P}_{k|k}$ from $\mathbf{z}_{1:k}$ that are (a) accurate, measured by the root-mean-square error (RMSE), (b) *consistent*, meaning $\mathbf{P}_{k|k}$ describes the actual error, and (c) *efficient*, meaning close to the posterior Cramér–Rao bound.

3 Measurement model and information geometry

3.1 A unified observation model

Under A2 every sensor produces one timestamp. TOA reports each timestamp minus the known emission instant. TDOA reports differences relative to a reference sensor, which cancels the unknown instant. Both are linear maps of the same timestamps:

$$\mathbf{z} = \mathbf{D} \frac{\mathbf{r}(\mathbf{p})}{c} + \mathbf{D}\mathbf{w}, \quad \mathbf{w} \sim \mathcal{N}(\mathbf{0}, \Sigma), \quad \Sigma = \text{diag}(\sigma_1^2, \dots, \sigma_S^2), \quad (1)$$

with $\mathbf{D} = \mathbf{I}_S$ (TOA) or, with sensor 1 as reference, $\mathbf{D} = [-\mathbf{1} \quad \mathbf{I}_{S-1}]$ (TDOA). Consequently

$$\mathbf{h}(\mathbf{p}) = \frac{\mathbf{D} \mathbf{r}(\mathbf{p})}{c}, \quad \mathbf{H}(\mathbf{p}) = \frac{\partial \mathbf{h}}{\partial \mathbf{p}} = \frac{\mathbf{D} \mathbf{U}(\mathbf{p})}{c}, \quad \mathbf{R} = \mathbf{D} \Sigma \mathbf{D}^\top. \quad (2)$$

For equal jitter σ_t this gives $\mathbf{R}_{\text{TOA}} = \sigma_t^2 \mathbf{I}_S$ and the well-known correlated form $\mathbf{R}_{\text{TDOA}} = \sigma_t^2 (\mathbf{I}_{S-1} + \mathbf{1}\mathbf{1}^\top)$. Here it follows from the structure of the measurement rather than being imposed. Heterogeneous jitter, as in Section 8.6, only changes Σ . Level sets of TOA are circles about $\mathbf{s}_i$, and those of TDOA are hyperbola branches with foci $\mathbf{s}_1, \mathbf{s}_i$. Three sensors in general position are needed for a planar fix in either mode.

3.2 What TDOA gives up

The Fisher information about the position carried by one epoch is $\mathbf{J} = \mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H}$ [13].

Proposition 1. For equal jitter σ_t ,

$$\mathbf{J}_{\text{TOA}} - \mathbf{J}_{\text{TDOA}} = \frac{1}{S c^2 \sigma_t^2} \left(\sum_i \mathbf{u}_i \right) \left(\sum_i \mathbf{u}_i \right)^\top. \quad (3)$$

Proof. By (2), $\mathbf{J} = \mathbf{U}^T \mathbf{D}^T (\mathbf{D} \mathbf{D}^T)^{-1} \mathbf{D} \mathbf{U} / (c^2 \sigma_t^2)$. For TOA the central factor is $\mathbf{I}_S$. For TDOA it is the orthogonal projector onto the row space of $\mathbf{D}$. Since $\mathbf{D} \mathbf{1} = \mathbf{0}$ and $\text{rank } \mathbf{D} = S - 1$, that row space is $\mathbf{1}^\perp$, and the projector equals $\mathbf{I}_S - \mathbf{1} \mathbf{1}^T / S$. The claim follows from $\mathbf{U}^T \mathbf{1} = \sum_i \mathbf{u}_i$. $\square$

Corollary 1. (a) TDOA is never more informative than TOA, and the loss has rank one. (b) The TDOA information does not depend on the choice of reference sensor. (c) TOA and TDOA are equally informative exactly when $\sum_i \mathbf{u}_i = \mathbf{0}$, for instance at the centre of any symmetric array.

The common explanation, that TOA's surplus measurement “informs only the emission time”, is thus inaccurate in setting A3. The surplus is information along the resultant of the line-of-sight vectors. It vanishes inside a surrounding array and grows as the target moves out of it.

3.3 Geometric dilution of precision

Writing $\mathbf{J} = \mathbf{G} / (c \sigma_t)^2$ with $\mathbf{G} = \mathbf{U}^T \mathbf{D}^T (\mathbf{D} \mathbf{D}^T)^{-1} \mathbf{D} \mathbf{U}$, the snapshot Cramér–Rao bound is

$$\sqrt{\mathbb{E} \|\hat{\mathbf{p}} - \mathbf{p}\|^2} \geq c \sigma_t \text{GDOP}(\mathbf{p}), \quad \text{GDOP}(\mathbf{p}) = \sqrt{\text{tr } \mathbf{G}^{-1}}. \quad (4)$$

Proposition 2. For $N \geq 3$ sensors equally spaced on a circle centred at the target, $\mathbf{G} = \frac{N}{2} \mathbf{I}_2$ and $\text{GDOP} = 2 / \sqrt{N}$ in both TOA and TDOA.

Proof. With $\mathbf{u}_i = (\cos \theta_i, \sin \theta_i)$ and $\theta_i = \theta_0 + 2\pi i / N$, the sums $\sum_i \cos 2\theta_i$, $\sum_i \sin 2\theta_i$ and $\sum_i \mathbf{u}_i$ vanish for $N \geq 3$. Hence $\mathbf{U}^T \mathbf{U} = \sum_i \mathbf{u}_i \mathbf{u}_i^T = \frac{N}{2} \mathbf{I}_2$, and by Proposition 1 the TDOA information coincides with it. $\square$

4 Dynamic models

The state is axis-interleaved, $\mathbf{x} = [p_x, p_y, v_x, v_y, \dots]^T$, and evolves as $\mathbf{x}_{k+1} = \mathbf{f}(\mathbf{x}_k) + \mathbf{q}_k$ with $\mathbf{q}_k \sim \mathcal{N}(\mathbf{0}, \mathbf{Q})$. Three of the four models are linear and are exact discretisations of one continuous-time family. Per axis, $d\mathbf{s} = \mathbf{A} \mathbf{s} dt + \mathbf{b} dW$ with white-noise intensity q (Table 1). For a sampling interval Δ ,

$$\Phi = e^{\mathbf{A}\Delta}, \quad \mathbf{Q}_{\text{axis}} = \int_0^\Delta e^{\mathbf{A}\tau} \mathbf{b} q \mathbf{b}^T e^{\mathbf{A}^T \tau} d\tau, \quad \mathbf{F} = \Phi \otimes \mathbf{I}_2, \quad \mathbf{Q} = \mathbf{Q}_{\text{axis}} \otimes \mathbf{I}_2. \quad (5)$$

The integral is evaluated exactly with Van Loan's block exponential [29]: $\exp\left(\begin{bmatrix} -\mathbf{A} & q\mathbf{b}\mathbf{b}^T \\ \mathbf{0} & \mathbf{A}^T \end{bmatrix} \Delta\right) = \begin{bmatrix} \cdot & \Phi^{-1} \mathbf{Q}_{\text{axis}} \\ \mathbf{0} & \Phi^T \end{bmatrix}$. For CV and CA this reproduces the textbook closed forms [19, Sec. 6.2–6.3]. For the Singer model [22] it avoids the cancellation that limits the hand-expanded closed form to about four significant digits at $\alpha\Delta = 0.005$. Because the family is consistent, Singer reduces to CA as $\alpha \rightarrow 0$ at fixed q .

