Delft University of Technology

Faculty of Electrical Engineering, Mathematics and Computer Science

MASTER OF SCIENCE THESIS

Distributed FMCW Radar Localization Under Residual Phase Errors

by

Floris Grootenboer

Student number: 5376505

Programme:	MSc Electrical Engineering
Track:	Wireless Communications and Sensing
Supervisor:	Prof. Olexander Yarovoy

November 2026

Preface

This thesis was completed as part of the Wireless Communications and Sensing track of the MSc programme in Electrical Engineering at Delft University of Technology. Throughout this journey, I have had the opportunity to significantly develop both my practical and theoretical skills, while exploring a field that has become increasingly fascinating to me.

First, I would like to express my sincere gratitude to my supervisors, Prof. Olexander Yarovoy and Rutkay Guneri. Both have been incredibly helpful in overcoming technical challenges and providing valuable guidance on the direction of this research. I am especially grateful for the time they have taken out of their busy schedules to support me throughout this process and help me work towards obtaining my Master of Science degree.

I would also like to thank the Microwave Sensing, Signals and Systems group for offering such interesting courses and excellent lecturers. The courses provided a perfect bridge between my background in physics and the field of radar engineering, allowing me to develop an understanding of radar systems that I could only have dreamed of a few years ago.

Lastly, I would like to thank my family and friends for always being there for me. Their support, encouragement, and the sense of stability they provide have been invaluable, especially during the more challenging moments of this journey.

Looking back, this thesis has been a challenging but rewarding conclusion to my time at TU Delft, and I am grateful to everyone who has been part of that journey.

Floris Grootenboer
Delft, November 2026

Abstract

Combining spatially separated frequency-modulated continuous-wave (FMCW) radar nodes can improve angular resolution beyond that of a single module. Coherent use of this distributed aperture requires carrier-phase calibration as well as timing and frequency synchronization. This thesis evaluates the localization sensitivity of two independently clocked 77 GHz radar nodes forming two monostatic and two reciprocal bistatic apertures.

A common processing framework compares Bartlett beamforming, independent FOCUSS, Block-FOCUSS, multi-task Bayesian compressive sensing (MT-BCS), and coherent FOCUSS, alongside coherent and non-coherent Cartesian backpropagation. The contribution is an estimator-specific assessment linking residual phase structure after synchronization to angular accuracy, source-count recovery, and relative-separation success. Controlled simulations distinguish intra-aperture manifold distortion, inter-aperture block rotations, and phase evolution within a coherent processing interval (CPI).

In the static-phase studies, MT-BCS gives the most consistent sparse close-pair performance without requiring a common carrier phase. Coherent processing produces finer ideal spatial responses but is more sensitive to calibration and baseline-dependent ambiguity. In single-target CPI tests with fixed waveform settings, peak-tracked coherent FOCUSS with a frozen reference retains angular root-mean-square error (RMSE) below 0.05° for all tested residual frequency offsets through 100 Hz. Strong random phase walks and incompatible target/reference Doppler readouts can degrade coherent estimates; accurate dominant-peak angles can coexist with incorrect source counts.

The main physical campaign uses a 25 cm radar baseline, nine target separations, and 72 frame-level cases. Some estimators recover relative spacing, but none achieves absolute-angle success: exactly two targets with both bearing errors within 2°. Later 14.03- and 19.5 cm-baseline campaigns compare processing without compensation against a separate spherical reference. Reference compensation remains limited by geometric uncertainty, while scene-assisted correction is an in-sample diagnostic. A measurement-matched simulator reproduces selected range–Doppler structures but not the measured estimator rankings. Together, these results identify when residual phase permits useful distributed-aperture localization and why a narrow coherent response alone is insufficient evidence of accurate physical localization.

Contents

Preface **i**

Nomenclature **vi**

1 Introduction **1**
1.1 Research Problem and Gap 2
1.2 Research Objective and Questions 3
1.3 Contributions 4
1.4 Thesis Outline 5

2 Literature Review **7**
2.1 Synchronization in Distributed FMCW Radar 7
2.1.1 Coarse Synchronization 8
2.1.2 Treatment of Residual Phase Error in Prior Work 8
2.1.3 Fine Synchronization 10
2.2 Direction-of-Arrival and Localization Methods 11
2.2.1 Conventional Spectral Estimators 11
2.2.2 Subspace-Based Estimators 12
2.2.3 Maximum-Likelihood Estimation 12
2.2.4 Sparse Reconstruction and Shared Support 12
2.3 Coherent and Non-Coherent Fusion 13
2.4 Synthesis and Research Gap 14

3 Theoretical Background and Signal Model **16**
3.1 System Model 16
3.2 Signal Model 18
3.3 Residual Phase-Error Model 20
3.3.1 Intra-Aperture Phase Error 20
3.3.2 Inter-Aperture Phase Error 21
3.3.3 Within-CPI Phase Evolution 22

4 Processing and Estimation Methodology **23**
4.1 Input Data and Channel Organization 23

4.2 Receive-Side Synchronization and Phase Stabilization 24
4.2.1 Coarse Bistatic Range–Doppler Alignment 24
4.2.2 Fine Synchronization and Inter-Aperture Phase Stabilization 25
4.3 Localization Preprocessing and Geometric Association 26
4.3.1 Processing Framework 26
4.3.2 Range-Resolved Static Snapshots 27
4.3.3 Geometric Parameterization 29
4.3.4 Aperture-Local Localization and Target Association 29
4.3.5 Angular Search Grid and Steering-Dictionary Construction 30
4.4 Direction-of-Arrival Estimators 33
4.4.1 Bartlett Beamforming 33
4.4.2 FOCUSS 34
4.4.3 Block-FOCUSS 35
4.4.4 Coherent Fine-Grid FOCUSS 36
4.4.5 Bayesian Compressive Sensing 37
4.4.6 Truth-Blind Support and Model-Order Selection 38
4.4.7 Numerical Settings 40
4.5 Cartesian Backpropagation Localization 41
4.5.1 Non-Coherent Backpropagation 41
4.5.2 Coherent Backpropagation 42

5 Simulation and Experimental Evaluation Methodology 45
5.1 Simulation Architecture 45
5.2 Target Geometry and IF Signal Model 46
5.3 Simulation Impairment Model 48
5.3.1 Inter-Radar Frequency and Phase Errors 48
5.3.2 Additional Nonidealities 49
5.4 Complete Receiver Data and Simulation Validation 50
5.5 Physical Measurement Setup and Data Acquisition 51
5.6 Evaluation Geometries and Performance Metrics 54
5.6.1 Single-Target Scene 54
5.6.2 Same-Cell Two-Target Scene 55
5.7 Evaluation Protocols 56
5.7.1 Ideal Controls and Estimator Baselines 56
5.7.2 Baseline Ambiguity Comparison 57
5.7.3 Residual Phase-Sensitivity Protocol 57
5.7.4 Within-CPI Phase-Evolution Study 58
5.7.5 Physical Frame Evaluation and Phase Calibration 60
5.7.6 Backpropagation and Measurement-Matched Validation 60
5.8 Ground Truth, Model Scope, and Evaluation Limits 60
5.9 Evaluation Methodology Summary 61

6 Results and Discussion **63**
6.1 Ideal Angular Resolution of the Distributed Aperture 63
6.2 Baseline-Dependent Coherent and Non-Coherent Ambiguity 65
6.3 Sensitivity to Chirp-Invariant Residual Phase Errors 66
6.3.1 Single-Target Phase Error Sensitivity 67
6.3.2 Two-Target Pair Phase Error Sensitivity 67
6.4 Sensitivity to Phase Evolution Within the CPI 68
6.5 Physical Close-Pair Measurements and Phase Quality 72
6.5.1 Target-Count and Relative-Separation Results 72
6.5.2 Residual Phase Quality in the Measured Scenes 74
6.6 Cartesian Backpropagation: Ideal and Physical Localization 77
6.7 Simulation-to-Measurement Comparison 79
6.8 Discussion 81

7 Conclusions and Future Work **85**
7.1 Summary of Main Findings 85
7.2 Answers to the Research Questions 86
7.2.1 RQ1: Synchronization and Residual Errors 86
7.2.2 RQ2: Estimator Sensitivity to Residual Phase Errors 87
7.2.3 RQ3: Conditions for Coherent Localization 88
7.2.4 RQ4: Simulation-to-Measurement Transfer 88
7.3 Research Contributions 89
7.4 Limitations 90
7.5 Future Work 90

A Implementation and Reproducibility Details **96**
A.1 Common Acquisition Parameters 96
A.2 Study-Specific Processing Configurations 97
A.3 Dataset and Evaluation Inventory 98
A.4 Reproducibility and Data Provenance 98
A.5 Supplementary Within-CPI Results 100
A.5.1 Fixed Range–Doppler Cells 100
A.5.2 Moving-Target Control 101
A.6 Additional Estimator Diagnostics 102
A.6.1 Measured and Synthetic Pair Comparison 102
A.6.2 Single-Target Model-Order Recovery 102
A.6.3 Physical Angular Error Relative to the Nominal Geometry 103

Nomenclature

Symbols

c Speed of light in vacuum.

f_c Radar carrier frequency.

