Triple-Satellite Geolocation from Low Earth Orbit in a Multi-Emitter GNSS Interference Environment

A Parametric System Analysis

J.L.F. Raithel

Delft University of Technology

TU Delft

nlr Accelerating the future of aerospace

Triple-Satellite Geolocation from Low Earth Orbit in a Multi-Emitter GNSS Interference Environment

A Parametric System Analysis

by

J.L.F. Raithel

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Friday May 29, 2026 at 2:00 PM.

Student number:	5129648	
Project duration:	September 15, 2025 – May 29, 2026	
Thesis committee:	Ir. M.R. Menken,	NLR, Daily Supervisor
	Dr. ir. E.J.O. Schrama,	TU Delft, Responsible Supervisor
	Dr. ir. M.J. Heiligers	TU Delft, Chair
	Dr. ir. D. Dirkx	TU Delft, Examiner

Cover: Modified artist impressions of HawkEye 360 satellites flying in tight clusters of three geolocating GNSS emitters

An electronic version of this thesis is available at `http://repository.tudelft.nl/`.

TU Delft

nlr Accelerating the future of aerospace

Preface

These pages mark the culmination of my thesis project at the Royal Netherlands Aerospace Centre (NLR). What continues to amaze me about GNSS is the sheer improbability of it all. A satellite is launched on top of millions of kilograms of explosive fuel, placed into orbit at roughly 20,000 km altitude, and from there transmits a signal so weak that receiving it has been compared to observing the light of a single light bulb in Amsterdam from Gibraltar, some 2,000 km away. Yet this signal underpins systems that enable the modern world. Working on such technology, so often taken for granted, has been a privilege.

This privilege has brought me into contact with many people, each with their own expertise and perspective. Their input has contributed greatly to both my professional and personal development, and I am grateful for the opportunity to have learned from them.

In particular, I would like to thank my supervisors at NLR and TU Delft. I want to thank Maarten Menken for all the time and effort he put into making this thesis a success. I especially appreciated his ability to cut through the noise and keep the focus on what mattered most. I would also like to thank Ernst Schrama for his patience and helpful feedback throughout the process. Our biweekly meetings became more enjoyable over time, which made it slightly unfortunate that a thesis, by definition, has a finite duration.

I would also like to thank my colleagues at NLR for their genuine interest in my work and for the enjoyable atmosphere they created every day. I would like to wish all the best of luck with their own fascinating work.

Altogether, I am much obliged.

J.L.F. Raithel
Amsterdam, May 2026

"Everything should be kept as simple as possible, but no simpler."

Abstract

GNSS jammers can disrupt critical positioning, navigation, and timing services. Low Earth orbit satellites offer a way to locate these emitters over large areas. This thesis studies passive geolocation of stationary terrestrial GNSS jammers with a three-satellite LEO system.

Two geolocation approaches are compared. The first estimates emitter positions directly from received I/Q data. The second first extracts receiver-differenced observables and then estimates position from FDOA measurements. These approaches are referred to as direct and indirect geolocation. A parametric simulation framework is developed to model the satellite formation, jammer signals, receiver data, and main error sources. A CRLB analysis is used to study the design space.

The results show that both approaches can achieve sub-kilometer per-emitter accuracy. In Scenario 1, the direct method reaches a mean error of 502 m, while the indirect method reaches 143 m. In Scenario 2, the corresponding values are 517 m and 292 m. The indirect method also provides precision estimates through 95% confidence ellipses. Its main advantage is computational cost. It is about 9.0×10^2 times faster per snapshot in the reported implementation.

The results also show that a tight formation is preferred. It maximizes the shared field of view while keeping enough geometric diversity. The FDOA loop-closure constraint is found to be essential. It rejects false candidates and makes indirect geolocation practical in a multi-emitter setting.

This thesis concludes that FDOA-based indirect geolocation with a tight three-satellite LEO formation is the most suitable option for low-latency wide-area GNSS interference geolocation within the considered scope. Future work should validate the method with more scenarios or real satellite data.

Contents

List of Figures

List of Tables

Nomenclature

Abbreviations

Abbreviation	Definition
CAF	Cross-Ambiguity Function
CRLB	Cramér–Rao Lower Bound
DFDOA	Differential Frequency Difference of Arrival
DRMS	Distance Root Mean Square
ECEF	Earth-Centered Earth-Fixed
ECI	Earth-Centered Inertial
FDOA	Frequency Difference of Arrival
FoV	Field of View
GMST	Greenwich Mean Sidereal Time
GNSS	Global Navigation Satellite System
I/Q	In-phase and Quadrature
LEO	Low Earth Orbit
LOS	Line of Sight
SNR	Signal-to-Noise Ratio
TCXO	Temperature-Compensated Crystal Oscillator
TDOA	Time Difference of Arrival
WGS-84	World Geodetic System 1984

Symbols

Symbol	Definition	Unit
Geometry, states, and timing		
a	Along-track separation between satellites	[m]
b	Cross-track separation between satellites	[m]
B	Signal bandwidth	[Hz]
c	Speed of light	[m/s]
e	Orbital eccentricity	[-]
f_c	Carrier frequency	[Hz]
f_s	Sampling frequency	[Hz]
h	Orbital altitude or emitter height, depending on context	[m]
M	Mean anomaly	[rad]
N_E	Number of emitters	[-]
N_s	Number of samples in one snapshot	[-]
R_E	Earth radius	[m]
$\mathbf{r}_{E_i}$	Position vector of emitter i	[m]
$\hat{\mathbf{r}}_E$	Estimated emitter position	[m]
$\mathbf{r}_{R_j}(t)$	Position vector of receiver j	[m]
$\dot{\mathbf{r}}_{R_j}(t)$	Velocity vector of receiver j	[m/s]
$\ddot{\mathbf{r}}_{R_j}(t)$	Acceleration vector of receiver j	[m/s^2]
T_{int}	Coherent integration time	[s]
T_{obs}	Observation time	[s]
T_s	Sampling period	[s]

Symbol	Definition	Unit
$\Delta_{ij}(t)$	Relative position vector from receiver j to emitter i	[m]
$\theta(t)$	Greenwich Mean Sidereal Time angle	[rad]
λ_i	Longitude of emitter i	[rad]
μ	Earth gravitational parameter	[m^3/s^2]
$\rho_{ij}(t)$	Range between emitter i and receiver j	[m]
$\dot{\rho}_{ij}(t)$	Range rate	[m/s]
$\ddot{\rho}_{ij}(t)$	Range acceleration	[m/s^2]
ϕ_i	Latitude of emitter i	[rad]
$\hat{\boldsymbol{\ell}}_{ij}(t)$	Line-of-sight unit vector	[-]
ω	Argument of perigee	[rad]
$\boldsymbol{\omega}_E$	Earth rotation vector	[rad/s]
Ω	Right ascension of the ascending node	[rad]
Ω_E	Earth rotation rate	[rad/s]
Signals and clock terms		
$x_i(t)$	Signal transmitted by emitter i	[-]
$y_j(t)$	Received signal at receiver j	[-]
$\tilde{y}_j[n]$	Quantized complex-baseband sample sequence	[-]
δt_{E_i}	Emitter clock offset	[s]
$\delta t_{R_j}^{\text{res}}(t)$	Residual receiver clock offset	[s]
$\delta \dot{t}_{E_i}$	Emitter clock drift	[-]
$\delta \dot{t}_{R_j}^{\text{res}}(t)$	Residual receiver clock drift	[-]
σ^2	Noise variance per complex sample	[-]
$\tau_{ij}^{g}(t)$	Geometric propagation delay	[s]
$\tilde{\tau}_{ij}(t)$	Effective propagation delay	[s]
$\phi_{ij}(t)$	Received phase	[rad]
Observables and operators		
$\Delta q_{i,jk}(t)$	Pairwise receiver difference of a quantity q	[varies]
$\Delta\tau_{i,jk}(t)$	Time Difference of Arrival (TDOA)	[s]
$\Delta f_{i,jk}(t)$	Frequency Difference of Arrival (FDOA)	[Hz]
$\Delta \dot{f}_{i,jk}(t)$	Differential Frequency Difference of Arrival (DFDOA)	[Hz/s]
$\mathbf{P}_{ij}^{\perp}(t)$	Orthogonal projector perpendicular to $\hat{\boldsymbol{\ell}}_{ij}(t)$	[-]
Indices		
i	Emitter index	[-]
j, k	Receiver indices	[-]
n	Sample index	[-]

1

Introduction

Low Earth orbit (LEO) receivers enable wide-area passive geolocation of terrestrial GNSS jammers. Prior work has shown that this is feasible with one and two satellites. The main open question is therefore not feasibility, but design. Which processing approach is most suitable when observations are imperfect and latency matters?

This work studies that question for a triple-satellite LEO system. It compares two geolocation approaches under the same scenario assumptions. The first estimates emitter position directly from the received I/Q data. The second first extracts receiver-differenced observables and then estimates position from those measurements. These approaches are referred to in this work as direct and indirect geolocation, respectively. The core question is how they behave in the same three-satellite, multi-emitter setting.

1.1. Background and Motivation

Global Navigation Satellite Systems (GNSS), including GPS, GLONASS, Galileo, and BeiDou, are central to modern positioning, navigation, and timing [21]. Dependence on GNSS continues to grow. So does the impact of disruption [30, 7].

Intentional interference is a major threat to GNSS availability and trustworthiness [1]. The two main forms are jamming and spoofing [15]. Jamming blocks authentic GNSS signals by transmitting interference in the relevant bands. Spoofing transmits counterfeit GNSS-like signals to mislead a receiver's position, velocity, or time estimate [35, 36]. Both matter in practice. This work focuses on terrestrial GNSS jammers.

Fast jammer geolocation supports monitoring, attribution, and mitigation [8, 4, 13]. Interference can sometimes be inferred from anomalies in GNSS-dependent services such as AIS or ADS-B. Such information is usually too coarse for localization [25]. Terrestrial and airborne sensors can provide finer estimates, but only over limited areas [2, 40, 33]. Wide-area geolocation therefore requires a different observation architecture.

LEO receivers provide that architecture. They combine wide-area access with strong Doppler dynamics from high orbital velocity. They can also receive authentic GNSS signals for time synchronization and orbit determination. These advantages have motivated operational and commercial systems such as those of Spire Global and HawkEye 360 [31, 11]. An overview of the operational scenario is shown in Figure 1.1.

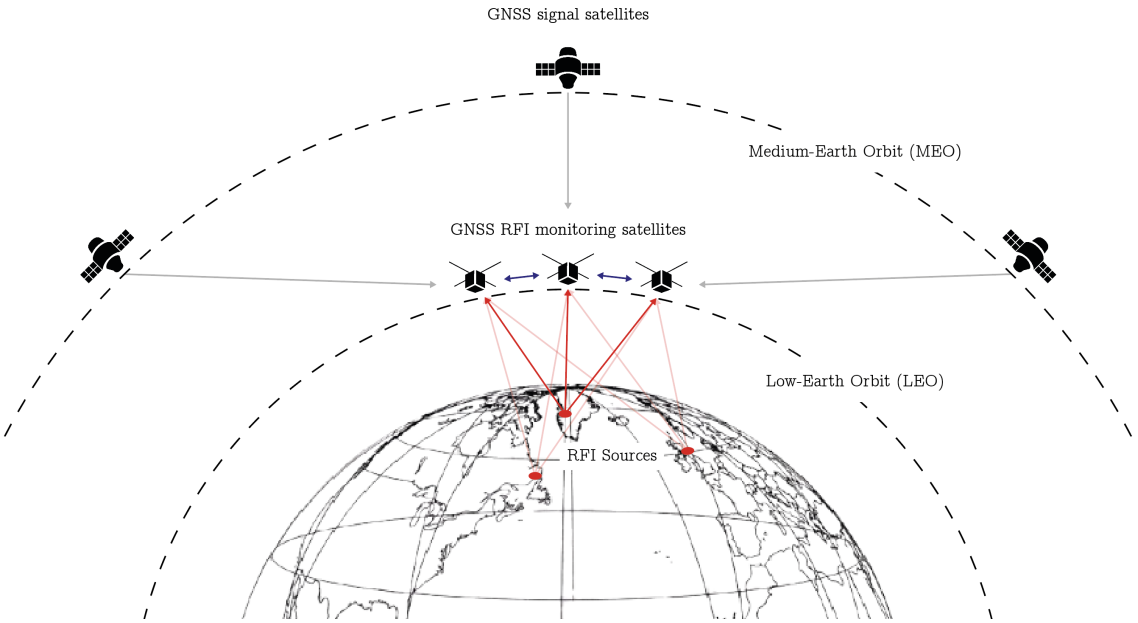

Figure 1.1: GNSS interference geolocation scenario using LEO satellites. The LEO satellites receive signals from both intentional RFI sources and authentic GNSS satellites, while inter-satellite communication enables cooperative detection and geolocation.

1.2. Related Work and Research Gap

LEO-based GNSS interference geolocation has been studied with one satellite, two satellites, and, to a much lesser extent, larger formations. Depending on system design, these methods use observables such as Doppler, Doppler rate, pseudorange, Time Difference of Arrival (TDOA), and Frequency Difference of Arrival (FDOA) [20, 29, 27, 16, 43, 12, 11, 15, 10, 6, 30].

Table 1.1: Overview of representative literature on LEO-based GNSS interference geolocation

Study	Year	Type	Sats.	Observables	Scenario	Performance	Limitations
[16]	2020	Exp.	1	Doppler, Doppler rate	Single-jammer	Position error: < 10 km	Limited samples; low-SNR, 2-bit data
[29]	2021	Exp.	1	Doppler, Doppler rate	Single-jammer	95% SMA: 220 m	Unknown ground truth; clock errors
[27]	2025	Exp.	1	FDOA	Single-jammer	Position error: 40 km (lat.), 70 km (lon.)	Short observation window; noisy observables
[12]	2025	Exp.	1	Doppler	Single-spoofer	Position error: 68 m; 95% SMA: 6.7 km	Clock errors; spoofer-specific
[43]	2024	Sim.	1	Doppler, Doppler rate	Single-jammer	Position error: < 1 km	Assumed observables
[11]	2023	Exp.	2	FDOA, TDOA	Multi-jammer	95% SMA: 620 m	Unknown ground truth; high computational cost
[15]	2024	Sim.	2	FDOA, TDOA	Single-jammer	Position error: 204 m	Assumed observables
[10]	2025	Sim.	2	FDOA, TDOA	Multi-jammer	Unknown	High computational cost

The studies in Table 1.1 are not directly comparable. They differ in geometry, scenario complexity, observation duration, observable definitions, error assumptions, and performance metrics. The table is therefore used to show how the field has developed and which limitations recur. In the Type column, Exp. denotes experimental work based on measured data, while Sim. denotes simulation-based work.

1.2.1. Single-Satellite and Dual-Satellite Geolocation

The literature can be grouped by receiver architecture and by how the received data are converted into a position estimate. Some methods operate on extracted observables such as Doppler, Doppler rate, TDOA, or FDOA. Others evaluate candidate emitter locations more directly from the received signal data. This distinction becomes important when comparing accuracy and latency, because observable extraction, data association, and position-domain search have different computational costs and error mechanisms.

The first conclusion from the literature is feasibility. Single-satellite studies showed that terrestrial GNSS jammers can be geolocated from LEO using Doppler, Doppler rate, or FDOA under controlled conditions [29, 27, 16, 43]. Their main limitation is fundamental. With one receiver, geometry is weak. Performance depends strongly on signal assumptions, clock behavior, and observation time.

The second conclusion is that two satellites improve the problem substantially. Dual-satellite systems enable receiver-differenced observables such as TDOA and FDOA. They improve geometry and reduce sensitivity to some nuisance terms [11, 15, 10]. They also make it possible to compare methods that use the received I/Q data directly with methods that first extract TDOA and FDOA.

1.2.2. Research Gap

The open problem begins at the next step. A third satellite adds another baseline. It also adds more geometry and more measurement combinations. Open GNSS-specific literature does not yet show clearly how a three-satellite system changes the comparison between direct and indirect geolocation.

This gap remains for three reasons. First, GNSS-specific work still focuses mainly on single-satellite and dual-satellite systems. Open work on triple-satellite architectures is therefore limited. Second, many studies rely on opportunistic datasets or operational systems. They show feasibility, but they are less suited to controlled architecture comparison. Geometry, observation strategy, signal assumptions, and processing methods vary across studies. Third, latency is rarely treated as a primary design driver, even though a highly accurate result may have limited operational value if it arrives too late.

HawkEye 360 has operated triple-satellite formations since 2016, but published work has focused mainly on formation design rather than GNSS-specific jammer geolocation performance [31]. A more recent comparison between single-satellite and triple-satellite systems found that three satellites are preferable when higher precision is required for the same cumulative measurement time [26]. That result is relevant, but it does not answer the design question addressed here. The open question is how different geolocation approaches compare when the same three-satellite scenario, signal environment, and latency constraints are used. This work therefore compares direct and indirect geolocation under common assumptions.

1.3. Research Objective, Questions, and Contributions

This section defines the aim and scope of the research. It states the research objective, the main research question, and the sub-questions. It then summarizes the main contributions of this work.

1.3.1. Research Objective

The objective of this work is:

To characterize and compare the latency–accuracy trade-off of direct and indirect geolocation for stationary terrestrial GNSS jammers in a triple-satellite LEO system, and to determine how formation and processing choices affect system suitability.

1.3.2. Research Question

This objective is addressed through the following main research question:

How do direct and indirect geolocation compare in a triple-satellite LEO system for stationary terrestrial GNSS jammers, and how do formation and processing choices affect the accuracy–latency trade-off?

This question is broken down into the following sub-questions:

- **RQ1:** How do formation geometry, observable set, coherent integration time, and observation duration affect the fundamental accuracy limits of triple-satellite LEO geolocation of stationary terrestrial GNSS jammers?
- **RQ2:** How do jammer type, receiver bandwidth, quantization, and dominant error sources affect the extraction and reliability of receiver-differenced observables in a triple-satellite system?
- **RQ3:** How do direct and indirect geolocation compare in accuracy, precision, observation time, and latency under common triple-satellite scenario assumptions?
- **RQ4:** Which combinations of formation and processing choices are most suitable for timely and accurate triple-satellite LEO geolocation of stationary terrestrial GNSS jammers?

Sub-kilometer accuracy is the benchmark used in this work. It is consistent with the open literature on LEO-based GNSS interference geolocation [29, 11, 15]. It is also operationally meaningful for low-latency intervention. A larger error region can still be useful for cueing follow-up assets.

1.3.3. Contributions of this work

This work makes five main contributions:

- A **CRLB-based characterization** of the triple-satellite LEO jammer geolocation design space. It shows how formation geometry, observable choice, integration time, and observation duration shape the fundamental accuracy limits.
- A **derived and implemented end-to-end simulation framework** for systematic analysis of LEO-based geolocation of terrestrial emitters under representative operating conditions.
- The **identification and exploitation of an FDOA loop-closure property**. In a three-receiver closed loop, the pairwise FDOAs must sum to zero. This is key to extending indirect geolocation beyond two receivers and to improving robustness in challenging environments.
- A **systematic comparison of direct and indirect geolocation** under common assumptions. This enables a fair assessment of their relative performance, strengths, and limitations.
- The **identification of latency bottlenecks** in LEO receiver geolocation systems, together with **architectural recommendations** for timely and accurate terrestrial emitter localization.

1.4. Scope and Limitations

This work focuses on passive geolocation of stationary terrestrial GNSS jammers. It does not address spoofers, moving emitters, or non-GNSS interferers. The jammer classes considered are tone, chirp, and modulated-signal emitters. This allows direct comparison with prior work [10] while still covering representative interference types.

The analysis focuses on a triple-satellite architecture and on the geolocation segment itself. It therefore considers in-pass performance rather than constellation-level access, revisit time, tasking latency, dissemination to end users, or broader mission design.

Within that scope, the study models the dominant first-order drivers of short-duration geolocation performance. These include signal structure, receiver and clock errors, atmospheric effects, and the processing chain used for observable extraction and geolocation. Sensitivity analyses quantify how these factors affect performance.

Several effects are excluded. They are expected to be of lower order over the observation intervals considered, strongly scenario-specific, or outside the segment being assessed. These include higher-order orbital perturbations, multipath, detailed antenna gain variation, detailed inter-satellite communication and relative-navigation implementation, spoofers, non-GNSS emitters, moving sources, and broader mission-level coverage considerations. These simplifications keep the problem tractable while preserving the dominant mechanisms relevant to the latency–accuracy trade-off.

1.5. Methodological Approach

First, this work characterizes the fundamental performance limits of triple-satellite geolocation using the Cramér-Rao Lower Bound (CRLB). This provides a first-order map of the design space. It also identifies

promising combinations of satellite formation, observable set, integration time, and observation duration.

Second, it develops a parametric end-to-end system model in Python. The model links representative jammer signals to the raw I/Q samples collected by the satellite receivers. It includes the signal, receiver, geometry, and error sources needed for consistent evaluation.

Third, it uses this framework to compare direct and indirect geolocation for a triple-satellite system under common scenario assumptions. The direct approach evaluates candidate emitter locations from the received I/Q data. The indirect approach first extracts receiver-differenced observables, associates them across the three receivers, and then solves for emitter position. Both approaches are evaluated on the same latency and accuracy basis.

The result is a controlled architecture-level assessment of when each concept is preferable and which design choices best support timely and accurate passive geolocation of stationary terrestrial GNSS jammers with three LEO satellites.

2

Methodology

This chapter presents the methodology used to analyse and compare direct and indirect geolocation of stationary terrestrial GNSS jammers with a triple-satellite LEO system. It first formulates the estimation problem and the underlying signal, measurement, and observable models. It then introduces the two geolocation approaches, the common comparison framework, the performance analysis used to assess the latency–accuracy trade-off, and the parametric simulation framework used for end-to-end evaluation.

2.1. Problem Formulation

This work considers passive geolocation of stationary terrestrial emitters from synchronized IQ data collected by three moving satellite receivers. The receivers are denoted by R_j, $j \in \{1, 2, 3\}$, and the emitters by E_i, $i = 1, \ldots, N_E$, where N_E is unknown. The objective is to estimate the two-dimensional ground position of each emitter from the three recorded complex-baseband datasets.

The receiver system is fully passive: the emitters do not cooperate, and their number is not known a priori. Each satellite is assumed to carry a GNSS receiver and to be synchronized to a common system time. Receiver position, velocity, and acceleration are assumed known exactly, so receiver-state uncertainty is neglected in the baseline formulation.

Unless stated otherwise, all position vectors in this chapter are expressed in Earth-Centered Earth-Fixed (ECEF) coordinates. Thus, $\boldsymbol{r}_{E_i}$ denotes the position of emitter E_i, $\boldsymbol{r}_{R_j}(t)$ the position of receiver R_j, and $\boldsymbol{r}_E$ a generic emitter position when the emitter index is not required.

2.1.1. Geometry

The multi-emitter, multi-receiver geometry considered in this work is shown in Figure 2.1. Terrestrial emitters are denoted by E_i, and spaceborne receivers by R_j.

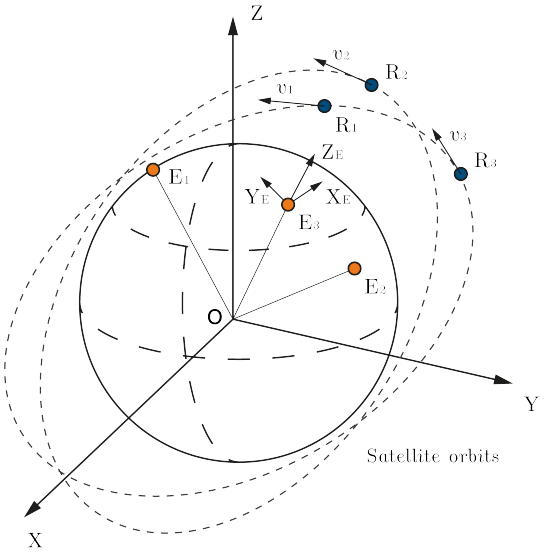

Figure 2.1: Geometry for multi-emitter and multi-receiver positioning

For emitter E_i and receiver R_j, define the relative position, range, and line-of-sight (LOS) unit vector as

$$\boldsymbol{\Delta}_{ij}(t) = \boldsymbol{r}_{E_i} - \boldsymbol{r}_{R_j}(t), \qquad \rho_{ij}(t) = \|\boldsymbol{\Delta}_{ij}(t)\|, \qquad \hat{\boldsymbol{\ell}}_{ij}(t) = \frac{\boldsymbol{\Delta}_{ij}(t)}{\rho_{ij}(t)}. \tag{2.1}$$

Because the receivers move with respect to Earth-fixed emitters, each receiver observes a distinct propagation path to the same emitter. These geometric differences generate receiver-dependent delays and Doppler shifts, which are the basis of passive geolocation.

2.1.2. Satellite Formation

The three receivers are assumed to fly in a formation inspired by the HawkEye 360 concept [31]. The formation consists of a triad in a common near-circular reference orbit and is parameterized by the orbital altitude h, the reference inclination i_0, the along-track separation a, and the cross-track separation b. In addition, the formation is anchored by a reference epoch t_0 and a prescribed ground target location, which determine where the triad is placed along the orbit at initialization. A schematic overview is shown in Figure 2.2.

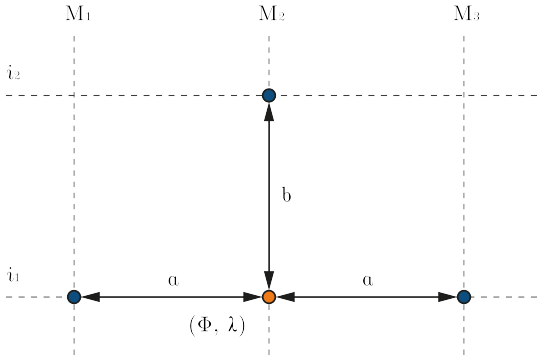

Figure 2.2: Overview of the satellite formation considered in this work

The reference orbit is assumed circular, with

$$r_0 = R_E + h, \qquad e = 0, \qquad \omega = 0, \tag{2.2}$$

where r_0 is the orbital radius, R_E the Earth radius, e the eccentricity, and ω the argument of perigee. At epoch t_0, the middle satellite is chosen such that its nadir direction is aligned with the prescribed

ground target. Let u denote its argument of latitude. For a circular orbit, the mean anomaly equals the argument of latitude, so the middle satellite uses

$$M_2 = u. \tag{2.3}$$

The two outer satellites lie in the same orbital plane and are placed symmetrically ahead of and behind the middle satellite. A small arc-length offset along the orbit corresponds approximately to a mean-anomaly offset, so that the along-track spacing a is realized as

$$\Delta M_{\frac{1}{2}} = \frac{a/2}{r_0}, \qquad M_1 = u + \Delta M_{\frac{1}{2}}, \qquad M_3 = u - \Delta M_{\frac{1}{2}}. \tag{2.4}$$

A non-collinear geometry is created by assigning the middle satellite a small inclination offset Δi relative to the two outer satellites. The resulting cross-track separation is approximately

$$b \approx r_0 \, |\sin u| \, \Delta i, \qquad \Delta i = \frac{b}{r_0 \, |\sin u|}. \tag{2.5}$$

The orbital inclinations are therefore

$$i_1 = i_3 = i_0, \qquad i_2 = i_0 + \Delta i. \tag{2.6}$$

Consequently, the initialized triad is completely specified by the six quantities

$$(h,\ i_0,\ a,\ b,\ t_0,\ \boldsymbol{r}_{E,\mathrm{ref}}), \tag{2.7}$$

where $\boldsymbol{r}_{E,\mathrm{ref}}$ denotes the reference target location used to place the formation. The inclination offset causes the middle satellite to oscillate relative to the two outer satellites, producing a time-varying cross-track baseline. This avoids a purely collinear receiver configuration and provides the spatial diversity required for passive geolocation from receiver-differenced observables. The influence of a and b on geolocation performance is analyzed in later chapters.

