

Article

Direct Determination of a Non-Cooperative LEO Satellite's Orbital State from Ground-Station Signals: A Cramér–Rao Analysis

Anthony J. Weiss 

School of Electrical Engineering, Tel Aviv University, Tel Aviv 6997801, Israel; tony@tauex.tau.ac.il

Abstract

This paper considers passive determination of the orbital state of a non-cooperative low-Earth-orbit satellite from radio signals received at several ground stations. Reversing the roles of direct geolocation in which a satellite with known ephemeris locates a ground emitter, the satellite's six-dimensional epoch state is estimated in a single step from the received data. Because the trajectory is a deterministic function of this state, every candidate is a valid orbit. The Cramér–Rao bound is derived; it factors into signal, geometry and orbital-dynamics terms, and a closed-form Keplerian state-transition block is obtained by implicit differentiation of the universal Kepler equation. The delay model includes the light-time correction, which leaves the bound essentially unchanged but removes a bias of more than a hundred meters. The unknown emission epoch is treated as a nuisance; it roughly doubles the three-station bound, whereas four stations recover the known-epoch accuracy. Simulations validate the dynamics implementation and demonstrate post-acquisition efficiency under a Gaussian delay/Doppler surrogate of the known-waveform model. Under an independent-cell model of acquisition failures, an ideal exhaustive search predicts a threshold about 6 dB below that of a two-step baseline with hard per-window peak decisions. The paper also quantifies how accuracy scales with arc length, station count and station layout.

Keywords: direct position determination; orbital-state estimation; Cramér–Rao bound; LEO satellites; passive geolocation; non-cooperative emitters; two-body dynamics; state-transition matrix; space situational awareness

1. Introduction

The rapid growth of low-Earth-orbit (LEO) constellations has made independent non-cooperative determination of satellite orbits from their radio emissions a problem of practical interest for space situational awareness, interference geolocation, and spectrum monitoring. High-precision orbit determination (POD) can achieve centimeter-level accuracy [1–3] but relies on cooperative signals (e.g., DORIS, GNSS, and satellite laser ranging) and a two-step processing chain: range/range-rate observables are extracted per station and then an orbit is fitted to them. Two-step processing can lose information through intermediate observable extraction, particularly when the extracted observables become unreliable at low SNR, where they cannot be reliably extracted [4,5]. The same two-step structure underlies classical TDOA/FDOA emitter geolocation in which delay and Doppler differences are first measured by cross-correlation or cross-ambiguity processing [6,7] and the source is then located algebraically or by iterative fitting [8,9] and the recent passive



Academic Editor: Alberto Gotta

Received: 4 September 2026

Revised: 24 September 2026

Accepted: 2 October 2026

Published: 6 October 2026

Copyright: © 2026 by the author.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

RF tracking of satellites by synchronized ground-receiver networks in which TDOA and FDOA measurements feed an initial orbit determination [10,11].

Direct position determination (DPD)—also called direct geolocation—removes the intermediate stage by mapping raw multi-receiver I/Q data straight to position through a single cost function [4,5,12]; related single-step formulations include TDOA-based direct positioning with its maximum-likelihood estimator and CRB [13], the joint detection and localization of intermittent emitters [14], and direct position estimation in GNSS receivers [15]. DPD is known to outperform two-step methods at low SNR and short data captures [4,12,13]. Recent work has applied DPD from LEO to geolocate terrestrial emitters [16], and Doppler and Doppler-rate measurements from a single LEO satellite have been used to geolocate ground emitters [17]. In all of this literature, however, the satellite trajectory is assumed known from ephemeris; the satellite is the sensor, and the ground emitter is the unknown.

A separate and rapidly growing body of literature treats non-cooperative LEO satellites as signals of opportunity for positioning, navigation and timing, and there the satellite orbit is a nuisance that must be determined or corrected [18,19]. Doppler-based opportunistic positioning from LEO downlinks, including differential schemes and the analysis of their baseline dependence, is developed in [20–25]; orbit determination and ephemeris tracking of non-cooperative satellites from Doppler observables extracted by a receiver at a known position, including the simultaneous tracking-and-navigation framework and its machine learning variants, in [26–31]; and the modeling and compensation of the timing and spatial ephemeris errors of such satellites in [32,33]. All of these are two-step: per-station delay and/or Doppler observables are extracted from the signal and an orbit or a receiver state is then fitted to them, with the extraction usually relying on knowledge of the signal structure (a beacon, a pilot, or a modulation format). The present work differs in three respects, which are stated here so that the novelty is not misread: (i) the estimand is the six-dimensional orbital state of the emitter with the receivers at known positions—the inverse of the receiver-localization problem—so that the satellite is treated exactly as a moving emitter in direct geolocation; (ii) the estimation is single-step, from the received samples, with the orbital dynamics built into the parameterization rather than imposed on extracted observables; and (iii) the Cramér–Rao bound is derived in a form that separates signal, geometry and dynamics, with the dynamics entering through a closed-form Keplerian state-transition block. The first two are what make the low-SNR behavior different from that of the two-step chain; the third is what makes the accuracy predictable. Direct processing of raw data (the DPD principle) and the inversion of roles are the setting rather than the novelty in themselves; the principal contributions are items (ii) and (iii): the incorporation of the orbital dynamics into the single-step parameterization, and the resulting factorized orbital-state CRB with its closed-form dynamics term.

This paper reverses the roles of the direct-geolocation configuration of [4,5,16,17]: instead of a satellite with known ephemeris locating an unknown ground emitter, ground stations at known positions determine the unknown six-dimensional orbital state of a non-cooperative satellite emitter. The contributions are:

1. A single-step estimator that determines the orbital state directly from the received signals in which the orbital-dynamics manifold is intrinsic to the parameterization. Every candidate is a valid orbit, so the entire trajectory, rather than a sequence of per-epoch positions, is searched through six orbital parameters (plus nuisances where present) independent of the arc length and the number of samples with no constraint penalty (Section 3);
2. A closed-form Cramér–Rao bound for the orbital state, with the light-time correction included in the delay model, that factors into signal, geometry, and dynamics terms,

the dynamics entering as the two-body state-transition matrix (STM), with the Keplerian STM position block derived in closed form by implicit differentiation of the universal Kepler equation (Sections 4 and 5);

3. Numerical validation of the derivation and known-waveform estimator efficiency, an observability analysis with a nondimensionalized state over several station layouts and passes, and a model-based comparison of the acquisition (wrong-peak) thresholds of the direct estimator and a two-step baseline that commits to hard per-window peak decisions (Section 7).

The approach is distinct from angle-only non-cooperative orbit determination (which uses optical bearings and two-step filtering) and POD (cooperative two-step); it is the direct RF single-step counterpart aimed at the weak and uncooperative regime.

2. Signal and Geometry Model

Let the unknown orbital state at epoch t_0 be

$$\mathbf{x} = \begin{bmatrix} \mathbf{r}_0 \\ \mathbf{v}_0 \end{bmatrix} \in \mathbb{R}^6, \quad (1)$$

the satellite position $\mathbf{r}_0$ and velocity $\mathbf{v}_0$ at the epoch t_0 , expressed in an Earth-centered inertial (ECI) frame. The ECI realization used throughout is the simplest one that is adequate for a single pass of a few minutes: it accounts for Earth's rotation about its polar axis at the sidereal rate $\omega_\oplus$ and neglects precession, nutation and polar motion, whose effects over the observation intervals considered are far below the accuracies of interest. For processing real data with published ephemerides, the full transformation between the terrestrial (ITRF) and celestial (GCRF) frames with the IAU precession–nutation and polar-motion models would be used [2,3]. Both the satellite and the ground stations are expressed in this frame; a station fixed in Earth-centered Earth-fixed (ECEF) coordinates therefore has a time-varying ECI position $\mathbf{q}_m(t) = \mathbf{R}_\oplus(t) \mathbf{q}_m^{\text{ECEF}}$, where $\mathbf{R}_\oplus(t)$ is the rotation about the polar axis by the angle $\omega_\oplus(t - t_0)$, and a corresponding ECI velocity $\dot{\mathbf{q}}_m(t) = \omega_\oplus \times \mathbf{q}_m(t)$ of a few hundred m/s, which enters the Doppler through the relative range rate $\dot{\tau}_m = \frac{1}{c} \hat{\mathbf{u}}_m^\top (\mathbf{v}_s - \dot{\mathbf{q}}_m)$, with $\mathbf{v}_s(t)$ the satellite velocity and $\hat{\mathbf{u}}_m(t)$ the unit line-of-sight vector defined below, and is retained throughout. Under the adopted dynamical model the satellite position $\mathbf{p}_s(t)$ is a deterministic function of the epoch state (for two-body motion in closed form; for perturbed motion by numerical integration of the equations of motion, which is deterministic once the force model is fixed). Stochastic accelerations are not modeled here; when process noise is modeled, the appropriate bound is the posterior (Bayesian) CRB of nonlinear filtering [34,35],

$$\mathbf{p}_s(t) = \boldsymbol{\phi}(\mathbf{x}, t), \quad \boldsymbol{\phi}(\mathbf{x}, t_0) = \mathbf{r}_0. \quad (2)$$

Ground station $m = 1, \dots, M$ receives the baseband signal

$$y_m(t) = b_m s(t - \tau_m(t)) e^{-j2\pi f_c \tau_m(t)} + w_m(t), \quad w_m \sim \mathcal{CN}(0, N_0), \quad (3)$$

where $s(\cdot)$ is the (initially known) complex envelope, f_c the carrier, b_m a (constant) complex gain, and c the speed of light. For a non-cooperative emitter the emission epoch is not known to the receivers, so the envelope in (3) is more generally $s(t - \tau_m(t) - \delta t_{\text{tx}})$, with δt_{tx} an unknown time offset common to all stations. Like the absolute carrier phase absorbed in b_m , δt_{tx} is not separable from the station clock offsets (Section 4); with the stations synchronized to a common reference it is a single common nuisance, and the delay observable then carries only time-difference information across stations while the Doppler and Doppler rate are unaffected. For clarity of exposition δt_{tx} is suppressed

in (3)–(9) and restored as a nuisance in Section 4. In this constant-gain model the geometric dependence is represented through the time-varying propagation delay, defined implicitly by the light-time equation

$$\tau_m(t) = \frac{1}{c} \|\boldsymbol{\phi}(\mathbf{x}, t - \tau_m(t)) - \mathbf{q}_m(t)\|, \quad (4)$$

i.e., the signal received at t was emitted at the retarded epoch $t_e = t - \tau_m(t)$ from the satellite position at that epoch. The propagation time τ_m , termed the one-way light time in orbit determination [2,36] (propagation time has also been modeled explicitly in LEO Doppler positioning [25]) (as opposed to the two-way light time of active ranging), is about 2–3.5 ms at the ranges of the observed arc, during which the satellite moves 14–25 m, and Section 7 shows that evaluating the satellite position at the reception time instead would leave the bound essentially unchanged but bias the estimate by more than a hundred meters. For given $\mathbf{x}$ and t , (4) is an implicit equation for the delay $\tau_m(t)$; it is solved by fixed-point iteration, starting from the simultaneous-time value $\tau_m^{(0)} = \|\boldsymbol{\phi}(\mathbf{x}, t) - \mathbf{q}_m(t)\|/c$ and iterating $\tau_m^{(k+1)} = \|\boldsymbol{\phi}(\mathbf{x}, t - \tau_m^{(k)}) - \mathbf{q}_m(t)\|/c$. Because the iteration contracts by the factor $|\hat{\mathbf{u}}_m^\top \mathbf{v}_s|/c \lesssim 2.5 \times 10^{-5}$ per step, two iterations reach machine precision. Equation (4) is the entire coupling: raw signal $\rightarrow$ time-varying delay $\rightarrow$ orbit map $\rightarrow$ epoch state. Its time derivatives carry the Doppler ($\dot{\tau}_m$) and Doppler-rate ($\ddot{\tau}_m$), both large for LEO and both informative. For a non-cooperative emitter the waveform s may be unknown; the known-waveform case is developed first for clarity, and the unknown-signal (covariance/QML) generalization is given in Section 6. The principal symbols are collected in Table 1.