λ Carrier wavelength, $\lambda = c/f_c$.

B FMCW sweep bandwidth.

μ FMCW chirp slope.

T FMCW chirp duration.

T_c Chirp repetition interval.

f_s ADC sampling frequency.

N_{ADC} Number of fast-time samples per chirp.

N_c Number of chirps per processed sequence.

$\mathcal{P}$ Set of available monostatic and bistatic apertures.

$\mathbf{p}$ Candidate or true target position in the global Cartesian frame.

$d_{p,ij,k}$ Propagation-path length for aperture p, TX element i, RX element j, and target k.

$\mathbf{y}_p$ Complex spatial snapshot of aperture p.

$\mathbf{A}_p$ Steering dictionary of aperture p.

$\mathbf{x}_p$ Sparse angular coefficient vector of aperture p.

θ_{g} Target bearing in the global coordinate system.

β_p Common residual phase rotation of aperture block p.

$\delta\phi_{p,q}$ Wrapped channel-phase residual.

$\eta_{\phi,p}$ Normalized coherent-projection metric.

$\Delta\theta$ Angular separation between a target pair.

Abbreviations

ADC Analogue-to-digital converter.

AR(1) First-order autoregressive process.

BCS Bayesian compressive sensing.

BP Backpropagation.

CA-CFAR Cell-averaging constant-false-alarm-rate detector.

CFAR Constant false alarm rate.

CPI Coherent processing interval.

CUT Cell under test.

DCA1000 Texas Instruments raw-data capture board.

DoA Direction of arrival.

eBIC Extended Bayesian information criterion.

FFT Fast Fourier transform.

FMCW Frequency-modulated continuous wave.

ESPRIT Estimation of signal parameters via rotational invariance techniques.

FOCUSS FOCal Underdetermined System Solver.

GNSS Global navigation satellite system.

IF Intermediate-frequency signal.

LO Local oscillator.

MIMO Multiple-input multiple-output.

MMV Multiple-measurement vector.

MT-BCS Multi-task Bayesian compressive sensing.

MUSIC Multiple signal classification.

MVDR Minimum-variance distortionless response.

OS-CFAR Ordered-statistic constant-false-alarm-rate detector.

PRI Pulse repetition interval.

RCS Radar cross section.

RD Range–Doppler.

RF Radio frequency.

RMS Root mean square.

RMSE Root-mean-square error.

RX Receive or receiver channel.

SNR Signal-to-noise ratio.

TDMA Time-division multiple access.

TX Transmit or transmitter channel.

ULA Uniform linear array.

Chapter 1

Introduction

Frequency-modulated continuous-wave (FMCW) radar is widely used for environmental sensing because it provides accurate range and radial-velocity information and remains operational under low lighting and water influenced weather conditions that challenge optical sensors [1, 2]. Angular resolution, however, remains limited by the physical aperture of a radar module. This limitation becomes most consequential when two reflectors overlap in the same range–Doppler region: targets that are already separated in range or velocity can be processed independently, whereas overlapping responses must be distinguished from their spatial signatures alone. Such close-pair scenes form the central angular-resolution test in this thesis.

For a uniformly weighted linear aperture, a representative angular-resolution relation is

$$\Delta\theta \approx \frac{0.886\lambda}{D}, \tag{1.1}$$

where λ is the wavelength and D is the aperture length. At approximately 77 GHz, λ is about 3.9 mm, while the available aperture remains constrained by the radar hardware and its installation. For the radar modules used in this work, approximately 15° is treated as a practical, configuration-specific resolution benchmark rather than a universal radar limit. The physical experiments vary the lateral distance between two targets that remain at approximately the same down-range. That commanded distance is used to reproduce the scene, but the relevant resolution variable is the corresponding expected angular separation $\Delta\theta$. Pairs below the 15° benchmark test whether a distributed aperture provides useful spatial information beyond that of one module.

Spatially separated radar nodes offer a larger effective aperture through multiple monostatic and reciprocal bistatic observations [3]. In principle, their measurements can be combined to produce a narrower angular response without enlarging a single radar module. That potential gain is conditional: independent nodes introduce timing and frequency mismatch, channel-dependent manifold errors, and uncertain phase rotations between aperture blocks. Receive-side correction can align target responses in range and Doppler, but such alignment does not demonstrate that the apertures share the carrier-phase origin required for coherent fusion.

Two spatial residual-error classes are consequently distinguished throughout the thesis. Intra-aperture errors distort the element-to-element phase progression and therefore corrupt the local steering manifold. Inter-aperture errors rotate otherwise consistent aperture blocks relative to one another. A method that tolerates an unknown common block phase is not necessarily robust to distorted local manifolds. Conversely, fully coherent processing can exploit distributed-baseline phase and achieve a narrow nominal response only when channel calibration, geometry, and inter-aperture phase are mutually consistent. Local and shared-support methods avoid part of this requirement, but they do not necessarily recover the complete coherent-aperture gain.

1.1 Research Problem and Gap

Distributed FMCW radar research has established the individual components required for coherent localization. Receive-side synchronization methods estimate timing and frequency mismatch, phase-calibration methods exploit hardware links, array redundancy, reference paths, or stationary scene scatterers, and distributed direction-of-arrival (DoA) methods combine spatial information from separated apertures [4, 5, 6, 7]. Related work also demonstrates coherent, shared-support, and non-coherent processing strategies under particular signal and calibration models [8, 9]. These studies establish that residual phase uncertainty limits coherent processing, but they commonly focus on the synchronization accuracy itself, a prescribed phase-error model, or the performance of one estimator architecture.

A remaining practical gap is the interface between these research directions. The residual phase quality delivered by a synchronization and calibration chain is not commonly translated into estimator-specific operating limits for distributed close-pair localization. It is therefore unclear when the achieved phase quality is sufficient for fully coherent processing, when local or shared-support fusion remains reliable, and when processing that avoids a common carrier-phase assumption is the safer alternative. Range–Doppler alignment, inter-aperture phase consistency, model-order recovery, relative pair separation, and correct absolute localization are related but distinct properties; success in one does not establish the others.

Residual phase error is also not one interchangeable scalar quantity. Intra-aperture error distorts the local array manifold, an inter-aperture block rotation invalidates the phase relation assumed when separated apertures are stacked coherently, and time-varying phase across the chirps of a coherent processing interval can reduce coherent integration and modify the range–Doppler response [10, 11]. These mechanisms must therefore be distinguished rather than represented only by one overall phase-error variance.

This thesis addresses the gap for a specified two-node 77 GHz FMCW configuration. Existing synchronization, phase-calibration, and DoA methods are combined in one processing and evaluation framework. Controlled studies isolate residual phase mechanisms and compare local, shared-support, fully coherent, and non-coherent estimator architec-

tures; physical measurements then relate the resulting tolerances to the phase quality and localization behaviour achieved in practice. Model-order recovery, relative target separation, and absolute angular or Cartesian localization are evaluated separately to identify which capability is lost as coherence deteriorates. Static spatial-error experiments address single targets and close pairs; the within-CPI extension uses a single-target scene.

The principal contribution is therefore not a new synchronization or DoA algorithm. It is an estimator-specific residual-phase error budget that connects the phase quality achieved after practical synchronization and calibration to the localization capabilities that can still be supported. The numerical thresholds remain dependent on the evaluated hardware, baseline, scene, and processing settings; the broader contribution is the evaluation methodology and the resulting evidence about which estimator assumptions remain credible under different residual-error mechanisms.

The research problem is summarized as follows:

> Under which residual intra-aperture, inter-aperture, and within-CPI phase conditions do existing synchronization and calibration procedures provide sufficient phase quality for reliable distributed close-pair localization, and how do these operating limits depend on the coherence assumptions of the selected estimator?

1.2 Research Objective and Questions

The objective is to implement and evaluate a common processing chain that accepts measured or synthetic dechirped intermediate-frequency (IF) data, corrects dominant receive-side synchronization errors, forms associated multi-aperture target snapshots, and applies localization methods with different coherence assumptions. The investigation combines a configurable signal-level simulator, controlled residual-phase sweeps, exact-path geometric modeling, multiple DoA estimator families, Cartesian backpropagation, and physical close-pair measurements. Figure 1.1 summarizes this research framework.

The work is guided by four research questions:

1. **Synchronization and residual errors:** To what extent can receive-side correction align monostatic and bistatic measurements from independent FMCW nodes, and what intra-aperture and inter-aperture phase residuals remain?