2.1.3. Emitter State

Each emitter E_i is parameterized in geodetic coordinates by latitude ϕ_i, longitude λ_i, and height h_i,

$$\boldsymbol{r}_{E_i,\mathrm{geo}} = \begin{bmatrix} \phi_i \\ \lambda_i \\ h_i \end{bmatrix}. \tag{2.8}$$

Although only two-dimensional ground geolocation is considered, the full geodetic state is introduced for consistency with the general model.

Using the WGS-84 ellipsoid [37], define

$$e^2 = f(2 - f), \qquad N(\phi_i) = \frac{a}{\sqrt{1 - e^2 \sin^2 \phi_i}}, \tag{2.9}$$

where a is the semi-major axis and f the flattening. The corresponding ECEF position is

$$\boldsymbol{r}_{E_i} = \begin{bmatrix} \big(N(\phi_i) + h_i\big) \cos\phi_i \cos\lambda_i \\ \big(N(\phi_i) + h_i\big) \cos\phi_i \sin\lambda_i \\ \big(N(\phi_i)(1 - e^2) + h_i\big) \sin\phi_i \end{bmatrix}. \tag{2.10}$$

Throughout this work, emitters are assumed stationary over the observation interval,

$$\dot{\boldsymbol{r}}_{E_i} = \boldsymbol{0}, \qquad \ddot{\boldsymbol{r}}_{E_i} = \boldsymbol{0}. \tag{2.11}$$

2.1.4. Receiver State

Each satellite receiver R_j is parameterized at epoch t_0 by the classical orbital elements [38]

$$(a_j, e_j, i_j, \Omega_j, \omega_j, M_{0,j}), \qquad j = 1, 2, 3. \tag{2.12}$$

The receiver state is first propagated in the Earth-Centered Inertial (ECI) frame with a Keplerian two-body model, yielding

$$\boldsymbol{r}_{R_j,\mathrm{ECI}}(t), \qquad \dot{\boldsymbol{r}}_{R_j,\mathrm{ECI}}(t), \qquad \ddot{\boldsymbol{r}}_{R_j,\mathrm{ECI}}(t). \tag{2.13}$$

The ECI position is rotated to ECEF with the Greenwich Mean Sidereal Time (GMST) angle $\theta(t)$,

$$\boldsymbol{r}_{R_j}(t) = \boldsymbol{R}_3\big(\theta(t)\big)\, \boldsymbol{r}_{R_j,\mathrm{ECI}}(t). \tag{2.14}$$

The corresponding ECEF velocity is

$$\dot{\boldsymbol{r}}_{R_j}(t) = \boldsymbol{R}_3\big(\theta(t)\big)\, \dot{\boldsymbol{r}}_{R_j,\mathrm{ECI}}(t) - \boldsymbol{\omega}_E \times \boldsymbol{r}_{R_j}(t), \tag{2.15}$$

where

$$\boldsymbol{\omega}_E = \begin{bmatrix} 0 \\ 0 \\ \Omega_E \end{bmatrix} \tag{2.16}$$

and Ω_E is the Earth rotation rate.

Receiver position, velocity, and acceleration are treated as deterministic inputs. The full transformation from orbital elements to ECI and ECEF state vectors is given in Appendix A.

2.1.5. Estimation Objective

For a given observation interval, let

$$\tilde{y}_j[n], \qquad j = 1, 2, 3 \tag{2.17}$$

denote the complex-baseband IQ samples recorded by the three receivers. The estimation problem is to infer, from these synchronized datasets and the known receiver trajectories, the ground position of each stationary emitter within the common field of view of the three-satellite system.

The problem is difficult for three reasons: the number of emitters is unknown, the emitters do not cooperate, and each receiver observes a superposition of delayed, Doppler-shifted, and noisy emitter contributions. The following sections introduce the signal model, the observables, and the estimation methods used to simulate and solve this problem.

2.2. Signal and Measurement Model

This section defines the signal and measurement model used throughouts. It first specifies the jammer signal classes considered and then describes how these signals are observed by the moving satellite receivers in complex baseband, including the effects of propagation, receiver clock behaviour, noise, and quantization. Finally, it identifies the relevant error sources and clarifies which terms cancel under receiver differencing, which are modeled explicitly, and which are excluded from the baseline analysis.

2.2.1. Emitter Signal Model

A signal transmitted by emitter E_i is modeled as

$$x_i(t) = s_i(t) e^{j2\pi f_c t}, \tag{2.18}$$

where $s_i(t)$ is the complex envelope, f_c is the carrier frequency, and the envelope bandwidth satisfies $B \ll f_c$. This analytic signal model is appropriate for GNSS-band signals (1164–1610) MHz [11]. As stated in section 1.4, three jammer types are considered: single-frequency, sweep, and modulated-signal jammers.

Single-frequency Jammer

A single-frequency jammer is modeled at baseband offset f_0 as

$$s_i(t) = e^{j2\pi f_0 t}. \tag{2.19}$$

Sweep Jammer
A sweep jammer is modeled as a chirp with instantaneous baseband frequency

$$f(t) = f_0 + kt, \tag{2.20}$$

where k is the sweep rate. The corresponding baseband signal is

$$s_i(t) = \exp\left(j2\pi \left(f_0 t + \tfrac{1}{2}kt^2\right)\right). \tag{2.21}$$

Modulated-Signal Jammer
A modulated-signal jammer is modeled by a periodic chip sequence $b[k] \in \{-1, +1\}$ with chip period T_c and code length N_c. With

$$k(t) = \left\lfloor \frac{t}{T_c} \right\rfloor \bmod N_c, \tag{2.22}$$

the unit-power baseband waveform is

$$s_i(t) = b[k(t)]. \tag{2.23}$$

2.2.2. Receiver Model

For a stationary emitter, the relative kinematics are

$$\dot{\boldsymbol{\Delta}}_{ij}(t) = -\dot{\boldsymbol{r}}_{R_j}(t), \qquad \ddot{\boldsymbol{\Delta}}_{ij}(t) = -\ddot{\boldsymbol{r}}_{R_j}(t) \tag{2.24}$$

The range rate is therefore

$$\dot{\rho}_{ij}(t) = \hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t) = -\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{r}}_{R_j}(t) \tag{2.25}$$

Define the geometric delay and geometric Doppler as

$$\tau_{ij}^{\mathrm{g}}(t) = \frac{\rho_{ij}(t)}{c}, \qquad f_{ij}^{\mathrm{g}}(t) = -\frac{f_c}{c}\dot{\rho}_{ij}(t) \tag{2.26}$$

The received contribution of emitter E_i at receiver R_j is modeled as

$$y_{ij}(t) = a_{ij}(t)\, s_i\left(t - \tilde{\tau}_{ij}(t)\right) e^{j\phi_{ij}(t)} \tag{2.27}$$

where $a_{ij}(t)$ is a time-varying amplitude, $\tilde{\tau}_{ij}(t)$ the effective delay, and $\phi_{ij}(t)$ the accumulated phase. Each receiver is assumed synchronized to GNSS time at the start of every snapshot. The deterministic receiver clock offset and the constant receiver frequency offset are therefore assumed compensated at $t = 0$. The remaining uncompensated receiver clock error is modeled as a residual intra-snapshot stochastic drift process. Let $\delta\dot{t}_{R_j}^{\mathrm{res}}(t)$ denote the residual fractional frequency error of receiver R_j. At the sampling rate $f_s = 1/T_s$, it is modeled as a discrete-time random walk [12],

$$\delta\dot{t}_{R_j}^{\mathrm{res}}[n] = \delta\dot{t}_{R_j}^{\mathrm{res}}[n-1] + \eta_j[n], \qquad \eta_j[n] \sim \mathcal{N}(0, q_{i,j}), \tag{2.28}$$

with

$$q_{i,j} = 2\pi^2 h_{-2,j}\, T_s, \qquad \delta\dot{t}_{R_j}^{\mathrm{res}}[0] = 0 \tag{2.29}$$

where $h_{-2,j}$ is the random-walk frequency-noise coefficient of receiver R_j. The corresponding residual receiver clock offset evolves as

$$\delta t_{R_j}^{\mathrm{res}}[n] = \delta t_{R_j}^{\mathrm{res}}[n-1] + \delta\dot{t}_{R_j}^{\mathrm{res}}[n]\, T_s, \qquad \delta t_{R_j}^{\mathrm{res}}[0] = 0 \tag{2.30}$$

After generation, both sequences are re-centered so that $\delta\dot{t}_{R_j}^{\mathrm{res}} = 0$ and $\delta t_{R_j}^{\mathrm{res}} = 0$ at the snapshot center epoch, enforcing the assumption that the deterministic offset and rate are compensated at that instant. The effective delay is modeled as [35]

$$\tilde{\tau}_{ij}(t) = \tau_{ij}^{\mathrm{g}}(t) + \left(\delta t_{E_i} - \delta t_{R_j}^{\mathrm{res}}(t)\right) + \frac{I_{ij}(t) + T_{ij}(t)}{c} + \left(d_{E_i} + d_{R_j}\right) + \epsilon_{\mathrm{mp},ij}(t) + \epsilon_{\tau,ij}(t) \tag{2.31}$$

where δt_{E_i} is the emitter clock offset, $I_{ij}(t)$ and $T_{ij}(t)$ are ionospheric and tropospheric delays, d_{E_i} and d_{R_j} are hardware group delays, and the remaining terms model multipath and delay-model error. The effective frequency shift is modeled as

$$\tilde{f}_{ij}(t) = f_{ij}^{\mathrm{g}}(t) + f_c\left(\delta\dot{t}_{E_i} - \delta\dot{t}_{R_j}^{\mathrm{res}}(t)\right) - \frac{f_c}{c}\left(\epsilon_{\mathrm{mp}',ij}(t) + \epsilon_{f,ij}(t)\right) \tag{2.32}$$

where $\delta\dot{t}_{E_i}$ is the emitter clock drift. By construction, no uncompensated constant receiver frequency-offset term remains at the snapshot center; only the residual stochastic intra-snapshot receiver drift is retained. The phase is related to the instantaneous frequency through

$$\phi_{ij}(t) = 2\pi \int_0^t \tilde{f}_{ij}(\xi)\, d\xi \tag{2.33}$$

The initial phase $\phi_{ij,0}$ is set to zero for all emitter–receiver pairs. This is adequate when geolocation relies on frequency-domain observables, since FDOA is independent of absolute phase.

Complex-Baseband Sampling and Quantization

A receiver observes the superposition of all visible emitters plus additive noise [35],

$$y_j(t) = \sum_{i=1}^{N_E} y_{ij}(t) + n_j(t) \tag{2.34}$$

where $n_j(t)$ is complex receiver noise. A snapshot of duration T is sampled at rate f_s, giving $N_s = f_s T$ samples on the centered grid

$$T_s = \frac{1}{f_s}, \qquad t_n = \left(n - \frac{N_s - 1}{2}\right) T_s, \qquad n = 0, \ldots, N_s - 1 \tag{2.35}$$

Centering the grid at the snapshot epoch simplifies the modeling of intra-snapshot dynamics. Within a snapshot, two classes of time-varying quantities are distinguished. Orbital quantities, receiver position, velocity, and the resulting geometric delay, Doppler, and amplitude, are evaluated at uniformly spaced block centers and linearly interpolated to each sample time t_n. The residual receiver clock process, being a stochastic random walk, is generated directly at the sampling rate without interpolation. The sampled signal is

$$y_j[n] = \sum_{i=1}^{N_E} a_{ij}[n]\, s_i\left(t_n - \tilde{\tau}_{ij}[n]\right) e^{j\phi_{ij}(t_n)} + n_j[n] \tag{2.36}$$

The accumulated phase at sample n is computed by trapezoidal integration of the effective frequency,

$$\phi_{ij}(t_n) = 2\pi\, T_s \sum_{m=1}^{n} \frac{\tilde{f}_{ij}[m] + \tilde{f}_{ij}[m-1]}{2}, \qquad \phi_{ij}(t_0) = 0 \tag{2.37}$$

The discrete-time noise is modeled as circularly symmetric complex Gaussian white noise,

$$n_j[n] \sim \mathcal{CN}(0, \sigma^2) \tag{2.38}$$

with variance σ^2 per complex sample. To model finite A/D resolution, the in-phase and quadrature channels are quantized separately,

$$\tilde{y}_j[n] = Q_b\left(\Re\{y_j[n]\}\right) + j\, Q_b\left(\Im\{y_j[n]\}\right) \tag{2.39}$$

where $Q_b(\cdot)$ is a b-bit scalar quantizer. The sequence $\tilde{y}_j[n]$ is therefore a b-bit, f_s-sample/s complex-baseband IQ snapshot of length N_s.

2.2.3. Error Sources

Table 2.1 summarizes the error sources for a single emitter–receiver link. In this work, however, geolocation is based on receiver-differenced observables. Emitter-specific terms therefore cancel, common terms are substantially reduced, and receiver-specific terms remain [35, 11]. TDOA, FDOA, and DFDOA are all retained as candidate observables in the analytical framework. The implemented indirect estimator uses only FDOA operationally, but DFDOA is still considered because Doppler-rate is often used jointly with Doppler in the literature.

Table 2.1: Candidate geolocation error sources, their dominant character, affected observable, and typical literature magnitude

Error source	Dominant character	Affected observable	Typical magnitude in literature
Receiver position	Random	Both	3–3.34 m [41, 23]
Receiver velocity	Random	FDOA	0.02 m/s [23]
Receiver clock synchronization offset	Random	TDOA	< 1 ns [28]
Receiver clock drift	Random	FDOA	< 1 Hz [9]
Receiver group delay	Systematic	TDOA	< 10 ns [17]
Emitter clock offset	Random	TDOA	-
Emitter frequency offset	Random	FDOA	-
Emitter clock drift	Random	FDOA	-
Emitter group delay	Random	TDOA	-
Ionospheric delay	Mostly systematic	TDOA	10–50 m [22]
Tropospheric delay	Mostly systematic	TDOA	< 10 m [3]
Multipath	Mixed	Both	-
Sagnac / Earth-rotation correction	Systematic	TDOA	-
Antenna phase-center offset	Systematic	TDOA	-
Terrain / elevation uncertainty	Random	Indirect	26 m [11]
Thermal noise	Random	Both	-

This leaves a smaller set of residual error sources that are relevant in principle for FDOA-based geolocation. Not all of them are modeled. Errors that depend strongly on mission-specific navigation performance or local propagation environment are excluded, so that the comparison between geolocation estimators is not blurred by external scenario assumptions.

Table 2.2: Residual error sources relevant to the baseline FDOA model and their treatment

Error source	Modeled	Treatment in this work
Receiver position error	Excluded	Perfect receiver position knowledge assumed
Receiver velocity error	Excluded	Perfect receiver velocity knowledge assumed
Receiver clock drift	Included	Gaussian random-walk process
Multipath	Excluded	Highly scenario-dependent and difficult to parameterize generically
Thermal noise	Included	Circularly symmetric complex Gaussian white noise

2.3. System Observables

Different receivers observe the same emitter with different delays and frequency shifts because their geometries differ. Receiver differencing produces observables that suppress emitter-specific error terms and are therefore better suited for geolocation. This work considers TDOA, FDOA, and DFDOA as candidate observables. For any receiver-dependent quantity $q_{ij}(t)$, define the pairwise difference

$$\Delta q_{i,jk}(t) = q_{ij}(t) - q_{ik}(t) \tag{2.40}$$

2.3.1. Time Difference of Arrival (TDOA)

For emitter E_i observed by receivers R_j and R_k, the time difference of arrival is

$$\Delta\tau_{i,jk}(t) = \tilde{\tau}_{ij}(t) - \tilde{\tau}_{ik}(t) \tag{2.41}$$

Substituting the delay model gives

$$\Delta\tau_{i,jk}(t) = \frac{\rho_{ij}(t) - \rho_{ik}(t)}{c} + \left(\delta t^{\mathrm{res}}_{R_k}(t) - \delta t^{\mathrm{res}}_{R_j}(t)\right) + \frac{\Delta I_{i,jk}(t) + \Delta T_{i,jk}(t)}{c} + \left(d_{R_j} - d_{R_k}\right) + \Delta\epsilon_{\tau,i,jk}(t) \tag{2.42}$$

with

$$\Delta\epsilon_{\tau,i,jk}(t) = \epsilon_{\mathrm{mp},ij}(t) - \epsilon_{\mathrm{mp},ik}(t) + \epsilon_{\tau,ij}(t) - \epsilon_{\tau,ik}(t) \tag{2.43}$$

Receiver differencing removes all emitter-specific delay terms, including the emitter clock offset δt_{E_i} and emitter group delay d_{E_i}. The remaining TDOA depends primarily on geometric range difference, together with residual receiver clock, hardware, atmospheric, and multipath terms.

For periodic or cyclostationary signals, TDOA may be ambiguous up to an integer multiple of the signal period,

$$\Delta\tau_{i,jk}(t) = \Delta\tau^{\star}_{i,jk}(t) + mT_p, \qquad m \in \mathbb{Z}, \tag{2.44}$$

where T_p is the signal period. Treating the error terms as additive and position-independent, the gradient of the single-receiver delay with respect to the emitter position $\boldsymbol{r}_E$ is

$$\nabla\tau_{ij}(t) = \frac{1}{c}\hat{\boldsymbol{\ell}}^{\mathsf{T}}_{ij}(t). \tag{2.45}$$

Hence,

$$\nabla\Delta\tau_{i,jk}(t) = \nabla\tau_{ij}(t) - \nabla\tau_{ik}(t). \tag{2.46}$$

The full derivation is given in section B.1.

2.3.2. Frequency Difference of Arrival (FDOA)

For emitter E_i observed by receivers R_j and R_k, the frequency difference of arrival is

$$\Delta f_{i,jk}(t) = \tilde{f}_{ij}(t) - \tilde{f}_{ik}(t). \tag{2.47}$$

Substituting the frequency-shift model gives

$$\Delta f_{i,jk}(t) = -\frac{f_c}{c}\left[\dot{\rho}_{ij}(t) - \dot{\rho}_{ik}(t) + c\left(\delta\dot{t}^{\mathrm{res}}_{R_j}(t) - \delta\dot{t}^{\mathrm{res}}_{R_k}(t)\right) + \Delta\epsilon_{f,i,jk}(t)\right], \tag{2.48}$$

where

$$\Delta\epsilon_{f,i,jk}(t) = \epsilon_{\mathrm{mp}',ij}(t) - \epsilon_{\mathrm{mp}',ik}(t) + \epsilon_{f,ij}(t) - \epsilon_{f,ik}(t). \tag{2.49}$$

As in the TDOA case, receiver differencing removes the emitter clock-drift term $\delta\dot{t}_{E_i}$. The remaining FDOA depends mainly on geometric range-rate difference and is particularly informative for LEO geolocation because orbital velocities are large.

For compactness, define the orthogonal projector

$$\boldsymbol{P}^{\perp}_{ij}(t) = \boldsymbol{I} - \hat{\boldsymbol{\ell}}_{ij}(t)\hat{\boldsymbol{\ell}}^{\mathsf{T}}_{ij}(t). \tag{2.50}$$

Then the gradient of the range rate is

$$\nabla\dot{\rho}_{ij}(t) = \left[\frac{1}{\rho_{ij}(t)}\,\boldsymbol{P}^{\perp}_{ij}(t)\,\dot{\boldsymbol{\Delta}}_{ij}(t)\right]^{\mathsf{T}}, \tag{2.51}$$

and the FDOA gradient becomes

$$\nabla\Delta f_{i,jk}(t) = -\frac{f_c}{c}\left(\nabla\dot{\rho}_{ij}(t) - \nabla\dot{\rho}_{ik}(t)\right). \tag{2.52}$$

The full derivation is given in section B.2.

2.3.3. Differential Frequency Difference of Arrival (DFDOA)

For emitter E_i observed by receivers R_j and R_k, the differential frequency-difference rate is

$$\Delta\dot{f}_{i,jk}(t) = \dot{\tilde{f}}_{ij}(t) - \dot{\tilde{f}}_{ik}(t). \tag{2.53}$$

DFDOA is included because Doppler-rate is often considered jointly with Doppler in the geolocation literature.

For a stationary emitter, the range acceleration is

$$\ddot{\rho}_{ij}(t) = \hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\ddot{\boldsymbol{\Delta}}_{ij}(t) + \frac{\left\|\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right\|^2}{\rho_{ij}(t)}. \tag{2.54}$$

DFDOA depends primarily on the difference in these range accelerations and may provide additional geometric information.

The corresponding gradient is

$$\nabla\ddot{\rho}_{ij}(t) = \left[\frac{\boldsymbol{P}_{ij}^{\perp}(t)\ddot{\boldsymbol{\Delta}}_{ij}(t)}{\rho_{ij}(t)} - \frac{2\left(\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right)\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t) + \left\|\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right\|^2\hat{\boldsymbol{\ell}}_{ij}(t)}{\rho_{ij}^2(t)}\right]^{\mathsf{T}}. \tag{2.55}$$

Hence,

$$\nabla\Delta\dot{f}_{i,jk}(t) = -\frac{f_c}{c}\left(\nabla\ddot{\rho}_{ij}(t) - \nabla\ddot{\rho}_{ik}(t)\right). \tag{2.56}$$

The full derivation is given in section B.3.

2.4. Geolocation Estimation Methods

This section presents the two geolocation approaches considered in this work and the framework used to compare them fairly. Direct geolocation estimates emitter position by evaluating signal consistency directly in the position domain, whereas indirect geolocation first extracts and associates differential observables before solving for position. Together, these methods represent two fundamentally different ways of converting the same synchronized IQ measurements into location estimates, with different trade-offs in information use, computational cost, and latency.

2.4.1. Direct Geolocation

As a benchmark for the proposed indirect method, this work also considers direct geolocation based on position-domain correlation [40, 10]. Instead of first estimating explicit TDOA or FDOA observables, direct geolocation evaluates signal alignment directly for each hypothesized emitter location.

A two-dimensional search grid is defined over latitude and longitude. The candidate emitter height is fixed to zero. Let

$$\mathcal{G} = \{\boldsymbol{r}_E\} \tag{2.57}$$

denote the set of candidate ground positions, where each $\boldsymbol{r}_E \in \mathcal{G}$ is obtained by mapping the WGS-84 grid point $(\phi, \lambda, 0)$ to ECEF coordinates.

For each candidate $\boldsymbol{r}_E \in \mathcal{G}$ and receiver pair (R_j, R_k), the known receiver trajectories yield the predicted differential delay and Doppler histories,

$$\Delta\tau_{jk}(\boldsymbol{r}_E, t_n), \qquad \Delta f_{jk}(\boldsymbol{r}_E, t_n), \qquad n = 0, \ldots, N_s - 1. \tag{2.58}$$

The associated accumulated differential phase is

$$\Delta\phi_{jk}(\boldsymbol{r}_E, n) = 2\pi T_s \sum_{m=0}^{n} \Delta f_{jk}(\boldsymbol{r}_E, t_m). \tag{2.59}$$

Let $\tilde{y}_k(t)$ denote an interpolant of the sampled sequence $\tilde{y}_k[n]$. The delay-compensated version of receiver R_k is

$$\tilde{y}_k^{(\boldsymbol{r}_E)}[n] = \tilde{y}_k\big(t_n - \Delta\tau_{jk}(\boldsymbol{r}_E, t_n)\big)\,. \tag{2.60}$$

The pairwise position-domain score is

$$S_{jk}(\boldsymbol{r}_E) = \left| \sum_{n=0}^{N_s-1} \tilde{y}_j[n] \left(\tilde{y}_k^{(\boldsymbol{r}_E)}[n]\right)^* e^{-j\Delta\phi_{jk}(\boldsymbol{r}_E, n)} \right|. \tag{2.61}$$

For three receivers, the total score is

$$S(\boldsymbol{r}_E) = S_{12}(\boldsymbol{r}_E) + S_{13}(\boldsymbol{r}_E) + S_{23}(\boldsymbol{r}_E). \tag{2.62}$$

If multiple snapshots are available, non-coherent integration is performed as

$$S_{\text{tot}}(\boldsymbol{r}_E) = \sum_{m=1}^{M} S^{(m)}(\boldsymbol{r}_E), \tag{2.63}$$

where $S^{(m)}(\boldsymbol{r}_E)$ is the score map for snapshot m. The direct geolocation estimate is

$$\hat{\boldsymbol{r}}_E = \arg\max_{\boldsymbol{r}_E \in \mathcal{G}} S_{\text{tot}}(\boldsymbol{r}_E). \tag{2.64}$$

Figure 2.3 shows an example single-snapshot score map. The dominant peak indicates the most likely emitter location, while the ridge structures reflect TDOA- and FDOA-consistent candidate locations.

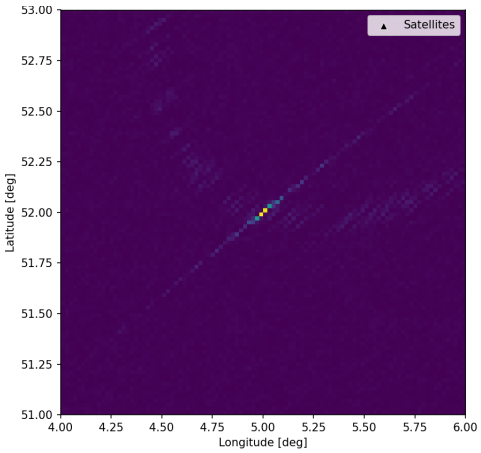

Figure 2.3: Example single-snapshot position-domain score map for direct geolocation

The main advantage of direct geolocation is that it avoids separate observable-extraction and data-association stages. Its main disadvantage is computational cost, since the score must be evaluated at every grid point.