Table 1. Principal notation. Vectors are in the ECI frame of Section 2 unless stated otherwise.

Symbol	Meaning
$t_0; \mathbf{x} = [\mathbf{r}_0^\top, \mathbf{v}_0^\top]^\top$	epoch; orbital state at t_0 (satellite position $\mathbf{r}_0$ and velocity $\mathbf{v}_0$)
$\boldsymbol{\phi}(\mathbf{x}, t)$	orbit propagator: satellite position at time t generated by the epoch state $\mathbf{x}$
$\mathbf{p}_s(t), \mathbf{v}_s(t)$	satellite position and velocity at time t
$\mathbf{r}_0, \mathbf{v}_0, v_{r0}$	$\ \mathbf{r}_0\ $, $\ \mathbf{v}_0\ $, and radial velocity $\mathbf{r}_0^\top \mathbf{v}_0 / r_0$ at epoch (Section 5)
$\mu, \omega_\oplus$	Earth's gravitational parameter; Earth rotation rate
$\mathbf{q}_m^{\text{ECEF}}, \mathbf{q}_m(t), \dot{\mathbf{q}}_m(t)$	position of station m in ECEF; its ECI position and velocity
$\mathbf{R}_\oplus(t)$	rotation from ECEF to ECI about the polar axis by $\omega_\oplus(t - t_0)$
M	number of ground stations
$\tau_m(t), t_e$	propagation delay to station m at reception time t ; emission epoch $t_e = t - \tau_m(t)$
$\mathbf{d}_m, \hat{\mathbf{u}}_m$	line-of-sight vector $\boldsymbol{\phi}(\mathbf{x}, t_e) - \mathbf{q}_m(t)$ and its unit vector
$s(t), f_c, B$	complex envelope; carrier frequency; rms bandwidth
b_m, N_0	complex gain of station m ; noise power spectral density
δt_{tx}	unknown emission-time offset (common to all stations)
$\delta f_m, \hat{f}_{\text{tx}}$	identifiable per-station frequency offset; common transmitter drift
$\mathbf{J}, \mathbf{C}$	Fisher information matrix; Cramér–Rao bound $\mathbf{J}^{-1}$
$\boldsymbol{\Phi}(t, t_0), \boldsymbol{\Phi}_r(t, t_0)$	state-transition matrix $\partial(\mathbf{p}_s, \mathbf{v}_s)/\partial \mathbf{x}$ and its position block $\partial \mathbf{p}_s / \partial \mathbf{x}$
$\chi, \alpha, C(z), S(z)$	universal anomaly; reciprocal semi-major axis; Stumpff functions
ρ, ρ_s	pass SNR per station; per-sample SNR (ρ/N per raw sample, ρ/N_s in the surrogate (16))
N, N_s, T	raw samples per station; delay/Doppler samples of the numerical surrogate; arc duration
K, T_{sub}	number of sub-windows of the two-step baseline; sub-window duration T/K
$\mathbf{a}(\mathbf{x}), \mathbf{R}(\mathbf{x})$	manifold vector and spatial covariance of the unknown-waveform model (Section 6)

3. The Direct Estimator

Stacking the received data, the single-step maximum-likelihood (ML) estimate of the orbital state is obtained by concentrating out the unknown per-station complex gains b_m . Because each b_m is an independent nuisance, its ML value at a hypothesized state is $\hat{b}_m = z_m(\mathbf{x})/E_s$, where $z_m(\mathbf{x}) = \int y_m(t) \hat{s}^*(t - \tau_m(\mathbf{x}, t)) e^{+j2\pi f_c \tau_m(\mathbf{x}, t)} dt$ is the station- m matched-filter output, $\hat{s}$ the known (or reconstructed) waveform and E_s its energy. Substituting $\hat{b}_m$ back yields the

concentrated objective, which is an *incoherent* (non-coherent) combination across stations. The unknown per-station phases preclude coherent summation:

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x}} \frac{1}{E_s} \sum_{m=1}^M \left| \int y_m(t) \hat{s}^*(t - \tau_m(\mathbf{x}, t)) e^{+j2\pi f_c \tau_m(\mathbf{x}, t)} dt \right|^2. \quad (5)$$

(If the relative inter-station gains are known or calibrated, the gain-weighted coherent sum $|\sum_m b_m^* z_m(\mathbf{x})|^2$ is instead optimal; the conservative unknown-gain case is treated here.), with $\tau_m(\mathbf{x}, t)$ given by (4). The key structural point is that (5) is optimized over the six-dimensional epoch state $\mathbf{x}$, and, through (2), every point of the search space is a physically valid orbit. Unlike constrained ground-emitter DPD, where the Earth-surface constraint is imposed separately (and the bound is a constrained CRB [37,38]), here the orbital-dynamics manifold is built into the parameterization: the whole trajectory, i.e., $3N$ positions for N samples, is searched through the six epoch-state parameters regardless of the arc length, and no constraint penalty is required. (Nuisances that cannot be concentrated in closed form, such as the frequency offsets and the emission-time offset of Section 4, add dimensions to the search or are optimized in an inner loop.) Initialization from a coarse two-line element set or a single-pass coarse fix confines the search to a small neighborhood, after which Gauss–Newton refinement converges rapidly (Section 7). The complete procedure is summarized in Algorithm 1.

Algorithm 1 Direct Orbital-State Estimation

The estimator is implemented in the following steps.

1. *Inputs.* Sampled baseband data y_m of the M stations over the arc; station ECEF coordinates; a common time reference for the stations; the known or structured waveform s ; a prior epoch state $\bar{\mathbf{x}}$ (e.g., propagated from a two-line element set) with a search box of $\pm\Delta r$ per position component and $\pm\Delta v$ per velocity component.
 2. *Evaluation of the objective for a candidate $\mathbf{x}$.* Propagate $\phi(\mathbf{x}, t)$ over the arc (universal variables, Section 5); rotate the stations to ECI; solve the light-time Equation (4) for every station and sample by two fixed-point iterations; form the matched-filter outputs $z_m(\mathbf{x})$ with the delay and carrier phase compensated; return $\sum_m |z_m(\mathbf{x})|^2$. Nuisances that cannot be concentrated in closed form (δt_{tx} and, where present, δf_m , $\dot{f}_{\text{tx}}$) are maximized in an inner loop. One evaluation costs $O(MN)$ operations, N being the number of samples per station.
 3. *Acquisition.* Maximize the objective over the search box. An exhaustive search at the resolution of the likelihood is impractical (Section 7); the search is organized hierarchically, first by a Doppler-only search on short sub-windows that reduces the box and then over the full arc.
 4. *Refinement.* Starting from the acquired cell, apply Gauss–Newton (scoring) iterations to $(\mathbf{x}, \boldsymbol{\eta})$, with the Jacobian obtained from (8) and the closed-form Φ_r , until the position update is below 1 mm (at most 12 iterations).
 5. *Output.* The estimate $\hat{\mathbf{x}}$ and its predicted covariance, the Schur-complemented CRB of Section 4 evaluated at $\hat{\mathbf{x}}$.
-

The Monte Carlo experiments of Section 7 exercise step 4 (from a single start after acquisition), and the acquisition of step 3 enters through the failure model of the threshold comparison.

4. Cramér–Rao Bound

For the deterministic model $\mu_m(t; \mathbf{x}) = b_m s(t - \tau_m) e^{-j2\pi f_c \tau_m}$ in white $\mathcal{CN}(0, N_0)$ noise, independent across stations, the Fisher information matrix (FIM) is [39]

$$[J]_{ij} = \frac{2}{N_0} \operatorname{Re} \sum_{m=1}^M \int \left(\frac{\partial \mu_m}{\partial x_i} \right)^* \frac{\partial \mu_m}{\partial x_j} dt. \quad (6)$$

This is the conditional FIM for fixed (known) complex gains; the unknown-gain case is handled by the Schur complement of the augmented FIM (nuisance discussion below). The problem reduces to the sensitivity $\partial\mu_m/\partial\mathbf{x}$, which is built in two layers.

4.1. Layer 1: Signal Sensitivity to Its Own Delay

Differentiating μ_m with respect to the scalar delay,

$$g_m(t) \equiv \frac{\partial\mu_m}{\partial\tau_m} = b_m [-\dot{s}(t - \tau_m) - j2\pi f_c s(t - \tau_m)] e^{-j2\pi f_c \tau_m}. \quad (7)$$

The carrier term $-j2\pi f_c s$ provides phase-level range sensitivity; the envelope term $-\dot{s}$ provides group-delay (bandwidth) sensitivity. All signal dependence is contained in g_m .