2. **Estimator sensitivity:** How do Bartlett, independent sparse, shared-support, and coherent localization methods respond to controlled residual phase perturbations, particularly for close pairs above and below the approximately 15° radar-specific benchmark?

3. **Conditions for coherent localization:** Under which tested conditions do reference-based phase calibration and geometric accuracy support coherent fusion, and how does coherent localization compare with methods that do not require a common carrier phase?

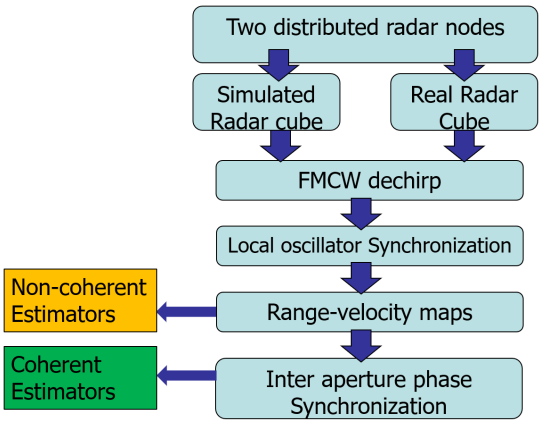


Figure 1.1: Research framework linking signal generation or acquisition, synchronization, target-snapshot extraction, localization, and phase-sensitivity evaluation.

4. **Simulated versus physical performance:** To what extent do the controlled and measurement-matched simulations predict physical target-count, relative-separation, and absolute angular or Cartesian-localization outcomes, and which modeled or unmodeled limitations explain the discrepancies?

These questions relate phase-error tolerance to the estimator, baseline, scene, and scoring criterion. The physical study distinguishes relative-separation recovery from absolute-angle accuracy to assess the practical value of coherent processing.

1.3 Contributions

The main contribution of this thesis is an estimator-specific assessment of the residual phase quality required for distributed close-pair localization. Rather than proposing a new synchronization or direction-of-arrival (DoA) algorithm, the work connects the phase quality remaining after practical synchronization and calibration to the localization capabilities that can still be supported.

More specifically, the thesis makes the following contributions:

1. **Estimator-specific residual-phase limits.** A controlled sensitivity study determines how residual intra-aperture manifold errors, inter-aperture block-phase errors, and phase evolution across the chirps of a coherent processing interval affect estimator architectures with different coherence assumptions. The static comparison includes local processing, shared-support reconstruction, fully coherent fusion, and non-coherent Cartesian localization. The within-CPI extension separately sweeps residual frequency offset and random phase walk after synchronization and before the Doppler FFT. It compares the five DoA methods under four paired RD-readout and reference-calibration protocols while waveform and CPI timing remain fixed.

2. **Separation of error mechanisms and localization claims.** The evaluation distinguishes phase distortion within an aperture, phase rotation between apertures,

and phase evolution over a CPI. It also separates correct target-count estimation, accurate relative target separation, and correct absolute angular or Cartesian localization. This prevents apparent target resolution from being interpreted automatically as accurate localization.

3. **Connection between controlled phase sensitivity and physical measurements.** Intra- and inter-aperture phase quality measured in the physical radar scenes is compared with the corresponding controlled spatial-error studies. This links the theoretical benefit of distributed coherent processing to the synchronization, calibration, geometry, and phase-reference accuracy achieved in the experimental setup.

4. **Comparison of phase-reference strategies.** The physical evaluation distinguishes scene-assisted phase correction, in which the phase reference is estimated from the target region itself, from compensation transferred from an independent spherical reference reflector. This exposes the difference between improving coherence for a known scene and demonstrating calibration that transfers to independently located targets.

5. **An integrated evaluation framework.** A common processing chain is implemented for two monostatic and two reciprocal bistatic apertures. It combines receive-side synchronization, range–Doppler processing, exact-path geometric association, multiple DoA estimators, and coherent and non-coherent Cartesian backpropagation. A configurable dechirped-IF simulator and measurement-matched room model support controlled comparison with physical data.

The reported numerical thresholds are specific to the evaluated hardware, baselines, scenes, and processing settings. The broader added value lies in the methodology for determining whether the residual phase quality achieved by a distributed radar system is compatible with the assumptions of a selected localization method.

1.4 Thesis Outline

The thesis is organized into seven chapters. Chapter 2 provides a critical review of distributed-radar synchronization, phase calibration, DoA estimation, and coherent and non-coherent fusion. It compares the assumptions and limitations of existing approaches and establishes the research gap addressed by this work.

Chapter 3, *Theoretical Background and Signal Model*, introduces the two-node radar architecture, monostatic and bistatic propagation geometry, the FMCW signal and array models, and the mathematical distinction between intra-aperture errors, inter-aperture errors, and phase evolution over a CPI.

Chapter 4, *Processing and Estimation Methodology*, describes the receive-side synchronization and phase-stabilization procedures, range–Doppler processing, snapshot extraction, geometric association, and steering-dictionary construction. It then presents the

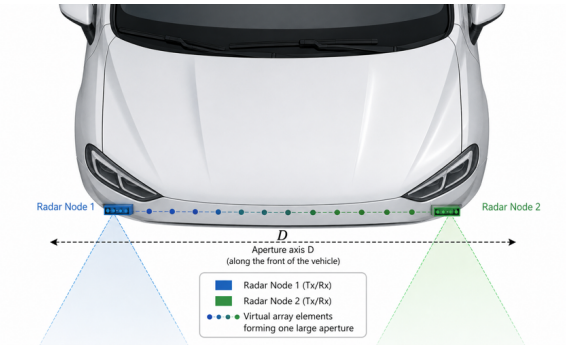


Figure 1.2: Illustration showing the idea of utilizing two distributed radar nodes to increase the aperture size. In this example it is mounted on a car, a common application.

implemented DoA estimators, truth-blind model-order selection, and coherent and non-coherent Cartesian backpropagation.

Chapter 5, *Simulation and Experimental Evaluation Methodology*, details the dechirped-IF simulator and its modeled nonidealities, the physical measurement setup, and the evaluation geometries. It also defines the performance metrics and protocols used for ideal controls, baseline-dependent ambiguity analysis, residual-phase sensitivity sweeps, physical measurements, and simulation-to-measurement comparison.

Chapter 6, *Results and Discussion*, presents the controlled simulation and physical measurement outcomes. It examines ideal close-pair resolution, coherent and non-coherent baseline ambiguity, sensitivity to residual phase errors, measured target-separation and phase-quality results, Cartesian backpropagation, and the extent to which measurement-matched simulations reproduce the physical observations. The discussion relates these findings to the assumptions and limitations of the methods reviewed in Chapter 2.

Chapter 7, *Conclusions and Future Work*, summarizes the findings, answers the four research questions from Section 1.2, identifies the main contributions and limitations, and outlines improvements in calibration, geometric modeling, simulation fidelity, and phase-quality-informed estimator selection. Supplementary implementation settings and reproducibility details are provided in the appendix.

Chapter 2

Literature Review

The angular resolution of an automotive FMCW radar is constrained by its available antenna aperture. This limitation is particularly relevant when nearby scatterers occupy the same range–Doppler region and cannot first be separated by range or velocity processing [1]. Spatially separated radar nodes offer an alternative to enlarging one module: their monostatic and bistatic observations can provide a larger effective aperture and multiple perspectives of the scene [12]. The resulting improvement, however, is conditional on the synchronization accuracy, array calibration, and propagation model required by the chosen localization method.

This chapter examines the resulting trade-off rather than ranking methods by ideal angular resolution alone. Coherent processing can exploit complex phase relations across the distributed aperture but requires an accurate inter-node phase and geometric reference. Non-coherent and shared-support methods relax that requirement, while retaining their own assumptions about local array calibration, target association, and scene sparsity [6, 9]. Moreover, a narrow angular spectrum or a correct estimate of target count does not in itself establish accurate absolute target localization. The review first contrasts synchronization strategies, then relates residual phase errors to the assumptions of localization methods, and finally identifies the specific evaluation gap addressed by this thesis.

2.1 Synchronization in Distributed FMCW Radar

Distributed-radar synchronization comprises several distinct requirements. Temporal alignment concerns the relative timing of samples, chirps, and frames; frequency alignment addresses oscillator differences and their drift; and carrier-phase coherence concerns the complex phase relationships needed to combine distributed channels coherently. These requirements are related but not interchangeable [11]. In particular, aligning corresponding detections in range and Doppler does not demonstrate that the measured phases agree with a common distributed-array steering model.

At approximately 77 GHz, the carrier wavelength is about 3.9 mm. Consequently, small errors in effective antenna position, reference-target position, or propagation-path

length can produce substantial phase errors. Calibration and near-field modeling therefore matter more for coherent aperture combination than for fusion of independently estimated powers or target hypotheses [13, 14]. A useful assessment must also distinguish channel-relative phase errors within an aperture from an unknown common phase rotation between apertures.