2.4.2. Indirect Geolocation

Indirect geolocation first extracts differential observables from the received IQ data and then estimates emitter position from those observables. This reduces computational cost relative to direct geolocation but introduces three additional challenges: observable extraction, per-snapshot candidate association, and temporal association across snapshots. In the passive multi-receiver, multi-emitter setting considered here, these stages are non-trivial and form a central part of the methodological contribution.

Although section 2.3 defines TDOA, FDOA, and DFDOA, the proposed indirect estimator uses only FDOA operationally. DFDOA is retained only as a candidate observable in the analytical framework. Many GNSS jammers are cyclostationary and can produce ambiguous delay peaks, so avoiding explicit

TDOA extraction prevents this ambiguity from propagating into later stages. In addition, FDOA resolution improves with coherent integration time, whereas TDOA resolution is fundamentally limited by signal bandwidth (subsection 2.5.1). Finally, the high orbital velocities of LEO receivers create substantial Doppler diversity, making FDOA particularly informative for space-based geolocation.

Cross-Ambiguity Function (CAF)
For each receiver pair (R_j, R_k), the cross-ambiguity function is evaluated over differential delay τ and differential frequency f [42],

$$\mathcal{A}_{jk}(\tau, f) = \left| \sum_{n=0}^{N_s-1} \tilde{y}_j[n]\, \tilde{y}_k^*(t_n - \tau)\, e^{-j2\pi f t_n} \right|, \tag{2.65}$$

where $\tilde{y}_k(t)$ again denotes an interpolant of the sampled sequence. Within one coherent processing interval, τ and f are assumed constant. Under this approximation, the CAF forms a map in the TDOA–FDOA plane whose dominant peaks correspond to candidate observables. The CAF can be evaluated efficiently with the FFT.

To suppress physically impossible peaks, the search is restricted to the region implied by the receiver geometry. Let

$$\mathcal{S}_{\text{FoV}} = \bigcap_{j=1}^{3} \mathcal{V}_j \tag{2.66}$$

denote the shared field of view, where $\mathcal{V}_j$ is the visible ground region of receiver R_j. For a hypothetical emitter at $\boldsymbol{r}_E$,

$$\Delta\tau_{jk}(\boldsymbol{r}_E) = \frac{\rho_j(\boldsymbol{r}_E) - \rho_k(\boldsymbol{r}_E)}{c}, \qquad \Delta f_{jk}(\boldsymbol{r}_E) = -\frac{f_c}{c}\left(\dot{\rho}_j(\boldsymbol{r}_E) - \dot{\rho}_k(\boldsymbol{r}_E)\right). \tag{2.67}$$

The valid search intervals are

$$\Delta\tau_{jk} \in \left[\min_{\boldsymbol{r}_E \in \mathcal{S}_{\text{FoV}}} \Delta\tau_{jk}(\boldsymbol{r}_E),\ \max_{\boldsymbol{r}_E \in \mathcal{S}_{\text{FoV}}} \Delta\tau_{jk}(\boldsymbol{r}_E) \right], \tag{2.68}$$

$$\Delta f_{jk} \in \left[\min_{\boldsymbol{r}_E \in \mathcal{S}_{\text{FoV}}} \Delta f_{jk}(\boldsymbol{r}_E),\ \max_{\boldsymbol{r}_E \in \mathcal{S}_{\text{FoV}}} \Delta f_{jk}(\boldsymbol{r}_E) \right]. \tag{2.69}$$

An example single-snapshot CAF is shown in Figure 2.4. This particular CAF contains a single symmetric sweep jammer, which appears as two diagonal ridges intersecting at the true FDOA and TDOA.

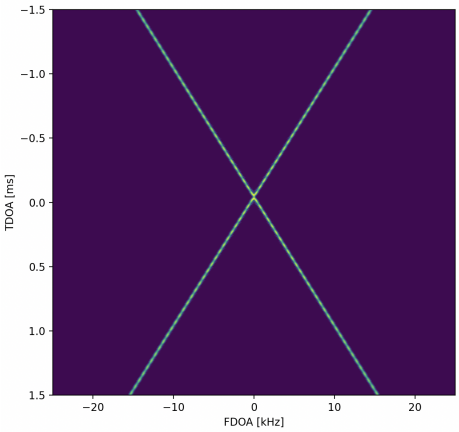

Figure 2.4: Example single-snapshot TDOA–FDOA CAF map for indirect geolocation

When multiple emitters are present, the CAF contains self-terms, cross-terms, and noise [5],

$$\mathcal{A}_{jk}(\tau, f) \approx \left| \sum_{i=1}^{N_E} \mathcal{A}_{jk}^{(i)}(\tau, f) + \sum_{i \neq \ell} C_{jk}^{(i,\ell)}(\tau, f) + \mathcal{N}_{jk}(\tau, f) \right|. \tag{2.70}$$

Strong emitters can mask weak ones, and cyclostationary signals can produce repeated CAF structures, complicating extraction.

FDOA Candidate Extraction
For each receiver pair (R_j, R_k), FDOA candidates are extracted from the CAF

$$\mathcal{A}_{jk}(\tau, f). \tag{2.71}$$

The output is a finite candidate set

$$\mathcal{F}_{jk} = \left\{\hat{\Delta f}_{jk}^{(p)}\right\}_{p=1}^{N_{jk}}, \tag{2.72}$$

which contains the promising FDOA values for that receiver pair. The extraction is performed in four steps.

1) Peak detection A noise floor is estimated. Two-dimensional local maxima are then detected subject to a minimum level above the noise floor and a minimum local prominence. This yields a raw peak set

$$\mathcal{P}_{jk}^{\text{raw}} = \left\{(f^{(m)}, \tau^{(m)}, \gamma^{(m)})\right\}_{m=1}^{M_{jk}}, \tag{2.73}$$

where $f^{(m)}$, $\tau^{(m)}$, and $\gamma^{(m)}$ denote frequency, delay, and SNR.

2) Separation of diagonal structures The CAF map is searched for diagonal line segments,

$$\mathcal{L}_{jk} = \{L_\ell\}_{\ell=1}^{N_L}. \tag{2.74}$$

Each raw peak is classified by its distance to the nearest detected diagonal. Peaks close to a diagonal are removed from the ordinary peak set and stored with that diagonal, while the remaining peaks are retained as point-like peaks.

3) Conversion to FDOA candidates The surviving point-like peaks are merged along the frequency axis, retaining only the strongest representative within a small FDOA neighbourhood. This produces the set

$$\mathcal{F}_{jk}^{\text{peak}}. \tag{2.75}$$

In parallel, three additional candidate types are extracted:

$$\mathcal{F}_{jk}^{\text{vert}}, \qquad \mathcal{F}_{jk}^{\text{int}}, \qquad \mathcal{F}_{jk}^{\text{single}}, \tag{2.76}$$

where $\mathcal{F}_{jk}^{\text{vert}}$ contains candidates from strong near-vertical CAF columns, $\mathcal{F}_{jk}^{\text{int}}$ candidates from intersections of diagonals with opposite slope, and $\mathcal{F}_{jk}^{\text{single}}$ one representative candidate for each unmatched diagonal. In all cases, only the FDOA coordinate is retained.

4) Final candidate set The final candidate set is

$$\mathcal{F}_{jk} = \mathcal{F}_{jk}^{\text{vert}} \cup \mathcal{F}_{jk}^{\text{peak}} \cup \mathcal{F}_{jk}^{\text{int}} \cup \mathcal{F}_{jk}^{\text{single}}. \tag{2.77}$$

Candidates that remain closer than a prescribed FDOA tolerance are merged.

Besides reducing the CAF to a compact set of plausible FDOA values, this procedure also preserves the origin of each candidate (vertical, point-like, diagonal-intersection, or single-diagonal). This origin can be exploited later as auxiliary information for emitter signal-type classification, which is an additional benefit of the indirect geolocation approach.

FDOA Candidate Association

With more than two receivers, each pair may produce multiple FDOA candidates, so it must be determined which candidates belong to the same emitter at a given snapshot. Pairwise differences generated by a common emitter satisfy a zero cyclic-sum condition along any closed receiver loop. While the exploitation of such consistency conditions is known in other fields, including microphone-array localization, its application to data association in passive multi-satellite geolocation is, to the best of the author's knowledge, novel [44, 32]. A proof is given in Appendix C. For three receivers that observe the same emitter, the TDOA and FDOA observables between each pair satisfy

$$\Delta\tau_{12} + \Delta\tau_{23} + \Delta\tau_{31} = 0, \qquad \Delta f_{12} + \Delta f_{23} + \Delta f_{31} = 0. \tag{2.78}$$

Only the FDOA closure relation is used operationally. Let

$$\mathbf{b} = \left(\hat{\Delta f}_{12}^{(p)},\ \hat{\Delta f}_{23}^{(q)},\ \hat{\Delta f}_{31}^{(r)}\right) \tag{2.79}$$

denote a candidate triplet. Its closure residual is

$$\eta_f(\mathbf{b}) = \left|\hat{\Delta f}_{12}^{(p)} + \hat{\Delta f}_{23}^{(q)} + \hat{\Delta f}_{31}^{(r)}\right|. \tag{2.80}$$

The triplet is accepted if $\eta_f(\mathbf{b}) \le \epsilon_f$, where ϵ_f accounts for noise and finite CAF resolution. Accepted triplets are treated as per-snapshot emitter observations. The principle extends to $N > 3$ receivers by enforcing zero closure on the independent cycles of the receiver graph. Figure 2.5 illustrates the three-receiver case.

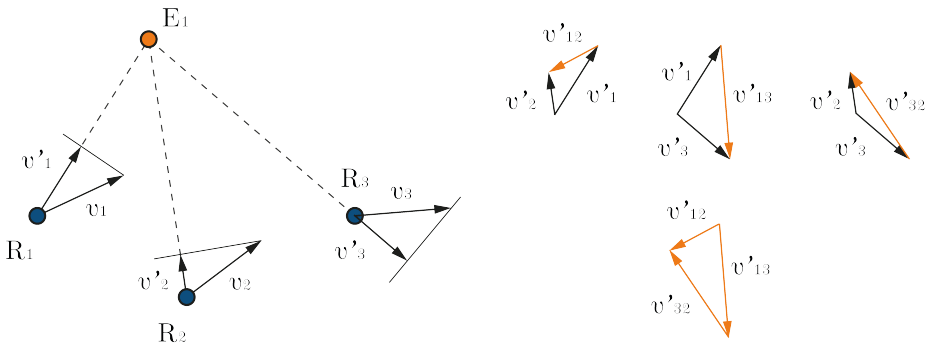

(a) Observed frequency shift f_i' **(b)** Closed-loop frequency consistency

Figure 2.5: Zero cyclic-sum condition for observed frequency

The same principle scales directly to more than three receivers. For $N_R > 3$, the receivers form a graph whose edges are pairwise FDOA candidate sets. Any candidate assignment that originates from one emitter must satisfy the zero-sum constraint on every closed receiver loop,

$$\sum_{(j,k)\in C_\ell} \hat{\Delta f}_{jk} = 0, \tag{2.81}$$

up to the tolerance ϵ_f. Thus, multiple independent loops can be checked in the same snapshot. This provides additional consistency tests without changing the underlying association principle. Figure 2.6 shows a six-receiver example with two possible closed loops.

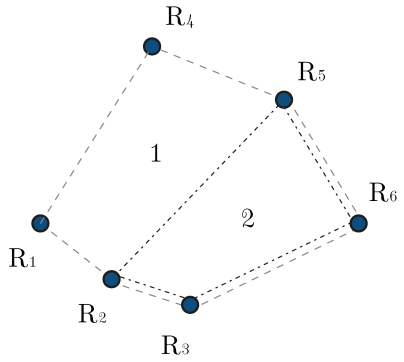

Figure 2.6: Six-receiver example with two closed FDOA consistency loops

FDOA Triplet Stacking

After per-snapshot association, each accepted triplet provides one three-pair FDOA observation at epoch t_m. These triplets must then be linked across snapshots to determine which observations originate from the same emitter over time.

Because the receiver geometry evolves smoothly, the pairwise FDOA history of a stationary emitter is also smooth. This property is used only for temporal association. In addition, the candidate origin retained during FDOA extraction (for example point-like, vertical, diagonal-intersection, or single-diagonal) can be used as auxiliary information during this matching stage. The scenarios considered here contains four emitters with distinct signal types, which makes temporal association comparatively simple. For each receiver pair, a cubic polynomial

$$p_{jk}(t) = a_{0,jk} + a_{1,jk}t + a_{2,jk}t^2 + a_{3,jk}t^3 \tag{2.82}$$

is fitted to the FDOA samples already assigned to a tentative emitter track. New triplets are then associated with existing tracks by consistency with these predicted FDOA values. The polynomial is therefore used only as a bookkeeping device for temporal track formation; it is not part of the final geolocation measurement model.

Once the temporal associations have been established, the raw FDOA measurements are stacked directly. Choosing the middle satellite R_2 as reference and retaining only an independent subset, the measurement vector for one emitter track observed at epochs $t_{m_1}, \ldots, t_{m_L}$ is

$$\mathbf{z} = \begin{bmatrix} \hat{\Delta f}_{12}(t_{m_1}) \\ \hat{\Delta f}_{23}(t_{m_1}) \\ \vdots \\ \hat{\Delta f}_{12}(t_{m_L}) \\ \hat{\Delta f}_{23}(t_{m_L}) \end{bmatrix}. \tag{2.83}$$

Only an independent subset is stacked, since the third receiver pair is redundant once closed-loop consistency has been enforced.

Nonlinear Position Estimation

The emitter position is estimated from the stacked raw FDOA measurements with a Gauss–Newton least-squares solver. For an emitter track observed at epochs $t_{m_1}, \ldots, t_{m_L}$, define the modelled measurement

vector

$$\mathbf{h}(\boldsymbol{r}_E) = \begin{bmatrix} \Delta f_{12}(\boldsymbol{r}_E, t_{m_1}) \\ \Delta f_{23}(\boldsymbol{r}_E, t_{m_1}) \\ \vdots \\ \Delta f_{12}(\boldsymbol{r}_E, t_{m_L}) \\ \Delta f_{23}(\boldsymbol{r}_E, t_{m_L}) \end{bmatrix}. \tag{2.84}$$

The estimate is obtained as

$$\hat{\boldsymbol{r}}_E = \arg\min_{\boldsymbol{r}_E \in \mathcal{S}_0} \|\mathbf{z} - \mathbf{h}(\boldsymbol{r}_E)\|_2^2, \tag{2.85}$$

where $\mathcal{S}_0$ is the Earth surface under the zero-height assumption.

The Gauss–Newton update at iteration i is

$$\boldsymbol{r}_E^{(i+1)} = \boldsymbol{r}_E^{(i)} + \left(\mathbf{H}^\mathsf{T}\mathbf{H}\right)^{-1} \mathbf{H}^\mathsf{T} \left(\mathbf{z} - \mathbf{h}\left(\boldsymbol{r}_E^{(i)}\right)\right), \tag{2.86}$$

where

$$\mathbf{H}\left(\boldsymbol{r}_E^{(i)}\right) = \begin{bmatrix} \nabla\Delta f_{12}(t_{m_1}) \\ \nabla\Delta f_{23}(t_{m_1}) \\ \vdots \\ \nabla\Delta f_{12}(t_{m_L}) \\ \nabla\Delta f_{23}(t_{m_L}) \end{bmatrix}_{\boldsymbol{r}_E = \boldsymbol{r}_E^{(i)}}. \tag{2.87}$$

After each iteration, the updated estimate is projected back to $\mathcal{S}_0$ to enforce the zero-height constraint.

2.4.3. Comparison Framework

The comparison is carried out in two stages.

First, the common physical setting is selected with the approach-independent analysis of section 2.5. This also screens the candidate observable sets.

Second, the two methods are compared end-to-end on the same simulated IQ datasets, receiver trajectories, and emitter scenarios. The two methods do not use the same information. Direct geolocation exploits joint TDOA/FDOA information through position-domain correlation. Indirect geolocation uses only FDOA after CAF-based extraction and association. A comparison at equal coherent integration time would therefore not be fair. It would favour the direct method in geolocation performance, because it uses additional TDOA information. It would also favour the indirect method in latency and processing cost, because the direct method would then be forced to use a longer integration time than necessary.

The coherent integration time is therefore chosen separately for the two approaches. The selected values make the FDOA-only indirect configuration and the T/FDOA direct configuration comparable in system-level performance at the chosen formation. This yields an apples-to-apples comparison.

The end-to-end comparison then isolates estimator-level differences under matched scenario conditions. These include extraction and association losses in the indirect chain, grid discretisation effects in the direct chain, and the computational cost of both approaches.

Two multi-emitter scenarios are used. The first reproduces a challenging configuration from the direct-geolocation literature. The second uses the same signal set but moves part of the emitter set farther from the formation centre, creating a more adverse geometry. The comparison is reported with three metrics: the minimum observation time required to localise all emitters, the final position error after the full observation window, and the wall-clock runtime.

2.5. System Performance Analysis

This section develops the approach-independent performance envelope: the geolocation accuracy that is fundamentally obtainable from the receiver geometry, signal properties, observation design, and observable set, irrespective of whether the final estimator is direct or indirect. Both approaches operate

on the same physical measurements and are therefore subject to the same limits. Estimator-specific effects, such as observable extraction errors, data-association failures, or search-grid discretisation, are not part of this analysis and are treated in chapter 3 through the comparison framework defined in subsection 2.4.3.

The accuracy depends on design variables chosen by the system designer and scenario parameters imposed by the emitter and environment. Table 2.3 lists both categories.

Table 2.3: Parameters affecting geolocation accuracy, grouped by category

Category	Parameter	Symbol	Influence
Design variables	Along-track separation	a	Geometric diversity and shared FoV
	Along-track / cross-track ratio	a/b	Cross-track geometric diversity and shared FoV
	Observable set	—	FIM rank and structure
	Total observation time	T_{obs}	Number of accumulated snapshots
	Snapshot interval	Δt_{snap}	Geometry change between snapshots
	Coherent integration time	T_{int}	Per-snapshot observable variance
Scenario parameters	Emitter transmit power	P_{tx}	Received SNR
	Signal bandwidth	B_c	TDOA variance (via $\bar{\beta}_2$)
	Carrier frequency	f_c	Path loss and Doppler scaling
	Receiver noise bandwidth	B_r	Noise power and observable variances

The CRLB links the scenario parameters and integration time analytically at the single-snapshot level (subsection 2.5.1). A physical upper bound on the integration time is derived in subsection 2.5.2. The remaining design variables affect how information accumulates across snapshots (subsection 2.5.3). Finally, subsection 2.5.4 defines a scalar coverage metric for systematic formation comparison.

2.5.1. Cramér–Rao Lower Bound

The CRLB provides the minimum achievable variance for an unbiased estimator of each observable. Following [18] for stationary constant-envelope signals and a rectangular integration window, the bounds are stated directly. These are consistent with Stein-type bounds for TDOA and FDOA [34] and additionally cover DFDOA. Related formulations appear in [15, 39, 19, 14, 24]. For receiver pair (R_j, R_k),

$$\sigma^2_{\tau,jk} \geq \frac{1}{\bar{\beta}_2^2 \, B_r \, T_{\text{int}} \, \text{SNR}_{\text{eff}}^{(jk)}}, \tag{2.88}$$

$$\sigma^2_{\dot{f},jk} \geq \frac{3}{\pi^2 \, B_r \, T_{\text{int}}^3 \, \text{SNR}_{\text{eff}}^{(jk)}}, \tag{2.89}$$

$$\sigma^2_{\dot{f},jk} \geq \frac{45}{\pi^2 \, B_r \, T_{\text{int}}^5 \, \text{SNR}_{\text{eff}}^{(jk)}}, \tag{2.90}$$

where $\bar{\beta}_2$ is the RMS signal bandwidth, B_r the receiver noise bandwidth, T_{int} the coherent integration time, and $\text{SNR}_{\text{eff}}^{(jk)}$ the effective pairwise SNR. The effective SNR includes the cross-correlation noise floor [34],

$$\frac{1}{\text{SNR}_{\text{eff}}^{(jk)}} = \frac{1}{2}\left(\frac{1}{\text{SNR}_{R_j}} + \frac{1}{\text{SNR}_{R_k}} + \frac{1}{\text{SNR}_{R_j}\,\text{SNR}_{R_k}}\right). \tag{2.91}$$

Each single-receiver SNR follows from a free-space link budget,

$$\text{SNR}_{R_j} = \frac{P_{\text{tx}} \, G_{\text{tx}} \, G_{\text{rx}}}{k_B \, T_{\text{sys}} \, F \, B_r} \left(\frac{c}{4\pi \rho_{ij} f_c}\right)^2, \tag{2.92}$$

where P_{tx} is the emitter transmit power, G_{tx} and G_{rx} antenna gains, k_B Boltzmann's constant, T_{sys} the system noise temperature, and F the noise figure. Since ρ_{ij} varies across the field of view, the CRLB variances are position-dependent.

TDOA variance is governed by the signal bandwidth $\bar{\beta}_2$ and improves as $T_{\text{int}}^{-1/2}$. FDOA variance is independent of signal bandwidth and improves as $T_{\text{int}}^{-3/2}$; DFDOA as $T_{\text{int}}^{-5/2}$. For narrow-bandwidth emitters, frequency-domain observables therefore benefit disproportionately from longer integration, while TDOA accuracy remains limited by $\bar{\beta}_2$.

2.5.2. Integration Time Constraints

Longer integration reduces the CRLB variances but also increases the computational cost of the CAF. Beyond computational considerations, there is a physical ceiling on T_{int} imposed by the assumption that both the differential delay and the differential Doppler remain approximately constant over the integration window.

Differential Delay Drift

The TDOA between receivers R_j and R_k evolves as

$$\Delta\tau_{jk}(t) \approx \Delta\tau_{jk}(t_0) + \Delta\dot{\tau}_{jk}(t_0)\,(t - t_0), \tag{2.93}$$

where $\Delta\dot{\tau}_{jk} = (\dot{\rho}_j - \dot{\rho}_k)/c$. Over a symmetric window of duration T_{int}, the maximum delay drift is

$$|\delta\tau|_{\max} = \frac{|\Delta\dot{\tau}_{jk}(t_0)|}{2}\,T_{\text{int}}. \tag{2.94}$$

Coherent integration requires $|\delta\tau|_{\max} \ll 1/B_c$, which yields

$$T_{\text{int}} \ll \frac{2}{B_c\,|\Delta\dot{\tau}_{jk}(t_0)|}. \tag{2.95}$$

For typical LEO range-rate differences of a few km/s and a signal bandwidth of 1 MHz, this bound is in the order of seconds and is therefore not the binding constraint.

Differential Doppler Drift

The FDOA evolves as

$$\Delta f_{jk}(t) \approx \Delta f_{jk}(t_0) + \Delta\dot{f}_{jk}(t_0)\,(t - t_0). \tag{2.96}$$

The CAF compensates only the constant component. The residual frequency error is

$$\delta f(t) = \Delta\dot{f}_{jk}(t_0)\,(t - t_0), \tag{2.97}$$

with maximum value

$$|\delta f|_{\max} = \frac{|\Delta\dot{f}_{jk}(t_0)|}{2}\,T_{\text{int}}. \tag{2.98}$$

This residual smears the CAF peak in frequency. Coherent integration remains effective as long as $|\delta f|_{\max}$ is small compared to the frequency resolution $1/T_{\text{int}}$. Requiring $|\delta f|_{\max} \ll 1/T_{\text{int}}$ gives

$$T_{\text{int}} \ll \sqrt{\frac{2}{|\Delta\dot{f}_{jk}(t_0)|}}. \tag{2.99}$$

Since $|\Delta\dot{f}_{jk}|$ increases with satellite separation, the maximum useful integration time decreases for wider formations. For the separations considered in this work, this Doppler-drift constraint is more restrictive than the delay-drift constraint and sets the practical ceiling on T_{int}. The integration time is therefore treated as a fixed design choice satisfying Equation 2.99.

2.5.3. Information Accumulation

Over an observation window T_{obs}, multiple snapshots are collected at intervals Δt_{snap}, yielding $M = \lfloor T_{\text{obs}}/\Delta t_{\text{snap}} \rfloor + 1$ epochs. At epoch t_m, the stacked measurement vector is

$$\mathbf{z}_m = \begin{bmatrix} z_{1,m} & \cdots & z_{N,m} \end{bmatrix}^{\mathsf{T}}, \tag{2.100}$$

where each entry is a TDOA, FDOA, or DFDOA observable from one receiver pair. For the indirect method considered in this work, each entry is a pairwise FDOA observable. The measurement covariance is

$$\mathbf{R}_m = \mathrm{diag}\left(\sigma_{1,m}^2, \ldots, \sigma_{N,m}^2\right), \tag{2.101}$$

with each entry set to the CRLB variance at the geometry of epoch t_m. The Jacobian is

$$\mathbf{G}_m(\boldsymbol{r}_E) = \begin{bmatrix} \nabla z_{1,m}(\boldsymbol{r}_E) \\ \vdots \\ \nabla z_{N,m}(\boldsymbol{r}_E) \end{bmatrix}, \tag{2.102}$$

and the Fisher information matrix accumulates additively,

$$\mathbf{J}_{\mathrm{tot}}(\boldsymbol{r}_E) = \sum_{m=1}^{M} \mathbf{G}_m^{\mathsf{T}}(\boldsymbol{r}_E)\, \mathbf{R}_m^{-1}\, \mathbf{G}_m(\boldsymbol{r}_E). \tag{2.103}$$

The CRLB covariance is $\mathbf{P}_{\mathrm{CRLB,ECEF}} = \mathbf{J}_{\mathrm{tot}}^{-1}$, transformed to the local ENU frame via the standard rotation [21]. The horizontal accuracy metric is

$$\mathrm{DRMS} = \sqrt{\mathbf{P}_{\mathrm{CRLB,ENU}}(1,1) + \mathbf{P}_{\mathrm{CRLB,ENU}}(2,2)}. \tag{2.104}$$

The information gain from each additional snapshot depends on how much the receiver geometry has changed. If the geometry evolves slowly relative to Δt_{snap}, successive Jacobians are nearly collinear and additional snapshots contribute limited new information. The effectiveness of a given snapshot interval is therefore coupled to the orbital velocity and formation geometry, and cannot be determined from the single-snapshot CRLB alone.

The total observation time controls how many geometrically distinct snapshots are sampled. With sufficient T_{obs}, even a single observable type can accumulate enough geometric diversity to satisfy an accuracy requirement across the full shared FoV. The question is whether the required T_{obs} is operationally acceptable, or whether combining observable types can achieve comparable accuracy in a shorter window. These trade-offs are evaluated in chapter 3.

2.5.4. Coverage Metric

The DRMS is a pointwise bound. To compare formations at the system level, it must be combined with the spatial extent over which geolocation is feasible.

Shared Field of View

The shared field-of-view model used here is geometric. Each satellite antenna is assumed to point nadir and to have an isotropic gain pattern within the visible region. The receiver footprint is therefore determined only by the Earth geometry and the minimum elevation angle $\varepsilon_{\min}$.