Table 1. The linear motion models as one continuous-time family (per axis) and their default parameters.

Model	axis state $\mathbf{s}$	dynamics	intensity q	default parameters
CV (white-noise acceleration)	(p, v)	$\dot{v} = w$	σ_a^2	$\sigma_a = 1 \text{ m s}^{-3/2}$
CA (white-noise jerk)	(p, v, a)	$\dot{a} = w$	σ_j^2	$\sigma_j = 0.5 \text{ m s}^{-5/2}$
Singer (Gauss–Markov acceleration)	(p, v, a)	$\dot{a} = -\alpha a + w$	$2\alpha\sigma_m^2$	$\alpha = 0.1 \text{ s}^{-1}$, $\sigma_m = 5 \text{ m/s}^2$

Coordinated turn. With state $[p_x, p_y, v_x, v_y, \omega]^T$ the velocity rotates by $\omega\Delta$ per step [19, Sec. 11.7], [21]. Let $s(\omega) = \sin(\omega\Delta)/\omega$ and $\gamma(\omega) = (1 - \cos \omega\Delta)/\omega$. Then

$$\mathbf{f}(\mathbf{x}) = \begin{bmatrix} p_x + s v_x - \gamma v_y \\ p_y + \gamma v_x + s v_y \\ \cos(\omega\Delta) v_x - \sin(\omega\Delta) v_y \\ \sin(\omega\Delta) v_x + \cos(\omega\Delta) v_y \\ \omega \end{bmatrix}, \quad \frac{\partial \mathbf{f}}{\partial \omega} = \begin{bmatrix} s' v_x - \gamma' v_y \\ \gamma' v_x + s' v_y \\ -\Delta(\sin(\omega\Delta) v_x + \cos(\omega\Delta) v_y) \\ \Delta(\cos(\omega\Delta) v_x - \sin(\omega\Delta) v_y) \\ 1 \end{bmatrix}, \quad (6)$$

with $s' = (\Delta \cos \omega \Delta - s)/\omega$ and $\gamma' = (\Delta \sin \omega \Delta - \gamma)/\omega$. These ratios are 0/0 at $\omega = 0$. For $|u| < 10^{-4}$, where $u = \omega \Delta$, they are replaced by the series $s \approx \Delta(1 - u^2/6)$, $\gamma \approx \Delta u(1 - u^2/12)/2$, $s' \approx -\Delta^2 u/3$ and $\gamma' \approx \Delta^2(1 - u^2/4)/2$, so $\mathbf{f}$ and its Jacobian remain smooth through straight flight. Process noise is white acceleration on the velocity (intensity σ_a^2 , default 0.5) plus a random walk on ω (intensity σ_ω^2 , default $0.02 \text{ rad s}^{-3/2}$).

5 Estimation

5.1 Single-epoch maximum-likelihood fix

A position fix is needed to seed the track and as a reference for the tracker's gain. It is computed in four steps.

Closed-form seeds. For TOA, subtracting the squared-range equation of sensor 1 eliminates $\|\mathbf{p}\|^2$ and leaves the linear system $2(\mathbf{s}_i - \mathbf{s}_1)^\top \mathbf{p} = \|\mathbf{s}_i\|^2 - \|\mathbf{s}_1\|^2 - r_i^2 + r_1^2$ for $i = 2, \dots, S$. For TDOA, Chan and Ho [11] introduce $R_1 = \|\mathbf{p} - \mathbf{s}_1\|$ and the range differences $d_i = cz_i$:

$$(\mathbf{s}_i - \mathbf{s}_1)^\top \mathbf{p} + d_i R_1 = \frac{1}{2}(\|\mathbf{s}_i\|^2 - \|\mathbf{s}_1\|^2 - d_i^2), \quad i = 2, \dots, S. \quad (7)$$

With $S \geq 4$ this is an over-determined linear system in $(\mathbf{p}, R_1)$. With $S = 3$ it is under-determined, and the unconstrained least-squares (minimum-norm) solution is in general far from the target. The constraint must be imposed explicitly.

Lemma 1. *Let $S = 3$ with non-collinear sensors and noise-free TDOA data. Then (7) gives $\mathbf{p}(R_1) = \mathbf{a} - \mathbf{b}R_1$, and every position consistent with the data equals $\mathbf{p}(R_1)$ for a root of*

$$(\mathbf{b}^\top \mathbf{b} - 1)R_1^2 - 2\mathbf{b}^\top(\mathbf{a} - \mathbf{s}_1)R_1 + \|\mathbf{a} - \mathbf{s}_1\|^2 = 0 \quad (8)$$

satisfying $R_1 \geq 0$ and $R_1 + d_i \geq 0$. At most two such positions exist.

Proof. With non-collinear sensors, the 2×2 matrix with rows $(\mathbf{s}_i - \mathbf{s}_1)^\top$ is invertible, which gives the affine form. Consistency requires $\|\mathbf{p} - \mathbf{s}_1\| = R_1$, and substituting the affine form yields (8). Conversely, a root makes the squared range-difference equations hold, and the sign conditions recover $r_i = R_1 + d_i$. A quadratic has at most two roots. $\square$

The ambiguity is therefore intrinsic to the data. No single-epoch estimator can resolve it, and a prior is required. When the closed forms are degenerate (collinear sensors), a coarse grid search supplies well-separated seeds instead.

Gauss–Newton refinement. Each seed is refined by weighted Gauss–Newton [8],

$$\mathbf{p} \leftarrow \mathbf{p} + \lambda (\mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H})^{-1} \mathbf{H}^\top \mathbf{R}^{-1} (\mathbf{z} - \mathbf{h}(\mathbf{p})), \quad (9)$$

with λ halved until the weighted cost decreases. This yields the maximum-likelihood fix and its Cramér–Rao covariance $\mathbf{C} = (\mathbf{H}^\top \mathbf{R}^{-1} \mathbf{H})^{-1}$.

Disambiguation. Candidates whose cost lies within 2 of the minimum are equally plausible. Examples are the two roots of Lemma 1 and the mirror images of a linear array. Among them the one nearest to a prior hint, normally the previous fix, is selected.