Synchronization strategies can be grouped according to their reference information. Shared clock or radio-frequency (RF) references directly constrain inter-node errors but increase hardware coupling [11]. Wireless reference paths and signal-domain approaches retain more deployment flexibility but require identifiable reference observations or sufficient information in the received measurements [15]. Global navigation satellite system (GNSS)-disciplined oscillators provide another route to time- and frequency-reference transfer, although the degree of carrier-phase coherence achieved at the radar carrier must be evaluated separately rather than inferred from timing synchronization alone [16]. The following subsections distinguish coarse receive-side alignment from the additional constraints needed for coherent phase calibration.

2.1.1 Coarse Synchronization

Vecchi, Güneri, and Yarovoy develop a model-based post-processing synchronization algorithm for cooperative FMCW radars operating with independent oscillators [5]. The approach relates timing and frequency mismatch to distortions observed in the bistatic beat signals and uses reference information in the received data to estimate compensating corrections. Its principal advantage over a shared-reference architecture is that recorded radar data can be corrected without physically distributing one local oscillator.

The distinction between observable alignment and carrier-phase coherence remains essential. A correction that reduces range displacement, migration, or slow-time drift does not by itself verify a stable inter-aperture carrier phase for every scene point. The usefulness of a target-derived reference also depends on its detectability and on the accuracy of the assumed signal and propagation model. Thus, the correction adopted in this thesis is treated as a receive-side synchronization stage, not as proof of complete coherent calibration. Its implementation and frequency-correction expression are given in Section 4.2.1, including (4.2).

2.1.2 Treatment of Residual Phase Error in Prior Work

The literature does not use one universal definition of phase error. Instead, the adopted model reflects the question being studied. Synchronization research commonly characterizes the residual phase produced by a particular reference-transfer algorithm. Yang and Blum, for example, derive statistics of the synchronization error for master-slave, round-trip, and consensus schemes [4]. Their analysis clearly states that it assumes that the synchronization interval is sufficiently short for oscillator phase noise to be negligible. Practical FMCW-network studies similarly report the accuracy or coherence achieved by a

correction or calibration procedure [5, 6, 15], but such a value does not by itself determine whether a specific localization estimator remains valid.

Theoretical localization studies often replace the synchronization procedure by a statistical mismatch model. Godrich and Haimovich treat the radar phase offsets as random unknown nuisance parameters and derive a hybrid Cramér–Rao bound (HCRB), showing a localization-error floor governed by the phase-error variance [17]. Rojhani, Greco, and Gini similarly assign independent Gaussian errors to distributed transmitters and receivers and evaluate location and velocity bounds as a function of their variance in non-Gaussian clutter [18]. These results establish fundamental sensitivity, but a scalar variance and a performance bound do not distinguish how the error enters a practical array snapshot or how different estimator architectures fail.

Other studies evaluate the processing architecture rather than only a bound. Li *et al.* model timing, frequency, and phase synchronization errors and compare coherent, hybrid, and non-coherent distributed-MIMO detectors [19]. Their finding that phase error directly affects phase-combining detectors supports an assumption-based comparison, although the reported metrics are detection probabilities rather than close-pair angular localization. Estimator-side approaches instead incorporate mismatch into a particular reconstruction. Sun, Chen, and Zou alternate between sparse multitarget localization and estimation of a static phase mismatch for each transmitter–receiver pair [20]. Chen *et al.* formulate an atomic-norm DoA estimator with sensor-dependent gain and phase errors [21]. The former addresses distributed block mismatch and the latter local manifold error, but both primarily demonstrate the proposed robust estimator under simulated mismatch.

Waveform-level FMCW studies introduce a further distinction: phase error can evolve during the measurement rather than remain one constant offset. Siddiq *et al.* model oscillator phase noise through its spectrum and analyze delay-dependent range correlation and spectral spreading [10]. Tschapek *et al.* show that systematic phase distortion should be separated from colored stochastic phase noise, because simplified white-noise models can produce optimistic performance predictions [22]. Relative frequency error also produces a deterministic phase progression across chirps. Its accumulated effect depends on chirp repetition and CPI duration, while the conversion of timing and frequency imperfections into the dechirped signal also depends on waveform parameters such as bandwidth and chirp duration.

These models motivate the decomposition used in this thesis. A channel-dependent intra-aperture perturbation isolates corruption of the local steering manifold. A common inter-aperture rotation isolates loss of the phase reference between otherwise internally consistent aperture blocks. A within-CPI perturbation isolates chirp-to-chirp phase evolution that cannot be represented by one constant block rotation. The controlled studies do not claim that any one of these abstractions is a complete oscillator model; their purpose is to identify which estimator assumption is violated. The tested spatial phase levels are subsequently related to intra- and inter-aperture residuals measured after the physical synchronization and calibration chain; the within-CPI histories remain imposed controlled

perturbations.

The distinction leads to different estimator vulnerabilities. Local processing is insensitive to a common rotation of its entire complex snapshot but remains sensitive to channel-relative distortion. Shared-support reconstruction can combine aperture-specific evidence without requiring a common complex amplitude, provided that the local manifolds and geometric support mapping remain consistent. Fully coherent reconstruction additionally uses the relative carrier phase and therefore requires accurate geometry, channel calibration, and inter-aperture phase [9, 6]. Actual measurements can also contain phase-centre error, multipath, and target-dependent scattering [14, 15]; physical validation is consequently needed before attributing a failure to a single idealized component.

2.1.3 Fine Synchronization

Fine synchronization seeks phase information beyond coarse timing and frequency alignment. The main alternatives differ in how they make relative aperture phase observable: through redundant or reciprocal array measurements, a direct inter-node reference path, or reference scatterers in the scene. Each mechanism imposes different requirements on the hardware geometry, reference visibility, and transfer of calibration from the reference to the targets of interest [6, 15, 7].

Calibration from Array Structure

Redundant virtual-array elements can provide constraints between measurements that represent equivalent spatial responses under the assumed array and propagation model. Janoudi *et al.* investigate coherent processing of uncoupled radar networks and calibration from suitable virtual-array structure [6, 23]. Such approaches reduce dependence on an external reference reflector, but the usable constraints are set by the array geometry and by whether the compared channels observe sufficiently compatible signals. For widely separated bistatic paths, equality of nominal virtual positions alone must not be interpreted as a guarantee of identical exact propagation phase; the appropriate signal model and calibration conditions still matter.

Calibration from a Reference Path

An identifiable direct radar-to-radar signal can provide synchronization and geometry information without a wired RF reference. Fenske *et al.* investigate joint constellation estimation, coherent processing, and multiperspective imaging for uncoupled bistatic radar, demonstrating the importance of accurately modeling the available reference-path information [15]. Reference-target approaches offer a related but distinct mechanism: an active reference target can provide synchronization information [24], while Griffiths *et al.* use multiple reference signals from an urban scene, including direct-path and stationary-scatterer responses, to estimate and correct relative phase variations [7].

Table 2.1: Comparison of synchronization and phase-calibration approaches. The limitations summarize the cited signal-model and experimental assumptions rather than universal performance bounds.

Approach and source	**Required information**	**Main capability**	**Important limitation**
Receive-side correction [5]	Observable reference response	Reduces timing/frequency-induced distortions	Range–Doppler alignment is not proof of carrier coherence
Virtual-array calibration [6, 23]	Suitable inter-aperture structural constraints	Estimates relative aperture calibration terms	Depends on geometry and signal-model consistency
Direct-path synchronization [15]	Identifiable inter-node path	Supports phase-aware synchronization and geometry estimation	Depends on direct-path observability and modeling
Active reference target [24]	Known active reference response	Supplies reference information for coherent processing	Requires dedicated reference hardware and calibration
Scene-based multi-reference correction [7]	Multiple identifiable stationary/direct-path returns	Estimates phase variations from environmental references	Depends on reference stability and scene observability

These methods show why the choice of reference is part of the experimental evidence. A scene-derived phase adjustment can improve coherence for the scene from which it was estimated, but this does not alone establish transfer to independent targets or measurement geometries. An independently surveyed reference offers a more stringent transfer test, although reflector phase centre, multipath, and survey accuracy remain limiting factors. Table 2.1 compares the approaches by the information they require and the limitations that follow from it.

2.2 Direction-of-Arrival and Localization Methods

DoA methods differ in the spatial information they exploit and in the statistical or structural assumptions needed for high resolution. The distinction is especially relevant when multiple target responses are present in a single range–Doppler region. For distributed bistatic apertures, one physical target also need not have the same *local* arrival angle at each radar node. Any joint-support model must therefore express correspondence through compatible geometry or a suitable common target parameterization, rather than simply equating untransformed local DoA grids. The models used in this thesis are defined in Chapter 3, and the implemented estimators are specified in Chapter 4.