4.2. Layer 2: Delay Sensitivity to the Orbital State

By the chain rule and (4), with $\mathbf{d}_m(t) = \boldsymbol{\phi}(\mathbf{x}, t_e) - \mathbf{q}_m(t)$ evaluated at the retarded epoch $t_e = t - \tau_m(t)$ and line-of-sight unit vector $\hat{\mathbf{u}}_m = \mathbf{d}_m/\|\mathbf{d}_m\|$, implicit differentiation of the light-time equation gives

$$\frac{\partial\tau_m}{\partial\mathbf{x}} = \frac{1}{c} \frac{\hat{\mathbf{u}}_m(t)^\top \boldsymbol{\Phi}_r(t_e, t_0)}{1 + \hat{\mathbf{u}}_m(t)^\top \mathbf{v}_s(t_e)/c} \approx \frac{1}{c} \hat{\mathbf{u}}_m(t)^\top \boldsymbol{\Phi}_r(t_e, t_0), \quad (8)$$

where $\boldsymbol{\Phi}_r(t_e, t_0) \in \mathbb{R}^{3 \times 6}$ is the position block of the orbital state-transition matrix $\boldsymbol{\Phi}(t, t_0) = \partial\boldsymbol{\phi}/\partial\mathbf{x}$, and the denominator, which differs from unity by $|\hat{\mathbf{u}}_m^\top \mathbf{v}_s|/c \lesssim 2.5 \times 10^{-5}$ for LEO, is retained in the code and dropped in the expressions below for legibility; it changes the bound by less than 10^{-4} relative. This is the step that distinguishes the bound from a static-emitter DPD CRB: the line-of-sight sensitivity is projected through the orbital dynamics, so the bound depends on how the entire trajectory responds to a perturbation of the epoch state, not merely instantaneous position.

4.3. Assembled FIM

Combining (7) and (8), $\partial\mu_m/\partial\mathbf{x} = c^{-1}g_m(t) \hat{\mathbf{u}}_m(t)^\top \boldsymbol{\Phi}_r(t_e, t_0)$ with $t_e = t - \tau_m(t)$, and (6) becomes

$$\mathbf{J} = \frac{2}{N_0 c^2} \text{Re} \sum_{m=1}^M \int \boldsymbol{\Phi}_r^\top(t_e, t_0) \hat{\mathbf{u}}_m(t) |g_m(t)|^2 \hat{\mathbf{u}}_m(t)^\top \boldsymbol{\Phi}_r(t_e, t_0) dt \quad (9)$$

The integrand factors into three physically distinct contributors:

- *Signal*: $|g_m(t)|^2 = |b_m|^2 [(2\pi f_c)^2 |s|^2 + |\dot{s}|^2 + \text{cross}]$, carrying SNR, carrier, and bandwidth—how informative the waveform is per instant;
- *Geometry*: The rank-one dyad $\hat{\mathbf{u}}_m \hat{\mathbf{u}}_m^\top$, summed over stations and time, is the classical dilution-of-precision structure;
- *Dynamics*: $\boldsymbol{\Phi}_r(t, t_0)$, mapping instantaneous position sensitivity back to the epoch state.

The CRB is $\mathbf{C} = \mathbf{J}^{-1}$; position and velocity RMS bounds are $\sqrt{\text{tr}(\mathbf{C}_{rr})}$ and $\sqrt{\text{tr}(\mathbf{C}_{vv})}$. The one-line summary: the orbital-state CRB is the multi-station delay CRB projected through the orbital state-transition matrix.

4.4. Nuisance Parameters

Unknown complex gains $\{b_m\}$ and station clock offsets are handled by augmenting $\mathbf{x}$ and partitioning the FIM, $\mathbf{J} = \begin{bmatrix} \mathbf{J}_{xx} & \mathbf{J}_{x\eta} \\ \mathbf{J}_{\eta x} & \mathbf{J}_{\eta\eta} \end{bmatrix}$, with $\boldsymbol{\eta}$ the nuisances. The effective orbital-state bound is the Schur complement $\mathbf{C}_{xx} = (\mathbf{J}_{xx} - \mathbf{J}_{x\eta} \mathbf{J}_{\eta\eta}^{-1} \mathbf{J}_{\eta x})^{-1}$, which quantifies the accuracy loss due to the nuisances. Four classes matter for a non-cooperative satellite, and they enter the two observables differently.

Constant complex gains b_m affect neither delay nor Doppler geometry and are separable per station. The gain b_m is taken constant over the pass; in practice it varies with range, antenna pattern, attitude, and propagation, but, over a single short pass, this variation is neglected in the present analysis, and its impact is left for evaluation as part of the experimental extension.

Station clock errors must be split: a constant offset δt_m contaminates the delay observable, while a clock drift $\dot{\delta t}_m$ contaminates both delay (linearly in t) and, through its rate, the Doppler. Both are per-station and enter as additive terms in the respective observable gradients.

Emission epoch. A non-cooperative emitter does not report its transmission time, so each delay also contains the unknown common offset δt_{tx} of Section 2. As with the frequencies, δt_{tx} and the station clock offsets δt_m appear only through their sums, so they are not separately estimable. If the stations are synchronized to a common reference (e.g., GNSS-disciplined), δt_m is known and δt_{tx} is a single common nuisance: the delays then contribute time-difference (TDOA) information, and absolute range is lost. If the station clocks are independently unknown, each delay carries its own unknown offset and only its variation over the pass—integrated Doppler—remains, so the delay observable adds essentially no information beyond the Doppler. The results of this paper assume synchronized stations; δt_{tx} enters η and is removed by the Schur complement in the same way as the frequency nuisances. Its effect is largest for small M , where the $M - 1$ independent time differences are fewest. For the baseline geometry at $\rho = 15$ dB it raises the three-station position bound from 21 m (emission epoch known) to 41 m, whereas with four stations the bound is 7.0 m in both cases. When the frequency nuisances below are also estimated, the additional cost of δt_{tx} is small (about $76 \rightarrow 78$ m at $M = 3$) because the Doppler information that remains after the frequency nuisances are removed is then the dominant contribution. All the numerical results of Section 7 include δt_{tx} .

Carrier-frequency errors are, for a non-cooperative emitter, the most consequential nuisance and were omitted from earlier treatments. The received Doppler is not $-f_c \dot{\tau}_m(t)$ alone; it also contains the unknown satellite transmit-frequency offset Δf_{tx} (common to all stations), the transmit oscillator drift $\dot{f}_{tx}$ (common, linear in t), and a per-station receiver offset Δf_m . A subtlety of identifiability arises: the common offset Δf_{tx} and the per-station offsets Δf_m appear only through their sum in each station's Doppler, so they are not separately estimable; the identifiable set is obtained by eliminating Δf_{tx} and defining the combined per-station offset $\delta f_m = \Delta f_{tx} + \Delta f_m$ so that the estimated frequency nuisances are $[\delta f_1, \dots, \delta f_M, \dot{f}_{tx}]$. This gauge choice makes the nuisance Fisher block nonsingular by construction rather than relying on an implicit constraint. Appending these to η and taking the Schur complement gives the effective orbital-state CRB in the presence of unknown frequencies. The numerical examples of this subsection (Figures 1–3) use the scenario of Section 7: orbit, station coordinates, carrier, bandwidth, arc and observable model are listed there, and the same script produces all the figures. For the baseline geometry, with δt_{tx} estimated, this inflates the position CRB by a constant factor of about 1.9 ($41 \rightarrow 78$ m at 15 dB) across the SNR range (Figure 1): the frequency nuisances tax a fixed fraction of the Doppler information rather than introducing an SNR-dependent floor. Equivalently, a differential formulation—referencing each station's Doppler to a common one—removes the common terms at the cost of one observable.

The observability and frequency-nuisance analyses can be combined: does the minimum station count change when the identifiable frequency nuisances are estimated jointly with the orbital state? Figure 2 answers this by plotting the effective position CRB versus M with and without the frequency nuisances (Schur-complemented). The minimum useful count is unchanged— $M = 3$ still yields a finite fix (with the conditioning quantified in Section 7)—but with all nuisances estimated the bound is large at small M : 959 m at

$M = 2$, 78 m at $M = 3$, and then 17 m at $M = 4$ and only slowly lower beyond (14 and 12 m at $M = 5, 6$). The penalty factor itself is largest at $M = 2$ ($\approx 7\times$) and lies between $1.9\times$ and $2.6\times$ for $M = 3\text{--}6$; the fourth station, which helps to separate the common oscillator drift and per-station offsets from the orbital Doppler, gives the largest absolute gain. The practical implication for the tested regional geometry and pass is that three stations give a finite orbital fix while four or more provide robustness to both geometric conditioning and oscillator (frequency-nuisance) uncertainty; Section 7 shows that the station layout matters as much as the count.




Figure 1. Effect of unknown carrier-frequency offset and drift on the orbital-state CRB. Estimating the identifiable frequency nuisances (per-station offsets and a common drift) jointly with the orbital state via the Schur complement inflates the position CRB by a constant factor (≈ 1.9) that is independent of SNR. The emission-time offset δt_{tx} is estimated in both curves.




Figure 2. Effective position CRB versus number of stations with frequencies known and the identifiable frequency nuisances estimated (Schur complement). The minimum useful count is unchanged at $M = 3$; the nuisance penalty (annotated) is largest at $M = 2$ ($\approx 7\times$), and the fourth station gives the largest absolute gain (78 $\rightarrow$ 17 m). The emission-time offset is estimated in both curves.

Finally, the equal-SNR assumption is favorable, so range-dependent free-space weighting is also tested. Applying $|b_m|^2 \propto 1/R_m^2(t)$, with $R_m(t) = \|\mathbf{d}_m(t)\|$ the station-to-satellite range, at fixed total received energy (which isolates the effect of redistributing energy across stations and does not test the absolute path-loss penalty), Figure 3 shows the station-count curve is almost unchanged (e.g., 41 $\rightarrow$ 45 m at $M = 3$) because the stations lie at comparable ranges to a 543 km orbit, so $1/R_m^2$ varies little across the network. The station-count result is therefore only weakly sensitive to the redistribution of a fixed total received energy according to free-space range loss, and this holds even when range weighting and the frequency nuisances are combined ($M = 3$ remains finite, $M \geq 4$ robust).




Figure 3. Robustness of the station-count result to range-dependent free-space station weighting. Range-dependent SNR ($|b_m|^2 \propto 1/R_m^2$, fixed total energy) barely shifts the equal-SNR curve; the three/four-station conclusion survives, and $M = 3$ remains usable even with path loss and frequency nuisances combined.