Each receiver footprint above elevation $\varepsilon_{\min}$ is approximately circular at LEO altitude. The shared field of view is $\mathcal{S}_{\mathrm{FoV}} = \bigcap_{j=1}^{3} \mathcal{V}_j$. The overlap area of two equal circles of radius R separated by d is

$$A(d) = 2R^2 \arccos\left(\frac{d}{2R}\right) - \frac{d}{2}\sqrt{4R^2 - d^2}. \tag{2.105}$$

For $d \ll 2R$, this reduces to $A \approx \pi R^2 - 2Rd$, so the along-track FoV loss is approximately linear when the separation remains a small fraction of the footprint diameter. The third satellite, offset by cross-track distance b, intersects the already-narrowed two-circle overlap. As a grows, this overlap thins in the cross-track direction, so a given b removes a proportionally larger fraction of the remaining area. The formation is parameterised by a and a/b.

Valid Area
Each grid cell within $\mathcal{S}_{\text{FoV}}$ is classified as valid if $\text{DRMS}(\phi, \lambda) \leq \text{DRMS}_{\max}$. The valid area and coverage percentage are

$$A_{\text{valid}} = \sum_{\substack{(\phi,\lambda)\in\mathcal{S}_{\text{FoV}} \\ \text{DRMS}\leq\text{DRMS}_{\max}}} \Delta A(\phi, \lambda), \qquad C_{\text{valid}} = \frac{A_{\text{valid}}}{A_{\text{FoV}}} \times 100\%. \tag{2.106}$$

This metric captures the central formation design trade-off: wider separation improves geometric diversity and lowers DRMS, but reduces A_{FoV}. An optimal formation maximises A_{valid}.

The DRMS surface structure differs by observable type. For TDOA, the surface is approximately conical and the sub-threshold region depends primarily on a. For FDOA, the surface exhibits multiple local minima from the non-monotonic variation of projected receiver velocity, making the sub-threshold region sensitive to both a and b. The coverage metric is evaluated across the full formation parameter space in chapter 3.

2.5.5. Latency Decomposition
The end-to-end latency from emitter activation to result availability can be decomposed as

$$L_{\text{e2e}} = T_{\text{access}} + T_{\text{obs}} + T_{\text{xfer}} + T_{\text{proc}} + T_{\text{dl}}. \tag{2.107}$$

T_{access} is the time until the emitter enters the shared FoV of a cluster. It depends on the number of clusters, the orbital parameters, and the emitter latitude, and can range from minutes to hours for a single cluster. Wider satellite separations reduce the shared FoV (subsection 2.5.4), which on average increases T_{access} because a smaller ground area is covered per pass. Formation design therefore affects latency not only through accuracy but also through coverage.

T_{obs} is the observation time required to accumulate sufficient information. It depends on the formation geometry, observable set, snapshot interval, and integration time, as developed in subsection 2.5.3.

T_{xfer} is the time to consolidate IQ data from all three receivers at a single processing node. For ground-based processing, this is governed by downlink scheduling, data volume, and link capacity. For onboard processing, it depends on the inter-satellite link: optical ISLs support high data rates but add complexity; RF-based ISLs are simpler but bandwidth-limited.

T_{proc} is the algorithmic processing time. Direct geolocation scales with the number of grid points, receiver pairs, snapshots, and samples. Indirect geolocation distributes cost across CAF evaluation, candidate association, and a lightweight nonlinear solve, resulting in fundamentally different scaling behaviour.

T_{dl} is the time to deliver the result to an end user. For onboard processing this reduces to transmitting a compact message; for ground-based processing the result is already available.

Of these five components, three fall within the geolocation-segment scope defined in section 1.4: the shared FoV area, which drives T_{access} through per-pass coverage; T_{obs}, evaluated through the CRLB-based design space analysis; and T_{proc}, evaluated through wall-clock runtime comparison of the two geolocation approaches (subsection 2.4.3). The remaining components depend on constellation-level and implementation-level decisions outside this scope.

2.6. Parametric Simulation Framework
This section summarizes how the mathematical models from the previous sections are implemented in software. Since NLR is the beneficiary of this system-engineering tool, it is designed for modularity and supports parametric variation of the emitter scenario, satellite system configuration, signal parameters, and processing settings. The framework's pipeline structure first generates the IQ sampling data for a given scenario, as detailed in subsection 2.6.1. The generated data is then processed by two different chains. The direct geolocation branch is detailed in subsection 2.6.2. The indirect geolocation branch is detailed in subsection 2.6.3.

2.6.1. Data Generation Framework

The data-generation framework produces the synchronized IQ snapshots used by both geolocation methods. For each epoch t_m, it propagates the receiver states $\boldsymbol{r}_{R_j}(t_m)$ and $\dot{\boldsymbol{r}}_{R_j}(t_m)$, checks emitter visibility, and computes the link-dependent delay, Doppler, Doppler rate, and received C/N_0. These quantities are stored in `states.csv`.

For each visible link, an IQ snapshot of duration T_{int} is synthesized from the signal model of subsection 2.2.2. The framework evaluates the geometric dynamics within the snapshot, adds thermal noise and residual receiver clock drift, and optionally applies ADC quantization. The output is a consistent set of IQ files $\tilde{y}_j[n]$ and receiver-state records for downstream processing.

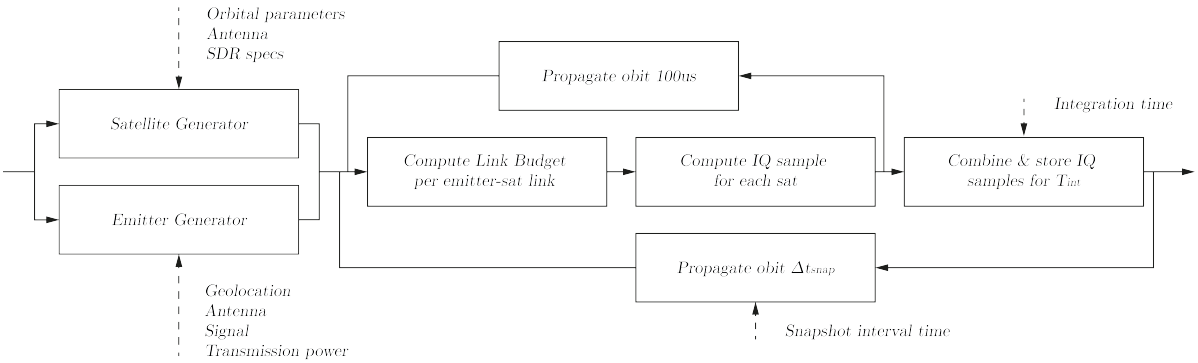

Figure 2.7: Overview of the IQ samples generation framework.

2.6.2. Direct Geolocation

In the implemented framework, direct geolocation is evaluated with a hierarchical search instead of a uniform full-resolution grid. The search starts from a coarse partition of $\mathcal{S}_0$, retains only the strongest cells, and recursively refines these branches until the target spatial resolution is reached. Nearby branch endpoints are then merged, after which a final local fine-grid evaluation is performed around each merged cluster. This reduces the number of evaluated candidate locations substantially while preserving the same position-domain score definition from subsection 2.4.1.

Algorithm 1 Direct Geolocation

Input: receiver IQ data $\tilde{y}_j^{(m)}[n]$, receiver states $\boldsymbol{r}_{R_j}(t_m)$ and $\dot{\boldsymbol{r}}_{R_j}(t_m)$, sampling frequency f_s, carrier frequency f_c
Output: emitter location estimates $\hat{\boldsymbol{r}}_{E,\ell}$

Partition $\mathcal{S}_0$ into a coarse set of search cells
Evaluate $S_{\text{tot}}(\boldsymbol{r}_E)$ at all coarse cell centres
Retain the significant top-K cells as initial branches
for all branches b **do**
    Set the current cell to the selected coarse cell
    while the cell size exceeds the target resolution **do**
        Subdivide the current cell into a finer grid of child cells
        Evaluate $S_{\text{tot}}(\boldsymbol{r}_E)$ at all child-cell centres
        Keep only the highest-scoring child cell
    end while
end for
Merge nearby branch endpoints by centroid-based spatial gating
for all merged clusters C_ℓ **do**
    Construct a local fine grid around the weighted centroid of C_ℓ
    Evaluate $S_{\text{tot}}(\boldsymbol{r}_E)$ on this local grid
    Set $\hat{\boldsymbol{r}}_{E,\ell}$ to the strongest local maximum
end for
return $\{\hat{\boldsymbol{r}}_{E,\ell}\}$

2.6.3. Indirect Geolocation

In the implemented framework, the indirect chain is divided into four explicit processing stages. First, each receiver pair produces snapshot-wise FDOA candidates from the gated CAF. Second, candidates

are associated across the three receivers using the FDOA closure constraint. Third, accepted observations are grouped over time to obtain the measurement vector $\mathbf{z}(t_0)$. Fourth, this vector is passed to the Gauss–Newton solver to estimate emitter position.

Algorithm 2 Indirect geolocation

Input: receiver IQ data $\tilde{y}_j^{(m)}[n]$, receiver states $\boldsymbol{r}_{R_j}(t_m)$ and $\dot{\boldsymbol{r}}_{R_j}(t_m)$, sampling frequency f_s, carrier frequency f_c
Output: emitter location estimates $\hat{\boldsymbol{r}}_E$

for each snapshot $m = 1, \ldots, M$ **do**
  for each receiver pair (R_j, R_k) **do**
    Compute the gated CAF $\mathcal{A}_{jk}^{(m)}(\tau, f)$
    Extract FDOA candidates $\mathcal{F}_{jk}^{(m)} = \{\hat{\Delta f}_{jk}^{(p)}\}$
  end for
  Form candidate triplets $\mathbf{b} = \left(\hat{\Delta f}_{12}^{(p)}, \hat{\Delta f}_{23}^{(q)}, \hat{\Delta f}_{31}^{(r)}\right)$
  Retain triplets satisfying the closure test $\eta_f(\mathbf{b}) \le \epsilon_f$
end for
Cluster accepted triplets over time into emitter tracks
for each emitter track **do**
  for each receiver pair (R_j, R_k) **do**
    Fit a cubic polynomial $p_{jk}(t)$ to the associated FDOA samples
  end for
  Form the measurement vector $\mathbf{z}(t_0)$
  Estimate $\hat{\boldsymbol{r}}_E$ with the Gauss–Newton method
end for
return the set of emitter location estimates $\hat{\boldsymbol{r}}_E$

3

Results

3.1. System Design Space Analysis

This section narrows the design space defined in Table 2.3 using the CRLB framework and coverage metric from section 2.5. The analysis proceeds in five steps. First, the coherent integration time is characterised as a function of the along-track baseline (subsection 3.1.1). Second, the shared field of view is quantified across the formation parameter space (subsection 3.1.2). Third, CRLB heatmaps illustrate the spatial accuracy distribution (subsection 3.1.3). Fourth, the valid area metric is evaluated for different observable sets, observation times, and integration times (subsection 3.1.4). Fifth, the selected formation and the integration time assignments for the direct and indirect approaches are stated (subsection 3.1.5).

The baseline parameter values are listed in Table 3.1.

Table 3.1: Baseline parameters for the design-space analysis

Category	Parameter	Symbol	Baseline value
Design Parameters	Along-track separation	a	100–1500 km (swept)
	Along-track / cross-track ratio	a/b	2–12 (swept)
	Total observation time	T_{obs}	10 s and 30 s
	Snapshot interval	Δt_{snap}	1 s
	Coherent integration time	T_{int}	25–800 ms (swept)
Scenario Parameters	Emitter transmit power	P_{tx}	10 W
	Signal bandwidth	B_c	1 MHz
	Carrier frequency	f_c	1575.42 MHz
	Receiver noise bandwidth	B_r	20 MHz
Evaluation criteria	Minimum elevation angle	ε_{min}	10°
	DRMS threshold	DRMS_{max}	250 m

The integration time and formation geometry are mutually dependent. The maximum useful T_{int} decreases with satellite separation (subsection 2.5.2), while the optimal separation depends on the per-snapshot accuracy set by T_{int}. This constraint, however, applies only to the indirect approach. As derived in subsection 2.5.2, the upper bound on T_{int} arises from the CAF assumption that the differential delay and differential Doppler remain constant within the integration window. Direct geolocation does not rely on this assumption: the position-domain score compensates the full delay and Doppler history for each candidate location, so no constant-observable approximation is required. The integration time is therefore a design constraint specific to CAF-based (indirect) processing.

This interdependence is resolved as follows. The integration time analysis in subsection 3.1.1 establishes how T_{int} must be chosen as a function of a for the indirect approach. The subsequent valid area analysis sweeps the formation geometry at both $T_{\text{int}} = 250$ ms and $T_{\text{int}} = 50$ ms. As will be shown, the optimal

formation lies at small a, where the T_{int} constraint is least restrictive. The two analyses are therefore consistent.

3.1.1. Integration Time

The coherent integration time is limited by differential Doppler drift within the integration window (subsection 2.5.2). This limit is specific to the CAF, which assumes constant differential delay and Doppler over each processing interval. Direct geolocation compensates the full delay–Doppler history per candidate position and is therefore not subject to this constraint.

To determine a suitable T_{int} for CAF-based processing, the truth-cell SNR is evaluated as a function of integration time for several along-track separations a. Only the two co-planar satellites are considered, as this pair exhibits the largest differential Doppler rate and thus imposes the most restrictive constraint. For each (a, T_{int}) combination, the CAF-based SNR is estimated over 10 Monte Carlo realisations with independent thermal-noise and receiver clock-drift samples. Figure 3.1 shows the mean SNR; shaded bands indicate the standard deviation.

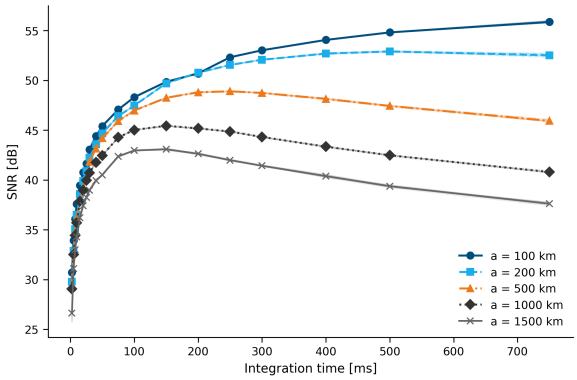

Figure 3.1: CAF truth-cell SNR as a function of coherent integration time for different along-track separations a

Three trends are visible. First, increasing a increases the differential Doppler rate and reduces the integration time at which the peak SNR occurs. Second, at small baselines, longer integration remains beneficial because thermal-noise reduction outweighs Doppler-induced decorrelation. Third, at large baselines the CAF peak smears at shorter integration times, causing the SNR to plateau and then decline. The spread across Monte Carlo runs is small relative to the overall trend, confirming that the limiting mechanism is differential Doppler drift rather than stochastic clock drift.

Based on this analysis, $T_{\text{int}} = 250\,\text{ms}$ is adopted as the baseline for the indirect approach. As shown in subsection 3.1.4, the most favourable formations occur near $a \approx 200\,\text{km}$, where this choice lies well within the useful integration range.

3.1.2. Shared Field of View

Figure 3.2 shows the shared FoV area across the formation parameter space. At 550 km altitude with $\varepsilon_{\text{min}} = 10$, each footprint has a diameter of approximately 3329 km.

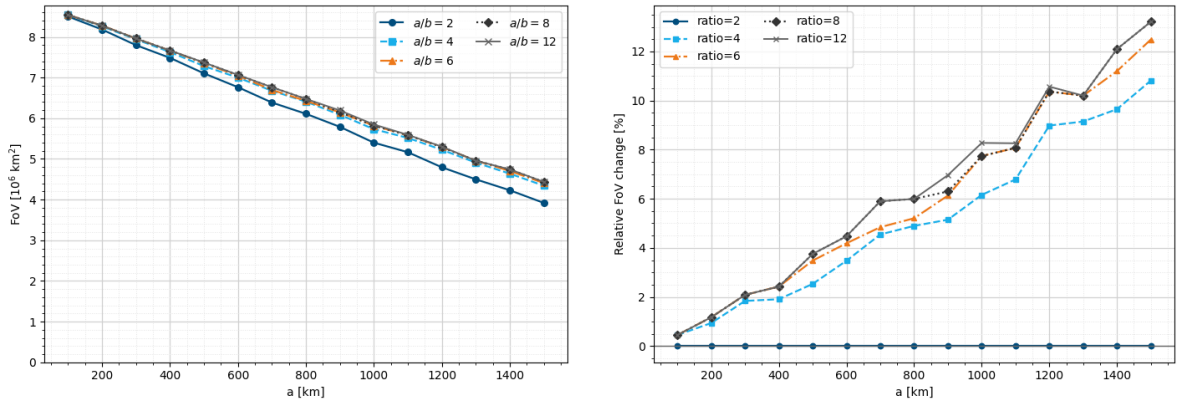

(a) Absolute shared FoV area

(b) Relative FoV reduction

Figure 3.2: Shared FoV area as a function of along-track separation a, parameterised by a/b.

The shared FoV decreases approximately linearly from 8.5×10^6 km^2 at $a = 100$ km to 4.4×10^6 km^2 at $a = 1500$ km—a reduction of 48%. This linearity is consistent with the linear-regime approximation from subsection 2.5.4: $a = 1500$ km is only 0.45 times the footprint diameter. At $a/b \geq 6$, all curves coincide; the cross-track offset has no measurable effect. At $a/b = 4$, a separation becomes visible. At $a/b = 2$, the additional FoV loss reaches approximately 13% at $a = 1500$ km relative to the two-satellite case. The cross-track baseline is therefore unconstrained from a FoV perspective provided $a/b \gtrsim 6$.

3.1.3. CRLB Analysis

Figure 3.3 shows the FDOA-based DRMS surface for an example formation ($a = 200$ km, $a/b = 8$, $T_{\text{obs}} = 30$ s, $T_{\text{int}} = 250$ ms).

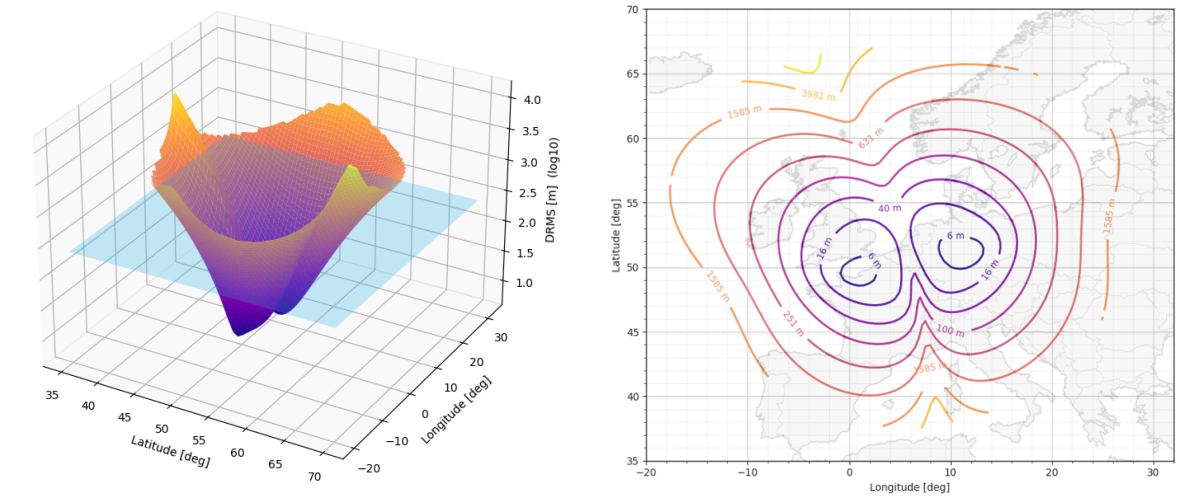

(a) 3-D surface with 250 m threshold plane

(b) Contour map

Figure 3.3: FDOA-based CRLB DRMS over the shared FoV ($a = 200$ km, $a/b = 8$, $T_{\text{obs}} = 30$ s, $T_{\text{int}} = 250$ ms).

Two distinct minima near 6 m DRMS appear symmetrically about the sub-satellite ground track, at approximately (51°N, 0°E) and (51°N, 10°E). The DRMS increases steeply toward the FoV edges, exceeding 1500 m beyond 60°N and below 43°N. This multi-lobed structure confirms the observable-dependent DRMS topology described in subsection 2.5.4: the FDOA surface is not conical, and the sub-threshold region depends on both a and b.

To compare the candidate observable sets more directly, Figure 3.4 plots the horizontal DRMS as a function of total observation time for four cases: FDOA only, TDOA only, joint T/FDOA, and joint

DFDOA/FDOA. The same example formation is used. Different coherent integration times are applied for different observable sets. The T/FDOA case uses $T_{\text{int}} = 125\,\text{ms}$, since the combined delay and Doppler information achieves comparable accuracy with a shorter coherent integration time. The other cases use $T_{\text{int}} = 250\,\text{ms}$. This reflects the observable-dependent integration-time limits from subsection 2.5.2.

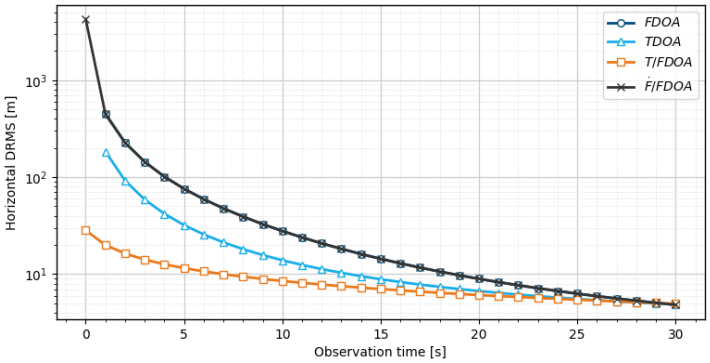

Figure 3.4: CRLB horizontal DRMS versus observation time for different observable sets. T/FDOA uses $T_{\text{int}} = 125\,\text{ms}$, whereas TDOA, FDOA and DFDOA/FDOA use $T_{\text{int}} = 250\,\text{ms}$.

Two conclusions follow from Figure 3.4. First, the joint T/FDOA case provides the fastest reduction in DRMS. At short observation times it clearly outperforms both TDOA-only and FDOA-only. This shows that combining delay and Doppler information is highly beneficial when only a limited observation window is available.

Second, the benefit of adding DFDOA to FDOA is negligible in the considered scenario. The DFDOA/FDOA curve is almost indistinguishable from the FDOA-only curve over nearly the full observation window, and both cases converge to essentially the same final accuracy. DFDOA is therefore retained in the analytical framework for completeness, but is not considered further in the indirect geolocation approach.

3.1.4. Valid Area Across the Formation Parameter Space

The valid area A_{valid} and coverage percentage C_{valid} are evaluated for FDOA, TDOA, and T/FDOA across the full $(a,\ a/b)$ parameter space. Results are presented for $T_{\text{obs}} = 10\,\text{s}$ and $T_{\text{obs}} = 30\,\text{s}$ to quantify the effect of observation time. For the 30 s case, two integration times are compared: $T_{\text{int}} = 250\,\text{ms}$ and $T_{\text{int}} = 50\,\text{ms}$. This comparison reveals the sensitivity of each observable set to integration time and satellite formation.

Short Observation Time, Long Integration Time ($T_{obs} = 10\,s$, $T_{int} = 250\,ms$)

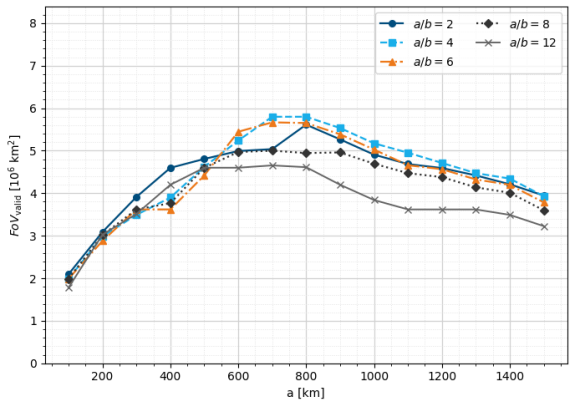

(a) FDOA: valid area

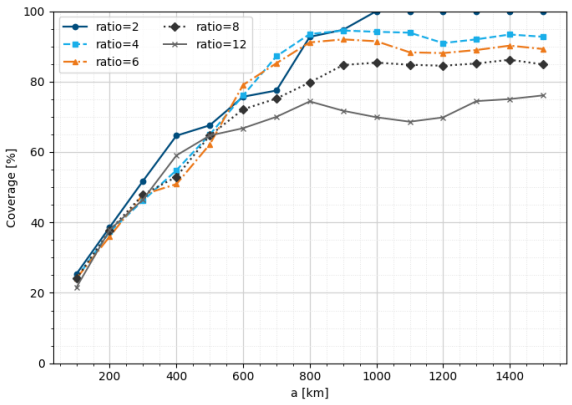

(b) FDOA: coverage

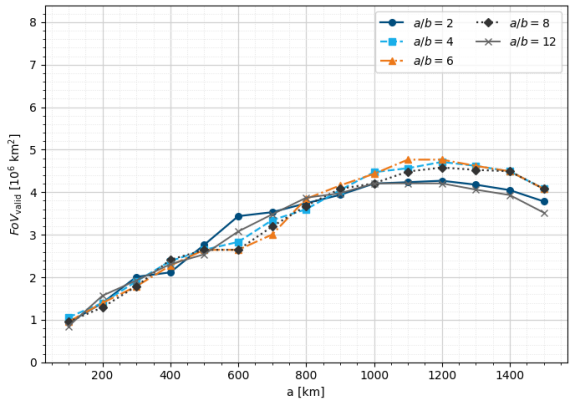

(c) TDOA: valid area

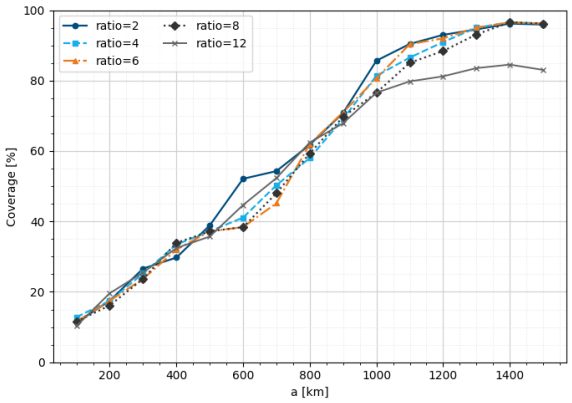

(d) TDOA: coverage

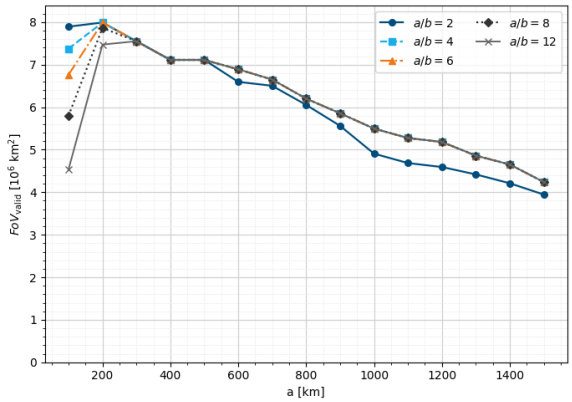

(e) T/FDOA: valid area

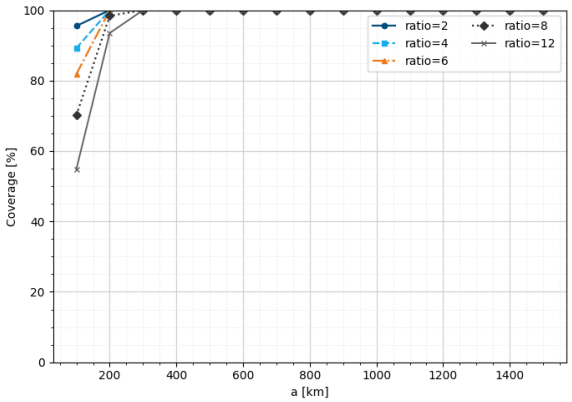

(f) T/FDOA: coverage

Figure 3.5: Valid area (left) and coverage percentage (right) for $T_{obs} = 10\,s$, $T_{int} = 250\,ms$.