2.2.1 Conventional Spectral Estimators

Bartlett beamforming provides an interpretable baseline by comparing measured spatial responses with steering vectors over candidate angles. With a fixed local aperture, its ability to separate closely spaced targets is restricted by the width and sidelobes of the array response [1]. Adaptive methods such as the minimum-variance distortionless re-

sponse (MVDR) beamformer can improve selectivity, but their performance depends on the estimated covariance and the accuracy of the assumed steering vector [25]. These dependencies are important for the limited-snapshot and potentially coherent target responses considered here.

2.2.2 Subspace-Based Estimators

Multiple signal classification (MUSIC) uses a signal/noise-subspace decomposition, whereas estimation of signal parameters via rotational invariance techniques (ESPRIT) exploits rotational or shift invariance in compatible array structures [26, 27]. Their potential for resolving nearby sources depends on conditions such as covariance quality, source-number identifiability, source coherence, and array structure. Such conditions are not automatically satisfied when one or a few complex snapshots are extracted from a radar range–Doppler region. These methods are consequently relevant state-of-the-art comparators, although they are not the principal implementations evaluated in this thesis.

2.2.3 Maximum-Likelihood Estimation

Maximum-likelihood estimation fits the observed data under an explicit statistical measurement model and may provide high accuracy when that model and its parameterization are correct [28]. Joint multi-target fitting, however, generally requires a nonlinear search over target parameters and can grow substantially in complexity when distributed geometry, nuisance phases, and unknown model order are included. The method is therefore discussed as a reference family rather than as an experimental estimator implemented in the present comparison.

2.2.4 Sparse Reconstruction and Shared Support

Sparse reconstruction assumes that comparatively few directions or candidate target locations explain the measured array response. This can provide concentrated solutions from limited snapshots, but the outcome still depends on steering-dictionary accuracy, noise, regularization, and the coherence between dictionary columns [29]. In particular, a fine angular grid does not remove off-grid bias or model mismatch automatically; off-grid Bayesian methods address this issue by explicitly modeling departures from a discretized grid [30].

FOCUSS promotes a sparse solution through iterative reweighting of an underdetermined inverse problem [31]. When applied to separate aperture snapshots, it can use their local phase structure without requiring calibrated inter-aperture carrier phase. This independence of block phase does not eliminate sensitivity to channel-level calibration errors, correlated dictionary atoms, initialization, or regularization. In this thesis, independent FOCUSS serves as a sparse local-processing comparator.

Block-FOCUSS combines aperture-specific coefficient vectors through a common-support constraint rather than a common complex amplitude. It can aggregate evidence

across apertures while allowing unknown block rotations, provided the tasks have been mapped to geometrically compatible support. The method sacrifices the direct use of distributed carrier-phase interference and may fail when visibility, range selection, multipath, or mismatched steering dictionaries make the supports inconsistent. Mo uses a two-stage distributed-radar framework with coarse association followed by coherent and non-coherent FOCUSS-based estimation, providing a closely related methodological starting point [9].

Coherent FOCUSS instead stacks phase-calibrated data against a unified distributed steering model. The extended baseline can then improve discrimination, but model mismatch, inter-aperture phase uncertainty, and spatial ambiguities may also create narrow yet incorrectly located peaks. Mo's results demonstrate the potential of this architecture under the models and conditions studied there; they do not establish that equivalent gains necessarily transfer to a differently calibrated physical setup [9]. The distinction between ideal resolution and defensible absolute localization is therefore central to the present evaluation.

Bayesian compressive sensing (BCS) estimates sparse angular contributions using probabilistic relevance parameters rather than only a fixed sparsity penalty. Carlin *et al.* investigate BCS strategies for DoA estimation, including the use of complex sensor measurements [32]. Multi-task formulations can share relevance across geometrically aligned aperture tasks while retaining task-specific complex amplitudes. This offers an alternative route to joint model-order evidence without coherent stacking, but still relies on a suitable noise model, grid, convergence procedure, and support-selection rule. The multi-task implementation and these choices are examined explicitly in the methodology of this thesis.

Table 2.2 summarizes the relevant assumptions, the entries follow the method references discussed above, particularly; [25, 26, 27, 31, 9, 32].

2.3 Coherent and Non-Coherent Fusion

The reviewed localization approaches are better viewed as a hierarchy of cross-aperture assumptions than as a binary choice. Power- or hypothesis-level fusion combines independently processed evidence without requiring a common carrier-phase origin [12]. Shared-support methods impose a stronger requirement: multiple apertures must support the same physical target hypotheses after geometric association, but their complex coefficients can remain aperture-specific [9]. Fully coherent approaches additionally require that the complex samples conform to a single phase-consistent distributed-array model [6].

The distinction matters because increasing the baseline creates both an opportunity and a vulnerability. Greater spatial separation can produce more discriminating phase variation, but coherent processing also becomes sensitive to phase-reference error, near-field geometry, and sparse-aperture ambiguities [14, 9]. Consequently, a valid comparison should separate correct model order, relative target spacing, and accurate absolute angle

Table 2.2: Assumption-based comparison of DoA and localization method families. Performance depends on aperture geometry, snapshots, signal coherence, and model accuracy.

Method family	**Information used**	**Main advantage**	**Critical condition or failure mode**
Bartlett	Local steering response	Simple reference spectrum	Local-aperture width and sidelobes
MVDR	Covariance and look-direction constraint	Adaptive selectivity	Covariance and manifold accuracy
MUSIC/ESPRIT	Subspaces or shift invariance	High resolution under suitable conditions	Snapshots, source coherence, and array structure
Maximum likelihood	Statistical data fit	Accurate inference under a correct model	Search complexity and nuisance parameters
Independent FOCUSS	Sparse aperture-local response	No cross-aperture phase required	Dictionary, regularization, and local phase errors
Block-FOCUSS	Geometrically shared sparse support	Combines aperture evidence without coherent stacking	Support correspondence and model mismatch
Coherent FOCUSS	Joint complex distributed response	Exploits distributed-baseline phase	Calibration, geometry, and ambiguous lobes
Multi-task BCS	Shared relevance, task-specific amplitudes	Joint support and model-order evidence	Grid, noise assumptions, and pruning

or position. These measures answer different questions and cannot be replaced by a single spectrum-width ranking.

2.4 Synthesis and Research Gap

The reviewed literature establishes the main components needed for distributed coherent localization, but the studies generally begin and end at different interfaces. Synchronization and calibration research estimates timing, frequency, geometry, or phase corrections and reports the remaining error or coherent gain [4, 5, 6, 15, 7]. Theoretical sensitivity studies begin with a prescribed phase-error distribution and report a bound on localization or velocity accuracy [17, 18]. Detection studies compare coherence assumptions under synchronization error [19], while phase-robust localization studies modify one estimator to absorb a particular mismatch model [20, 21]. FMCW phase-noise studies provide more realistic temporal and spectral models, but normally evaluate waveform-level range or frequency-estimation effects rather than distributed angular resolution [10, 22].

The unresolved issue is therefore not whether phase error matters, but how the residual phase quality delivered by a practical synchronization and calibration chain should be translated into localization claims for estimators with different coherence assumptions. A single RMS phase value does not reveal whether the local manifold, the relative phase between apertures, or the phase evolution over the CPI is responsible. Likewise, a correct target count or pair separation does not establish correct absolute localization. Within the reviewed literature, no study was identified that combines measured post-synchronization spatial phase quality, separately controlled intra-aperture, inter-aperture, and within-CPI phase errors, multiple estimator coherence assumptions, and separate close-pair scoring of

Table 2.3: Comparison of the residual-phase effects studied in prior work and the remaining estimator-evaluation question addressed by this thesis.

Literature line	**Phase-error treatment**	**Reported outcome**	**Open interface**
Synchronization and calibration [4, 5, 15]	Error estimated from a reference exchange, path, or scene	Residual error, alignment, or coherent gain	Translation to estimator-specific localization limits
Theoretical sensitivity [17, 18]	Random phase offsets with prescribed variance	HCRB for location or velocity	Achieved physical residuals and estimator failure modes
Architecture comparison [19]	Timing, frequency, and phase sync errors	Coherent, hybrid, and non-coherent detection	Close-pair DoA and absolute localization
Robust localization [20, 21]	Static block or sensor gain–phase mismatch	Performance of a proposed robust estimator	Common comparison across estimator assumptions
FMCW phase-noise modeling [10, 22]	Systematic distortion and colored stochastic phase evolution	Range correlation, spectral, and precision effects	Distributed-aperture estimator selection
This thesis	Controlled spatial/temporal perturbations; measured spatial residuals	Model order and angle; relative separation for close pairs	Connects achieved phase quality to supported localization claims

model order, relative separation, and absolute position in one distributed FMCW evaluation work.