5. Keplerian State-Transition Position Block

The dynamics term Φ_r admits a closed form for two-body motion. Universal-variable formulations of Keplerian motion and its state-transition matrix are classical [40,41]; the derivation below obtains the position block directly by implicit differentiation of the universal Kepler equation in a form that reuses the quantities already computed during propagation. Using the Lagrange coefficients, the position is expressed as $\mathbf{r}(t) = f \mathbf{r}_0 + g \mathbf{v}_0$, with coefficients f, g that depend on the epoch state, so

$$\frac{\partial \mathbf{r}}{\partial \mathbf{r}_0} = f \mathbf{I}_3 + \mathbf{r}_0 \partial_{\mathbf{r}_0} f + \mathbf{v}_0 \partial_{\mathbf{r}_0} g, \quad (10)$$

$$\frac{\partial \mathbf{r}}{\partial \mathbf{v}_0} = g \mathbf{I}_3 + \mathbf{r}_0 \partial_{\mathbf{v}_0} f + \mathbf{v}_0 \partial_{\mathbf{v}_0} g, \quad (11)$$

and $\Phi_r = [\partial \mathbf{r} / \partial \mathbf{r}_0 \mid \partial \mathbf{r} / \partial \mathbf{v}_0]$. In universal variables with anomaly χ , epoch radius $r_0 = \|\mathbf{r}_0\|$, epoch speed $v_0 = \|\mathbf{v}_0\|$, epoch radial velocity $v_{r0} = \mathbf{r}_0^\top \mathbf{v}_0 / r_0$, reciprocal semi-major axis $\alpha = 2/r_0 - v_0^2/\mu$, argument $z = \alpha \chi^2$, and Stumpff functions $C(z), S(z)$,

$$f = 1 - \frac{\chi^2}{r_0} C(z), \quad g = \Delta t - \frac{\chi^3}{\sqrt{\mu}} S(z). \quad (12)$$

The scalars f, g depend on the epoch state through the primitives $\{r_0, v_{r0}, \alpha\}$ and χ , which is defined implicitly by the universal Kepler equation $\mathcal{K}(\chi; r_0, v_{r0}, \alpha, \Delta t) = 0$. Implicit differentiation gives

$$\frac{\partial \chi}{\partial \mathbf{x}} = -\frac{1}{r(t)} \left(\frac{\partial \mathcal{K}}{\partial r_0} \frac{\partial r_0}{\partial \mathbf{x}} + \frac{\partial \mathcal{K}}{\partial v_{r0}} \frac{\partial v_{r0}}{\partial \mathbf{x}} + \frac{\partial \mathcal{K}}{\partial \alpha} \frac{\partial \alpha}{\partial \mathbf{x}} \right), \quad (13)$$

where the Jacobian denominator collapses to the instantaneous radius, $\partial \mathcal{K} / \partial \chi = r(t)$ —the defining property $dt/d\chi = r/\sqrt{\mu}$ of the universal anomaly. The primitive gradients are elementary,

$$\frac{\partial r_0}{\partial \mathbf{x}} = [\hat{\mathbf{r}}_0^\top \quad \mathbf{0}^\top], \quad \frac{\partial \alpha}{\partial \mathbf{x}} = \left[-\frac{2}{r_0^2} \hat{\mathbf{r}}_0^\top \quad -\frac{2}{\mu} \mathbf{v}_0^\top \right], \quad \frac{\partial v_{r0}}{\partial \mathbf{x}} = \left[\frac{\mathbf{v}_0^\top}{r_0} - \frac{\mathbf{r}_0^\top \mathbf{v}_0}{r_0^3} \mathbf{r}_0^\top \quad \frac{\mathbf{r}_0^\top}{r_0} \right], \quad (14)$$

and the coefficient gradients follow with $\partial z / \partial \mathbf{x} = \chi^2 \partial \alpha / \partial \mathbf{x} + 2\alpha \chi \partial \chi / \partial \mathbf{x}$ and the Stumpff derivatives $C'(z) = \frac{1}{2z}(1 - zS - 2C)$, $S'(z) = \frac{1}{2z}(C - 3S)$ (series near $z = 0$). Substituting yields Φ_r in closed form, every term a function of the epoch state, χ , and the Stumpff functions—quantities already computed during propagation. For J_2 or higher perturbations, Φ_r is obtained by integrating the variational equations; the CRB structure (9) is unchanged. Classical or equinoctial elements $\mathbf{e}_0$ are accommodated by the chain rule, $\partial \mathbf{r}(t) / \partial \mathbf{e}_0 = \Phi_r \partial \mathbf{x}_0 / \partial \mathbf{e}_0$, i.e., by post-multiplying Φ_r with the Jacobian of the Cartesian epoch state with respect to the elements.

6. Unknown-Waveform Case

When s is unknown, the ML estimator (5) is replaced by its quasi-ML/covariance form: the cost is the largest eigenvalue of the cross-station correlation matrix parameterized by $\mathbf{x}$, exactly as in delay-Doppler DPD for unknown signals [4], with the fixed emitter position replaced by the orbital state. It remains useful to establish the corresponding CRB for the unknown-waveform case since it is the conservative non-cooperative model; the estimator-efficiency demonstration of Section 7 uses the known-waveform concentrated ML, and the practical regime of the unknown-waveform case is quantified at the end of this section.

6.1. Proposition (Stochastic-Signal FIM)

Model the signal per frequency bin: in each bin the unknown-waveform coefficient is a zero-mean complex-Gaussian scalar of power σ_s^2 so that the per-bin spatial covariance is rank-one, $\mathbf{R}(\mathbf{x}) = \sigma_s^2 \mathbf{a}(\mathbf{x}) \mathbf{a}(\mathbf{x})^H + N_0 \mathbf{I}$, and the total information is the sum over bins, where $\mathbf{a}(\mathbf{x})$ is the (frequency-dependent) manifold vector whose parameter dependence is exactly that of the deterministic model. Because the orbital manifold varies over the pass, the snapshots do not share a common $\mathbf{a}(\mathbf{x})$: snapshot ℓ in bin n has its own manifold vector $\mathbf{a}_{n\ell}(\mathbf{x})$, evaluated at the time of that snapshot, and the Slepian-Bangs formula gives the total information as the sum over bins and snapshots, $[J]_{ij} = \sum_n \sum_{\ell=1}^{L_n} \text{tr}(\mathbf{R}_{n\ell}^{-1} \partial_i \mathbf{R}_{n\ell} \mathbf{R}_{n\ell}^{-1} \partial_j \mathbf{R}_{n\ell})$ with $\partial_i \mathbf{R}_{n\ell} = \sigma_s^2 (\partial_i \mathbf{a}_{n\ell} \mathbf{a}_{n\ell}^H + \mathbf{a}_{n\ell} \partial_i \mathbf{a}_{n\ell}^H)$. (When the manifold is constant over L snapshots the sum collapses to the familiar factor L ; in the time-varying case the factor is absorbed in the sum, which is why it does not appear explicitly in (15), a per-snapshot expression.) The manifold has one unit-modulus entry per station so that its norm $\|\mathbf{a}(\mathbf{x})\|^2 = M$ is constant in the parameter. Under this convention the rank-one signal covariance reduces, after applying the matrix-inversion lemma, to

$$[J]_{ij} = \frac{2\gamma^2}{M(1+\gamma)} \text{Re}\{\partial_i \mathbf{a}^H \mathbf{P}_a^\perp \partial_j \mathbf{a}\}, \quad \gamma = \frac{\sigma_s^2}{N_0} \|\mathbf{a}\|^2, \quad (15)$$

where $\mathbf{P}_a^\perp = \mathbf{I} - \mathbf{a}\mathbf{a}^H / \|\mathbf{a}\|^2$ and γ is the effective array SNR (the per-sample SNR $\rho_s = \sigma_s^2 / N_0$ scaled by the manifold norm, $\gamma = M\rho_s$). The manifold norm is taken constant in the parameter, so no additional amplitude-sensitivity term arises. This is identical in structure to the deterministic FIM (9), the same geometry and dynamics quadratic form $\partial_i \mathbf{a}^H \mathbf{P}_a^\perp \partial_j \mathbf{a}$, which, through (8), carries the $\hat{\mathbf{u}}_m^\top \Phi_r$ factor with an additional stochastic-signal information factor $\gamma / (1 + \gamma)$ relative to the corresponding deterministic unknown-amplitude model. That is, relative to the deterministic FIM with the same projected manifold quadratic form $J_{\text{det}} = (2\gamma / M) \text{Re}\{\partial_i \mathbf{a}^H \mathbf{P}_a^\perp \partial_j \mathbf{a}\}$ (i.e., after the unknown source amplitude has been concentrated out, which is what produces the projector $\mathbf{P}_a^\perp$), the stochastic FIM is reduced by the factor $J_{\text{stoch}} = \frac{\gamma}{1 + \gamma} J_{\text{det}}$, which approaches unity at high effective array SNR. This relationship refers to the projected unknown-amplitude deterministic FIM, not the conditional known-gain FIM of (9). (Equation (15) was verified numerically against the deterministic FIM; the ratio matches $\gamma / (1 + \gamma)$ to four significant figures across $\gamma \in [-10, 20]$ dB.)

6.2. Practical Regime of the Unknown-Waveform Case

The factor $\gamma / (1 + \gamma)$ is innocuous only when the *per-snapshot* array SNR γ is not small, and, for a wideband long capture, it is very small: with $N = BT$ samples per station and pass SNR ρ , the per-sample SNR is $\rho_s = \rho / N$ and $\gamma = M\rho_s$. For the scenario of Section 7 ($B = 2$ MHz, $T = 60$ s, $N \approx 1.2 \times 10^8$, $M = 3$) a pass SNR of 15 dB corresponds to $\rho_s \approx -66$ dB and $\gamma \approx 8 \times 10^{-7}$, so the unknown-waveform bound is about 10^3 times the deterministic one, and cross-station correlation of a Gaussian waveform is above its detection threshold only when $N\rho_s^2 \gg 1$, i.e., for pass SNRs of the order of 45–50 dB. The unknown-waveform direct estimator is therefore viable for strong narrowband or short captures (high ρ_s), whereas, at low-pass SNRs where the single-step advantage of Section 7 is predicted, the waveform must be known or at least structured (a known modulation format, a pilot, or a cyclostationary signature), which is the usual situation for a non-cooperative but standard-compliant downlink. The low-SNR results of this paper are accordingly stated for the known-waveform model, and the extension of the efficiency demonstration to structured-waveform estimators is left to future work.

The numerical non-cooperative signal model retains delay/Doppler geometry and does not exploit carrier-phase ranging. Although the signal model (3) contains the carrier-phase term $e^{-j2\pi f_c \tau_m(t)}$, the unknown constant complex gain b_m absorbs the absolute per-station carrier phase; the estimator therefore exploits the *time-varying* phase evolution (Doppler and Doppler rate), not absolute carrier-phase ranging. This is what the “delay/Doppler” description of the numerical observable model (Section 7) refers to.

7. Numerical Results

Table 2 collects the modeling assumptions, which are layered: a core idealized model, the two signal (waveform) models, the nuisance model, and the numerical scenario.

Table 3 maps each numerical experiment to the signal model, the nuisance assumptions and the bound against which it is compared; the observing geometry common to all the experiments is shown in Figure 4.

Table 2. Layered modeling assumptions.