FDOA. The valid area peaks at $a \approx 700$–$800\,km$ with $A_{valid} \approx 5.8 \times 10^6\,km^2$ ($a/b = 4$). Coverage spread across ratios is pronounced: $a/b = 2$ reaches 100% only at $a = 1500\,km$, while $a/b = 12$ saturates near 75%. This sensitivity is consistent with the multi-lobed FDOA DRMS structure. At 10 s, the few accumulated snapshots provide limited geometric diversity, requiring a larger baseline to compensate.

TDOA. The valid area peaks at $a \approx 1100$–$1200\,km$ with $A_{valid} \approx 4.8 \times 10^6\,km^2$ for multiple ratios. Coverage reaches 95% only at $a \geq 1300\,km$; $a/b = 12$ remains below 85%.

T/FDOA. The combined observable set peaks at $a \approx 200\,\text{km}$ with $A_{\text{valid}} \approx 8.0 \times 10^6\,\text{km}^2$. Coverage reaches 100% by $a \approx 300\,\text{km}$ for all ratios, with only $a/b = 12$ starting lower ($\approx 55\%$) at $a = 100\,\text{km}$. The complementary range and range-rate information brings the DRMS below 250 m across the entire shared FoV at small separations. T/FDOA is therefore the only observable set for which tight formations ($a \leq 300\,\text{km}$) achieve near-complete coverage at short observation times.

Long Observation Time, Long Integration Time ($T_{\text{obs}} = 30\,\text{s}$, $T_{\text{int}} = 250\,\text{ms}$)

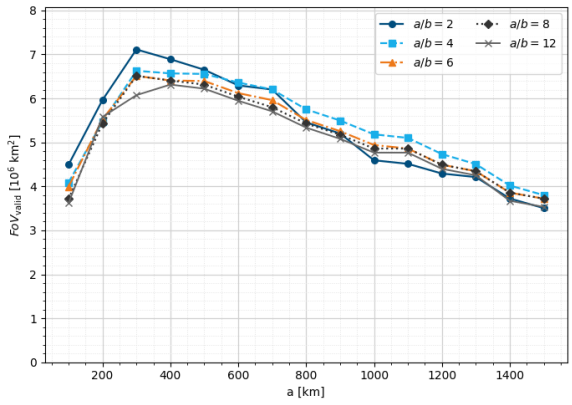

(a) FDOA: valid area

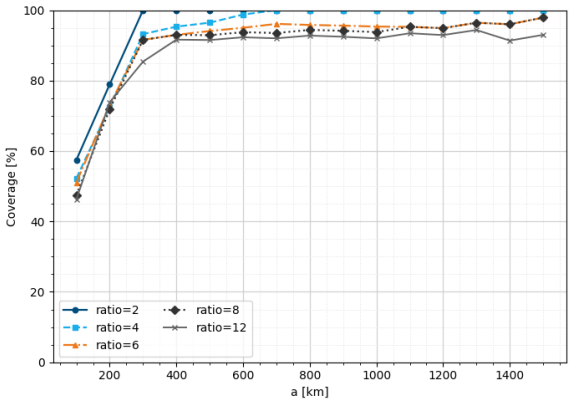

(b) FDOA: coverage

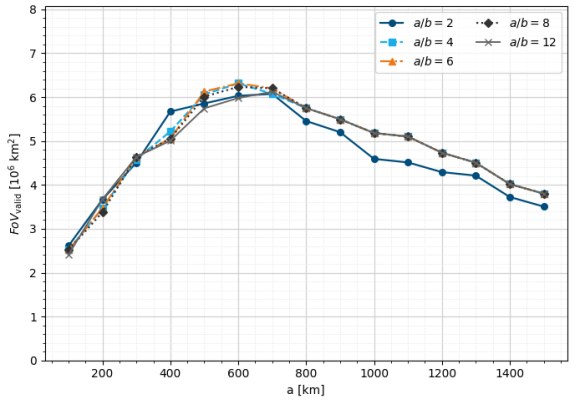

(c) TDOA: valid area

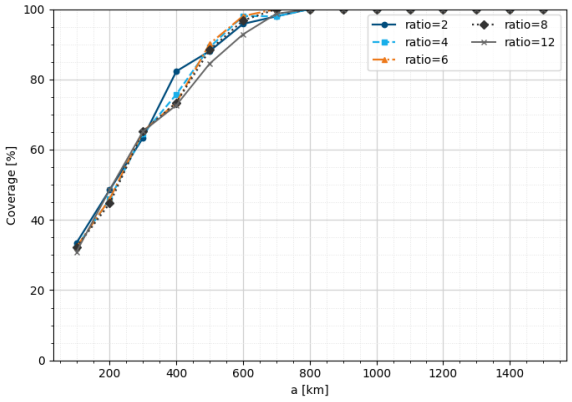

(d) TDOA: coverage

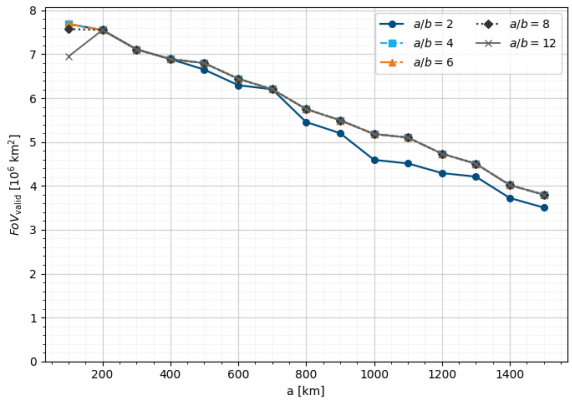

(e) T/FDOA: valid area

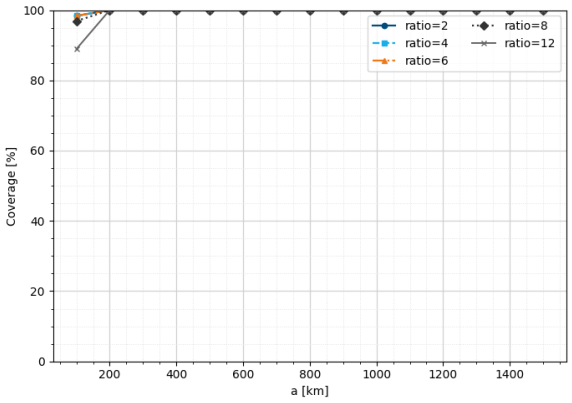

(f) T/FDOA: coverage

Figure 3.6: Valid area (left) and coverage percentage (right) for $T_{\text{obs}} = 30\,\text{s}$, $T_{\text{int}} = 250\,\text{ms}$.

FDOA. Tripling the observation time shifts the optimum from $a \approx 700\,\text{km}$ to $a \approx 300\,\text{km}$ and increases the peak valid area from 5.8×10^6 to $7.1 \times 10^6\,\text{km}^2$ ($a/b = 2$). Coverage reaches 100% by $a \approx 300\,\text{km}$ for

$a/b = 2$ and stabilises near 93% for higher ratios. The sensitivity to a/b persists; $a/b = 2$ becomes more advantageous at smaller baselines.

TDOA. The optimum shifts from $a \approx 1000\,\text{km}$ to $a \approx 600\,\text{km}$, with $A_{\text{valid}} \approx 6.2 \times 10^6\,\text{km}^2$. All a/b curves are nearly indistinguishable up to $a \approx 700\,\text{km}$, confirming that the conical DRMS surface is insensitive to the cross-track baseline. Coverage reaches 100% by $a \approx 800\,\text{km}$.

T/FDOA. The combined set yields $A_{\text{valid}} \approx 7.5 \times 10^6\,\text{km}^2$ at $a \approx 200\,\text{km}$, declining monotonically beyond. Coverage reaches 100% by $a \approx 200\,\text{km}$ for $a/b \leq 8$. These values are nearly identical to the 10 s results. The DRMS continues to improve with additional observation time, but the 250 m threshold is already satisfied across the entire shared FoV at $a \approx 200\,\text{km}$. Further accuracy gains therefore do not increase A_{valid}; widening the formation beyond this point only reduces the shared FoV.

Long Observation Time, Short Integration Time ($T_{\text{obs}} = 30\,\text{s}$, $T_{\text{int}} = 25\,\text{ms}$)

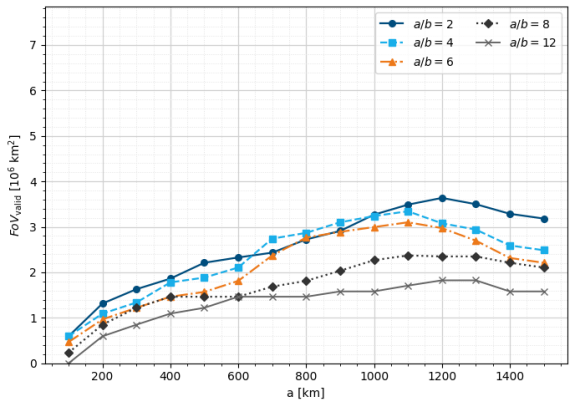

(a) FDOA: valid area

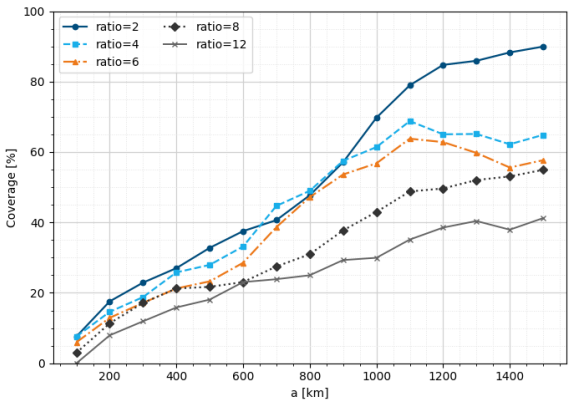

(b) FDOA: coverage

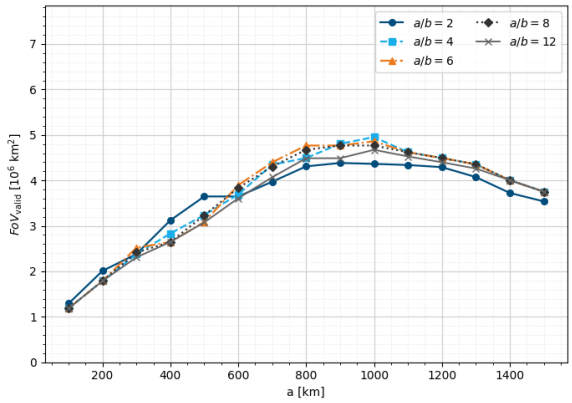

(c) TDOA: valid area

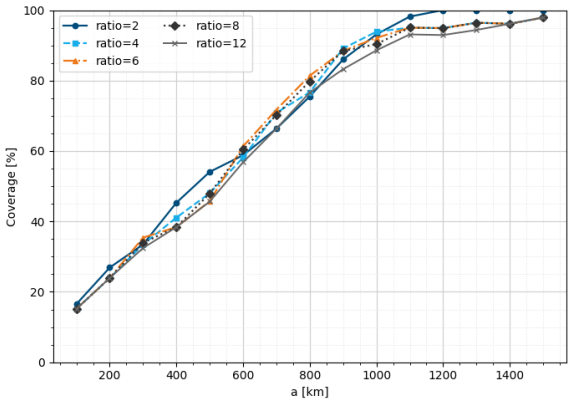

(d) TDOA: coverage

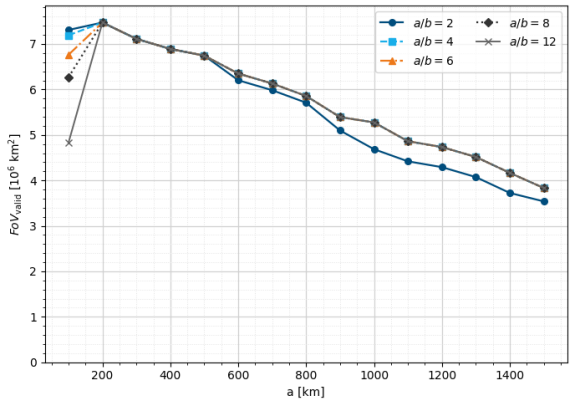

(e) T/FDOA: valid area

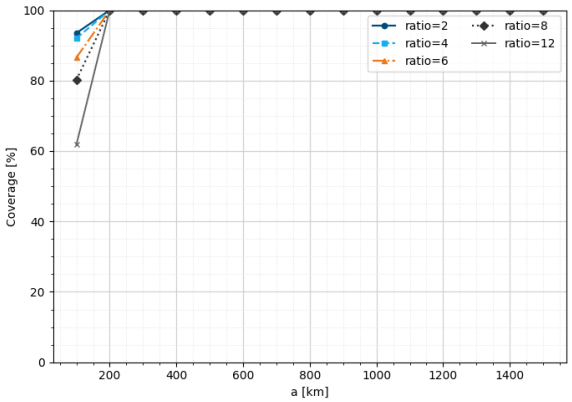

(f) T/FDOA: coverage

Figure 3.7: Valid area (left) and coverage percentage (right) for $T_{\text{obs}} = 30\,\text{s}$, $T_{\text{int}} = 25\,\text{ms}$.

FDOA. Reducing the integration time per snapshot from 250 ms to 25 ms impacts the valid area significantly. The peak valid area drops from $7.1 \times 10^6\,\text{km}^2$ to $3.6 \times 10^6\,\text{km}^2$ at $a \approx 1200\,\text{km}$ for $a/b = 2$. Coverage never reaches 100%, and the sensitivity to a/b is more pronounced. The valid area range from $1.8 \times 10^6\,\text{km}^2$ to $3.6 \times 10^6\,\text{km}^2$. This degradation is consistent with the T_{int}^{-3} scaling of the FDOA variance.

TDOA. TDOA performance degrades modestly relative to the 250 ms case. The impact of a/b on the valid area is less pronounced as well. The optimum shifts from $A_{\text{valid}} \approx 6 \times 10^6\,\text{km}^2$ at $a \approx 600\,\text{km}$ to

$A_{\text{valid}} \approx 4.8 \times 10^6\,\text{km}^2$ at $a \approx 900 - 1000\,\text{km}$. These results are consistent with the weaker T_{int}^{-1} scaling of the TDOA variance.

T/FDOA. The combined T/FDOA set at 25 ms achieves complete coverage at $a \approx 200\,\text{km}$ for all a/b. This analysis reveals the either a long observation time with short integration or a short observation time with long integration can achieve full coverage at tight formations if both TDOA and FDOA are exploited.

3.1.5. Summary and System Selection

Table 3.2 summarises the peak valid area across observable sets, observation times, and integration times.

Table 3.2: Peak valid area, optimal along-track separation, and coverage at optimum for each configuration

Observable	T_{obs}	T_{int}	Optimal a	A_{valid} [$\times 10^6\,\text{km}^2$]	C_{valid}
FDOA	10 s	250 ms	700 km	5.8	80%
FDOA	30 s	250 ms	300 km	7.1	95%
FDOA	30 s	25 ms	1200 km	3.6	85%
TDOA	10 s	250 ms	1100 km	4.8	95%
TDOA	30 s	250 ms	600 km	6.2	90%
TDOA	30 s	25 ms	900 km	4.8	90%
T/FDOA	10 s	250 ms	200 km	7.5	100%
T/FDOA	30 s	250 ms	200 km	7.5	100%
T/FDOA	30 s	25 ms	200 km	7.5	100%

Every observable set can achieve $C_{\text{valid}} \approx 100\%$. The sets differ in the combination of along-track separation and observation time required to reach it. TDOA-only requires $a \geq 800\,\text{km}$ at 30 s. FDOA-only requires $a \geq 300\,\text{km}$ at 30 s or $a \geq 1000\,\text{km}$ at 10 s. T/FDOA reaches full coverage at $a \approx 200\,\text{km}$ regardless of the considered observation times.

Since A_{valid} is fundamentally maximised at small a, the formation should be kept as tight as possible. A tight formation preserves the shared FoV and thereby minimises T_{access} (subsection 2.5.5). It also relaxes the integration time constraint (subsection 3.1.1), because the differential Doppler rate is smaller at short baselines. Coverage and integration time are therefore aligned rather than conflicting. The cross-track baseline b has negligible effect on the shared FoV for $a/b \gtrsim 6$. A ratio of $a/b = 8$ is selected as a compromise between geometric diversity and formation-keeping cost.

To ensure an apples-to-apples comparison, the integration time is chosen such that both approaches operate at comparable CRLB performance levels at the selected formation at the emitter geolocations. The indirect approach uses FDOA only at $T_{\text{int}} = 250\,\text{ms}$; the direct approach uses T/FDOA at $T_{\text{int}} = 125\,\text{ms}$. All other parameters are identical. Based on the above, the parameters for the end-to-end comparison are:

- Along-track separation: $a = 200\,\text{km}$
- Cross-track separation: $b = 25\,\text{km}$ $(a/b = 8)$
- Observation time: $T_{\text{obs}} = 30\,\text{s}$
- Snapshot interval: $\Delta t_{\text{snap}} = 1\,\text{s}$
- Indirect approach: FDOA, $T_{\text{int}} = 250\,\text{ms}$
- Direct approach: T/FDOA, $T_{\text{int}} = 125\,\text{ms}$

3.2. Comparison of Direct and Indirect Geolocation

This section compares the direct and indirect geolocation methods under the same end-to-end conditions defined in subsection 2.4.3. The selected formation and observation settings follow from the design-space analysis of section 3.1: $a = 200\,\text{km}$, $b = 25\,\text{km}$ $(a/b = 8)$, $T_{\text{obs}} = 30\,\text{s}$, and $\Delta t_{\text{snap}} = 1\,\text{s}$. The

indirect method uses FDOA-only processing with $T_{\text{int}} = 250\,\text{ms}$; the direct method uses position-domain correlation with $T_{\text{int}} = 125\,\text{ms}$.

Two scenarios are evaluated. Scenario 1 reproduces the challenging four-emitter configuration from [10]. Scenario 2 keeps the same signal set but moves two emitters farther from the formation centre, increasing off-nadir geometry and reducing measurement quality. The comparison is reported in four steps: scenario definition, direct geolocation results, indirect geolocation results, and final quantitative comparison.

All runtime measurements were obtained on a 64-bit Windows workstation equipped with an Intel(R) Core(TM) i5-10310U CPU @ 1.70 GHz and 16.0 GB RAM, using a CPU-only NumPy implementation.

3.2.1. Scenario Definition

The emitter definitions are listed in Table 3.3. Only the emitter locations differ between the two scenarios; the signal classes, carrier frequency, and transmit powers are unchanged. This isolates the effect of geometry from the effect of signal design.

Table 3.3: Emitter definitions used in the end-to-end comparison scenarios

Scenario	Emitter	Signal type	Latitude [deg]	Longitude [deg]	P_{tx} [W]	Signal-specific parameters
1	E_1	Single-frequency tone	51.0	8.0	30	–
	E_2	Asymmetric sweep jammer	51.0	5.0	30	$f_0 = -10\,\text{kHz}$, $k = 100\,\text{MHz/s}$, $T_{\text{sweep}} = 10\,\text{ms}$, sweep range = 1 MHz
	E_3	Symmetric sweep jammer	52.0	8.0	30	$f_0 = -10\,\text{kHz}$, $k = 25\,\text{MHz/s}$, $T_{\text{sweep}} = 10\,\text{ms}$, sweep range = 250 kHz
	E_4	BPSK jammer	52.0	5.0	30	$T_c = 1000\,\text{ns}$, B_k = PRN(1)
2	E_1	Single-frequency tone	45.0	15.0	30	–
	E_2	Asymmetric sweep jammer	51.0	5.0	30	$f_0 = -10\,\text{kHz}$, $k = 100\,\text{MHz/s}$, $T_{\text{sweep}} = 10\,\text{ms}$, sweep range = 1 MHz
	E_3	Symmetric sweep jammer	60.0	15.0	30	$f_0 = -10\,\text{kHz}$, $k = 25\,\text{MHz/s}$, $T_{\text{sweep}} = 10\,\text{ms}$, sweep range = 250 kHz
	E_4	BPSK jammer	52.0	5.0	30	$T_c = 1000\,\text{ns}$, B_k = PRN(1)

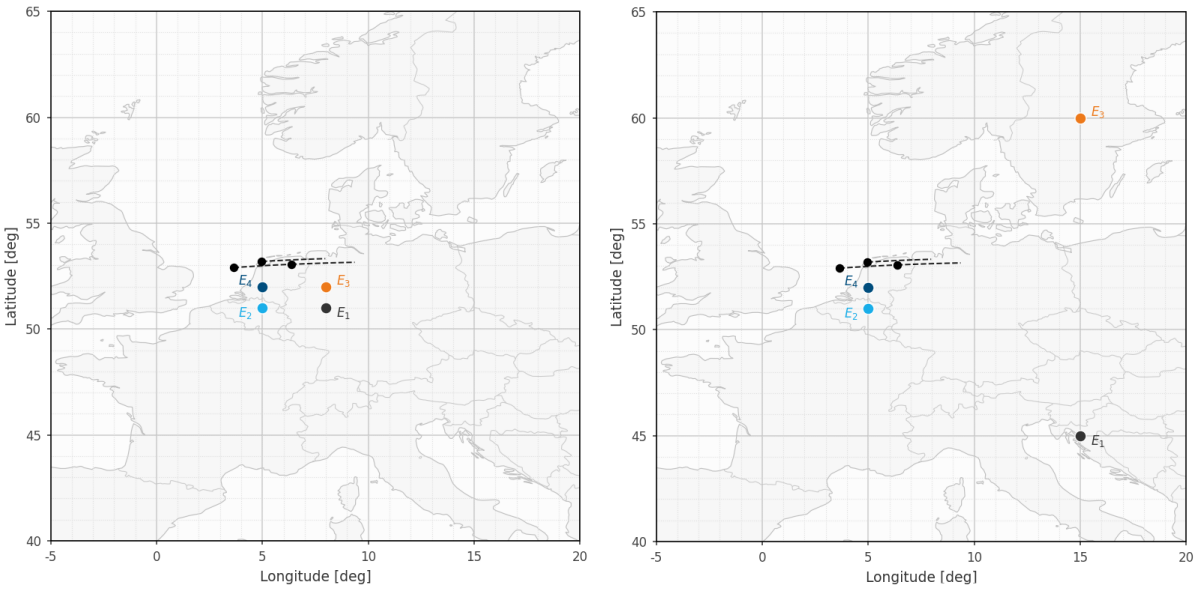

(a) Scenario 1 **(b)** Scenario 2

Figure 3.8: Emitter locations and receiver ground tracks for the two end-to-end scenarios.

Emitter Signals

The power spectral densities and spectrograms of the four emitter classes used in the comparison are shown in Figure 3.9. The set consists of a single-frequency tone jammer (E_1), an asymmetric sweep jammer (E_2), a symmetric sweep jammer (E_3), and a BPSK jammer (E_4).