Algorithm 1 Track initialisation from a window of M epochs**Require:** epochs $t_1 < \dots < t_M$, measurements $z_1, \dots, z_M$, prior hint $\mathbf{h}_0$

```

1: for  $i = 1, \dots, M$  do ▷ A: snapshot fixes
2:    $(\hat{\mathbf{p}}_i, C_i) \leftarrow \text{Fix}(z_i, \text{hint} = \hat{\mathbf{p}}_{i-1} \text{ or } \mathbf{h}_0)$ 
3: end for
4:  $\tau_i \leftarrow t_i - t_M, \quad \mathbf{H}_i \leftarrow [\mathbf{I}_2 \quad \tau_i \mathbf{I}_2]$ 
5:  $\boldsymbol{\theta} \leftarrow (\sum_i \mathbf{H}_i^\top C_i^{-1} \mathbf{H}_i)^{-1} \sum_i \mathbf{H}_i^\top C_i^{-1} \hat{\mathbf{p}}_i$  ▷ B: CV segment anchored at  $t_M$ 
6: repeat ▷ C: joint Gauss–Newton on raw data
7:    $\mathbf{p}_i \leftarrow \boldsymbol{\theta}_{1:2} + \tau_i \boldsymbol{\theta}_{3:4}, \quad \mathbf{J}_i \leftarrow \mathbf{H}(\mathbf{p}_i)[\mathbf{I}_2 \quad \tau_i \mathbf{I}_2]$ 
8:    $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} + (\sum_i \mathbf{J}_i^\top \mathbf{R}_i^{-1} \mathbf{J}_i)^{-1} \sum_i \mathbf{J}_i^\top \mathbf{R}_i^{-1} (z_i - \mathbf{h}(\mathbf{p}_i))$  ▷ with step halving
9: until converged
10:  $\mathbf{P}_\theta \leftarrow (\sum_i \mathbf{J}_i^\top \mathbf{R}_i^{-1} \mathbf{J}_i)^{-1}$ 
11: return  $\hat{\mathbf{x}}_{M|M} = [\boldsymbol{\theta}; \mathbf{0}], \quad \mathbf{P}_{M|M} = \text{blkdiag}(\mathbf{P}_\theta, \text{model prior})$  ▷ D: lift

```

5.2 Track initialiser

Algorithm 1 turns the first $M = 5$ epochs into the EKF prior, following the window least-squares recipe of the tracker-school notes [30] with two refinements. First, stage C refines the constant-velocity segment on the raw measurements, so the result is the maximum-likelihood estimate of the window. Second, the segment is parametrised at the *last* epoch, so the prior refers to t_M , exactly where the filter takes over. Parametrising at t_1 without propagation offsets the prior by $(M - 1)\Delta \mathbf{v}$. That is about 19 m at 47 m/s, enough to dominate the first seconds of tracking. Components that a 0.4 s window cannot observe (acceleration, turn rate) receive the model’s prior variance. Over 3000 simulated windows of the straight-line path the prior’s NEES averages 4.08 against an expected value of 4 (95 % Monte Carlo band [3.90, 4.10]).

5.3 Extended Kalman filter**Algorithm 2** One EKF cycle with Joseph-form update**Require:** $\hat{\mathbf{x}}, \mathbf{P}$ at t_{k-1} ; measurement $\mathbf{z}$ at t_k ; models $\mathbf{f}, \mathbf{h}$; noise $\mathbf{R}$ or $\mathbf{R}(\cdot)$

```

1:  $\mathbf{F} \leftarrow \partial \mathbf{f} / \partial \mathbf{x}|_{\hat{\mathbf{x}}}, \quad \hat{\mathbf{x}} \leftarrow \mathbf{f}(\hat{\mathbf{x}}), \quad \mathbf{P} \leftarrow \mathbf{F} \mathbf{P} \mathbf{F}^\top + \lambda \mathbf{Q}$  ▷ predict;  $\lambda = 1$  unless adapting
2:  $\mathbf{H} \leftarrow \mathbf{H}(\mathbf{E} \hat{\mathbf{x}}) \mathbf{E}, \quad \mathbf{v} \leftarrow \mathbf{z} - \mathbf{h}(\mathbf{E} \hat{\mathbf{x}}), \quad \mathbf{R} \leftarrow \mathbf{R}(\mathbf{E} \hat{\mathbf{x}})$  if map-aware
3:  $\mathbf{S} \leftarrow \mathbf{H} \mathbf{P} \mathbf{H}^\top + \mathbf{R} = \mathbf{L} \mathbf{L}^\top$  ▷ one Cholesky factorisation
4:  $\mathbf{K} \leftarrow \mathbf{P} \mathbf{H}^\top \mathbf{S}^{-1}$  ▷ two triangular solves
5:  $\hat{\mathbf{x}} \leftarrow \hat{\mathbf{x}} + \mathbf{K} \mathbf{v}$ 
6:  $\mathbf{P} \leftarrow (\mathbf{I} - \mathbf{K} \mathbf{H}) \mathbf{P} (\mathbf{I} - \mathbf{K} \mathbf{H})^\top + \mathbf{K} \mathbf{R} \mathbf{K}^\top$  ▷ Joseph form
7:  $\varepsilon_v \leftarrow \|\mathbf{L}^{-1} \mathbf{v}\|^2$ 

```

Algorithm 2 is the first-order EKF [19, Ch. 10], written once for any pair of motion and measurement models. The measurement Jacobian with respect to position is lifted to the state by the $2 \times n$ position selector $\mathbf{E}$. The covariance update uses the Joseph form, a sum of two positive semi-definite terms, which remains symmetric and positive definite under round-off in the gain [17, 18]. The innovation covariance is factorised once and serves the gain, the NIS and the likelihood. Because the measurement model is an argument of the update, the sensor set may change between epochs without reinitialisation.

5.4 Extensions

Map-aware measurement noise. Under A2, obstructions inflate per-sensor jitter but leave the geometry unchanged. Given an environment model $\sigma_i(\mathbf{p})$ (Section 8.6), the filter evaluates it at its

predicted position and uses $\mathbf{R}(\hat{\mathbf{p}}) = \mathbf{D} \text{diag}(\sigma_i^2(\hat{\mathbf{p}}))\mathbf{D}^\top$, which is (2) with a position-dependent Σ . The *nominal* filter assumes clear-sky jitter everywhere.

Adaptive process noise. A smooth motion model cannot follow an instantaneous turn, and its innovation then becomes implausibly large. When $\varepsilon_v > \chi_m^2(0.999)$ the cycle is repeated from the previous posterior with $\lambda\mathbf{Q}$, for $\lambda = 10, 10^2, \dots, 10^5$, and the first level whose innovation passes the gate is kept. This is a data-driven instance of adjustable-level process noise [19, Sec. 11.2], in the spirit of innovation-based adaptive filtering [24]. Under correct modelling the gate opens at one epoch in a thousand.