Layer	Assumptions
Core idealized model	free-space propagation; white Gaussian noise; two-body (Keplerian) orbit; known station positions
Known-waveform model	transmitted waveform known (deterministic)
Unknown-waveform model	zero-mean complex-Gaussian signal; geometric information via delay and Doppler; no carrier phase
Nuisance model	per-station carrier offset δf_m ; common transmitter drift $\dot{f}_{tx}$; common emission-time offset δt_{tx} (stations synchronized; station clock offsets and drift discussed in Section 4)
Station geometry	known station coordinates; Earth rotation included ($q_m(t)$, $\dot{q}_m$ in ECI); timing/frequency errors as nuisances
Numerical scenario	$M = 3$ stations; 543 km, 65° orbit; $f_c = 2.2$ GHz; $B = 2$ MHz rms; 60 s arc; equal per-station SNR; Doppler window rms duration $T_d = 1$ s

Table 3. Signal model, nuisance assumptions and bound used in each numerical experiment. “Deterministic” denotes the delay/Doppler surrogate of (9) defined in Section 7.

Experiment	Waveform/Gains	Nuisances	Bound
STM validation (Figure 5)	—	—	analytic vs. finite difference
CRB vs. SNR, efficiency (Figure 6)	known; unknown constant gains	δt_{tx} (frequencies known)	deterministic (9), gains concentrated
Arc length (Figure 7)	known; unknown constant gains	δt_{tx}	deterministic
Frequency nuisances (Figures 1 and 2)	known; unknown constant gains	δt_{tx} , δf_m , $\dot{f}_{tx}$ (Schur)	deterministic, Schur-complemented
Path loss (Figure 3)	known; range-weighted gains	$\delta t_{tx}/+$ frequencies	deterministic
Stations, observability (Figures 8–10)	known	δt_{tx}	deterministic; scaled FIM spectrum
Light time (Section 7)	known	δt_{tx}	deterministic, retarded vs. simultaneous
Two-step comparison (Figure 11)	known	δt_{tx}	deterministic
Unknown waveform (Section 6)	Gaussian, unknown	none	stochastic (15), analytic only

7.1. Scenario

The observing geometry is shown in Figure 4. A near-circular LEO orbit at 543 km altitude and inclination 65° is observed by $M = 3$ ground stations on Earth’s surface, at $(40^\circ, 0^\circ)$, $(41.5^\circ, 3^\circ)$, and $(38^\circ, -2.5^\circ)$ latitude/longitude, forming a regional network with inter-station baselines of 303 km, 310 km, and 611 km. The satellite transits the network over a 120 s window during which all three stations observe it above a 10° elevation mask (elevations from 22° to 88°) with complementary elevation histories (Figure 4b): each station reaches peak elevation at a different time, providing the line-of-sight diversity that conditions the estimation problem. All the results of this section were produced by one script from this pass (sub-satellite point 34° N, 4° W at the epoch $t_0 = 0$, heading along the 65° -inclination track, 60 s arc starting at $t = 30$ s) with the light-time delay model. The carrier is $f_c = 2.2$ GHz and the rms bandwidth $B = 2$ MHz; the pass-arc duration is 60 s unless varied. The scenario data of Figures 5–7 are those of this section: the same orbit (543 km, 65°), stations, carrier, bandwidth and arc; the STM validation of Figure 5 uses the same epoch state propagated over $\Delta t = 300$ s. All the numerical results include Earth rotation: station positions and velocities are transformed into the ECI frame as functions of time, and the station velocity is retained in the Doppler. Unless otherwise stated, the numerical CRBs below use the ideal frequency-known model; the carrier-frequency nuisance penalty is reported separately in Figure 1. All the results assume synchronized stations and estimate the common emission-time offset δt_{tx} of Section 4 jointly with the state (Schur complement for the bounds, an additional unknown in the estimators), so the delays contribute time-difference information. The numerical observable model is a delay/Doppler surrogate for a known or sufficiently structured non-cooperative

waveform. The unstructured unknown-Gaussian-waveform case of Section 6 is treated analytically only.




Figure 4. Observing geometry. (a) Earth-centered inertial (ECI) view of the LEO pass (solid), the three ground stations (triangles, in the station colors of panel (b)), and the station-to-satellite lines of sight at mid-pass (dashed). (b) Elevation angle of the satellite from each station over the pass; the staggered peaks provide the geometric diversity that conditions the Fisher information. Inter-station baselines are 303–611 km.

7.2. Numerical Observable Model

The numerical bounds are computed for a windowed delay/Doppler surrogate of (9). Each station's pass is divided into N_s samples; at each sample the delay and the Doppler are treated as independent Gaussian observables with per-sample information $2\rho_s(2\pi B)^2$ and $2\rho_s(2\pi T_d)^2$, i.e., variances

$$\sigma_\tau^2 = \frac{1}{2\rho_s(2\pi B)^2}, \quad \sigma_f^2 = \frac{1}{2\rho_s(2\pi T_d)^2}, \quad (16)$$

where $\rho_s = \rho/N_s$ is the per-sample SNR, B the rms bandwidth and T_d the effective rms duration of the window that produces each Doppler observable (T_d plays the role of B for frequency; σ_f is in Hz). The value $T_d = 1$ s is used throughout. Because the per-sample information scales as $1/N_s$, the total information does not depend on N_s . The surrogate treats the Doppler of successive windows as independent and so does not exploit phase continuity across windows; it is therefore conservative relative to the full-data information (9), which retains the coherent phase evolution over the arc.

7.3. STM Validation

Figure 5 validates the closed-form Keplerian state-transition position block of Section 5. Each of the six abscissae is one component of the epoch state x : indices 1–3 are the epoch position r_0 and indices 4–6 the epoch velocity v_0 . The ordinate is the Euclidean norm of the corresponding column of $\Phi_r(t, t_0)$, i.e., $\|\partial p_s(t)/\partial x_{0,i}\|$ —the sensitivity of the satellite position at time t to a unit perturbation of that epoch-state component. Two features are physically meaningful and both are reproduced exactly by the analytic expression: the velocity-sensitivity block (indices 4–6) has a characteristic scale of order Δt (here 300 s), as expected from first-order kinematic propagation, whereas the position-sensitivity block (indices 1–3) is $O(1)$. This comparison is unit-dependent (the position-state sensitivity $\partial r/\partial r_0$ is dimensionless, while the velocity-state sensitivity $\partial r/\partial v_0$ has units of time), so the raw column magnitudes should not be compared without normalization by representative

state scales. Over the propagation interval (here $\Delta t = 300$ s), the velocity-state sensitivity scales roughly as Δt . The analytic f, g derivation and an independent central finite difference of the propagated position agree to a relative error $\|\Phi_r^{\text{an}} - \Phi_r^{\text{fd}}\|_F / \|\Phi_r^{\text{fd}}\|_F \sim 10^{-9}$, the finite difference being taken with fixed steps of 1 m in each epoch-position component and 10^{-3} m/s in each epoch-velocity component, the propagation itself being the closed-form universal-variable solution (no ODE solver is involved; MATLAB R2025b, MathWorks, Natick, MA, USA), providing a numerical verification of the closed-form result. The finite-difference check requires no orbital-mechanics software and is the recommended verification before the bound is trusted.




Figure 5. Analytic vs. finite-difference position block of the Keplerian state-transition matrix; column norms coincide to relative error $\sim 10^{-9}$ (finite-difference steps of 1 m and 10^{-3} m/s).

7.4. CRB and Efficiency

Figure 6 shows the position (panel a) and velocity (panel b) CRB versus pass SNR on separate axes. Here the pass SNR ρ is defined as the total received signal-energy-to-noise ratio per station, integrated over the observation interval; per-sample SNR is ρ/N_s for N_s samples. The stations are taken to have equal SNRs (identical $|b_m|$) in the numerical scenario; this is a favorable choice since the stations span different ranges and elevation histories. A range-dependent (free-space path-loss) SNR would weight the stations unequally; its effect on the station-count result is quantified in Figure 3 and found to be small. Both bounds fall as $\text{SNR}^{-1/2}$ (an exact -1 dB/dB slope): 41 m and 0.15 m/s at $\rho = 15$ dB, reaching meter-level position and cm/s velocity accuracy at the upper end of the range for a single pass. Because the two bounds share this slope, their ratio is *constant in SNR* (here $\sigma_r/\sigma_v \approx 280$ s, an effective dynamic timescale set by the geometry and propagation interval); they are plotted separately so this fixed offset is not mistaken for coincidence. Monte Carlo runs of the direct estimator (Gauss–Newton on the delay/Doppler geometric parameters, circles and squares) lie on both bounds. This estimator is the ML estimator under the Gaussian delay/Doppler surrogate of the known-waveform model, used here as a tractable efficiency demonstration. Section 6 gives the corresponding stochastic (unknown-waveform) FIM, which differs from the corresponding projected deterministic FIM by the scalar $\gamma/(1 + \gamma)$ (a relation between those two manifold models, not a correction to be applied to the surrogate bound plotted here); demonstrating QML efficiency against that bound is left for a dedicated unknown-waveform Monte Carlo study.




Figure 6. Orbital-state CRB versus pass SNR, shown as (a) position and (b) velocity on separate axes. The direct known-waveform estimator, simulated after acquisition under the Gaussian delay/Doppler surrogate (Monte Carlo markers), attains both bounds. The ratio $\sigma_r/\sigma_v \approx 280$ s (an effective dynamic timescale) is constant in SNR; the panels are separated so this fixed offset is not read as overlap.

7.5. Arc Length

Figure 7 shows the position CRB versus pass-arc duration. Accuracy improves rapidly with the observed orbital arc: a longer arc both sweeps the line of sight over more directions and exposes more of the orbit's dynamic response, tightening Φ_r in (9). To separate the sources of this improvement, Figure 7 plots the position CRB versus arc under two normalizations: fixed total pass energy (holding ρ constant as T grows) and fixed received power (letting the integrated energy grow with T). The fixed-power curve improves steadily, as expected from the added integration gain. The fixed-energy curve is the more interesting one: with no additional received energy the bound still falls from 80 m at 10 s to 47 m at 30 s, 42 m at 45 s and 38 m at 90 s because lengthening the observation changes the line-of-sight geometry and the state-transition sensitivity; most of the gain is obtained within the first 30 s, after which the improvement is slow. Thus, at fixed total received energy, orbital motion increases the information about the epoch state through the time evolution of the observation manifold $\hat{\mathbf{u}}_m^\top(t) \Phi_r(t, t_0)$ —a feature of a moving orbital emitter that is absent from static-emitter DPD—but the gain is bounded and geometry-dependent; the integration gain of the fixed-power curve is what makes long arcs pay.