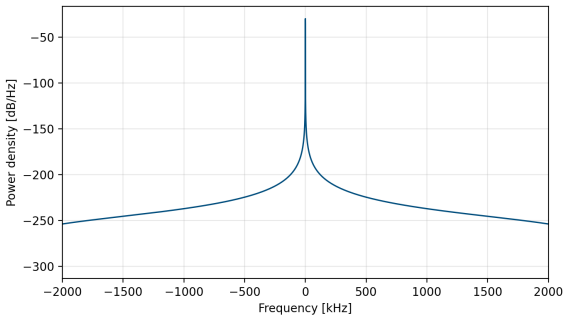

(a) Single-frequency tone jammer E_1: power spectral density

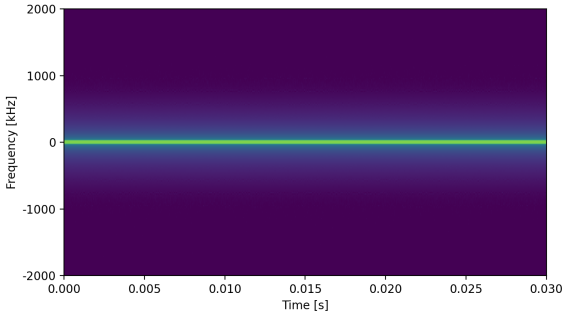

(b) Single-frequency tone jammer E_1: spectrogram

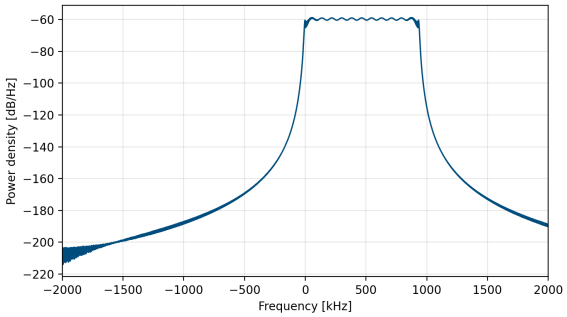

(c) Asymmetric sweep jammer E_2: power spectral density

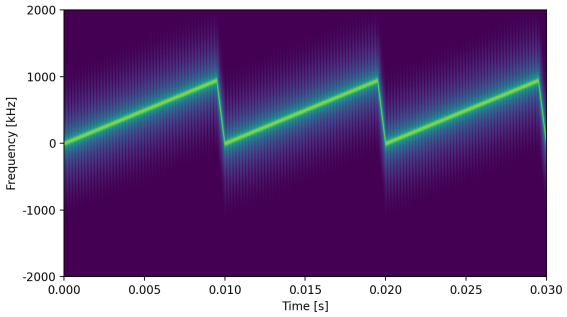

(d) Asymmetric sweep jammer E_2: spectrogram

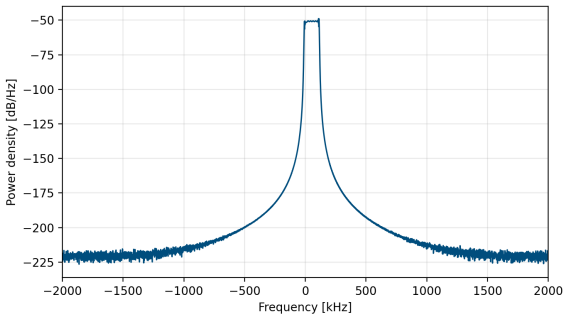

(e) Symmetric sweep jammer E_3: power spectral density

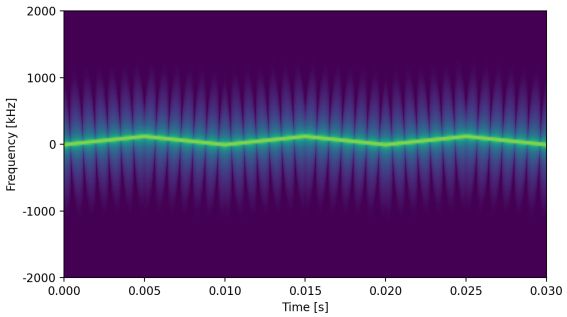

(f) Symmetric sweep jammer E_3: spectrogram

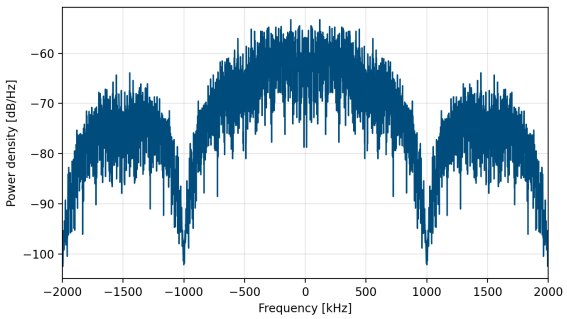

(g) BPSK jammer E_4: power spectral density

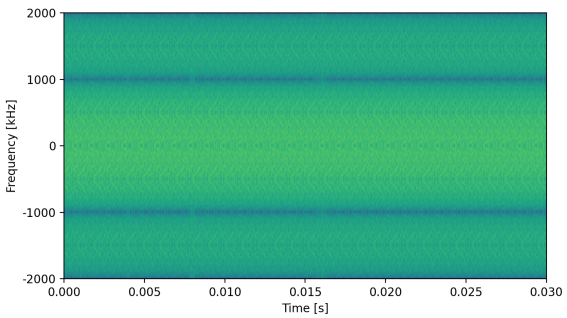

(h) BPSK jammer E_4: spectrogram

Figure 3.9: Power spectral densities (left) and spectrograms (right) of the four emitter classes used in both scenarios.

The tone signal is spectrally narrow and appears as a constant line in the spectrogram. The two sweep

signals occupy a wider bandwidth and show time-varying frequency content. The asymmetric and symmetric sweep cases differ in their sweep structure and spectral shape. The BPSK signal has a broader spectrum with a noise-like time-domain appearance caused by the code modulation.

These plots summarise the signal characteristics of the four emitters used in the end-to-end comparison.

3.2.2. Direct Geolocation Results

Direct geolocation operates directly in the position domain, bypassing a separate observable-extraction stage. The result is therefore a score map in the position domain, from which emitter locations are read as local maxima. In the implemented framework, the score is evaluated with the hierarchical search of subsection 2.4.1, so the reported accuracies are already the result of a coarse-to-fine refinement.

Grid Spacing Lower Bound for $T_{\text{int}} = 125\,\text{ms}$

[11] motivates a coarse-to-fine search strategy for wide-area surveillance, in which a coarse grid identifies promising emitter regions and subsequent refinements improve the location estimate. In the direct method, the coherent integration time T_{int} and the grid spacing are coupled: a longer T_{int} sharpens the position-domain peaks and improves spatial resolution, but also requires a finer grid to sample those peaks adequately.

In this work, the direct method is evaluated at a fixed coherent integration time of $T_{\text{int}} = 125\,\text{ms}$. This value was not optimised in the present subsection, and lower integration times were not investigated for the direct method. Instead, $T_{\text{int}} = 125\,\text{ms}$ is adopted from the CRLB-based comparison framework of subsection 3.1.5, where it provides a fair comparison with the indirect method.

For $T_{\text{int}} = 125\,\text{ms}$, several grid spacings were evaluated in Scenario 1. Figure 3.10 shows a set of representative score maps after non-coherent integration over $T_{\text{obs}} = 30\,\text{s}$. It is found that spacings coarser than 0.03° did not resolve all four emitters. The adopted lower bound is therefore 0.03°. Figure 3.10 also shows that the false-peak structures become more pronounced at coarser spacings.

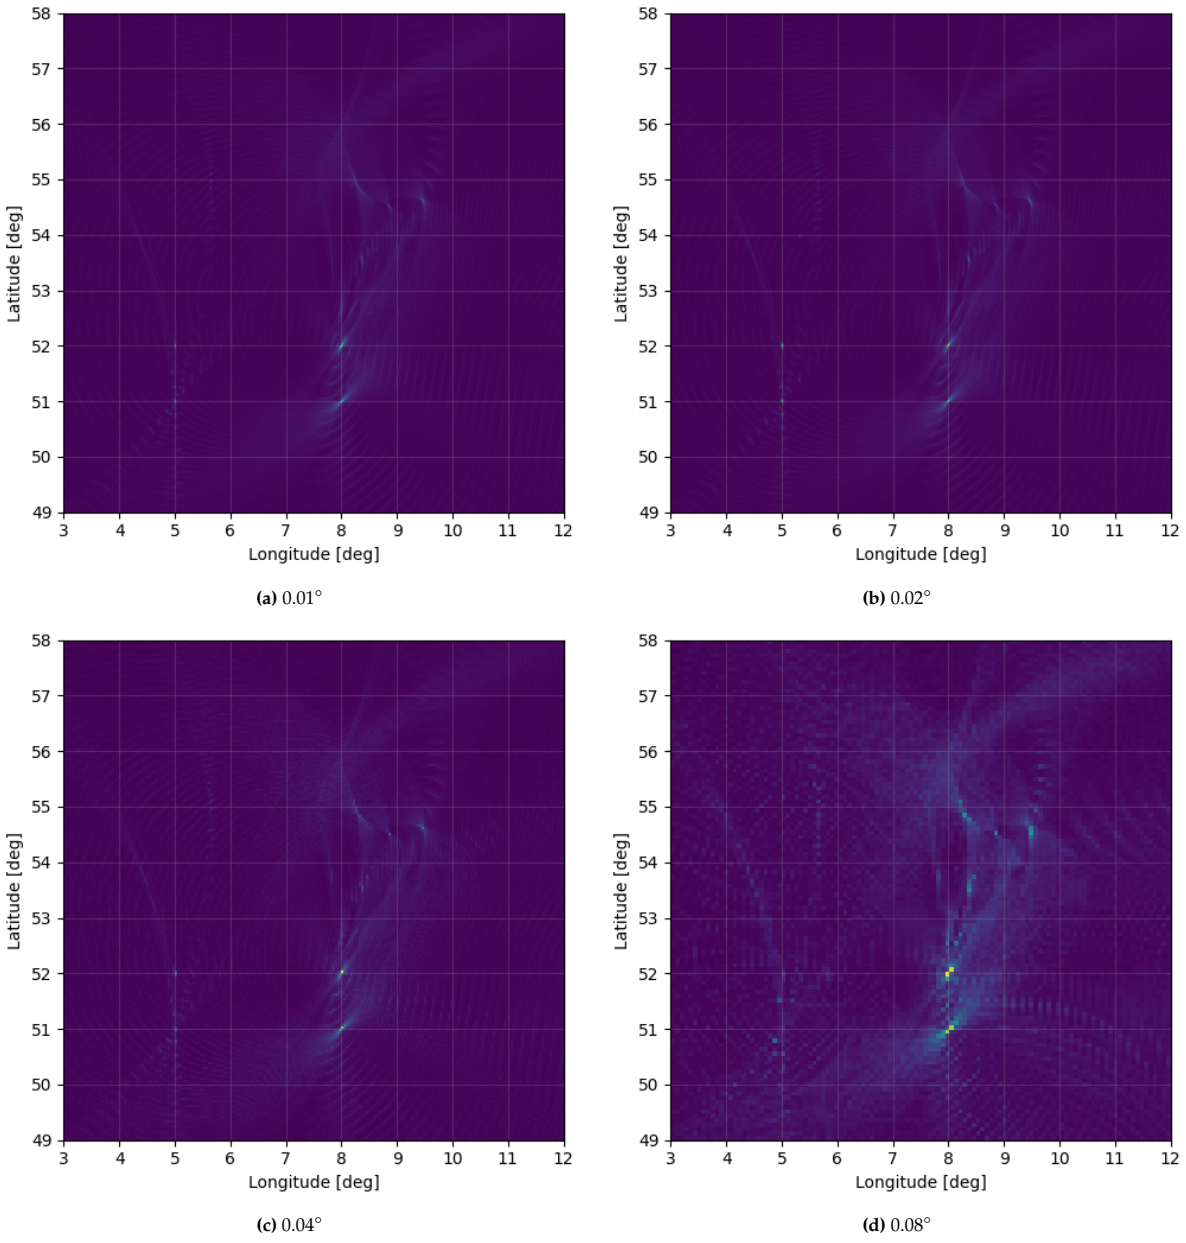

(a) 0.01°

(b) 0.02°

(c) 0.04°

(d) 0.08°

Figure 3.10: Representative position-domain score maps for Scenario 1 after non-coherent integration over $T_{\text{obs}} = 30\,\text{s}$ with $T_{\text{int}} = 125\,\text{ms}$, shown for different grid spacings. From these results, 0.03° is adopted as the coarsest admissible initial grid spacing.

Over the shared field of view of the selected formation, a uniform grid with spacing 0.03° corresponds to approximately 1.2×10^6 cells. Since this grid only provides kilometre-level spatial sampling, hierarchical refinement remains necessary to reach sub-kilometre accuracy.

Figure 3.11 illustrates the score structure for 0.03°. For a single snapshot, multiple false peaks are visible. After non-coherent integration over $T_{\text{obs}} = 30\,\text{s}$, the false-peak structure is largely suppressed and the true emitters appear as distinct local maxima.

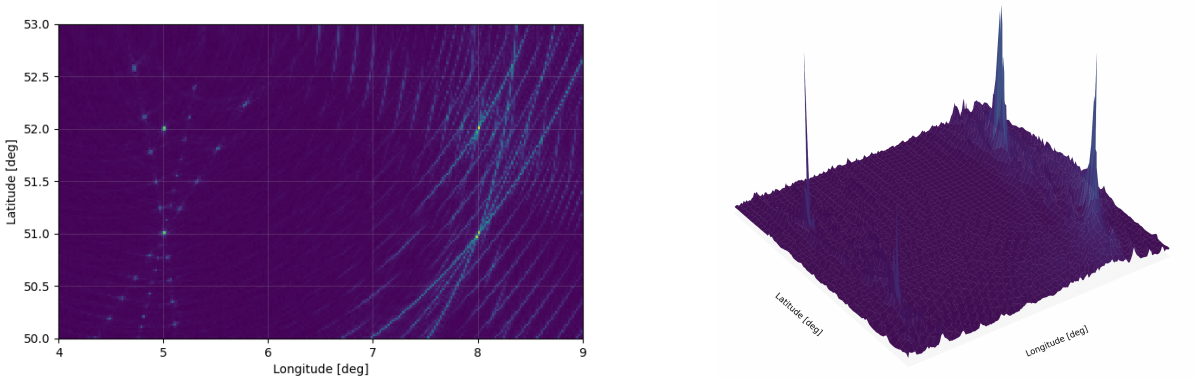

(a) $T_{\text{obs}} = 1\,\text{s}$ **(b)** 3D score map for $T_{\text{obs}} = 30\,\text{s}$

Figure 3.11: Additional visualisations of the position-domain score map for Scenario 1 with $T_{\text{int}} = 125\,\text{ms}$ and grid spacing 0.02°.

Hierarchical Search

Figure 3.12 shows a wide-area position-domain score map for scenario 1, evaluated with the hierarchical search method and a grid spacing of 0.03°. The right plot zooms in on the area around the single-frequency jammer E_1, which shows the hierarchical search refinement process at promising cells. The other emitters show similar behaviour. This illustrates the effectiveness of the hierarchical search strategy in high accuracy geolocation estimation at the lowest possible computational cost.

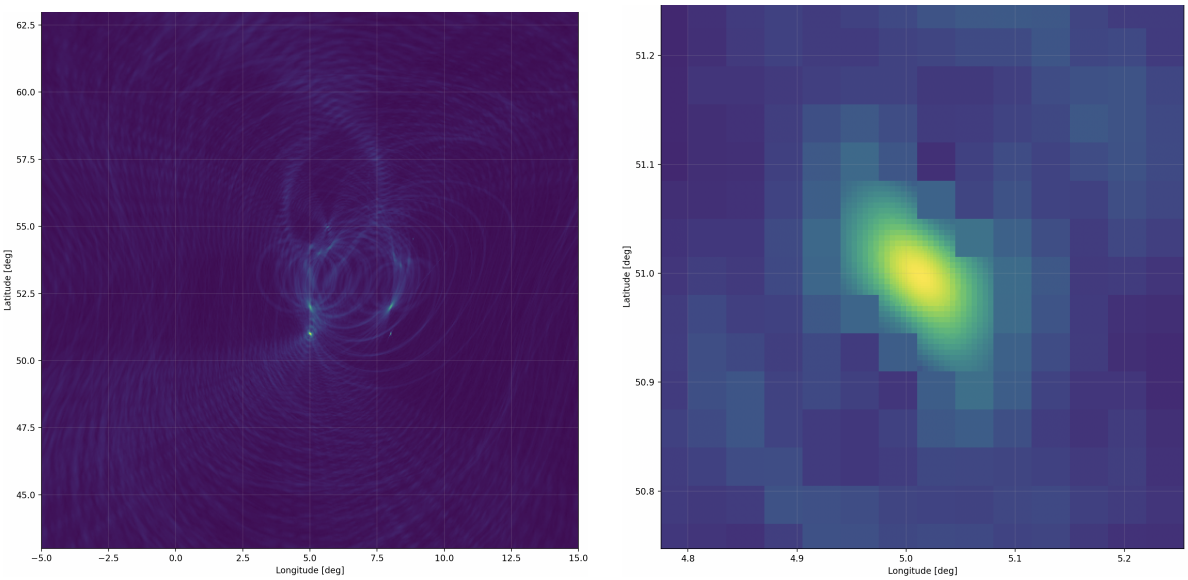

(a) Full field of view **(b)** Zoom on a single peak

Figure 3.12: Accumulated position-domain score maps for Scenario 1. The true emitter locations are marked by red crosses.

Observation Time Lower Bound

Unambiguously resolving all four emitters requires a minimum observation time $T_{\text{obs,min}}$. Asessment of $T_{\text{obs,min}}$ was performed by sweeping over T_{obs} and evaluating the position-domain score map. Once all four emitters were resolved as distinct local maxima, the minimum observation time was determined. The result is a $T_{\text{obs,min}} = 3\,\text{s}$ for Scenario 1 with $T_{\text{int}} = 125\,\text{ms}$ and grid spacing 0.03°. Scenario 2 has a higher observation time required of $T_{\text{obs,min}} = 15\,\text{s}$. This is expected due to the more challenging geometry for two of the four emitters, and the larger difference in SNR. Figure 3.13 shows the number of clusters detected in the position-domain score map for Scenario 2 as a function of T_{obs}. Scenario 1 shows a similar trend, but with a much faster convergence to four clusters.

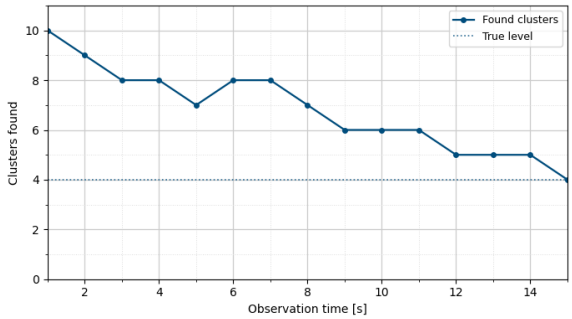

Figure 3.13: Number of clusters detected in the position-domain score map for Scenario 2 as a function of observation time T_{obs}. The minimum observation time $T_{\text{obs,min}} = 15\,\text{s}$ is reached when the four correct clusters are detected

Final Performance Metrics

For the direct method, the core runtime includes the position-domain score evaluation and the hierarchical refinement used to localise emitter peaks. The resulting performance summary is given in Table 3.4.

Table 3.4: Direct-geolocation performance for each emitter at $T_{\text{obs}} = 30\,\text{s}$. Accuracy is the absolute geolocation error, i.e., the distance to truth in meters. The runtime is the time to process one snapshot of $T_{\text{int}} = 125\,\text{ms}$ IQ samples across three receivers.

Scenario	$T_{\text{obs,min}}$ [s]	Emitter	Geo. error [m]	$t_{\text{core,snap}}$ [s]
1	3	E_1	450.1	274 140
		E_2	473.9	
		E_3	650.9	
		E_4	432.4	
2	15	E_1	518.4	
		E_2	648.8	
		E_3	472.6	
		E_4	426.2	

Scenario 1 reaches stable recovery after 3 s. At $T_{\text{obs}} = 30\,\text{s}$, the four per-emitter errors are 450.1 m, 473.9 m, 650.9 m, and 432.4 m for E_1 through E_4, respectively. Scenario 2 requires a longer observation time of 15 s. The corresponding per-emitter errors are 518.4 m, 648.8 m, 472.6 m, and 426.2 m. The core runtime is 274 140 s per snapshot.

3.2.3. Indirect Geolocation Results

The indirect method is reported in the same order as the processing chain of subsection 2.4.2: FDOA candidate extraction, three-receiver association by closure, FDOA track clustering, and nonlinear position estimation.

Observable Extraction

Figure 3.14 shows one representative CAF for the receiver pair (R_1, R_2) at the first snapshot. The true FDOA and TDOA for four emitters are marked with a red circle. Unfortunatelly, they do not all appear as isolated peaks. In particular, the sweep jammers form extended diagonal structures, so the extraction stage must convert structured CAF energy into discrete FDOA candidates. In this specific scenario, the asymmetric sweep jammer E_2 and symmetric sweep jammer E_3 may interfere with the BPSK jammer E_4, which can cause a missed detection in this receiver pair at this snapshot.

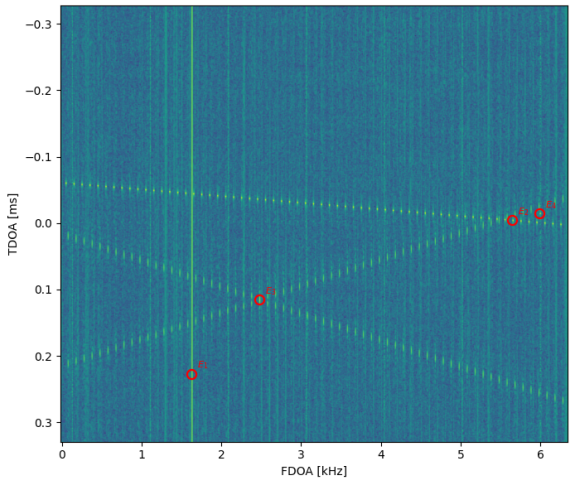

(a) Raw CAF, pair (R_1, R_2)

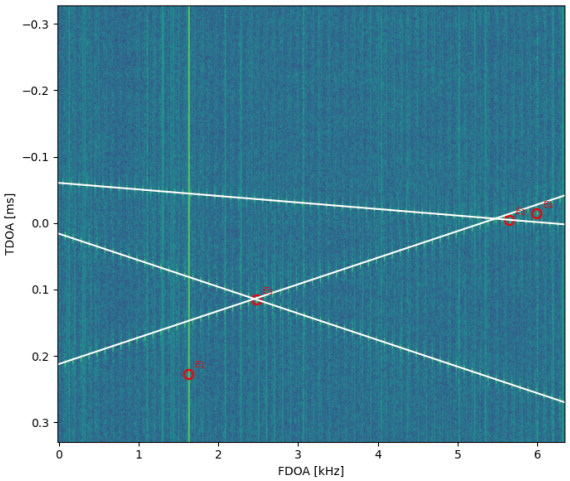

(b) Detected diagonal structures

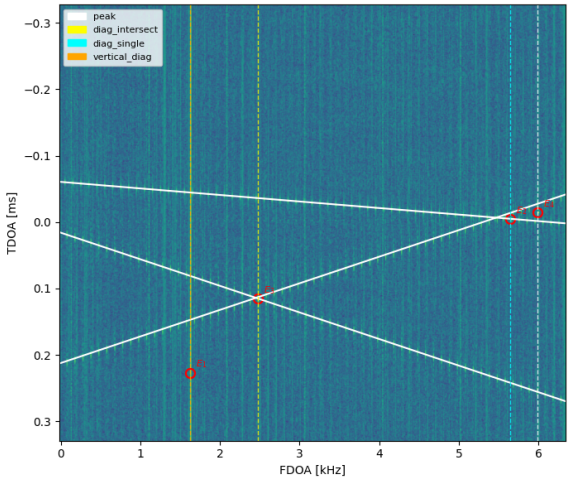

(c) Final retained FDOA candidates

Figure 3.14: Example of the indirect observable-extraction stage for pair (R_1, R_2) at the first snapshot.

Table 3.5 summarises the extraction performance over the full 30 s observation window. The pairwise rows correspond to CAF-based candidate extraction before closure. The final row for each scenario shows the result after enforcing the three-receiver FDOA cyclic-sum constraint.

Table 3.5: Indirect observable-extraction performance over the full observation window.

Scenario	Processing level	Candidates	Missed detection rate (MDR)	False-positive rate (FPR)
1	Pair (R_1, R_2)	155	6.45%	6.45%
	Pair (R_1, R_3)	141	13.71%	2.73%
	Pair (R_2, R_3)	155	3.23%	3.23%
	All receiver pairs	451	7.80%	4.19%
	Accepted triplets	100	24.19%	**0.00%**
2	Pair (R_1, R_2)	150	6.45%	2.52%
	Pair (R_1, R_3)	140	12.10%	0.00%
	Pair (R_2, R_3)	155	1.61%	1.61%
	All receiver pairs	445	6.72%	1.38%
	Accepted triplets	103	16.94%	**0.00%**

The pairwise CAF stage already recovers most true emitter FDOA's, with aggregate MDRs of 7.80% in Scenario 1 and 6.72% in Scenario 2. Its main limitation is not detection, but association. The three receiver pairs produce many candidate combinations, and in principle every candidate from (R_1, R_2) can be combined with every candidate from (R_1, R_3) and (R_2, R_3). The closure constraint is therefore essential, because it identifies the physically consistent triplets among this large combinatorial set.

This is reflected directly in Table 3.5. After closure, only 100 triplets remain in Scenario 1 and 103 in Scenario 2, while the false-positive rate drops to 0.00% in both cases. The closure step is thus the key mechanism of the indirect chain. It converts ambiguous pairwise detections into validated three-receiver hypotheses. This comes at the cost of a higher MDR, since some true candidates are also rejected, but the accepted triplets are highly reliable.

The extraction errors are signal-dependent. The asymmetric sweep jammer is the most difficult case because its CAF structure is asymmetric and does not yield a clear intersection point. The BPSK jammer produces a compact peak, which makes it easy to detect in isolation but sensitive to interference from neighbouring emitters.

FDOA Tracks and Nonlinear Position Estimation

After closure, the accepted triplets are linked over time and converted into emitter-wise FDOA histories. For each receiver pair, a cubic polynomial is fitted to the accepted measurements and evaluated at the reference epoch for the nonlinear solver. Figure 3.15 shows the raw FDOA estimates, the fitted curves, and the residual FDOA errors for the representative receiver pair (R_1, R_2).

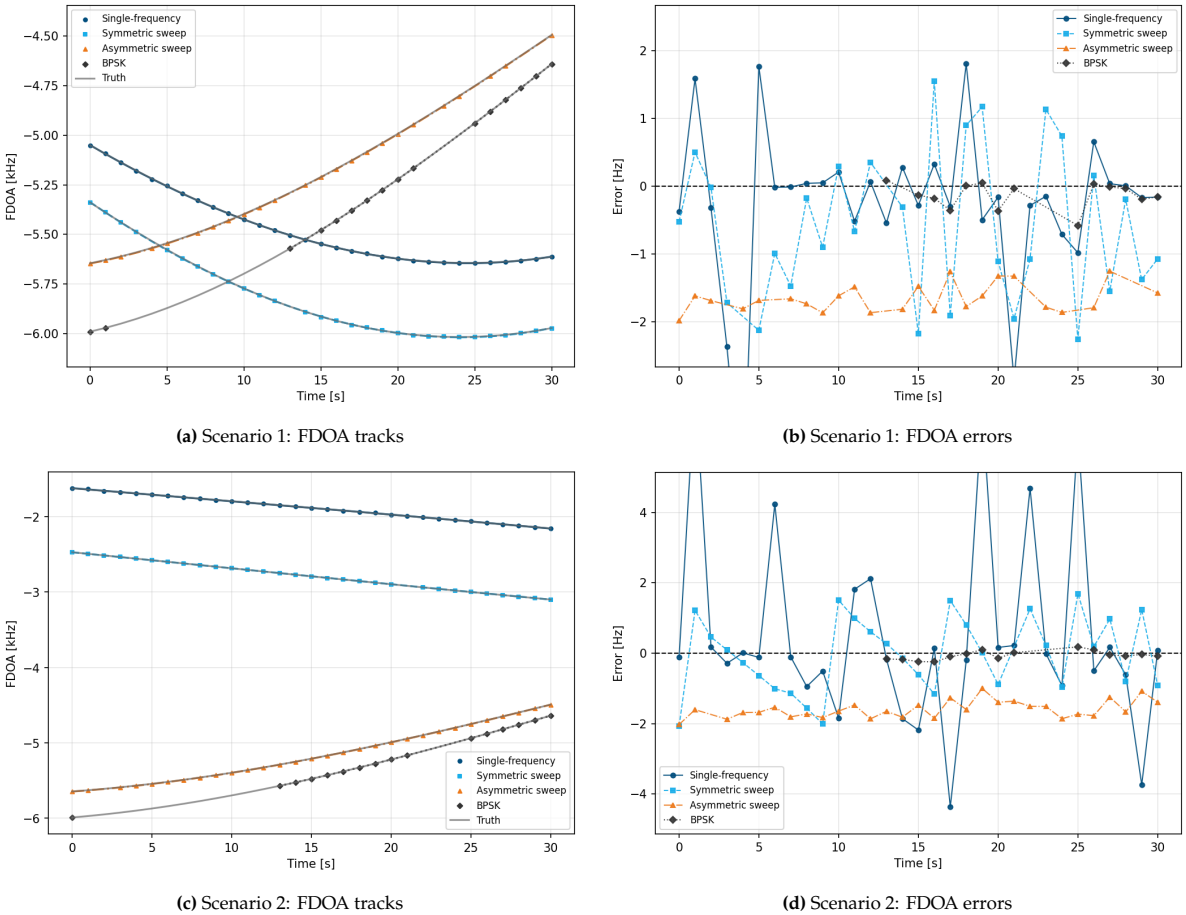

(a) Scenario 1: FDOA tracks

(b) Scenario 1: FDOA errors

(c) Scenario 2: FDOA tracks

(d) Scenario 2: FDOA errors

Figure 3.15: Snapshot-wise FDOA estimates, fitted FDOA curves, and FDOA errors for the receiver pair (R_1, R_2).