6 Performance theory

6.1 Consistency

With estimation error $\tilde{\mathbf{x}}_k$ and innovation $\mathbf{v}_k$, a correctly modelled filter has $\varepsilon_x(k) = \tilde{\mathbf{x}}_k^\top \mathbf{P}_{k|k}^{-1} \tilde{\mathbf{x}}_k \sim \chi_n^2$ and $\varepsilon_v(k) = \mathbf{v}_k^\top \mathbf{S}_k^{-1} \mathbf{v}_k \sim \chi_m^2$ [19, Sec. 5.4]. Averaged over N independent runs, $N\bar{\varepsilon}_x(k) \sim \chi_{Nn}^2$, so the average NEES (ANEES) must lie in

$$\bar{\varepsilon}_x(k) \in \frac{1}{N} [\chi_{Nn}^2(0.025), \chi_{Nn}^2(0.975)], \quad (10)$$

and likewise for the ANIS with m . Values above the band indicate over-confidence and values below it conservatism.

6.2 Posterior Cramér–Rao bound

For additive Gaussian noises the PCRLB information matrix obeys [26, 27]

$$\mathbf{J}_{k+1} = (\mathbf{Q} + \mathbf{F}_k \mathbf{J}_k^{-1} \mathbf{F}_k^\top)^{-1} + \mathbf{E}^\top \mathbf{H}_k^\top \mathbf{R}_k^{-1} \mathbf{H}_k \mathbf{E}, \quad (11)$$

with Jacobians evaluated on the true trajectory, and $\mathbf{E}[\tilde{\mathbf{x}}_k \tilde{\mathbf{x}}_k^\top] \succeq \mathbf{J}_k^{-1}$. The bound presumes that the truth follows the motion model. It is therefore used with truth sampled from that model.

6.3 Steady-state scaling

Proposition 3. Consider the n -fold integrator $\dot{\mathbf{s}} = \mathbf{A}\mathbf{s} + \mathbf{b}w$, where $\mathbf{A}$ is the upward shift matrix, $\mathbf{b} = \mathbf{e}_n$ and w is white with intensity q , observed through $y = \mathbf{e}_1^\top \mathbf{s} + v$, where v is white with intensity r . Let $\tau = (r/q)^{1/(2n)}$. The stabilising solution of the filter Riccati equation is

$$\mathbf{P} = \frac{r}{\tau} \mathbf{T} \tilde{\mathbf{P}}_n \mathbf{T}, \quad \mathbf{T} = \text{diag}(1, \tau^{-1}, \dots, \tau^{-(n-1)}), \quad (12)$$

where $\tilde{\mathbf{P}}_n$ is the solution for $q = r = 1$, a constant matrix depending only on n . In particular, $P_{jj} = (\tilde{P}_n)_{jj} q^{(2j-1)/(2n)} r^{(2n-2j+1)/(2n)}$, and for $n = 2$, $P_{11} = \sqrt{2} q^{1/4} r^{3/4}$.

Proof. The Riccati equation is $\mathbf{A}\mathbf{P} + \mathbf{P}\mathbf{A}^\top + q\mathbf{b}\mathbf{b}^\top - \mathbf{P}\mathbf{e}_1\mathbf{e}_1^\top\mathbf{P}/r = \mathbf{0}$. One checks that $\mathbf{A}\mathbf{T} = \tau^{-1}\mathbf{T}\mathbf{A}$, $\mathbf{T}\mathbf{e}_1 = \mathbf{e}_1$ and $\mathbf{T}\mathbf{b} = \tau^{-(n-1)}\mathbf{b}$. Substituting (12), multiplying by $\mathbf{T}^{-1}$ on both sides and dividing by r/τ^2 gives $\mathbf{A}\tilde{\mathbf{P}}_n + \tilde{\mathbf{P}}_n\mathbf{A}^\top + (q\tau^{2n}/r)\mathbf{b}\mathbf{b}^\top - \tilde{\mathbf{P}}_n\mathbf{e}_1\mathbf{e}_1^\top\tilde{\mathbf{P}}_n = \mathbf{0}$. Since $q\tau^{2n}/r = 1$, this is the normalised equation. $(\mathbf{A}, \mathbf{e}_1^\top)$ is observable and $(\mathbf{A}, \mathbf{b})$ controllable, so its stabilising solution is unique [31]. For $n = 2$, solving the three scalar equations gives $\tilde{\mathbf{P}}_2 = \begin{bmatrix} \sqrt{2} & 1 \\ 1 & \sqrt{2} \end{bmatrix}$. $\square$

When position fixes of per-axis variance σ_p^2 arrive every $\Delta \ll \tau$ seconds, they act as a continuous observation of intensity $r = \sigma_p^2 \Delta$ [32]. With $\sigma_p \propto c \sigma_t \text{GDOP}$:

Corollary 2. *In the regime $\Delta \ll \tau$, the steady-state errors scale as*

$$RMSE_p \propto (\sigma_t \text{GDOP})^{(2n-1)/(2n)}, \quad RMSE_v \propto (\sigma_t \text{GDOP})^{(2n-3)/(2n)}. \quad (13)$$

For CV ($n = 2$) the exponents are $3/4$ and $1/4$. For CA ($n = 3$) they are $5/6$ and $1/2$. On a ring of N sensors (Proposition 2) the position RMSE therefore scales as $N^{-3/8}$ for CV and $N^{-5/12}$ for CA.

A tracker thus gains less than proportionally from better clocks: halving σ_t improves a CV tracker by $2^{3/4} \approx 1.68$. The effective memory τ shrinks as the measurements improve. When τ approaches Δ the filter can no longer average, and the exponent tends to that of a snapshot fix, namely 1. This is the continuous-time counterpart of the tracking-index analysis of α - β trackers [28]. In the experiments the discrete Riccati equation is solved exactly to give quantitative predictions.

7 Experimental methodology

7.1 Simulation design

Unless stated otherwise: $\Delta = 0.1$ s, $\sigma_t = 10$ ns (3 m of range), TDOA with sensor 1 as reference, a square array inscribed in a circle of radius 2.5 km (all regular layouts share this circumradius), an initialisation window of $M = 5$ epochs, and the default parameters of Table 1. Comparisons across motion models use manoeuvre-tolerant tunings ($\sigma_a = 4$; $\sigma_j = 2$; CT $\sigma_a = 2$, $\sigma_\omega = 0.05$; Singer $\alpha = 0.2$, $\sigma_m = 8$).

7.2 Ground truth

Two kinds of truth serve different purposes. *Matched* truth is sampled from the filter's own stochastic motion model. Under it, the chi-square tests and the PCRLB hold exactly, so it is used to verify the implementation (RQ1) and to test the scaling laws (RQ3). *Designed* trajectories (Table 2) are deterministic paths that probe specific behaviour. On them a filter tuned for random accelerations is expected to be conservative (ANEES below n), and departures above n indicate model mismatch.

Table 2. Designed ground-truth trajectories.