Figure 7. Position CRB versus pass-arc duration under two normalizations: fixed total pass energy ($\rho = 15$ dB held constant as T grows) and fixed received power (ρ growing with T). The fixed-energy curve isolates the information gained purely from changing geometry and dynamics, most of it within the first 30 s; the fixed-power curve adds the ordinary integration gain.

7.6. Station Geometry

A single station yields a numerically rank-deficient Fisher information for the measurement model, arc and geometry studied here: the geometry factor $\hat{\mathbf{u}}_m \hat{\mathbf{u}}_m^\top$ in (9) is then a sum of nearly collinear rank-one dyads whose time-integrated contribution produces an effectively rank-deficient Fisher information for the tested geometry. Each additional suitably separated station—observing the pass from a distinct line-of-sight history (Figure 4b)—fills this factor and conditions $\mathbf{J}$, sharply reducing the bound. In the three-station geometry of Figure 4 the Fisher information has full rank throughout the pass; its conditioning is quantified below.

Figure 8 quantifies this dependence. Panel (a) shows the position CRB and the scaled FIM condition number as stations are added from a regional pool of six; panel (b) shows the CRB versus SNR for each station count. Two regimes appear. With one or two stations the Fisher information is rank-deficient or nearly so and the state is effectively unobservable, so each added station lowers the bound by orders of magnitude—an observability-limited regime driven by geometric diversity, not SNR. Beyond about four stations further stations give the ordinary $\mathcal{O}(M^{-1/2})$ improvement—an SNR-limited regime. Three stations produce a full-rank FIM in this geometry; four or more improve the conditioning. Because the state mixes position and velocity, the condition number and the singular values of $\mathbf{J}$ in SI units depend on the choice of units; the observability statements below are therefore made on the nondimensionalized information $\tilde{\mathbf{J}} = \mathbf{D}\mathbf{J}\mathbf{D}$, $\mathbf{D} = \text{diag}(a \mathbf{I}_3, v_{\text{circ}} \mathbf{I}_3)$ with a the semi-major axis and v_{circ} the circular speed as characteristic scales (a Jacobi scaling by $\text{diag}(\mathbf{J})^{-1/2}$ gives the same qualitative picture), and exact rank deficiency is distinguished from numerical conditioning. The transition identifies the minimum network size that yields an observable well-conditioned orbital fix for the tested geometry.




Figure 8. Accuracy vs. number of ground stations. (a) Position CRB (left axis) and condition number of the nondimensionalized information (right axis) vs. M at $\rho = 15$ dB; shaded bands mark the observability-limited ($M \lesssim 4$) and SNR-limited ($M \gtrsim 4$) regimes. (b) Position CRB vs. SNR for $M = 2, 3, 4$; a single station leaves the state unobservable.

Because the line-of-sight direction rotates during a pass, rank deficiency cannot be inferred from the instantaneous geometry factor $\hat{\mathbf{u}}_m \hat{\mathbf{u}}_m^\top$ alone; it must be established from the time-integrated Fisher information. Figure 9a shows the six singular values of the scaled information $\tilde{\mathbf{J}}$ for $M = 1, \dots, 4$: a single station leaves two modes far below the numerical floor (numerical rank deficiency for this measurement model, arc and geometry; the statement is not a generic single-station Doppler result), two stations leave one marginal

mode (smallest normalized singular value $\sim 10^{-7}$), and three stations give full rank with a scaled condition number of about 5×10^5 , reduced by more than an order of magnitude (to 3×10^4) with a fourth station. Figure 9b confirms this is not a short-arc artifact: the smallest scaled singular value for $M = 1$ remains below 10^{-12} across the full range of pass durations, so the rotation of the line of sight does not by itself restore single-station observability; for $M = 2$ and 3 it grows with the arc.




Figure 9. Observability from the time-integrated Fisher information. (a) Normalized singular values of the nondimensionalized information $\tilde{J}$ (state scaled by a and v_{circ}) for $M = 1, \dots, 4$ in descending order of the index k ; each index is a direction (singular vector) in the scaled state space, a combination of position and velocity components rather than a single component. Shaded: below 10^{-12} , numerically unobservable. (b) Smallest normalized singular value versus pass-arc duration for $M = 1, 2, 3$.

7.7. Station Layout and Pass

The three stations of Figure 4 lie nearly on a line, which is the least favorable layout for a given count. Figure 10 repeats the observability and bound analysis for an equilateral triangle of 300 km side, an equilateral triangle of 600 km side, and a 300 km square with four stations on the same pass. At $M = 3$ and $\rho = 15$ dB the position bound is 41 m for the near-collinear layout, 13 m for the 300 km triangle and 5 m for the 600 km triangle; the square gives 7 m with three of its stations and 6 m with all four, and a fourth station at the center of either triangle adds comparatively little ($12.8 \rightarrow 12.3$ and $5.2 \rightarrow 4.1$ m) because it repeats lines of sight the vertices already provide. The unknown emission epoch costs most in the near-collinear layout ($21 \rightarrow 41$ m) and little in the well-separated ones. The scaled condition number improves correspondingly, from 5×10^5 to 5×10^4 and 9×10^3 . The single-station rank deficiency (two modes below the numerical floor) and the marginal two-station mode (smallest normalized singular value $\sim 10^{-7}$) are the same in every layout. Two further passes over the baseline network (a descending pass and a north-westward pass, not shown) give three-station bounds of 18 and 23 m, so the baseline pass of Figure 4 is the least favorable of the three and the three-station results lie within a factor of about 2.3 of one another; with the emission epoch known the spread would be within a factor of 1.5 (15–22 m). Well-separated non-collinear stations are therefore worth as much as an additional station, which is the familiar dilution-of-precision conclusion; the count statements above are for the tested regional layouts and are not universal.




Figure 10. Sensitivity of observability to the station layout. (a) Normalized singular values of the nondimensionalized information for $M = 1, \dots, 4$ on the baseline layout plus a fourth station, indexed as in Figure 9a. (b) Position CRB at $\rho = 15$ dB versus station count for four layouts: the near-collinear baseline with a fourth station added, equilateral triangles of 300 and 600 km sides with a fourth station at their center, and a 300 km square.

7.8. Light-Time Correction

With the retarded-epoch delay model (4) the bounds of Figures 6–10 change by less than 10^{-4} relative to the simultaneous-time approximation because the sensitivity (8) differs only by the factor $1/(1 + \hat{\mathbf{u}}_m^\top \mathbf{v}_s/c)$ and by the evaluation of the line of sight and of the transition matrix at the retarded epoch.

The approximation is nevertheless not admissible for the estimator: fitting the simultaneous-time model to noise-free observables generated with the light-time model on the reproduced pass leaves a pseudo-true bias of 276 m in position (27 m along track) and 1.04 m/s in velocity—more than six times the 15 dB bound and far above the meter-level bounds at high SNR. All the estimator results in this paper use the light-time model; the light time at the start of the arc is 2.6 ms, during which the satellite moves 20 m.

7.9. Estimator Validation

To confirm that the CRB is operationally meaningful—not merely an asymptotic limit—the direct estimator is exercised in a Monte Carlo experiment under the Gaussian delay/Doppler surrogate of (16). In each trial, and for every station and every one of the 60 time samples of the arc, a delay and a Doppler measurement are drawn from the Gaussian distributions (16) about the light-time-model values, with a common unknown offset δt_{tx} added to all delays; the six-dimensional orbital state and δt_{tx} are then estimated jointly by weighted Gauss–Newton (at most 12 iterations) from a state perturbed by 50 m and 0.5 m/s per component. This is the ML estimator under the surrogate model, which stands in for the concentrated ML of Equation (5) after acquisition, and it exercises the geometry and dynamics layers of the bound, the light-time model and the solver; it does not simulate the raw-sample likelihood surface or its acquisition (see below). Across the signal-to-noise range investigated the post-acquisition root-mean-square error is consistent with the deterministic CRB within the Monte Carlo uncertainty of 15–30 trials per point (Figure 6): 120, 45, 17 and 4.7 m at 6, 15, 24 and 33 dB against bounds of 115, 41, 14.5 and 5.2 m with no gross errors. This also serves as a consistency check on the derivation.

Acquisition. The Monte Carlo starts after acquisition. A raw-sample search of the direct likelihood over a prior box of the size of a two-line-element error, ± 2 km in each position component and ± 2 m/s in each velocity component, at the resolution of the likelihood (about $c/B = 150$ m in range and $\lambda/T \approx 2$ mm/s in radial velocity, with $\lambda = c/f_c$ the wavelength, for a 60 s coherent arc) comprises $N_{\text{cell}}^{\text{dir}} \approx 10^{14}$ cells and is not practical by exhaustive search; in practice it has to be organized hierarchically—a coarse Doppler-only search on short sub-windows followed by refinement—as in the acquisition stage of the two-step chain. Because the surrogate likelihood used in the Monte Carlo is unimodal, it cannot produce acquisition errors by itself; the next paragraph therefore models the acquisition failures of both estimators explicitly with the same independent-cell argument.

7.10. Model-Based Threshold Comparison with Two-Step Processing

A two-step estimator first extracts, per station and per sub-window, a delay and a Doppler observable from the received data, then fits the orbital state to those observables. Above a signal-to-noise threshold this is asymptotically equivalent to the direct estimator and attains the same CRB. Below threshold, however, the per-sub-window observable extraction—a peak search over the delay/Doppler cross-ambiguity surface—begins to select noise peaks, injecting gross outliers that the orbit fit cannot repair. This effect is quantified here for both estimators with a model of the wrong-peak probability rather than a raw-sample simulation; the threshold locations are derived as follows. Within each sub-window of duration $T_{\text{sub}} = T/K$, the two-step estimator locates a delay and a Doppler by searching the cross-ambiguity function (CAF) over an a priori (coarse-ephemeris) box of extent $\Delta\tau$ in delay and Δf in Doppler. The CAF is sampled at its resolution, $1/B$ in delay and $1/T_{\text{sub}}$ in Doppler. The independent-cell approximation used below applies when cells are sampled at approximately the ambiguity resolution; residual cell correlation can be accounted for approximately through an effective number of independent cells. The number of cells searched is

$$N_{\text{cell}} = (\Delta\tau B)(\Delta f T_{\text{sub}}). \quad (17)$$

For a matched CAF [7] in complex Gaussian noise, each non-signal cell has an exponentially distributed magnitude-squared, so a single noise cell exceeds the signal cell with probability $e^{-\rho_{\text{sub}}}$, where $\rho_{\text{sub}} = \rho/K$ is the per-sub-window SNR. The probability that the peak is picked incorrectly is, under a high-SNR (fixed-signal-cell) approximation in which the signal-cell statistic is treated as deterministic,

$$P_{\text{wrong}} = 1 - (1 - e^{-\rho_{\text{sub}}})^{N_{\text{cell}}-1} \approx N_{\text{cell}} e^{-\rho_{\text{sub}}}, \quad (18)$$