In both scenarios, the fitted curves follow the overall truth trend for all four emitters. The residual errors are nevertheless emitter-dependent. The BPSK jammer produces the smallest and most stable

residuals, while the asymmetric sweep jammer shows the largest deviations. The off-centre geometry in Scenario 2 also increases the spread of the residuals. The main result of Figure 3.15 is therefore that the temporal fitting stage preserves the large-scale FDOA evolution, but does not remove emitter-dependent bias introduced in the extraction stage.

These fitted FDOA histories are then passed to the Gauss–Newton solver. The solver is evaluated cumulatively: at each timestamp, all accepted measurements up to that time are included. Figure 3.16 shows the resulting geolocation error as a function of cumulative observation time.

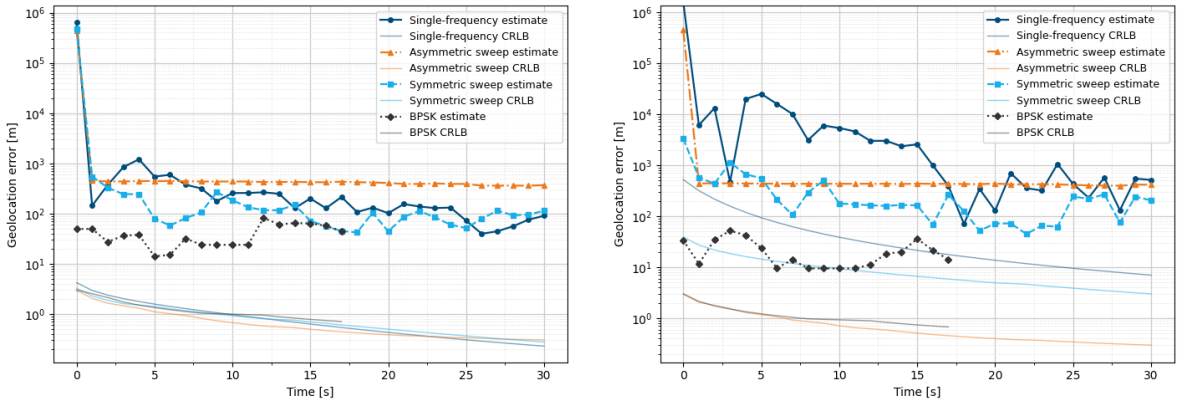

(a) Scenario 1 **(b)** Scenario 2

Figure 3.16: Indirect geolocation error as a function of cumulative observation time for the four emitters in both scenarios.

The geolocation error is large at short observation times and decreases as more measurements are accumulated. In Scenario 1, all emitters converge rapidly and reach a stable error level within a few seconds. In Scenario 2, convergence is slower and the final error level is higher, especially for the emitters moved farther from the formation centre. The figure therefore shows that the indirect solution benefits directly from temporal accumulation, but that the achievable error remains emitter- and geometry-dependent.

The same convergence behaviour can be interpreted from the optimisation surface. Figure 3.17 shows the weighted least-squares cost surface for the representative emitter E_2 at short and long observation times.

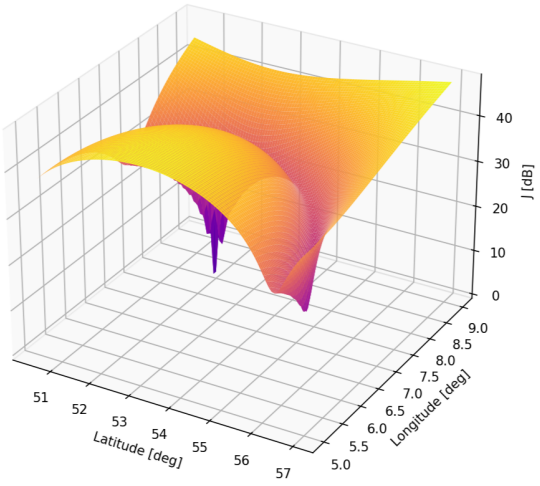

(a) $T_{\text{obs}} = 2\,\text{s}$: surface

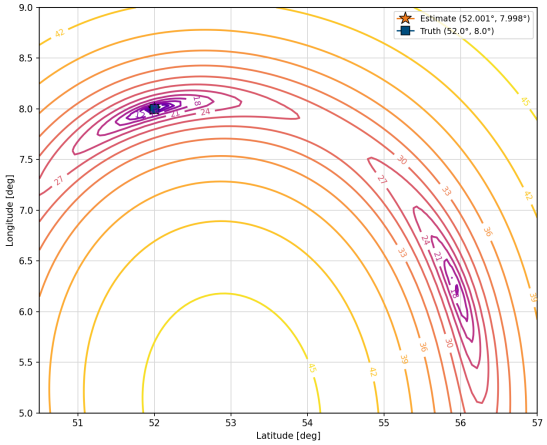

(b) $T_{\text{obs}} = 2\,\text{s}$: contour

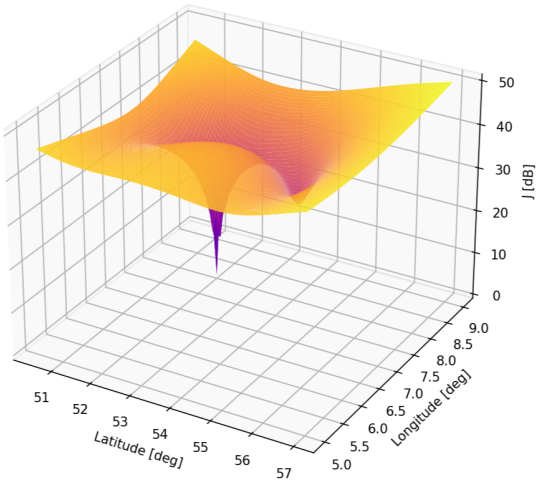

(c) $T_{\text{obs}} = 30\,\text{s}$: surface

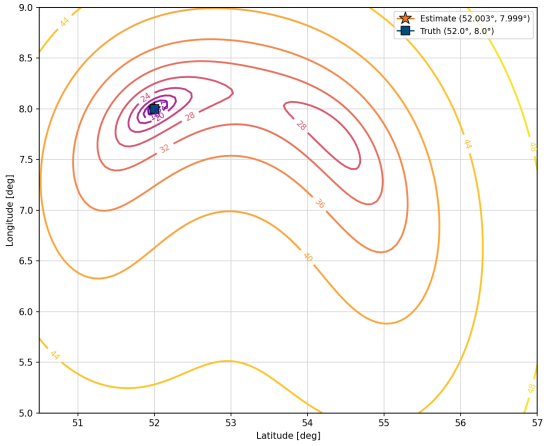

(d) $T_{\text{obs}} = 30\,\text{s}$: contour

Figure 3.17: Weighted least-squares cost surface for emitter E_2 at short and long observation time.

At $T_{\text{obs}} = 2\,\text{s}$, the cost surface already shows a well-defined structure with two narrow local minima, so the solution is largely determined even at short observation time. At $T_{\text{obs}} = 30\,\text{s}$, the true minimum becomes much more dominant relative to the false minimum, which makes the solution more stable. The main effect of longer observation time is therefore variance reduction through improved FDOA averaging and geometric diversity. Any bias in the extracted FDOA tracks remains, however, so the geolocation error approaches a plateau rather than converging to zero.

The final output of the indirect geolocation approach is shown in Figure 3.18, which plots the estimated emitter positions and their associated 95% confidence ellipses for Scenario 1.

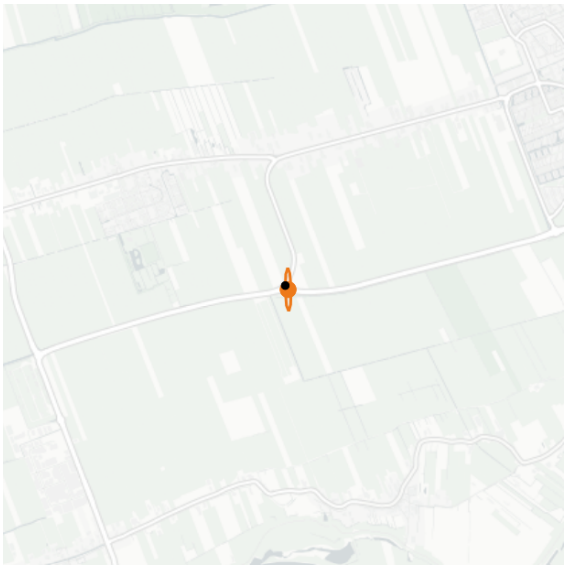

(a) BPSK jammer E_4

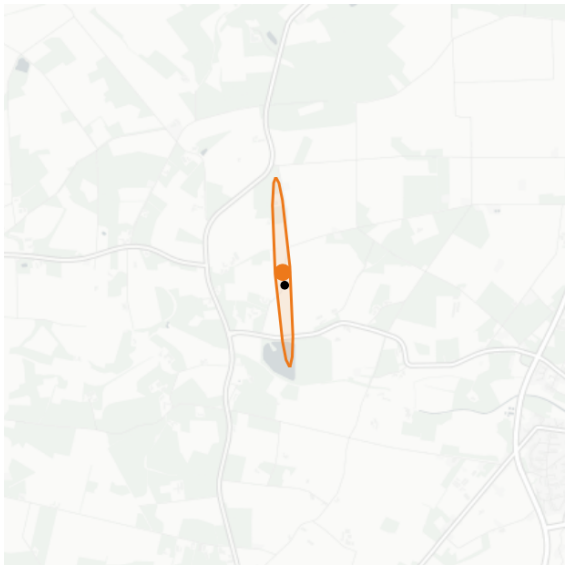

(b) Symmetric sweep jammer E_3

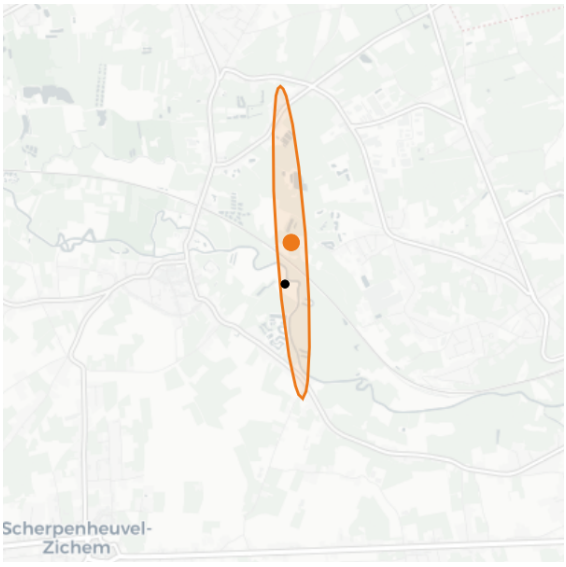

(c) Asymmetric chirp jammer E_2

(d) Constant frequency jammer E_1

Figure 3.18: Final geolocation estimates for Scenario 1. The true emitter geolocation is shown as a black dot, the estimated geolocation as an orange dot, and the 95% confidence ellipse in orange.

All four emitters are recovered in Scenario 1, but the estimate quality is not equal. The BPSK jammer E_4 gives the smallest error and the most compact confidence ellipse, with a geolocation error of 40 m and ellipse axes $a = 190$ m, $b = 30$ m. The asymmetric sweep jammer E_2 gives the largest error and the widest confidence ellipse, with a geolocation error of 339 m and ellipse axes $a = 1760$ m, $b = 170$ m. The tone jammer E_1 and symmetric sweep jammer E_3 lie between these cases. This is consistent with the emitter-dependent FDOA extraction behaviour discussed above.

Final Performance Metrics

For the indirect method, the core runtime includes FDOA candidate extraction, three-receiver association by closure, FDOA track clustering, and nonlinear position estimation. The resulting performance summary is given in Table 3.6.

Table 3.6: Indirect-geolocation performance for each emitter at $T_{obs} = 30\,s$. Accuracy is the absolute geolocation error (distance to truth) in meters. Precision is the 95% confidence ellipse semi-major and semi-minor axes in meters. The runtime is the time to process one snapshot of $\Delta t_{snap} = 1\,s$ IQ data across three receivers.

Scenario	$T_{obs,min}$ [s]	Emitter	Geo. error [m]	95% conf. ellipse [m]	$t_{core,snap}$ [s]
1	2	E_1	88	$a = 1130$, $b = 90$	303.24
		E_2	339	$a = 1760$, $b = 170$	
		E_3	105	$a = 840$, $b = 70$	
		E_4	40	$a = 190$, $b = 30$	
2	16	E_1	514	$a = 4640$, $b = 560$	
		E_2	225	$a = 1230$, $b = 143$	
		E_3	414	$a = 2250$, $b = 461$	
		E_4	15	$a = 85$, $b = 19$	

Scenario 1 reaches stable recovery after 2 s. At $T_{obs} = 30\,s$, the four per-emitter errors are 88 m, 339 m, 105 m, and 40 m for E_1 through E_4, respectively. The associated confidence ellipses show the same emitter dependence: E_4 has the most compact ellipse, while E_2 has the largest semi-major axis. In Scenario 2, the required observation time increases to 16 s. The per-emitter errors are 514 m, 225 m, 414 m, and 15 m. The confidence ellipses are also larger for the off-centre emitters, especially E_1. The core runtime is 303.24 s per snapshot and is dominated by the CAF stage.

3.2.4. Comparison Summary

The results show that both approaches can achieve sub-kilometre geolocation from LEO under the selected formation and observation settings. In Scenario 1, the direct method reaches a mean per-emitter error of 502 m, with errors ranging from 432 m to 651 m. The indirect method gives a lower mean error of 143 m, with errors ranging from 40 m to 339 m. The required observation times are similar, with $T_{obs,min} = 3\,s$ for the direct method and 2 s for the indirect method.

In Scenario 2, the direct method requires 15 s and reaches a mean per-emitter error of 517 m. The indirect method requires 16 s and reaches a mean per-emitter error of 292 m. The indirect method therefore remains more accurate on average, although the difference is emitter-dependent. Its confidence ellipses show that precision also depends strongly on geometry and signal type, with the largest uncertainty occurring for the single-frequency jammer E_1.

The runtime difference is much larger than the accuracy difference. The direct method requires 274 140 s per snapshot, whereas the indirect method requires 303.24 s per snapshot. The indirect method is therefore approximately 9.0×10^2 times faster in the reported implementation.

Overall, the direct method remains attractive because it is conceptually simple and does not depend on a separate observable-extraction stage. The indirect method is more complex, but in the two scenarios considered here it gives lower average error, provides confidence-ellipse information, and retains a very large computational-efficiency advantage.

3.3. Implementation Validation

Certain aspects of the implementation require explicit validation to ensure that the observed behaviour supports reliable conclusions. Other aspects are already validated implicitly through the results presented earlier in this chapter. This section therefore focuses on the explicit validation of the Jacobian expressions, the nonlinear solver, and the clock model.

3.3.1. Observable Jacobian Validation

The CRLB analysis in section 3.1 and the nonlinear solver in Equation 2.4.2 both depend on analytical Jacobians of the TDOA, FDOA, and DFDOA observables. These Jacobians were validated by comparing them with central finite-difference approximations.

The validation used the selected three-satellite formation over a 20 s window with a 2 s step size, giving 11 epochs. At each epoch, 40 random emitter locations were generated within the visible region. For each location, both independent receiver pairs were evaluated. This results in $11 \times 40 \times 2 = 880$ Jacobian

comparisons per observable. The finite-difference step size was set to $h = 10^{-3}$ km.

The relative Jacobian error is defined as

$$\varepsilon_{\text{rel}} = \frac{\|\mathbf{g}_{\text{ana}} - \mathbf{g}_{\text{num}}\|_2}{\|\mathbf{g}_{\text{num}}\|_2}, \tag{3.1}$$

where $\mathbf{g}_{\text{ana}}$ and $\mathbf{g}_{\text{num}}$ are the analytical and numerical Jacobians, respectively.

Figure 3.19 shows the empirical CDF of the relative Jacobian error for TDOA, FDOA, and DFDOA.

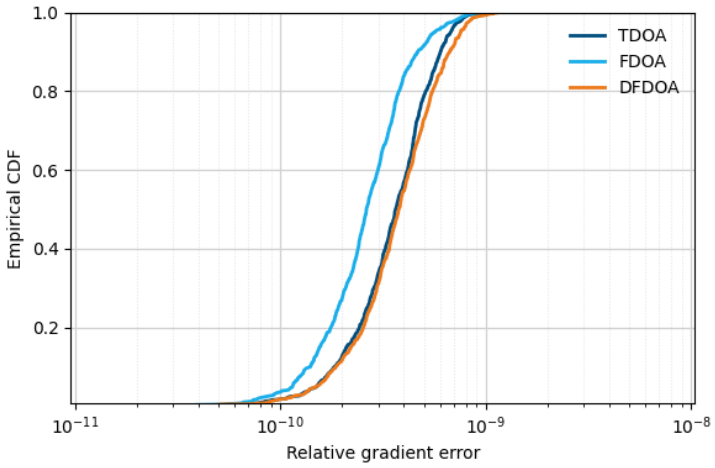

Figure 3.19: Empirical CDF of the relative Jacobian error for TDOA, FDOA, and DFDOA.

All three curves are concentrated between 10^{-10} and 10^{-9}, which indicates agreement at numerical-precision level between the analytical and finite-difference Jacobians.

3.3.2. Gauss–Newton Solver Validation

The nonlinear measurement model and Gauss-Newton solver were validated in a controlled single-emitter test case. A stationary emitter was placed at the known location (52°N, 5°E). Receiver states and Doppler shifts were generated from the satellite trajectories at a single epoch. The resulting receiver positions, velocities, and Doppler shifts are listed in Table 3.7. The derived FDOA measurement vector is given in Table 3.8.

Table 3.7: Receiver states and Doppler shifts for the single-emitter solver validation case

R	$r_{R,x}$ [km]	$r_{R,y}$ [km]	$r_{R,z}$ [km]	$\dot{r}_{R,x}$ [km/s]	$\dot{r}_{R,y}$ [km/s]	$\dot{r}_{R,z}$ [km/s]	f_D [kHz]
i	4059.9	954.2	5532.3	-1.5	7.1	-0.1	26.1
j	4053.1	340.6	5608.5	-1.2	7.2	0.4	0.7
k	4251.6	-227.0	5465.5	-0.8	7.2	0.9	-25.6

Table 3.8: FDOA measurement vector for the single-emitter solver validation case

Component	Value [kHz]
$z_{ij} = \Delta f_D(i,j) = f_{D,i} - f_{D,j}$	25.4
$z_{ik} = \Delta f_D(i,k) = f_{D,i} - f_{D,k}$	51.8
$z_{jk} = \Delta f_D(j,k) = f_{D,j} - f_{D,k}$	26.3

The weighted least-squares cost function was then evaluated over a latitude-longitude grid. The resulting surface and contour map are shown in Figure 3.20.

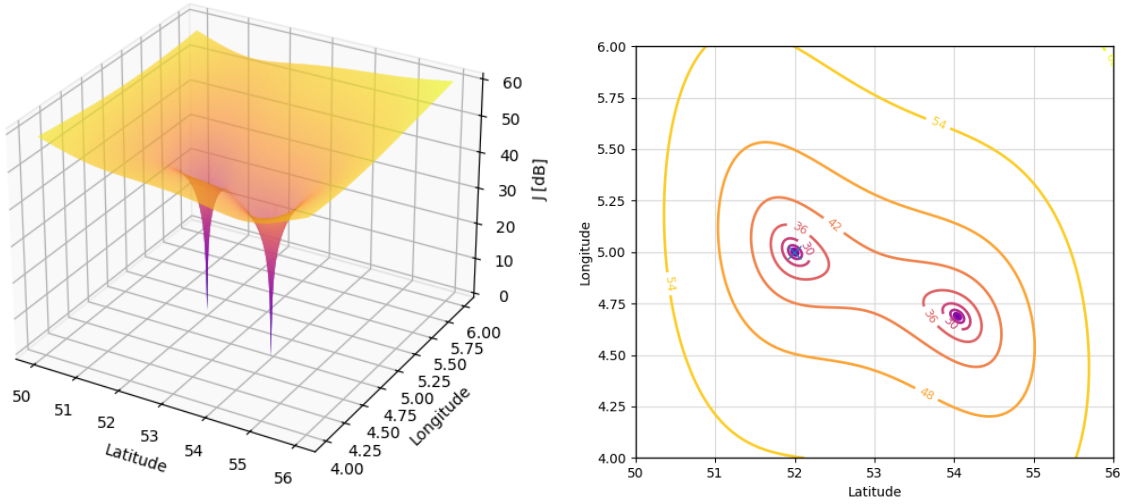

(a) Surface view **(b)** Contour view

Figure 3.20: Weighted least-squares cost surface for the single-emitter FDOA validation case.

The global minimum coincides with the true emitter location. This confirms that the nonlinear measurement model is implemented correctly and that the Gauss-Newton solver converges to the correct solution when exact single-epoch FDOA measurements are used. It also shows that, for the considered three-receiver geometry, the single-emitter solution is unique without requiring temporal accumulation.

3.3.3. Clock Error Validation

Figure 3.21 shows the simulated residual receiver clock offset $\delta t_{R_j}^{\mathrm{res}}(t)$ and residual fractional frequency error $\delta \dot{t}_{R_j}^{\mathrm{res}}(t)$ for the TCXO clock model. In accordance with subsection 2.2.2, both sequences are re-centered to zero at the snapshot midpoint, which here occurs at $t = 0.5\,\mathrm{s}$. This represents GNSS time synchronization between receivers at the reference epoch, after which only the stochastic intra-snapshot residual remains.

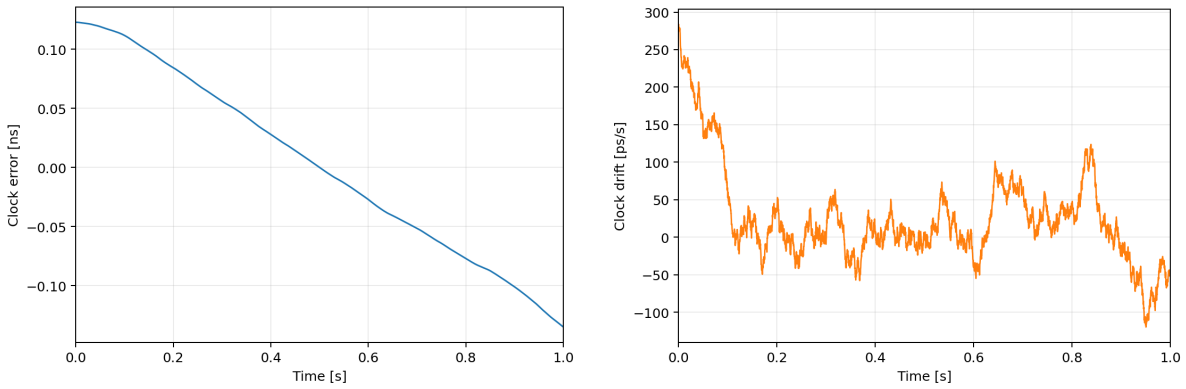

(a) Residual clock offset **(b)** Residual fractional frequency error

Figure 3.21: Residual TCXO clock offset and fractional frequency error over the validation window, re-centered to zero at the synchronization epoch $t = 0.5\,\mathrm{s}$.

The residual clock offset reaches $0.12\,\mathrm{ns}$ at $t = 0\,\mathrm{s}$ and $-0.135\,\mathrm{ns}$ at $t = 1\,\mathrm{s}$. The residual fractional frequency error reaches $275\,\mathrm{ps/s}$ at $t = 0\,\mathrm{s}$ and $-50\,\mathrm{ps/s}$ at $t = 1\,\mathrm{s}$. These values are consistent with a realistic TCXO-based receiver clock in LEO after midpoint synchronization, and therefore confirm that the simulated residual clock behaviour is of the expected order of magnitude for the model adopted in this work.

4

Conclusion and Future Work

4.1. Conclusion

This work investigated passive geolocation of stationary terrestrial GNSS jammers with a three-satellite LEO system. It developed a consistent end-to-end simulation and processing framework, analysed the system design space, and compared direct and indirect geolocation under a common comparison framework.

The main result is that, for the triple-satellite scenarios considered here, indirect geolocation offers the most attractive latency–accuracy trade-off. Both direct and indirect geolocation achieved sub-kilometer per-emitter accuracy with the designed system, and both required observation times of the same order to resolve the emitters. In Scenario 1, the direct method reached a mean per-emitter error of 502 m, while the indirect method reached 143 m. In Scenario 2, the direct method reached 517 m, while the indirect method reached 292 m. The decisive difference was computational cost: in the reported implementation, the indirect method was approximately 9.0×10^2 times faster per snapshot while also achieving lower average localisation error in both evaluated scenarios.

A second main result is that a tight satellite formation is preferred for wide-area surveillance. Small along-track separation preserves the shared field of view and still provides sufficient geometric diversity for sub-kilometer geolocation within short observation times. For the configurations considered in this work, the most favourable operating region lies at small baselines, with the selected design point at $a = 200\,\text{km}$, $b = 25\,\text{km}$, and $a/b = 8$.

A third main result is that the FDOA closure constraint is the key enabler for practical indirect geolocation with three receivers. The pairwise CAF stage produced many plausible candidates, but the closed-loop consistency test reduced the false-positive rate to 0.00% in both scenarios. The importance of the third satellite therefore lies not only in added geometry, but also in the added redundancy for robust candidate association. This makes indirect geolocation much more viable in a multi-emitter setting than it would be for a single or dual-receiver system.