Trajectory	definition	duration
straight line	$\mathbf{p}_0 = (-1500, -800)$ m, $\mathbf{v} = (40, 25)$ m/s	30 s
constant acceleration	$\mathbf{p}_0 = (-1000, -1000)$ m, $\mathbf{v}_0 = (30, 0)$ m/s, $\mathbf{a} = (0, 1.5)$ m/s ²	30 s
coordinated turn	circle of radius 600 m, $\omega = 0.08$ rad/s (48 m/s)	40 s
zigzag	four legs at 60 m/s, 124° instantaneous turns	100 s
figure-eight	$\mathbf{p} = A(\sin \omega t, \frac{1}{2} \sin 2\omega t)$, $A = 800$ m, $\omega = 0.1$ rad/s	62.8 s

7.3 Statistical analysis

Each configuration is repeated N times with independent random streams spawned from a fixed seed. N is between 30 and 100 and is given with each result. The run, not the epoch, is the independent unit. Every RMSE is therefore reported with a 95 % percentile-bootstrap interval over runs (2000 resamples) [33]. Comparisons between two filters (TOA/TDOA, fixed/adaptive, nominal/map-aware) use *common random numbers*: run i of both filters sees the same truth and the same noise. Their RMSE ratio is bootstrapped pairwise, which removes most of the run-to-run variance from the comparison. Consistency is judged by the band (10), efficiency by RMSE/PCRLB, and scaling exponents by least squares on logarithms, with 95 % intervals of ± 1.96 standard errors.

7.4 Reproducibility

All experiments are implemented in the open `math_tracker` package. Running `python -m math_tracker reproduce` regenerates every figure, table and quoted number in this paper from fixed seeds, in about three minutes on a laptop. Section A summarises the software.

8 Results

8.1 RQ1: consistency and efficiency

Table 3 and Figure 1 report each filter on matched truth (100 runs of 30 s). All four filters have ANEES and ANIS within a few per cent of n and m , inside their Monte Carlo bands. Their RMSE equals the posterior bound throughout the run, including the initial transient: RMSE/PCRLB lies between 0.967 and 1.004, and the bound falls inside the RMSE’s 95 % interval for 4 of the 4 models. Figure 2 shows the per-epoch picture for CA. A single run scatters over the wide one-run band, while the 100-run average stays within the narrow Monte Carlo band in 94 % of the epochs. *Answer to RQ1*: at kilometre ranges and metre-level noise the first-order EKF is consistent and statistically efficient. Its linearisation error is negligible.

Table 3. Consistency and efficiency on matched truth (TDOA, $m = 3$). RMSE: mean $\pm$ half-width of the 95 % run-bootstrap interval. Bands: the 95 % Monte Carlo interval of the average.

Model	n	RMSE (m)	PCRLB (m)	RMSE/PCRLB	ANEES	95% band	ANIS	95% band
CV	4	1.27 ± 0.02	1.26	1.004 ± 0.017	4.03	[3.46, 4.57]	3.02	[2.54, 3.50]
CA	6	1.32 ± 0.02	1.33	0.995 ± 0.017	5.88	[5.34, 6.70]	3.01	[2.54, 3.50]
CT	5	1.26 ± 0.02	1.26	0.999 ± 0.019	5.00	[4.40, 5.64]	3.02	[2.54, 3.50]
Singer	6	1.65 ± 0.07	1.71	0.967 ± 0.039	5.86	[5.34, 6.70]	3.01	[2.54, 3.50]




Figure 1. Monte Carlo position RMSE (solid) and posterior Cramér–Rao bound (dotted) on matched truth. The Singer truth wanders furthest from the array, so its bound rises late in the run.

Designed trajectories. Table 4 and Figures 3 and 4 show each designed path tracked by its matching model (50 runs). On the smooth paths the RMSE is 1.1–1.3 m, well below the 3 m range noise, and ANEES is below n , the expected conservatism on noise-free paths. The zigzag violates every smooth model, and its ANEES exposes severe over-confidence at the turns (Section 8.5).




Figure 2. NEES and NIS of the CA filter. Light band: 95 % interval for a single run. Dark band: for the 100-run average.

Table 4. Designed trajectories tracked by the matching model (50 runs; metres).

Trajectory	Model	RMSE	final	peak	vel. RMSE	ANEES (n)	ANIS (m)
Straight line	CV	1.17 ± 0.03	1.05	2.97	1.36	2.19 (4)	2.97 (3)
Constant acceleration	CA	1.26 ± 0.03	1.11	2.72	1.64	3.38 (6)	2.99 (3)
Coordinated turn	CT	1.15 ± 0.02	1.12	2.59	1.43	2.52 (5)	2.97 (3)
Zigzag	Singer	3.89 ± 0.03	1.45	20.62	14.16	289.59 (6)	6.59 (3)

8.2 RQ2: TOA versus TDOA

Figure 5 maps the GDOP ratio of the two modes for the square array. As Corollary 1 predicts, it is exactly 1 near the centre and grows as the line-of-sight vectors align outside the array. On the constant-acceleration path, which stays inside the array, the paired RMSE ratio TDOA/TOA is 1.01 [1.01, 1.02] (100 runs). On an outbound path the two modes agree until the target crosses the sensor line. Over the final 5 s the ratio reaches 2.06 [1.91, 2.22] (Table 5). *Answer to RQ2:* TDOA costs nothing while the sensors surround the target, and its penalty grows with the resultant $\|\sum_i \mathbf{u}_i\|$, exactly as Proposition 1 states.

Table 5. TOA versus TDOA (100 paired runs).

Path	Meas.	RMSE (m)	final (m)	vel. RMSE (m/s)	ANEES (n)	ANIS (m)
inside	TOA	1.23 ± 0.02	1.14	1.59	3.31 (6)	3.97 (4)
inside	TDOA	1.25 ± 0.02	1.14	1.61	3.32 (6)	2.97 (3)
outbound	TOA	1.10 ± 0.02	1.21	1.11	2.13 (4)	3.95 (4)
outbound	TDOA	1.42 ± 0.04	2.50	1.14	2.11 (4)	2.95 (3)

8.3 RQ3: scaling with timing noise and sensor count

Figure 6 sweeps σ_t over three decades on a six-sensor ring with matched CV truth (30 runs per point). The snapshot bound grows with slope 1.000, as (4) requires. The tracker’s position error grows with slope 0.749 ± 0.003 above 10 ns. The exact discrete Riccati solution gives 0.764, and Corollary 2 gives the continuous limit $3/4$. The velocity error grows with slope 0.239 ± 0.003 , against 0.257 and $1/4$. Below 10 ns the position exponent steepens to 0.832 ± 0.017 , as τ approaches Δ . The Riccati prediction, which uses no fitted parameter, matches the Monte Carlo RMSE within 3 % for position and 6 % for velocity




Figure 3. Designed trajectories and EKF estimates (one run). The cross marks the initialiser's output at the end of the 0.4 s window.




Figure 4. One CA run: plan view; per-axis error with the filter's $\pm 3\sigma$ envelope; NEES and NIS against their 95 % bands.




Figure 5. Left: GDOP ratio TDOA/TOA for the square array, with the inside (solid) and outbound (dashed) paths. Right: RMSE on the outbound path (mean and 5–95 % range of single runs).

over the whole sweep. The EKF stays within 3 % of the PCRLB. The tracker’s gain over a single fix rises from $1.1\times$ at 0.1 ns to $4.3\times$ at 100 ns.