which is the outlier mechanism analyzed for frequency estimation in [42] and, for maximum-likelihood estimators in general, in [43,44]. It crosses order unity at the per-sub-window threshold $\rho_{\text{sub}}^* = \ln N_{\text{cell}}$, i.e., at a total-pass threshold

$$\rho^* = K \ln N_{\text{cell}}. \quad (19)$$

For the present parameters ($B = 2$ MHz, $T = 60$ s, $K = 4$, and a $2 \text{ km} \times 2 \text{ kHz}$ coarse-ephemeris box) this gives $N_{\text{cell}} \approx 4 \times 10^5$ and $\rho^* \approx 17$ dB. This onset of order-unity wrong-pick probability is denoted $\rho_{\text{onset}} \approx 17$ dB; a practical low-error operating point ($P_{\text{wrong}} \leq 10^{-2}$) lies about 1.5 dB higher, $\rho_{1\%} \approx 18.5$ dB. A two-step estimator can, of course, also extract a *third* observable per sub-window—the Doppler rate—to use more of the information that the single-step estimator exploits. This does not remove the threshold; it raises it. The search becomes three-dimensional, adding a Doppler-rate factor $\Delta\dot{f} T_{\text{sub}}^2/2$ to (17), so N_{cell} grows by several orders of magnitude (to $\sim 9 \times 10^9$ for the same parameters with a Doppler-rate search span $\Delta\dot{f} = 200$ Hz/s, which covers the LEO Doppler-rate

range at 2.2 GHz) and, by (19), the threshold rises to $\rho^* \approx 20$ dB. The extra observable improves the above-threshold accuracy (the fit uses more information) but worsens the detection SNR required to obtain it—a direct expression of the short-window/long-window tension noted above, now sharpened because resolving frequency curvature demands still longer windows. Figure 11 shows both two-step variants. (No claim is made that the three extracted observables retain all of the information in the raw likelihood; the higher-order baseline is the more complete of the two two-step schemes tested, not an information-equivalent one.) Below ρ^* , under a uniform-wrong-pick model, a wrong selection can inject range errors of order $c \Delta\tau$, which map into large orbital-state position errors depending on the geometry; above it the picks are reliable and the two-step estimate attains the CRB.

The direct estimator is subject to the same mechanism, but it makes a single peak decision on the concentrated statistic (5) of the whole arc and all stations over the $N_{\text{cell}}^{\text{dir}} \approx 10^{14}$ state cells of the two-line-element box. A noise cell of that statistic, an incoherent sum over M stations, is $\Gamma(M, 1)$ -distributed. In the same fixed-signal-cell approximation it exceeds the signal level $M\rho$ with probability

$$p_1^{\text{dir}} = e^{-M\rho} \sum_{k=0}^{M-1} \frac{(M\rho)^k}{k!}, \quad p_{\text{wrong}}^{\text{dir}} = 1 - (1 - p_1^{\text{dir}})^{N_{\text{cell}}^{\text{dir}} - 1}. \quad (20)$$

For $M = 3$ this crosses order unity at $\rho_{\text{dir}}^* \approx 11$ dB, and $p_{\text{wrong}}^{\text{dir}} \leq 10^{-2}$ at 11.6 dB: the far larger number of cells is more than offset by using the energy of the whole arc and all the stations in one decision, whereas the two-step estimator decides per station on a quarter of the energy. The direct threshold of (20) assumes an exhaustive search; a hierarchical acquisition reintroduces intermediate decisions at its coarse stage, so ρ_{dir}^* is a best case, and the practical advantage lies between zero and the model value.

Both estimators estimate δt_{tx} jointly with the state. In the Monte Carlo experiment of Figure 11 both are assigned acquisition failures from these models: the two-step per-sub-window extraction fails with probability (18) (a wrong pick being uniform in the search box) before the orbit fit, and the direct estimator fails with probability (20) (a wrong pick being uniform in the state box, which leaves an error of the order of kilometers); otherwise each proceeds to its orbit fit. The experiment reproduces the model thresholds and the convergence to the CRB above them: the probability of a gross error (more than ten times the bound) is one up to $\rho = 15$ dB, 0.56 at 18 dB and zero from 21 dB for the delay/Doppler baseline and one up to 18 dB and zero from 21 dB for the Doppler-rate variant. For the direct estimator every trial at 6 and 9 dB is a wrong pick, with errors of 1–2 km, and there are no gross errors from 12 dB. Above threshold the root-mean-square error follows the bound (within the Monte Carlo scatter of 15 trials per point). Because a two-step wrong pick produces errors of hundreds of kilometers, the two-step RMSE below threshold is dominated by the failures (each capped at 10^6 m, the order of the orbital error that the search-box extent maps to) and is reported together with the gross-error probability, which shows the mechanism. Under this common model the direct estimator's threshold lies about 6 dB below that of the delay/Doppler two-step baseline and about 8.5 dB below that of the Doppler-rate variant. The advantage is thus an extension of usable operation into the weak-signal regime, not a lower asymptote: the two approaches coincide where observables are reliably extractable. This advantage is predicted by the failure model, not demonstrated by a simulation of the raw-sample likelihood and its acquisition, which is left to future work. This comparison is against a two-step architecture that commits to a hard peak decision per sub-window; a more elaborate scheme retaining soft likelihoods or multiple candidate peaks could lower its threshold, so the precise claim is that direct processing avoids the threshold associated with hard intermediate peak decisions and replaces it with the lower threshold of a single decision on the full-arc statistic. A structural point underlies

this: two-step processing is caught between using short sub-windows, needed to track the rapidly varying Doppler, and long ones, needed for reliable detection—a tension the single-step estimator avoids at the level of its final decision, although a practical hierarchical acquisition reintroduces sub-window processing at its coarse stage.




Figure 11. Model-based threshold comparison for a single emitter. Both estimators are assigned acquisition (wrong-peak) failures from the same independent-cell model: the direct estimator by (20), one decision over the two-line-element state box; the two-step estimators by (18), one decision per station and sub-window. (a) Root-mean-square position error over all trials versus pass SNR (two-step failures capped at 10^6 m): CRB (line); direct estimator (circles); two-step with delay and Doppler (dashed, diamonds) and the Doppler rate additionally extracted (dotted, triangles). Vertical lines: model thresholds, ≈ 11 dB (direct, dash-dotted), ≈ 17 and ≈ 20 dB (two-step, Equation (19)). (b) Probability of a gross error (more than ten times the bound). Above its threshold each estimator coincides with the CRB. The failures are imposed by the model; the underlying surrogate likelihood is unimodal.

7.11. Summary of Accuracy

Table 4 collects the principal quantitative results for the scenario of this section.

Table 4. Summary of the orbital-state accuracy obtained in this section (pass SNR $\rho = 15$ dB per station, 60 s arc, and synchronized stations with the emission-time offset estimated unless stated otherwise). Position values are $\sqrt{\text{tr } C_{rr}}$.

Case	Result
Baseline network, $M = 3$, frequencies known: position/velocity CRB	41 m/0.15 m/s
Direct estimator after acquisition, Monte Carlo RMSE, same case	45 m
Same, $\rho = 33$ dB: CRB/Monte Carlo RMSE	5.2 m/4.7 m
Frequency offsets and drift also estimated, $M = 3/4/6$	78/17/12 m
Baseline network plus a fourth station, frequencies known	7.0 m
Layouts, $M = 3$: 300 km triangle/600 km triangle/300 km square	13/5/7 m
Arc length (fixed energy): 10 s/45 s/90 s	80/42/38 m
Bias of the simultaneous-time (no light-time) model	276 m, 1.04 m/s
Acquisition threshold (model): direct/two-step/two-step with Doppler rate	11/17/20 dB

8. Conclusions

This paper formulated the single-step direct determination of a non-cooperative LEO satellite's orbital state from multi-station RF signals, inverting the usual direct-geolocation roles so that the satellite is the unknown emitter. Because the trajectory is a deterministic function of the epoch state, the estimator searches only over valid orbits, with the orbital-dynamics manifold intrinsic to the parameterization. The orbital-state Cramér–Rao

bound was derived and shown to factor into signal, geometry, and dynamics terms, the last being the two-body state-transition matrix for which a closed-form position block was given via implicit differentiation of the universal Kepler equation. The delay model evaluates the satellite position at the emission epoch (light-time correction), which leaves the bound essentially unchanged but removes a bias of more than a hundred meters. The unknown emission epoch of a non-cooperative emitter enters as a common time-offset nuisance so that, with synchronized stations, the delays carry time-difference information; this roughly doubles the three-station bound on the near-collinear baseline, whereas four stations recover the known-epoch accuracy. Simulations validate the analytical derivation under the stated model assumptions, demonstrate post-acquisition efficiency against the deterministic CRB under a Gaussian delay/Doppler surrogate of the known-waveform model, and, under a common model of acquisition failures, place the direct estimator's threshold about 6 dB below that of a two-step baseline that commits to hard per-window peak decisions; this advantage is model-based, an upper limit for a practical hierarchically acquired direct estimator, and remains to be confirmed by a raw-I/Q simulation; more elaborate two-step schemes with soft or multiple-hypothesis processing may narrow the gap. The stochastic unknown-waveform FIM retains the same geometry/dynamics structure with an information factor $\gamma/(1 + \gamma)$ in the per-sample array SNR, which confines the unknown-waveform estimator to high per-sample SNR; at low-pass SNRs where the single-step advantage is predicted, the waveform must be known or structured. The station-count conclusions were established on a nondimensionalized information matrix for the regional layouts and passes tested, and well-separated non-collinear stations were found to be worth as much as an added station. The limitations include the computational cost of the six-dimensional search (mitigated by coarse initialization and the smoothness of the orbit), the assumption of Keplerian dynamics in the closed-form bound (relaxed numerically via the variational equations), and a single-emitter model. The most consequential extension is to the multiple-emitter case—for example, a satellite downlink coincident in frequency and time with a co-channel terrestrial interferer, a realistic non-cooperative scenario—where the single-step approach may provide a further signal-to-noise-independent advantage by avoiding the measurement association that two-step processing requires; this hypothesis remains to be quantified in future multi-emitter analysis. Future work will also address maneuvering (non-Keplerian) satellites—linking to thrust estimation for non-cooperative targets—and experimental validation with synchronized multi-station raw-I/Q capture. Public LEO I/Q datasets exist, but the coherent multi-station recordings that are required for experimental validation are not available in the present study and define a concrete next step.

Funding: This research received no external funding.