Within the final indirect chain, raw FDOA measurements proved sufficient for geolocation. Neither TDOA nor FDOA-rate information was required in the final estimator. Excluding both reduced model and processing complexity without changing the main conclusions of this work. The indirect method also provided a precision estimate through the 95% confidence ellipse. These ellipses showed that precision is strongly emitter- and geometry-dependent, with the largest uncertainty occurring for the asymmetric chirp jammer.

Overall, this work provides a system-level understanding of passive terrestrial-emitter geolocation from LEO with three receivers. It shows how the main design choices interact. These include formation geometry, observation time, integration time, observable set, and estimation concept. It identifies which parameters drive performance most strongly and where the main trade-offs arise. It also establishes a consistent framework for applying and comparing direct and indirect geolocation in a triple-satellite setting. Within the considered scope, the results point to tight formations and FDOA-based indirect

processing as an attractive design direction for low-latency wide-area surveillance. More broadly, this work provides guidance for the design and assessment of future LEO geolocation systems.

4.2. Future Work

The next steps for this work can be grouped into four themes. The first priority is validation. The second is to broaden the scenario framework. The third is to improve the processing methods. The fourth is to extend the analysis from the geolocation segment to the full mission level.

4.2.1. Validation with Higher-Fidelity Data

The most important next step is external validation. This work is simulation-based, so the developed methods should next be tested on real satellite data. If such data are not yet available, a high-fidelity GNSS interference simulator is the most relevant intermediate step. This would test whether the observed performance persists under more realistic signal distortions, hardware behaviour, and propagation conditions.

A second validation step is broader statistical testing. This work used two controlled end-to-end scenarios, which was sufficient for a clear comparison, but not for a full statistical characterisation. Future work should therefore evaluate both direct and indirect geolocation over many more scenarios with varying emitter locations, powers, signal types, and interference combinations. This is needed to quantify detection performance, localisation accuracy, and robustness more systematically.

4.2.2. Extension of the Scenario and Signal Framework

The simulation framework can be extended in several directions. First, the emitter library should be broadened beyond tones, linear sweeps, and BPSK-like signals. Relevant additions are nonlinear sweeps, pulse jammers, and broadband noise. This would improve the representativeness of the framework and test whether the present conclusions remain valid for a wider signal set.

Second, the environmental and hardware models can be made more realistic. Useful extensions are detailed antenna gain patterns, more advanced clock models, and more sophisticated propagation effects. These additions are important to determine which conclusions are robust system properties and which are sensitive to modelling assumptions.

4.2.3. Algorithmic Development

The clearest algorithmic extension is the CAF observable-extraction stage. In this work, extraction is rule-based and depends on signal structure. A promising next step is to apply machine learning to CAF interpretation. The developed simulation framework can generate large labelled CAF datasets, since the true observables are known by construction. This makes the framework well suited for supervised learning. A learned extractor could simplify the indirect chain and improve robustness in more complex signal scenes.

A second extension is hybrid processing. The results of this work show that direct and indirect geolocation each have clear strengths. Direct geolocation avoids early commitment to observables. Indirect geolocation is far more efficient computationally. A hybrid concept could therefore use indirect geolocation for rapid wide-area search and reserve direct geolocation for local refinement or ambiguous cases.

A third topic is motion and temporal design. Future work should extend the framework to moving emitters such as vehicles or drones. In addition, the effect of larger snapshot intervals should be studied. This work used $\Delta t_{\text{snap}} = 1\,\text{s}$, but longer intervals may provide similar performance at substantially lower processing cost.

4.2.4. System- and Mission-Level Extension

This work focused on the geolocation segment. A natural next step is to extend the analysis to the supporting system architecture. This includes inter-satellite data exchange, synchronization maintenance, onboard versus ground processing, and downlink design, for example via relay satellites. Such work is required to translate geolocation performance into true end-to-end latency.

Finally, the analysis should be extended from one three-satellite cluster to constellation scale. This includes access time, revisit behaviour, tasking strategy, and constellation-level geolocation KPIs. It would connect the local design insights of this work to operational surveillance performance. ef

References

[1] European Union Aviation Safety Agency. *Global Navigation Satellite System (GNSS) Outages and Alterations*. Accessed September 22, 2025. 2025. URL: https://www.easa.europa.eu/en/domains/air-operations/global-navigation-satellite-system-outages-and-alterations.

[2] Alon Amar and Anthony J Weiss. "Localization of narrowband radio emitters based on Doppler frequency shifts". In: *IEEE Transactions on signal processing* 56.11 (2008), pp. 5500–5508.

[3] Michael Bevis et al. "GPS meteorology: Mapping zenith wet delays onto precipitable water". In: *Journal of Applied Meteorology and Climatology* 33.3 (1994), pp. 379–386.

[4] Sriramya Bhamidipati and Grace Xingxin Gao. "GPS multireceiver joint direct time estimation and spoofer localization". In: *IEEE Transactions on Aerospace and Electronic Systems* 55.4 (2018), pp. 1907–1919.

[5] Jahshan Ahmed Bhatti. "Sensor deception detection and radio-frequency emitter localization". PhD thesis. 2015.

[6] Daniel CaJacob et al. "Geolocation of RF emitters with a formation-flying cluster of three microsatellites". In: (2016).

[7] Liang Chen et al. "Robustness, security and privacy in location-based services for future IoT: A survey". In: *Ieee Access* 5 (2017), pp. 8956–8977.

[8] Qidong Chen et al. "Localization of Multiple GNSS Interference Sources Based on Target Detection in C/N0 Distribution Maps". In: *Remote Sensing* 17.15 (2025), p. 2627.

[9] Li Cheng et al. "Structure and performance analysis of signal acquisition and Doppler tracking in LEO augmented GNSS receiver". In: *Sensors* 21.2 (2021), p. 525.

[10] Jacob S Clements and Zachary L Clements. "GPU-accelerated Direct Geolocation of GNSS Interference". In: *arXiv preprint arXiv:2508.06672* (2025).

[11] Zachary Clements, Todd E Humphreys, and Patrick Ellis. "Dual-satellite geolocation of terrestrial GNSS jammers from low Earth orbit". In: *2023 IEEE/ION Position, Location and Navigation Symposium (PLANS)*. IEEE. 2023, pp. 458–469.

[12] Zachary L Clements et al. "Single-satellite-based geolocation of broadcast GNSS spoofers from low Earth orbit". In: *arXiv preprint arXiv:2503.17791* (2025).

[13] Andrew G Dempster and Ediz Cetin. "Interference localization for satellite navigation systems". In: *Proceedings of the IEEE* 104.6 (2016), pp. 1318–1326.

[14] HU Dexiu et al. "Joint TDOA, FDOA and differential Doppler rate estimation: Method and its performance analysis". In: *Chinese Journal of Aeronautics* 31.1 (2018), pp. 137–147.

[15] Nikolas Duetsch et al. "GNSS interference detection and geolocation from LEO satellites–satellite formation and payload design specific considerations and their impact on the detection sensitivity and geolocation accuracy". In: *Proceedings of the 37th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2024)*. 2024, pp. 737–768.

[16] Patrick Ellis, Donald Van Rheeden, and Farid Dowla. "Use of Doppler and Doppler rate for RF geolocation using a single LEO satellite". In: *IEEE Access* 8 (2020), pp. 12907–12920.

[17] Simon Heine, Christian Hofmann, and Andreas Knopp. "End-to-end latency evaluation of LEO satellite IoT systems". In: *IET Conference Proceedings CP873*. Vol. 2023. 48. IET. 2023, pp. 216–221.

[18] Dexiu Hu et al. "CRLB for joint estimation of TDOA, phase, FDOA, and Doppler rate". In: *The Journal of Engineering* 2019.21 (2019), pp. 7628–7631.

[19] LI Jinzhou, GUO Fucheng, et al. "A linear-correction least-squares approach for geolocation using FDOA measurements only". In: *Chinese Journal of Aeronautics* 25.5 (2012), pp. 709–714.

[20] Ashkan Kalantari et al. “Frequency of arrival-based interference localization using a single satellite”. In: *2016 8th Advanced Satellite Multimedia Systems Conference and the 14th Signal Processing for Space Communications Workshop (ASMS/SPSC)*. IEEE. 2016, pp. 1–6.

[21] Elliott D Kaplan and Christopher Hegarty. *Understanding GPS/GNSS: principles and applications*. Artech house, 2017.

[22] John A Klobuchar. “Ionospheric time-delay algorithm for single-frequency GPS users”. In: *IEEE Transactions on aerospace and electronic systems* 3 (1987), pp. 325–331.

[23] Daero Lee and Soon Sik Hwang. “Enhanced Real-Time Onboard Orbit Determination of LEO Satellites Using GPS Navigation Solutions with Signal Transit Time Correction”. In: *Aerospace* 12.6 (2025), p. 508.

[24] Zhixin Liu et al. “An algebraic method for moving source localization using TDOA, FDOA, and differential Doppler rate measurements with receiver location errors”. In: *EURASIP Journal on Advances in Signal Processing* 2019.1 (2019), p. 25.

[25] Zixi Liu, Sherman Lo, and Todd Walter. “GNSS interference source localization using ADS-B data”. In: *Proceedings of the 2022 International Technical Meeting of The Institute of Navigation*. 2022, pp. 158–167.

[26] Pengrui Mao et al. “GNSS interference monitoring based on LEO: suitability study of single-satellite and triple-satellite formation systems”. In: *GPS Solutions* 30.1 (2026), p. 44.

[27] Francesco Menzione et al. “Interference Monitoring from Low Earth Orbit: The OPS-SAT Experiment”. In: *Engineering Proceedings* 88.1 (2025), p. 8.

[28] Oliver Montenbruck et al. “GNSS-based synchronization and monitoring of LEO-PNT onboard time”. In: *Advances in Space Research* (2025).

[29] Matthew J Murrian et al. “First results from three years of GNSS interference monitoring from low Earth orbit”. In: *Navigation* 68.4 (2021), pp. 673–685.

[30] Jia Qiao et al. “A survey of GNSS interference monitoring technologies”. In: *Frontiers in Physics* 11 (2023), p. 1133316.

[31] Karan Sarda et al. “Making the invisible visible: Precision RF-emitter geolocation from space by the hawkeye 360 pathfinder mission”. In: (2018).

[32] Jan Scheuing and Bin Yang. “Disambiguation of TDOA estimates in multi-path multi-source environments (DATEMM)”. In: *2006 IEEE International Conference on Acoustics Speech and Signal Processing Proceedings*. Vol. 4. IEEE. 2006, pp. IV–IV.

[33] Marco Spanghero et al. “GNSS jammer localization and identification with airborne commercial GNSS receivers”. In: *IEEE Transactions on Information Forensics and Security* (2025).

[34] Seymour Stein. “Algorithms for ambiguity function processing”. In: *IEEE Transactions on Acoustics, Speech, and Signal Processing* 29.3 (2003), pp. 588–599.

[35] Peter JG Teunissen, Oliver Montenbruck, et al. *Springer handbook of global navigation satellite systems*. Vol. 10. Springer, 2017.

[36] Micaela Troglia Gamba et al. “GNSS radio frequency interference monitoring from LEO satellites: An in-laboratory prototype”. In: *Sensors* 24.2 (2024), p. 508.

[37] Hans Van der Marel. “Reference systems for surveying and mapping”. In: *Delft University of Technology. http://gnss1. tudelft. nl/pub/vdmarel/reader/CTB3310_RefSystems_1-2a_online. pdf* (2014).

[38] Karel F Wakker. “Fundamentals of astrodynamics”. In: *TU Delft Repository, Delft* (2015), pp. 604–612.

[39] Daniel Weinzierl, Christian A Hofmann, and Andreas Knopp. “Blind Geolocation of RF-Signals with LEO Satellite Formations”. In: *MILCOM 2023-2023 IEEE Military Communications Conference (MILCOM)*. IEEE. 2023, pp. 365–370.

[40] Anthony J Weiss. “Direct geolocation of wideband emitters based on delay and Doppler”. In: *IEEE Transactions on Signal Processing* 59.6 (2011), pp. 2513–2521.

[41] Anthony S Williams. *Expected Position Error for an Onboard Satellite GPS Receiver*. Tech. rep. 2015.

[42] Christopher L Yatrakis. *Computing the cross-ambiguity function: A Review*. State University of New York at Binghamton, 2005.

[43] Tasneem Yousif et al. "A novel GNSS RF interference detection and geolocation algorithm for LEO satellites". In: *Proceedings of the 37th International Technical Meeting of the Satellite Division of The Institute of Navigation (ION GNSS+ 2024)*. 2024, pp. 3375–3389.

[44] Cecilia Maria Zannini et al. "Improved TDOA disambiguation techniques for sound source localization in reverberant environments". In: *Proceedings of 2010 IEEE International Symposium on Circuits and Systems*. IEEE. 2010, pp. 2666–2669.

A

Receiver-state Frame Transformations

This appendix summarizes the propagation of each satellite receiver state from the classical orbital elements to Earth-Centered Inertial (ECI) and Earth-Centered Earth-Fixed (ECEF) coordinates.

Each receiver R_i is parameterized at epoch t_0 by

$$(a, e, i, \Omega, \omega, M_0), \tag{A.1}$$

where a is the semi-major axis, e the eccentricity, i the inclination, Ω the right ascension of the ascending node, ω the argument of perigee, and M_0 the mean anomaly at epoch. Let μ denote the gravitational parameter and define

$$\Delta t = t - t_0. \tag{A.2}$$

The mean motion and mean anomaly at time t are

$$n = \sqrt{\frac{\mu}{a^3}}, \qquad M(t) = M_0 + n\Delta t. \tag{A.3}$$

The eccentric anomaly E follows from Kepler's equation,

$$M = E - e \sin E, \tag{A.4}$$

which is solved numerically. The true anomaly is then obtained as

$$\nu = 2 \arctan\left(\sqrt{\frac{1+e}{1-e}} \tan \frac{E}{2}\right). \tag{A.5}$$

Defining the semi-latus rectum

$$p = a(1 - e^2), \tag{A.6}$$

the position and velocity in the perifocal frame (PQW) are

$$\boldsymbol{r}_{\mathrm{PQW}} = \frac{p}{1 + e\cos\nu} \begin{bmatrix} \cos\nu \\ \sin\nu \\ 0 \end{bmatrix}, \qquad \dot{\boldsymbol{r}}_{\mathrm{PQW}} = \sqrt{\frac{\mu}{p}} \begin{bmatrix} -\sin\nu \\ e + \cos\nu \\ 0 \end{bmatrix}. \tag{A.7}$$

The rotation from the perifocal frame to the ECI frame is written as

$$\boldsymbol{Q} = \boldsymbol{R}_3(\Omega)\boldsymbol{R}_1(i)\boldsymbol{R}_3(\omega), \tag{A.8}$$

where

$$\boldsymbol{R}_1(\alpha) = \begin{bmatrix} 1 & 0 & 0 \\ 0 & \cos\alpha & \sin\alpha \\ 0 & -\sin\alpha & \cos\alpha \end{bmatrix}, \qquad \boldsymbol{R}_3(\alpha) = \begin{bmatrix} \cos\alpha & \sin\alpha & 0 \\ -\sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{bmatrix}. \tag{A.9}$$

The satellite state in ECI is therefore

$$\boldsymbol{r}_{R,\mathrm{ECI}} = \boldsymbol{Q}\,\boldsymbol{r}_{\mathrm{PQW}}, \qquad \dot{\boldsymbol{r}}_{R,\mathrm{ECI}} = \boldsymbol{Q}\,\dot{\boldsymbol{r}}_{\mathrm{PQW}}. \tag{A.10}$$

To express the receiver state in Earth-Centered Earth-Fixed coordinates, the ECI state is rotated about the Earth z-axis using the Greenwich Mean Sidereal Time angle $\theta(t)$,

$$\boldsymbol{r}_{R,\mathrm{ECEF}} = \boldsymbol{R}_3(\theta)\,\boldsymbol{r}_{R,\mathrm{ECI}}. \tag{A.11}$$

The corresponding ECEF velocity is

$$\dot{\boldsymbol{r}}_{R,\mathrm{ECEF}} = \boldsymbol{R}_3(\theta)\,\dot{\boldsymbol{r}}_{R,\mathrm{ECI}} - \boldsymbol{\omega}_E \times \boldsymbol{r}_{R,\mathrm{ECEF}}, \qquad \boldsymbol{\omega}_E = \begin{bmatrix} 0 \\ 0 \\ \Omega_E \end{bmatrix}, \tag{A.12}$$

where Ω_E is the Earth rotation rate. This correction accounts for the fact that the ECEF frame rotates with the Earth, whereas the propagated orbital state is first obtained in the inertial frame. This neglects smaller effects such as precession, nutation, and polar motion, which are not expected to significantly impact the results of this work.

B

Observable gradient derivations

This appendix derives the gradients of the geometric parts of the TDOA, FDOA, and DFDOA observables with respect to the emitter position $\boldsymbol{r}_E$. Receiver states are treated as known and independent of $\boldsymbol{r}_E$. Gradients of scalar quantities are written as row vectors.

B.1. TDOA gradient derivation

From chapter 2

$$\tau_{ij}^{\mathrm{g}}(t) = \frac{\rho_{ij}(t)}{c}, \qquad \rho_{ij}(t) = \|\boldsymbol{\Delta}_{ij}(t)\|, \qquad \boldsymbol{\Delta}_{ij}(t) = \boldsymbol{r}_E - \boldsymbol{r}_{R_j}(t). \tag{B.1}$$

Since $\boldsymbol{r}_{R_j}(t)$ is independent of $\boldsymbol{r}_E$,

$$\nabla\rho_{ij}(t) = \nabla\|\boldsymbol{\Delta}_{ij}(t)\| = \frac{\boldsymbol{\Delta}_{ij}^{\mathsf{T}}(t)}{\|\boldsymbol{\Delta}_{ij}(t)\|} = \hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t). \tag{B.2}$$

Therefore,

$$\nabla\tau_{ij}(t) = \nabla\tau_{ij}^{\mathrm{g}}(t) = \frac{1}{c}\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t), \tag{B.3}$$

where the non-geometric terms are treated as additive and position-independent in the local gradient calculation.

Hence, for the receiver pair (R_j, R_k),

$$\nabla\Delta\tau_{i,jk}(t) = \nabla\tau_{ij}(t) - \nabla\tau_{ik}(t). \tag{B.4}$$

B.2. FDOA gradient derivation

The geometric range rate is

$$\dot{\rho}_{ij}(t) = \frac{\boldsymbol{\Delta}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)}{\rho_{ij}(t)}. \tag{B.5}$$

Because the emitter is stationary, $\dot{\boldsymbol{\Delta}}_{ij}(t)$ is independent of $\boldsymbol{r}_E$. Differentiating with respect to $\boldsymbol{r}_E$ gives

$$\nabla\dot{\rho}_{ij}(t) = \frac{\dot{\boldsymbol{\Delta}}_{ij}^{\mathsf{T}}(t)}{\rho_{ij}(t)} - \frac{\boldsymbol{\Delta}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)}{\rho_{ij}^2(t)}\nabla\rho_{ij}(t). \tag{B.6}$$

Using

$$\nabla\rho_{ij}(t) = \hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t), \qquad \boldsymbol{\Delta}_{ij}(t) = \rho_{ij}(t)\hat{\boldsymbol{\ell}}_{ij}(t), \tag{B.7}$$

this becomes

$$\nabla\dot{\rho}_{ij}(t) = \frac{1}{\rho_{ij}(t)}\left[\dot{\boldsymbol{\Delta}}_{ij}(t) - \hat{\boldsymbol{\ell}}_{ij}(t)\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right]^{\mathsf{T}}. \tag{B.8}$$

Defining

$$\boldsymbol{P}_{ij}^{\perp}(t) = \boldsymbol{I} - \hat{\boldsymbol{\ell}}_{ij}(t)\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t), \tag{B.9}$$

yields

$$\nabla\dot{\rho}_{ij}(t) = \left[\frac{1}{\rho_{ij}(t)}\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right]^{\mathsf{T}}. \tag{B.10}$$

From the geometric Doppler relation,

$$f_{ij}^{\mathrm{g}}(t) = -\frac{f_c}{c}\dot{\rho}_{ij}(t), \tag{B.11}$$

and with non-geometric terms again treated as additive and position-independent,

$$\nabla\tilde{f}_{ij}(t) = -\frac{f_c}{c}\nabla\dot{\rho}_{ij}(t). \tag{B.12}$$

Therefore,

$$\nabla\Delta f_{i,jk}(t) = -\frac{f_c}{c}\left(\nabla\dot{\rho}_{ij}(t) - \nabla\dot{\rho}_{ik}(t)\right). \tag{B.13}$$

B.3. DFDOA gradient derivation

Starting from

$$\dot{\rho}_{ij}(t) = \hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t), \tag{B.14}$$

differentiate with respect to time:

$$\ddot{\rho}_{ij}(t) = \dot{\hat{\boldsymbol{\ell}}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t) + \hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\ddot{\boldsymbol{\Delta}}_{ij}(t). \tag{B.15}$$

Using

$$\hat{\boldsymbol{\ell}}_{ij}(t) = \frac{\boldsymbol{\Delta}_{ij}(t)}{\rho_{ij}(t)}, \tag{B.16}$$

its time derivative is

$$\dot{\hat{\boldsymbol{\ell}}}_{ij}(t) = \frac{1}{\rho_{ij}(t)}\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t). \tag{B.17}$$

Substitution gives

$$\ddot{\rho}_{ij}(t) = \hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\ddot{\boldsymbol{\Delta}}_{ij}(t) + \frac{\left\|\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right\|^2}{\rho_{ij}(t)}. \tag{B.18}$$

Now differentiate with respect to $\boldsymbol{r}_E$. For the first term,

$$\nabla\left(\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\ddot{\boldsymbol{\Delta}}_{ij}(t)\right) = \left[\frac{\boldsymbol{P}_{ij}^{\perp}(t)\ddot{\boldsymbol{\Delta}}_{ij}(t)}{\rho_{ij}(t)}\right]^{\mathsf{T}}, \tag{B.19}$$

since $\ddot{\boldsymbol{\Delta}}_{ij}(t)$ is independent of $\boldsymbol{r}_E$. For the second term, let

$$\alpha_{ij}(t) = \left\|\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right\|^2 = \|\dot{\boldsymbol{\Delta}}_{ij}(t)\|^2 - \left(\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right)^2. \tag{B.20}$$

Then

$$\nabla\left(\frac{\alpha_{ij}(t)}{\rho_{ij}(t)}\right) = \frac{1}{\rho_{ij}(t)}\nabla\alpha_{ij}(t) - \frac{\alpha_{ij}(t)}{\rho_{ij}^2(t)}\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t). \tag{B.21}$$

Because $\dot{\boldsymbol{\Delta}}_{ij}(t)$ is independent of $\boldsymbol{r}_E$,

$$\nabla\alpha_{ij}(t) = -\frac{2}{\rho_{ij}(t)}\left(\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right)\left(\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right)^{\mathsf{T}}. \tag{B.22}$$

Hence,

$$\nabla\left(\frac{\alpha_{ij}(t)}{\rho_{ij}(t)}\right) = \left[-\frac{2\left(\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right)\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t) + \left\|\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right\|^2\hat{\boldsymbol{\ell}}_{ij}(t)}{\rho_{ij}^2(t)}\right]^{\mathsf{T}}. \tag{B.23}$$

Combining both parts yields

$$\nabla\ddot{\rho}_{ij}(t) = \left[\frac{\boldsymbol{P}_{ij}^{\perp}(t)\ddot{\boldsymbol{\Delta}}_{ij}(t)}{\rho_{ij}(t)} - \frac{2\left(\hat{\boldsymbol{\ell}}_{ij}^{\mathsf{T}}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right)\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t) + \left\|\boldsymbol{P}_{ij}^{\perp}(t)\dot{\boldsymbol{\Delta}}_{ij}(t)\right\|^2\hat{\boldsymbol{\ell}}_{ij}(t)}{\rho_{ij}^2(t)}\right]^{\mathsf{T}}. \tag{B.24}$$

Finally, from

$$\dot{\tilde{f}}_{ij}(t) = -\frac{f_c}{c}\ddot{\rho}_{ij}(t) + \text{additive non-geometric terms}, \tag{B.25}$$

it follows that

$$\nabla\Delta\dot{f}_{i,jk}(t) = -\frac{f_c}{c}\left(\nabla\ddot{\rho}_{ij}(t) - \nabla\ddot{\rho}_{ik}(t)\right). \tag{B.26}$$

C

Zero Cyclic Sum Mathematical Theory

Let there be N nodes (receivers) 1, ..., N. Each node has some scalar quantity x_i associated with it. Define the pairwise differences between nodes (i, j) as

$$\Delta_{ij} = x_i - x_j \tag{C.1}$$

Where $\Delta_i j = -\Delta_{ji}$. Now pick any subset of nodes

$$S = \{x_1, x_2, ..., x_M\} \tag{C.2}$$

Now cycle through the subset and define the cycle sum As

$$\sum_{i=1}^{M} \Delta_{ij} = \Delta_{i_1,i_2} + \Delta_{i_2,i_3} + ... + \Delta_{i_{M-1},i_M} + \Delta_{i_M,i_1} \tag{C.3}$$

Then the condition holds

$$\sum_{i=1}^{M} \Delta_{ij} = (x_{i_1} - x_{i_2}) + (x_{i_2} - x_{i_3}) + ... + (x_{i_{M-1}} - x_{i_M}) + (x_{i_M} - x_{i_1}) = 0 \tag{C.4}$$

Hence for any subset of M nodes (receivers), and any closed loop on that subset, the sum of the pairwise differences along that loop is identically zero.

D

Thesis Planning

The full planning overview is shown on the next page.

Student Name	Joep Raithel		2025														2026																							
Student ID	5129648		38	39	40	41	42	43	44	45	46	47	48	49	50	51	52	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22
1	Milestones																																						
		Research proposal Review																																					
		Mid-term Review																																					
		Green light Review																																					
		Thesis Defence																																					
2	Research Proposal																																						
		Understand & Scope (read 10 papers)																																					
		Draft introduction section																																					
		Propose 3 research questions																																					
		Propose problem statement																																					
		Scope research																																					
		Draft final Research Proposal																																					
3	Build Simulation Framework																																						
		Simulate sats & emitters in Python																																					
		Simulate signal transfer in Python																																					
		Simulate signal geolocation in Python																																					
4	Explore geolocation approaches																																						
		Indirect Geolocation																																					
		Direct Geolocation																																					
5	Verification & Validation																																						
		Validate Signal Transmitter and Receiver Dynamics																																					
6	System Performance Analysis																																						
		Analyse runtime per geolocation approach																																					
		Exploration Global Parameter Sensitivity																																					
		Cramer-Rao Lower Bound analysis																																					
7	Finalize geolocation algorithms																																						
		Indirect Geolocation																																					
		Direct Geolocation																																					
7	Project conclusion																																						
		Synthesise main findings																																					
		Finalize thesis paper																																					
		Prepare Greenlight																																					
		Finalize Research portfolio																																					
		Prepare Defence																																					