Figure 6. Steady-state RMSE against timing noise (six-sensor ring, matched CV truth). Markers: Monte Carlo. Dotted: the steady-state Riccati prediction of Section 6.3.

Figure 7 varies the number of sensors on a ring along the constant-acceleration path (40 runs per point). The fitted exponents are -0.45 ± 0.04 (TOA) and -0.45 ± 0.04 (TDOA). They lie between the asymptotic tracker value $-5/12 \approx -0.42$ of Corollary 2 and the snapshot value $-1/2$, as expected at $\sigma_t = 10$ ns, where the noise exponent has begun to drift towards 1. *Answer to RQ3:* accuracy improves sub-linearly, as $\sigma_t^{3/4}$ and roughly $N^{-0.45}$ here, and the steady-state Riccati equation predicts the error quantitatively from the geometry and the noise level alone.

8.4 RQ4: sensor geometry

Figure 8 maps the TDOA GDOP of five layouts with a common circumradius. Table 6 tracks the constant-acceleration path with each (50 runs). The ring of six is best (1.04 m), followed by the square (1.26 m), the triangle (1.52 m), the random layout (2.23 m) and the line (2.56 m). The ranking is that of the path-averaged GDOP. The collinear array is usable off its axis, but it cannot distinguish a target from its mirror image. Here a one-off prior on the side of the line resolves this (Section 5.1). Without it the error equals twice the target’s distance from the line, kilometres in this geometry. *Answer to RQ4:* layouts of




Figure 7. Steady-state RMSE against the number of sensors on a ring (40 runs per point), between the tracker law $N^{-5/12}$ and the snapshot law $N^{-1/2}$.

equal footprint span a factor of 2.5 in RMSE, comparable to the largest algorithmic gain observed (RQ5). Degenerate layouts additionally require prior information.




Figure 8. TDOA GDOP of five layouts sharing a 2.5 km circumradius (logarithmic scale; contours at 1, 2 and 5).

Table 6. Layout comparison on the constant-acceleration path (50 runs).

Layout	S	GDOP (path mean)	RMSE (m)	final (m)	ANEES (n)	ANIS (m)
ring	6	0.84	1.04 ± 0.02	1.05	3.21 (6)	5.00 (5)
square	4	1.03	1.26 ± 0.03	1.11	3.38 (6)	2.99 (3)
triangle	3	1.27	1.52 ± 0.04	1.30	3.37 (6)	1.99 (2)
random	5	1.66	2.23 ± 0.08	0.89	3.37 (6)	4.01 (4)
line	4	2.43	2.56 ± 0.09	3.52	3.39 (6)	3.00 (3)

8.5 RQ5: motion models and manoeuvres

Figure 9 and Table 7 cross every model with every designed trajectory (40 runs). On the three smooth paths all models are within a factor 1.1 of the best, and all are conservative. There the level of process noise matters more than the structure of the model. On the zigzag every fixed model is over-confident at the turns: ANEES/ n ranges from 5.4 (CV) to 56.5 (CA), whose acceleration state keeps carrying the estimate past each turn.

NIS-gated adaptive process noise addresses this failure (Figure 10 and Table 8, 40 paired runs). The paired RMSE gain ranges from 1.6 to 2.6 across models (CT: 2.6 [2.6, 2.7]). The error spikes at the turns




Figure 9. Every motion model on every trajectory. Left: RMSE. Right: ANEES/ n , where red means overconfident and blue conservative.

Table 7. Position RMSE (m) of each model on each trajectory; best per row in bold.

Trajectory	CV	CA	CT	Singer
straight line	1.54	1.51	1.46	1.66
constant acceleration	1.53	1.47	1.44	1.63
coordinated turn	1.58	1.44	1.40	1.59
zigzag	2.89	4.07	4.14	2.83

shrink, and ANEES/ n for CT falls from 17.4 to 3.3. The gate opens in 0.3 % of the epochs, essentially only at the turns, and on the smooth constant-acceleration path the option costs 1.7 % in RMSE. *Answer to RQ5:* the motion model matters little on smooth targets. On abrupt manoeuvres a single adaptive filter recovers most of the loss without a multiple-model bank.




Figure 10. Zigzag with fixed and adaptive process noise. Left: CT filter over time (turns dotted). Right: RMSE of every model, fixed (grey) and adaptive (colour).

8.6 RQ6: the propagation environment

The environment model maps a scene to per-sensor jitter. A weather factor W scales the base noise, and each obstacle on the line of sight adds a penalty:

$$\sigma_i(\mathbf{p}) = W\sigma_{\text{base}} + \sum_{\text{walls} \cap \overline{\mathbf{p}s_i}} \pi_{\text{wall}} + \sum_{\text{trees} \cap \overline{\mathbf{p}s_i}} \pi_{\text{tree}} + \pi_{\text{crowd}} + \sum_{\text{base stations } b} \frac{P_b}{\max(d_{ib}, 1 \text{ m})}. \quad (14)$$

Table 8. Fixed against adaptive process noise on the zigzag (40 paired runs). Gain: paired RMSE ratio with 95 % interval.

Model	RMSE (m)			peak RMSE (m)		ANEES/ n	
	fixed	adaptive	gain	fixed	adaptive	fixed	adaptive
CV	2.89	1.57	1.85 ± 0.01	17.58	5.56	5.43	1.81
CA	4.07	1.75	2.32 ± 0.02	20.82	10.99	56.48	9.87
CT	4.14	1.57	2.64 ± 0.03	23.40	5.67	17.37	3.31
Singer	2.83	1.78	1.59 ± 0.01	15.52	8.68	16.39	4.27

The parameters (Table 9) are representative rather than calibrated (compare [15, 34]). In line with A2, obstructions add variance only. The NLOS bias that real obstructions also cause is outside the scope of this study.

Table 9. Environment model parameters (penalties add to the jitter, in ns).

Element	variants	effect
weather W	clear, light rain, snow, heavy rain, fog, storm	$\times 1, 1.5, 2.2, 2.5, 3, 5$
wall (per crossing)	thin / medium / thick (10/25/50 cm)	$+5 / 15 / 35$
tree (per crossing)	small / medium / large (crown 4/9/18 m)	$+1.5 / 4 / 8$
crowd	N people, N_{ph} phones in use	$+0.4N + 1.2N_{ph}$
base station	low / medium / high power	$+P_b/d, P_b = 1500 / 4000 / 9000 \text{ ns m}$

Environment ladder. A figure-eight is tracked in four increasingly degraded scenes (Figure 11 and Table 10, 30 paired runs). The storm multiplies the map-aware RMSE by 4.17, against $5^{5/6} \approx 3.8$ predicted by Corollary 2 for a five-fold increase in jitter on a CA tracker. Full clutter costs a factor 7.9. The nominal filter’s ANIS reaches 344.2 against an expected 3, so its covariance carries no information. The map-aware filter keeps ANIS at 3 in every scene and is also more accurate, by a paired factor of 1.22 [1.20, 1.24] in the cluttered scene.