Data Availability Statement: No observational data were collected in this study; all the results are simulations whose inputs are specified in Section 7 and Tables 2 and 3. The MATLAB (R2025b) script that implements the state-transition matrix, Fisher information, and CRB with the light-time model and direct estimator and two-step baseline and produces every figure and number quoted in Section 7 with a fixed random seed is available from the author upon reasonable request.

Conflicts of Interest: The author declares no conflicts of interest.

References

1. Tapley, B.D.; Schutz, B.E.; Born, G.H. *Statistical Orbit Determination*; Elsevier Academic Press: Amsterdam, The Netherlands, 2004.
2. Montenbruck, O.; Gill, E. *Satellite Orbits: Models, Methods and Applications*; Springer: Berlin/Heidelberg, Germany, 2000. [[CrossRef](#)]
3. Vallado, D.A. *Fundamentals of Astrodynamics and Applications*, 5th ed.; Microcosm Press: Hawthorne, CA, USA, 2022.

4. Weiss, A.J. Direct geolocation of wideband emitters based on delay and Doppler. *IEEE Trans. Signal Process.* **2011**, *59*, 2513–2521. [\[CrossRef\]](#)
5. Amar, A.; Weiss, A.J. Localization of narrowband radio emitters based on Doppler frequency shifts. *IEEE Trans. Signal Process.* **2008**, *56*, 5500–5508. [\[CrossRef\]](#)
6. Knapp, C.H.; Carter, G.C. The generalized correlation method for estimation of time delay. *IEEE Trans. Acoust. Speech Signal Process.* **1976**, *24*, 320–327. [\[CrossRef\]](#)
7. Stein, S. Algorithms for ambiguity function processing. *IEEE Trans. Acoust. Speech Signal Process.* **1981**, *29*, 588–599. [\[CrossRef\]](#)
8. Ho, K.C.; Chan, Y.T. Geolocation of a known altitude object from TDOA and FDOA measurements. *IEEE Trans. Aerosp. Electron. Syst.* **1997**, *33*, 770–783. [\[CrossRef\]](#) [\[PubMed\]](#)
9. Ho, K.C.; Xu, W. An accurate algebraic solution for moving source location using TDOA and FDOA measurements. *IEEE Trans. Signal Process.* **2004**, *52*, 2453–2463. [\[CrossRef\]](#)
10. Feuge-Miller, B.; Joplin, A.; York, J. Initial orbit determination from ambiguous TDOA and FDOA measurements of passive radio frequency signals. *J. Astronaut. Sci.* **2025**, *72*, 37. [\[CrossRef\]](#)
11. Schmidt, B.; Jäger, E.; Dietl, G. Orbit determination using difference of arrival measurements. *CEAS Space J.* **2025**. [\[CrossRef\]](#)
12. Weiss, A.J. Direct position determination of narrowband radio frequency transmitters. *IEEE Signal Process. Lett.* **2004**, *11*, 513–516. [\[CrossRef\]](#)
13. Vankayalapati, N.; Kay, S.; Ding, Q. TDOA based direct positioning maximum likelihood estimator and the Cramér–Rao bound. *IEEE Trans. Aerosp. Electron. Syst.* **2014**, *50*, 1616–1635. [\[CrossRef\]](#)
14. Oispuu, M.; Nickel, U. Direct detection and position determination of multiple sources with intermittent emission. *Signal Process.* **2010**, *90*, 3056–3064. [\[CrossRef\]](#)
15. Closas, P.; Fernández-Prades, C.; Fernández-Rubio, J.A. Maximum likelihood estimation of position in GNSS. *IEEE Signal Process. Lett.* **2007**, *14*, 359–362. [\[CrossRef\]](#)
16. Clements, Z.; Ellis, P.; Humphreys, T.E. Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit. In Proceedings of the 2023 IEEE/ION Position, Location and Navigation Symposium (PLANS), Monterey, CA, USA, 24–27 April 2023; pp. 458–469.
17. Ellis, P.; Rheeden, D.V.; Dowla, F. Use of Doppler and Doppler rate for RF geolocation using a single LEO satellite. *IEEE Access* **2020**, *8*, 12907–12920. [\[CrossRef\]](#)
18. Reid, T.G.R.; Neish, A.M.; Walter, T.; Enge, P.K. Broadband LEO constellations for navigation. *Navigation* **2018**, *65*, 205–220. [\[CrossRef\]](#)
19. Prol, F.S.; Ferre, R.M.; Saleem, Z.; Välisuo, P.; Pinell, C.; Lohan, E.S.; Elsanhoury, M.; Elmusrati, M.; Islam, S.; Çelikbilek, K.; et al. Position, navigation, and timing (PNT) through low Earth orbit (LEO) satellites: A survey on current status, challenges, and opportunities. *IEEE Access* **2022**, *10*, 83971–84002. [\[CrossRef\]](#)
20. Khalife, J.J.; Kassas, Z.M. Receiver design for Doppler positioning with LEO satellites. In Proceedings of the ICASSP 2019—2019 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP), Brighton, UK, 12–17 May 2019; pp. 5506–5510.
21. Psiaki, M.L. Navigation using carrier Doppler shift from a LEO constellation: TRANSIT on steroids. *Navigation* **2021**, *68*, 621–641. [\[CrossRef\]](#)
22. Neinavaie, M.; Khalife, J.; Kassas, Z.M. Acquisition, Doppler tracking, and positioning with Starlink LEO satellites: First results. *IEEE Trans. Aerosp. Electron. Syst.* **2022**, *58*, 2606–2610. [\[CrossRef\]](#)
23. Kozhaya, S.; Kassas, Z.M. A first look at the OneWeb LEO constellation: Beacons, beams, and positioning. *IEEE Trans. Aerosp. Electron. Syst.* **2024**, *60*, 7528–7534. [\[CrossRef\]](#)
24. Zhao, C.; Qin, H.; Wu, N.; Wang, D. Analysis of baseline impact on differential Doppler positioning and performance improvement method for LEO opportunistic navigation. *IEEE Trans. Instrum. Meas.* **2023**, *72*, 8501110. [\[CrossRef\]](#)
25. Liu, Q.; Reus-Bergas, A.; López-Salcedo, J.A.; Seco-Granados, G. An innovative outer-loop positioning framework for LEO Doppler observations considering propagation time. *Satell. Navig.* **2026**, *7*, 3. [\[CrossRef\]](#)
26. Kassas, Z.M.; Khairallah, N.; Kozhaya, S. Ad astra: Simultaneous tracking and navigation with megaconstellation LEO satellites. *IEEE Aerosp. Electron. Syst. Mag.* **2024**, *39*, 46–71. [\[CrossRef\]](#)
27. Mortlock, T.; Kassas, Z.M. Assessing machine learning for LEO satellite orbit determination in simultaneous tracking and navigation. In Proceedings of the 2021 IEEE Aerospace Conference, Big Sky, MT, USA, 6–13 March 2021; pp. 1–8.
28. Deng, R.; Qin, H.; Li, H.; Wang, D.; Lv, H. Noncooperative LEO satellite orbit determination based on single pass Doppler measurements. *IEEE Trans. Aerosp. Electron. Syst.* **2023**, *59*, 1096–1106. [\[CrossRef\]](#)
29. Deng, R.; Qin, H.; Zhang, Y. Non-cooperative LEO satellite orbit determination using single station for space-based opportunistic positioning. *Remote Sens.* **2024**, *16*, 912. [\[CrossRef\]](#)
30. Lei, M.; Liu, Y.; Yang, Z.; Li, H.; Fang, Z. Adaptive weighting-based multi-station Doppler orbit determination for noncooperative LEO satellites. *Sensors* **2026**, *26*, 5932. [\[CrossRef\]](#) [\[PubMed\]](#)

31. Allahviridi-Zadeh, A.; El-Mowafy, A.; Wang, K. Network-based precise orbit determination of broadband LEO satellites using Doppler-shift measurements. *GPS Solut.* **2026**, *30*, 16. [[CrossRef](#)]
32. Hayek, S.; Kassas, Z.M. Modeling and compensation of timing and spatial ephemeris errors of non-cooperative LEO satellites with application to PNT. *IEEE Trans. Aerosp. Electron. Syst.* **2025**, *61*, 5579–5593. [[CrossRef](#)]
33. Saroufim, J.; Kassas, Z.M. Ephemeris and timing error disambiguation enabling precise LEO PNT. *IEEE Trans. Aerosp. Electron. Syst.* **2025**, *61*, 6138–6153. [[CrossRef](#)]
34. Tichavský, P.; Muravchik, C.; Nehorai, A. Posterior Cramér–Rao bounds for discrete-time nonlinear filtering. *IEEE Trans. Signal Process.* **1998**, *46*, 1386–1396. [[CrossRef](#)]
35. Trees, H.L.V.; Bell, K.L. (Eds.) *Bayesian Bounds for Parameter Estimation and Nonlinear Filtering/Tracking*; Wiley-IEEE Press: Piscataway, NJ, USA, 2007.
36. Moyer, T.D. *Formulation for Observed and Computed Values of Deep Space Network Data Types for Navigation*; JPL Deep-Space Communications and Navigation Series; Wiley: Hoboken, NJ, USA, 2003. [[CrossRef](#)]
37. Gorman, J.D.; Hero, A.O. Lower bounds for parametric estimation with constraints. *IEEE Trans. Inf. Theory* **1990**, *36*, 1285–1301. [[CrossRef](#)]
38. Moore, T.J.; Sadler, B.M.; Kozick, R.J. Maximum-likelihood estimation, the Cramér–Rao bound, and the method of scoring with parameter constraints. *IEEE Trans. Signal Process.* **2008**, *56*, 895–908. [[CrossRef](#)]
39. Kay, S.M. *Fundamentals of Statistical Signal Processing, Volume I: Estimation Theory*; Prentice Hall: Upper Saddle River, NJ, USA, 1993.
40. Battin, R.H. *An Introduction to the Mathematics and Methods of Astrodynamics*, revised ed.; AIAA: Reston, VA, USA, 1999.
41. Shepperd, S.W. Universal Keplerian state transition matrix. *Celest. Mech.* **1985**, *35*, 129–144. [[CrossRef](#)]
42. Rife, D.C.; Boorstyn, R.R. Single tone parameter estimation from discrete-time observations. *IEEE Trans. Inf. Theory* **1974**, *20*, 591–598. [[CrossRef](#)]
43. Athley, F. Threshold region performance of maximum likelihood direction of arrival estimators. *IEEE Trans. Signal Process.* **2005**, *53*, 1359–1373. [[CrossRef](#)]
44. Richmond, C.D. Mean-squared error and threshold SNR prediction of maximum-likelihood signal parameter estimation with estimated colored noise covariances. *IEEE Trans. Inf. Theory* **2006**, *52*, 2146–2164. [[CrossRef](#)]

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