This gap motivates the selected approach. Existing receive-side synchronization and phase-reference procedures are retained rather than presented as new algorithms. Their output is evaluated through controlled error sweeps, a within-CPI phase-evolution study, and physical residual-phase diagnostics. Local, shared-support, fully coherent, and non-coherent estimators are then compared through one processing and scoring framework. The resulting contribution is an estimator-specific residual-phase error sensitivity: a structured connection between the phase quality achieved after synchronization and calibration and the localization performance that remains. The static studies include single targets and close pairs; the within-CPI comparison uses the five DoA methods in one single-target geometry with fixed waveform and CPI timing.

The scope remains deliberately configuration-specific. The numerical thresholds cannot be transferred unchanged to every radar network, and the thesis does not claim invention of the underlying synchronization or DoA algorithms. Its added value is the mechanism-resolved methodology and the experimental evidence used to determine when a nominal distributed-aperture resolution gain is supported by the available phase quality. The following chapters formalize the corresponding signal and error models and evaluate this question in controlled and physical scenes.

Chapter 3

Theoretical Background and Signal Model

This chapter establishes the theoretical foundations used by the processing and evaluation methods developed in the following chapters. Section 3.1 defines the two-node radar architecture, the monostatic and bistatic apertures, and the geometric assumptions relevant to distributed-array processing. Section 3.2 introduces the FMCW waveform, propagation and array response models, and the measurement representation used for localization. Finally, Section 3.3 distinguishes intra-aperture manifold distortion, inter-aperture block-phase offsets, and phase evolution over a CPI. This mechanism-resolved model provides the basis for relating residual phase quality to the assumptions of the different estimator architectures.

3.1 System Model

The radar setup consists of two spatially separated sensing nodes, denoted as M_1 and M_2, forming a distributed multiple-input multiple-output (MIMO) radar system. Each node is equipped with $M_{\mathrm{Tx}} = 2$ transmit (TX) antennas and $M_{\mathrm{Rx}} = 4$ receive (RX) antennas. Together, the two nodes form four effective apertures: two monostatic apertures, $M_1 \to M_1$ and $M_2 \to M_2$, and two bistatic apertures, $M_1 \to M_2$ and $M_2 \to M_1$.

Let $\mathbf{p}_t$ denote the position of a target. The distances from the target to radar nodes M_1 and M_2 are denoted by R_1 and R_2, respectively. The corresponding local angles with respect to the two radar nodes M_1 and M_2 are denoted by θ_1 and θ_2. These angles represent the relevant direction-of-departure and direction-of-arrival variables for the monostatic and bistatic channels. Each radar node operates using an FMCW waveform. Radar 1 operates at 77 GHz and Radar 2 at 77.005 GHz. This deliberate 5 MHz offset makes the two transmitted bands easy to distinguish at the receiver. The transmitted signal consists of linear chirps with bandwidth B and duration T_{chirp}, resulting in a chirp slope

$$\mu = \frac{B}{T_{\mathrm{chirp}}}. \tag{3.1}$$

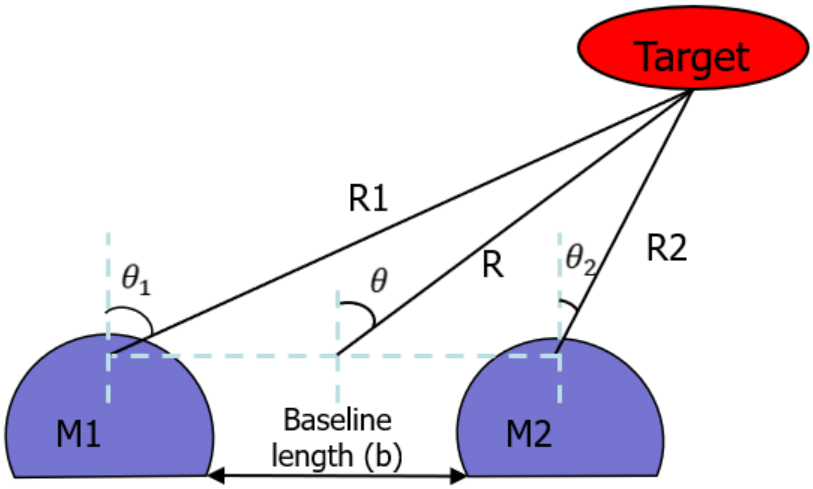


Figure 3.1: System configuration. The magnified part shows the individual antenna elements within one radar node.

The received signal is sampled in both fast time and slow time, with N_r samples per chirp and N_c chirps per coherent processing interval (CPI). Together with the spatial channels, these samples form a three-dimensional data cube used for range, Doppler, and angle estimation.

In a MIMO radar configuration, a virtual array can be constructed by combining all transmitter–receiver pairs. The effective position of each virtual element is expressed as

$$d_v(i, q) = d_{\text{tx}}(i) + d_{\text{rx}}(q), \tag{3.2}$$

where $d_{\text{tx}}(i)$ and $d_{\text{rx}}(q)$ represent the scalar array-axis locations of the i-th transmit and q-th receive antennas, respectively. This virtual array increases the effective aperture size and therefore improves the achievable angular resolution without requiring additional physical antennas. This virtual-position interpretation applies directly to colocated MIMO radar under the far-field approximation. For distributed bistatic apertures, the transmit and receive directions may differ, and the response is therefore modeled using separate transmit and receive steering vectors.

Accurate modeling of the received phase depends on whether the target lies in the near-field or far-field region of the array. The boundary between these regimes is commonly defined by the Fraunhofer distance

$$R_F = \frac{2D^2}{\lambda}, \tag{3.3}$$

where D denotes the aperture size and λ the wavelength. If the target distance satisfies $R \geq R_F$, the incoming wavefront can be approximated as planar. Under this assumption, the phase at the u-th array element can be written as

$$\phi_u \approx -\beta d_u \sin\theta, \tag{3.4}$$

with $\beta = \frac{2\pi}{\lambda}$ and d_u the element position.

For targets located within the near-field region, $R < R_F$, the curvature of the wavefront must be taken into account. The propagation distance from the target to the u-th element is then given by

$$r_u = \sqrt{R^2 + d_u^2 - 2Rd_u \sin\theta}, \tag{3.5}$$

and the corresponding phase becomes

$$\phi_u = -\beta r_u. \tag{3.6}$$

In distributed radar systems, this distinction is especially important because the effective aperture can be much larger than the aperture of an individual radar node. Small inaccuracies in distance modeling can therefore lead to significant phase errors, particularly at millimeter-wave frequencies such as 77 GHz. This becomes critical when coherent processing across multiple nodes is required.

To define a relative phase reference within each aperture, distances are referenced to the path through its first transmit and receive elements:

$$r_{\text{ref}} = \|\mathbf{t}_{p,1} - \mathbf{p}_t\| + \|\mathbf{r}_{p,1} - \mathbf{p}_t\|, \tag{3.7}$$

where $\mathbf{t}_{p,1}$ and $\mathbf{r}_{p,1}$ denote the position vectors of the first transmit and receive elements of aperture p. This operation removes one common phase within aperture p while preserving its spatial progression. It does not establish inter-aperture carrier coherence, which requires absolute-path steering and calibrated aperture-block phases.

3.2 Signal Model

The transmitted FMCW waveform is modeled as a linear frequency-modulated chirp:

$$s_{\text{tx}}(t) = \exp\left(j2\pi\left[f_c t + \frac{\mu}{2}t^2\right]\right), \quad 0 \le t < T_{\text{chirp}}. \tag{3.8}$$

This waveform enables simultaneous estimation of target range and velocity through frequency analysis of the received signal.

Assume K targets located at positions $\mathbf{p}_k(n)$ during the n-th chirp. For transmitter element i and receiver element q of aperture p, the propagation distance associated with target k is

$$L_{p,iq,k}(n) = \|\mathbf{p}_k(n) - \mathbf{t}_{p,i}\| + \|\mathbf{p}_k(n) - \mathbf{r}_{p,q}\|. \tag{3.9}$$

Here, $\mathbf{t}_{p,i}$ and $\mathbf{r}_{p,q}$ denote the position vectors of the transmit and receive antennas for aperture p, respectively. For monostatic apertures, the transmitter and receiver belong to the same radar node. For bistatic apertures, the transmitter and receiver are located on

different radar nodes. The corresponding propagation delay is

$$\tau_{p,iq,k}(n) = \frac{L_{p,iq,k}(n)}{c}. \tag{3.10}$$

The received signal is modeled as a superposition of delayed and scaled replicas of the transmitted waveform:

$$y_{p,iq}(t,n) = \sum_{k=1}^{K} \alpha_{p,iq,k} \exp\left(j2\pi \left[f_c \left(t - \tau_{p,iq,k}(n)\right) + \frac{\mu}{2}\left(t - \tau_{p,iq,k}(n)\right)^2\right]\right) + w_{p,iq}(t,n), \tag{3.11}$$

where $\alpha_{p,iq,k}$ represents the complex reflection coefficient of target k, and $w_{p,iq}(t,n)$ denotes additive noise. The coefficient $\alpha_{p,iq,k}$ includes the target response, propagation loss, and aperture-dependent amplitude and phase effects.