Table 10. Environment ladder (30 paired runs). $\bar{\sigma}_t$: mean true jitter. ANIS should be 3. Ratio: nominal over map-aware RMSE, with 95 % interval.

Scene	$\bar{\sigma}_t$ (ns)	nominal filter		map-aware filter		RMSE ratio
		RMSE (m)	ANIS	RMSE (m)	ANIS	
Clean	10.0	1.55	3.1	1.55	3.1	1.00 ± 0.00
Storm	50.0	7.10	73.5	6.46	3.1	1.10 ± 0.01
Storm + wall	65.6	9.27	135.6	7.71	3.1	1.20 ± 0.02
Storm + wall + clutter	106.8	14.97	344.2	12.28	3.1	1.22 ± 0.02




Figure 11. Environment ladder. Top: the four scenes (trees four times enlarged; crosses are base stations). Bottom: RMSE of the nominal and the map-aware filter.

Occlusion. All sensors stand west of a 3.8 km thick wall with a forest to the east, in heavy rain (Figure 12). Behind the wall every line of sight is obstructed, and the mean jitter rises from 25 to 60 ns. The map-aware filter's RMSE behind the wall is 4.7 [4.4, 5.0] times its value in front. That ratio combines the jitter increase, damped by the sub-linear response of Corollary 2, with the weaker geometry of a one-sided array. Behind the wall the nominal filter is 1.14 [1.10, 1.19] times worse again, with ANIS 108.4. *Answer to RQ6:* modelling the environment buys a moderate gain in accuracy and, more importantly, it restores consistency. The nominal filter's covariance cannot be used for gating, fusion or decision-making in cluttered scenes.




Figure 12. Occlusion scene (trees four times enlarged) and the RMSE of both filters; shaded intervals have the target behind the wall (30 paired runs).

9 Discussion

9.1 Synthesis

The results rank the levers available to a designer. Sensor geometry and manoeuvre handling are the two large ones, each worth a factor of about 2.5 here. Knowledge of the noise, through a map-aware $\mathbf{R}$, decides whether the reported uncertainty means anything in clutter. Timing precision and sensor count act only sub-linearly, through the exponents of Corollary 2. The structure of the motion model matters least on smooth targets. The choice between TOA and TDOA is decided by whether the emitter cooperates, and TDOA is free inside a surrounding array.

9.2 Relation to prior work

Proposition 1 is consistent with the known equivalence between TDOA and TOA with an unknown emission time treated as a nuisance parameter [13, 14]. It makes the loss explicit and geometric. The efficiency of the EKF at high signal-to-noise ratio agrees with the statistical efficiency of Chan–Ho-type estimators in the same regime [11]. Proposition 3 is the continuous-time counterpart of the tracking-index characterisation of fixed-gain trackers [28]. The adaptive scheme is a simple, data-driven alternative to interacting multiple models [23] for impulsive manoeuvres. It does not replace them for sustained regime changes.

9.3 Threats to validity

- *Simulation only.* No hardware data are used. Timing jitter is Gaussian and line-of-sight (A2). Real obstructions also bias ranges (NLOS) [15], and biases are not modelled.
- *Representative environment.* The penalties of Table 9 are plausible, not calibrated. Conclusions about the *mechanism*, consistency through a correct $\mathbf{R}$, are robust to them. Conclusions about magnitudes are not.
- *Tuning.* Motion-model comparisons depend on process-noise levels. We used one manoeuvre-tolerant tuning per model and did not optimise per trajectory.
- *Synchronisation.* Sensor clocks are assumed exact (A1). Clock offsets act like a common or per-sensor bias and would alter Proposition 1’s balance.
- *Sampling.* Monte Carlo sizes of 30–100 runs give the intervals reported. Effects smaller than those intervals are not claimed.

9.4 Future research

Natural continuations include joint estimation and mitigation of NLOS biases; interacting multiple-model and unscented or particle filters for short-range, strongly non-linear geometries [25, 26]; sensor placement as optimal experimental design, that is, maximising $\det \mathbf{G}$ along anticipated trajectories; estimation of clock offsets; multi-target tracking with data association [35]; and validation on an ultra-wideband testbed.

10 Conclusion

Casting TOA and TDOA as one observation, $\mathbf{z} = \mathbf{D}\mathbf{r}(\mathbf{p})/c + \mathbf{D}\mathbf{w}$, gives the measurement model, its correlated covariance and an exact account of TDOA’s information loss. With an initialiser that resolves intrinsic ambiguities and is anchored where the filter starts, a first-order EKF with Joseph updates is consistent and attains the posterior Cramér–Rao bound in this problem. Its steady-state accuracy follows the scaling law of Proposition 3, and the exact Riccati solution predicts it quantitatively. Two

inexpensive extensions address the main practical failure modes: adaptive process noise for abrupt manoeuvres, and a map-aware measurement covariance for cluttered environments.

Code and data availability. The `mtl_tracker` package, its 76 tests and the scripts that generate every figure, table and number in this paper accompany this report. See Section A for usage.

Acknowledgements. The author thanks Kadir Eken and Sudem Bulut for their contributions to the MTH 407 project. The window initialiser follows the least-squares and Kalman filter lecture notes of M. Altun (HAVELSAN tracker school) [30].

References

- [1] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 2nd ed. Boston, MA: Artech House, 2013.
- [2] F. Gustafsson and F. Gunnarsson, “Mobile positioning using wireless networks: Possibilities and fundamental limitations based on available wireless network measurements,” *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 41–53, 2005.
- [3] S. Gezici, Z. Tian, G. B. Giannakis, H. Kobayashi, A. F. Molisch, H. V. Poor, and Z. Sahinoglu, “Localization via ultra-wideband radios: A look at positioning aspects for future sensor networks,” *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 70–84, 2005.
- [4] J. A. del Peral-Rosado, R. Raulefs, J. A. López-Salcedo, and G. Seco-Granados, “Survey of cellular mobile radio localization methods: From 1G to 5G,” *IEEE Commun. Surveys Tuts.*, vol. 20, no. 2, pp. 1124–1148, 2018.
- [5] N. Patwari, J. N. Ash, S. Kyperountas, A. O. Hero, R. L. Moses, and N. S. Correal, “Locating the nodes: Cooperative localization in wireless sensor networks,” *IEEE Signal Process. Mag.*, vol. 22, no. 4, pp. 54–69, 2005.
- [6] N. J. Willis, *Bistatic Radar*, 2nd ed. Raleigh, NC: SciTech Publishing, 2005.
- [7] Bureau International des Poids et Mesures, “The international system of units (SI),” 9th edition, 2019.
- [8] W. H. Foy, “Position-location solutions by Taylor-series estimation,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-12, no. 2, pp. 187–194, 1976.
- [9] D. J. Torrieri, “Statistical theory of passive location systems,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-20, no. 2, pp. 183–198, 1984.
- [10] J. O. Smith and J. S. Abel, “Closed-form least-squares source location estimation from range-difference measurements,” *IEEE Trans. Acoust., Speech, Signal Process.*, vol. 35, no. 12, pp. 1661–1669, 1987.
- [11] Y. T. Chan and K. C. Ho, “A simple and efficient estimator for hyperbolic location,” *IEEE Trans. Signal Process.*, vol. 42, no. 8, pp. 1905–1915, 1994.
- [12] K. C. Ho and W. Xu, “An accurate algebraic solution for moving source location using TDOA and FDOA measurements,” *IEEE Trans. Signal Process.*, vol. 52, no. 9, pp. 2453–2463, 2004.
- [13] S. M. Kay, *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*. Upper Saddle River, NJ: Prentice Hall, 1993.
- [14] Y. Shen and M. Z. Win, “Fundamental limits of wideband localization—Part I: A general framework,” *IEEE Trans. Inf. Theory*, vol. 56, no. 10, pp. 4956–4980, 2010.
- [15] İ. Güvenç and C.-C. Chong, “A survey on TOA based wireless localization and NLOS mitigation techniques,” *IEEE Commun. Surveys Tuts.*, vol. 11, no. 3, pp. 107–124, 2009.
- [16] R. E. Kalman, “A new approach to linear filtering and prediction problems,” *Trans. ASME, J. Basic Eng.*, vol. 82, no. 1, pp. 35–45, 1960.
- [17] R. S. Bucy and P. D. Joseph, *Filtering for Stochastic Processes with Applications to Guidance*. New York: Interscience, 1968.

- [18] M. S. Grewal and A. P. Andrews, *Kalman Filtering: Theory and Practice Using MATLAB*, 3rd ed. Hoboken, NJ: Wiley, 2008.
- [19] Y. Bar-Shalom, X. R. Li, and T. Kirubarajan, *Estimation with Applications to Tracking and Navigation: Theory, Algorithms and Software*. New York: Wiley, 2001.
- [20] S. Blackman and R. Popoli, *Design and Analysis of Modern Tracking Systems*. Norwood, MA: Artech House, 1999.
- [21] X. R. Li and V. P. Jilkov, “Survey of maneuvering target tracking. Part I: Dynamic models,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. 39, no. 4, pp. 1333–1364, 2003.
- [22] R. A. Singer, “Estimating optimal tracking filter performance for manned maneuvering targets,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-6, no. 4, pp. 473–483, 1970.
- [23] H. A. P. Blom and Y. Bar-Shalom, “The interacting multiple model algorithm for systems with Markovian switching coefficients,” *IEEE Trans. Autom. Control*, vol. 33, no. 8, pp. 780–783, 1988.
- [24] R. K. Mehra, “On the identification of variances and adaptive Kalman filtering,” *IEEE Trans. Autom. Control*, vol. 15, no. 2, pp. 175–184, 1970.
- [25] S. J. Julier and J. K. Uhlmann, “Unscented filtering and nonlinear estimation,” *Proc. IEEE*, vol. 92, no. 3, pp. 401–422, 2004.
- [26] B. Ristic, S. Arulampalam, and N. Gordon, *Beyond the Kalman Filter: Particle Filters for Tracking Applications*. Boston, MA: Artech House, 2004.
- [27] P. Tichavský, C. H. Muravchik, and A. Nehorai, “Posterior Cramér–Rao bounds for discrete-time nonlinear filtering,” *IEEE Trans. Signal Process.*, vol. 46, no. 5, pp. 1386–1396, 1998.
- [28] P. R. Kalata, “The tracking index: A generalized parameter for α - β and α - β - γ target trackers,” *IEEE Trans. Aerosp. Electron. Syst.*, vol. AES-20, no. 2, pp. 174–182, 1984.
- [29] C. F. Van Loan, “Computing integrals involving the matrix exponential,” *IEEE Trans. Autom. Control*, vol. 23, no. 3, pp. 395–404, 1978.
- [30] M. Altun, “Linear least squares estimation; Kalman filter; extended Kalman filter,” Tracker-school lecture notes, HAVELSAN.
- [31] B. D. O. Anderson and J. B. Moore, *Optimal Filtering*. Englewood Cliffs, NJ: Prentice-Hall, 1979.
- [32] A. Gelb, Ed., *Applied Optimal Estimation*. Cambridge, MA: MIT Press, 1974.
- [33] B. Efron and R. J. Tibshirani, *An Introduction to the Bootstrap*. New York: Chapman & Hall, 1993.
- [34] ITU-R, “Recommendations P.526 (propagation by diffraction), P.833 (attenuation in vegetation) and P.838 (specific attenuation model for rain),” International Telecommunication Union.
- [35] R. P. S. Mahler, *Statistical Multisource-Multitarget Information Fusion*. Norwood, MA: Artech House, 2007.

A Software and reproducibility

The implementation (Figure 13) is a pure-Python package built on NumPy, SciPy and Matplotlib. Any motion model, providing $\mathbf{f}$, $\mathbf{F}$, $\mathbf{Q}$ and a position selector, combines with any measurement model, providing $\mathbf{h}$, $\mathbf{H}$ and $\mathbf{R}$. A new linear motion model is specified by its continuous-time $(\mathbf{A}, \mathbf{b}, q)$, and a new observation mode by its matrix $\mathbf{D}$. One streaming Tracker drives batch runs, Monte Carlo studies and the interactive simulator. Configurations are immutable, and random streams are spawned from a single seed, so a configuration determines its result exactly. The 76 tests cover Jacobians against finite differences, closed-form noise matrices against Van Loan, Proposition 1, the scaling of Proposition 3 via the Riccati equation, the solvers on ambiguous geometries, the initialiser’s alignment, Monte Carlo consistency and efficiency, the command line and a headless run of the GUI. A 300-epoch scenario runs in about 20 ms.




Figure 13. Layers of the `mth_tracker` package. Arrows point from a module to what it depends on. The models are independent of the estimator, the estimator is independent of the experiments, and the GUI reuses the batch Tracker.

<code>pip install -e .</code>	install the package (Python ≥ 3.10)
<code>python -m pytest</code>	run the 76 tests (about ten seconds)
<code>python -m mth_tracker reproduce</code>	regenerate every figure, table and number (about three minutes)
<code>python -m mth_tracker run --motion CT</code>	one scenario, with an overview figure
<code>python -m mth_tracker gui</code>	the interactive simulator (Figure 14)
<code>latexmk -pdf report.tex</code>	rebuild this document




Figure 14. Interactive simulator built on the same estimator. Sensors, walls, trees and base stations are placed with the mouse. Weather, crowd and timing noise are set live, and map-aware R and adaptive Q can be toggled. The panel reports error, reported uncertainty, NEES, NIS, process-noise level and GDOP, and a 95 % ellipse surrounds the estimate.