The baseband signal is obtained by mixing the received signal with the conjugate of the transmitted waveform:

$$z_{p,iq}(t,n) = y_{p,iq}(t,n) s^*_{\mathrm{tx}}(t). \tag{3.12}$$

This dechirping operation yields the beat signal used for subsequent range and Doppler processing. The beat frequency is directly related to the propagation delay and therefore to the target range, while the phase progression across chirps is used for Doppler estimation.

For angular processing, the spatial response of the array must be modeled. For a uniform linear array under the plane-wave assumption, the steering response for an incoming wave from angle θ is given element-wise by

$$a(\theta) = \exp(j\beta d \sin\theta), \tag{3.13}$$

where d denotes the antenna element position. In the MIMO case, the overall steering vector is constructed as the Kronecker product of the transmit and receive steering vectors:

$$a^{(p)}(\theta) = a^{(p)}_{\mathrm{Rx}}(\theta) \otimes a^{(p)}_{\mathrm{Tx}}(\theta). \tag{3.14}$$

This formulation applies directly to monostatic apertures, where the direction of departure and direction of arrival are identical under the far-field approximation.

For bistatic configurations, the transmit and receive angles generally differ. In that case, the steering vector becomes

$$a^{(p)}(\theta_{\mathrm{Tx}}, \theta_{\mathrm{Rx}}) = a^{(p)}_{\mathrm{Rx}}(\theta_{\mathrm{Rx}}) \otimes a^{(p)}_{\mathrm{Tx}}(\theta_{\mathrm{Tx}}), \tag{3.15}$$

where θ_{Tx} is the direction of departure at the transmitting node and θ_{Rx} is the direction of arrival at the receiving node.

After range–Doppler processing, a complex snapshot can be extracted from a selected

range–Doppler cell. The observation vector for aperture p can then be expressed as

$$\mathbf{y}^{(p)} = \mathbf{A}^{(p)}(\theta)\mathbf{x}^{(p)} + \mathbf{n}^{(p)}, \tag{3.16}$$

where $\mathbf{A}^{(p)}(\theta)$ is the aperture-specific steering matrix, $\mathbf{x}^{(p)}$ contains the complex target coefficients for aperture p on the chosen angular or geometrically aligned target grid, and $\mathbf{n}^{(p)}$ represents noise. The coefficient vectors may have common support after geometric alignment without sharing the same complex amplitudes or phases.

When the relative aperture phases are known or compensated and a common complex coefficient model is valid, the aperture observations can be stacked into the coherent multi-aperture model:

$$\mathbf{y} = \mathbf{A}(\theta)\mathbf{x} + \mathbf{n}. \tag{3.17}$$

This coherent model uses a single coefficient vector $\mathbf{x}$ across the distributed stack. In contrast, non-coherent and shared-support approaches retain aperture-dependent coefficients $\mathbf{x}^{(p)}$ and combine the resulting information without imposing a common carrier-phase origin. The distinction between these formulations is important when residual inter-aperture phase offsets are present.

Before parameter estimation, the raw radar measurements are organized into a three-dimensional signal cube. The first dimension consists of fast-time samples and is used for range estimation. The second dimension consists of slow-time samples over multiple chirps and is used for Doppler estimation. The third dimension consists of spatial channels and is used for angular processing. This signal cube is the input to the processing pipeline developed in the following chapters.

3.3 Residual Phase-Error Model

Residual phase disagreement is separated into intra-aperture, inter-aperture, and within-CPI components. The first two describe spatial disagreement in an extracted snapshot, whereas the third describes phase evolution over the chirps used to form that snapshot. For a single target observed by aperture p, the ideal spatial snapshot is

$$\mathbf{y}_p = \alpha_p \mathbf{a}_p(\theta), \tag{3.18}$$

where α_p is the aperture-dependent complex target amplitude and $\mathbf{a}_p(\theta)$ is the corresponding steering vector. The following spatial models isolate two limiting error structures; actual measurements can contain both simultaneously.

3.3.1 Intra-Aperture Phase Error

Channel-relative residuals within aperture p are represented by

$$\mathbf{y}_p^{\text{intra}} = \mathbf{D}_p \alpha_p \mathbf{a}_p(\theta), \tag{3.19}$$

with

$$\mathbf{D}_p = \mathrm{diag}\left(e^{j\delta\phi_{p,1}}, \ldots, e^{j\delta\phi_{p,N}}\right), \tag{3.20}$$

where $\delta\phi_{p,n}$ denotes the residual phase of channel n. Because these terms are not generally equal, they distort the relative phase progression that carries the local angular information. The measured response then no longer lies exactly on the assumed array manifold, as illustrated in Fig. 3.2.

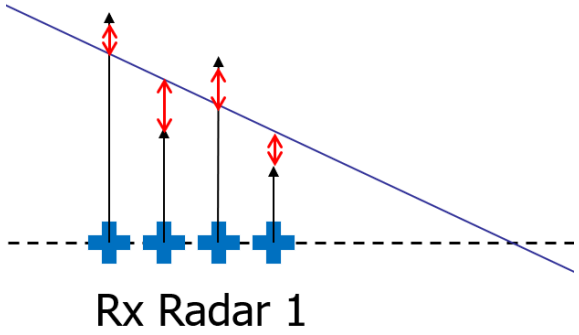


Figure 3.2: Illustration of an intra-aperture phase error. The sloped blue line represents the ideal phase progression, the black arrows the measured channel phases, and the red arrows the channel-relative residuals.

Fixed hardware offsets can be estimated by receive-channel calibration, but residual disagreement can remain because of noise, phase noise, imperfect calibration, antenna-position uncertainty, multipath, or mismatch between the assumed and actual scattering model. Since the local steering manifold itself is distorted, both aperture-local and distributed phase-based DoA methods can be affected.

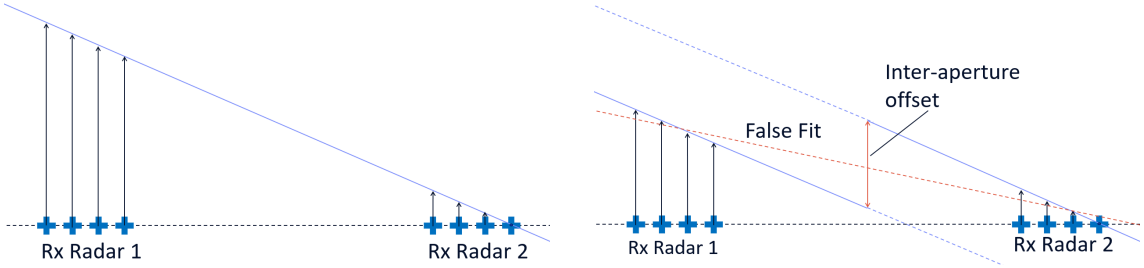
3.3.2 Inter-Aperture Phase Error

An inter-aperture residual instead rotates an otherwise valid local snapshot by an aperture-dependent common phase,

$$\mathbf{y}_p^{\mathrm{inter}} = e^{j\beta_p}\alpha_p\mathbf{a}_p(\theta), \tag{3.21}$$

where β_p is common to all channels in aperture p. The internal phase progression is preserved, so the aperture can still support a local angle estimate. However, snapshots with unknown relative block phases cannot be stacked as one calibrated distributed array because their complex relationship no longer matches the predicted global steering vector. Figure 3.3 shows the difference between the calibrated and rotated cases.

Estimating β_p establishes a common phase origin between aperture blocks but does not repair the channel-relative distortion in Eqs. (3.19)–(3.20). Conversely, accurate local calibration does not determine the relative carrier phase between independently clocked nodes. The decomposition is therefore useful for sensitivity analysis, although unknown geometry, target-dependent scattering, and multipath can violate the ideal assumption that the inter-aperture residual is exactly constant over every channel and scene location.



(a) Consistent phase progression across apertures.

(b) Unknown common phase rotation of one aperture block.

Figure 3.3: Inter-aperture phase consistency without an offset (a) and with an uncompensated block rotation (b).

3.3.3 Within-CPI Phase Evolution

A residual that changes over chirp index n is represented by

$$\mathbf{y}_p[n] = e^{\mathrm{j}\gamma_p[n]}\alpha_p\mathbf{a}_p(\theta), \tag{3.22}$$

where $\gamma_p[n]$ contains any phase evolution not explained by the target motion and nominal waveform model. A constant term is indistinguishable from an inter-aperture block rotation, while an approximately linear term behaves as a residual frequency or Doppler offset. Faster or stochastic variation reduces coherent integration and can broaden or displace the range–Doppler response [10, 22]. This third mechanism is therefore evaluated separately from the static spatial errors: the same phase spread can have different consequences depending on whether it occurs across channels, between aperture blocks, or over slow time.

Chapter 4

Processing and Estimation Methodology

This chapter describes the processing chain used to transform the raw monostatic and reciprocal bistatic measurements into target direction and Cartesian-position estimates. The geometry, FMCW signal representation, and residual phase-error models underlying the processing are established in Chapter 3. The present chapter focuses on the implemented operations: organization of received data, receive-side synchronization, target-snapshot extraction and association, steering-dictionary construction, DoA estimation, truth-blind support selection, and Cartesian backpropagation. The simulation and experimental configurations used to evaluate these operations are specified separately in Chapter 5.

The processing distinguishes between the four available apertures, M_1M_1, M_1M_2, M_2M_1, and M_2M_2. Coarse receive-side correction reduces bistatic timing- and frequency-induced distortions, but does not establish a common carrier phase. Accordingly, the estimator branches differ in how they combine aperture measurements: local and shared-support methods retain aperture-dependent complex coefficients, whereas fully coherent processing additionally relies on inter-aperture phase calibration and a consistent absolute propagation model. The final number of reported sources is selected from each estimator's output without using ground truth. These operations are not presented as new synchronization or localization algorithms. Their role is to provide one common chain in which the residual phase mechanisms can be measured or imposed and then related to the capabilities and coherence assumptions of each estimator branch.

4.1 Input Data and Channel Organization

The processing accepts complex radar samples organized into fast-time samples, slow-time chirps, and spatial channels. Two radar nodes provide the monostatic and reciprocal bistatic observations described in Chapter 3. For the azimuth processing considered here, two transmit antennas and four receive antennas per node form an eight-element local virtual azimuth array with $\lambda/2$ element spacing. The approximately 15° single-module

resolution benchmark used in the evaluation refers to close-target angular separation, not to single-target angular accuracy or the field of view. The physical hardware, frame-trigger arrangement, and measurement geometry are detailed in the experimental setup of Chapter 5.

For two-transmitter time-division multiple-access (TDMA) operation, the virtual-channel ordering used throughout the processing chain is

$$[\mathrm{TX1RX1}, \ldots, \mathrm{TX1RX4}, \mathrm{TX2RX1}, \ldots, \mathrm{TX2RX4}]. \tag{4.1}$$

The same ordering is used for measured and simulated data. The configuration-specific acquisition issues and frame handling are documented in Section A.1.

4.2 Receive-Side Synchronization and Phase Stabilization

The radar nodes are coarsely triggered but use independent oscillators. Synchronization therefore proceeds in two conceptually distinct stages. First, the dominant beat-frequency displacement, chirp-dependent migration, and slow-time phase drift are estimated from static reference responses and corrected on the received samples. Second, when coherent fusion is attempted, a reference response is used to estimate the remaining common phase rotation of each aperture block. The latter operation does not compensate arbitrary channel-dependent manifold errors; their implications for estimator performance are considered separately in the evaluation.

4.2.1 Coarse Bistatic Range–Doppler Alignment

The external trigger establishes a common frame start, but the two radar nodes retain independent carrier and sampling clocks. Their bistatic channels can therefore contain a constant beat-frequency displacement and a chirp-dependent drift. In the range–Doppler representation these terms appear as a range offset, range migration over slow time, and an associated Doppler-like phase progression.

Following the reference-assisted post-processing approach of Vecchi *et al.* [5], a strong static response that is visible in the relevant channels is tracked over chirp index n. Its measured migration is represented by the fast-time correction frequency

$$f_{\mathrm{corr}}(n) = f_0 + f_1 n, \tag{4.2}$$

where f_0 accounts for the dominant constant displacement and f_1 accounts for the approximately linear chirp-to-chirp drift. Multiplying the affected dechirped samples by the corresponding complex exponential removes the dominant frequency displacement before range and Doppler processing. The residual slow-time phase of the same static response is then used to estimate the remaining Doppler-like trend.

This operation is a receive-side alignment of the independently recorded data. It is intended to place corresponding static responses in compatible range–Doppler cells and to

stabilize their slow-time phase. It does not by itself establish carrier-level phase coherence between all four aperture blocks; that requires the additional phase treatment described in the following synchronization stages. Any chirp-dependent phase remaining after this correction is likewise not absorbed into a static inter-aperture offset. It is retained as a separate within-CPI residual mechanism in the controlled evaluation, consistent with the model in Section 3.3.3.

4.2.2 Fine Synchronization and Inter-Aperture Phase Stabilization

After the range–Doppler alignment, coherent processing still requires the complex aperture blocks to share a consistent carrier-phase reference. For a stationary reference response at known position $\mathbf{p}_{\mathrm{ref}}$, the expected propagation phase is computed with the same monostatic or bistatic path model used by the localization dictionary. The measured reference snapshot of aperture p is projected onto its expected steering vector, producing one complex block coefficient. Relative to a chosen reference aperture p_0, the residual block rotation is estimated as

$$\Delta\widehat{\beta}_p = \arg\left(\frac{\widehat{c}_p}{\widehat{c}_{p_0}}\right), \qquad \widetilde{\mathbf{y}}_p^{\mathrm{aligned}} = \widetilde{\mathbf{y}}_p \exp\left(-\mathrm{j}\Delta\widehat{\beta}_p\right), \tag{4.3}$$

where $\widehat{c}_p$ is the complex reference-projection coefficient after removal of the predicted geometric phase. The real-valued quantity $\Delta\widehat{\beta}_p$ is the relative phase angle obtained from the ratio of these coefficients, rather than the coefficient itself. The correction is applied as one common rotation per aperture block. It therefore addresses an inter-aperture phase offset; it cannot repair channel-relative errors within an aperture or a steering-model error that varies across channels.

Two reference choices are used in the measurement study and must be interpreted separately. In the original 25 cm-baseline experiment, the block phase is estimated from the same detected range–Doppler cluster that contains the evaluated target pair. This is a scene-conditioned or self-referenced correction: it is useful for testing the best coherence that the observed cluster can support, but it is not an independent calibration and must not be presented as one. In the later baseline experiments, a separate spherical reflector is used as the external reference. That procedure is operationally more independent of the target pair, but the predicted carrier phase then depends strongly on the surveyed reflector and radar positions.

At 77 GHz, a position error that is small on the scale of the room can still correspond to a substantial carrier-phase error. Consequently, reference-based compensation is applied only when the reference response is identifiable, and its validity is assessed from the residual phase diagnostics rather than assumed. The corrected snapshots are supplied to the coherent estimators described in Section 4.4 and Section 4.5; non-coherent and shared-support branches retain their own phase assumptions.

4.3 Localization Preprocessing and Geometric Association

Before applying the individual estimators, the detections from the available apertures are converted into a common geometric representation. This stage produces aperture-local position estimates, associates detections belonging to the same physical target, and determines a cluster centroid for each associated target.

The cluster centroid provides the initial position prior for the angular DoA search and the non-coherent Cartesian localization stage. The non-coherent Cartesian estimate subsequently initializes the fine coherent Cartesian search. Consequently, all estimators use the same target association while retaining the processing order used in the implementation.

4.3.1 Processing Framework

For an ordinary single-cell measurement, the input to the localization stage for aperture p is the complex spatial snapshot

$$\mathbf{y}_p = \mathbf{X}_p(r_p, d_p, :), \tag{4.4}$$

where (r_p, d_p) denotes the selected range–Doppler cell. The corresponding measured equivalent half-path range and Doppler estimate are denoted by R_p and v_p, respectively. For 1TX operation, $\mathbf{y}_p$ contains the four receive-channel samples of the selected aperture. For 2TX TDMA operation, it contains eight virtual-channel samples grouped according to the two active transmitters.

Because the transmitter groups in 2TX TDMA operation are measured on consecutive chirps, the later transmitter block contains an additional motion-induced phase. For transmitter group i, the implemented compensation is

$$\mathbf{y}_{p,i} \leftarrow \mathbf{y}_{p,i} \exp\left[-\mathrm{j}2\pi f_{D,p}(i-1)T_c\right], \tag{4.5}$$

where T_c is the chirp repetition interval and $f_{D,p}$ is the Doppler frequency associated with aperture p. No TDMA compensation is required for 1TX measurements.

The localization preprocessing consists of three stages:

1. an aperture-local position is estimated for each selected range–Doppler detection;
2. geometrically consistent aperture detections are associated into target clusters;
3. the cluster centroid is used to centre the target-specific angular search grid and the non-coherent Cartesian search grid.

The non-coherent Cartesian estimate obtained later is then used to centre the fine coherent Cartesian search.

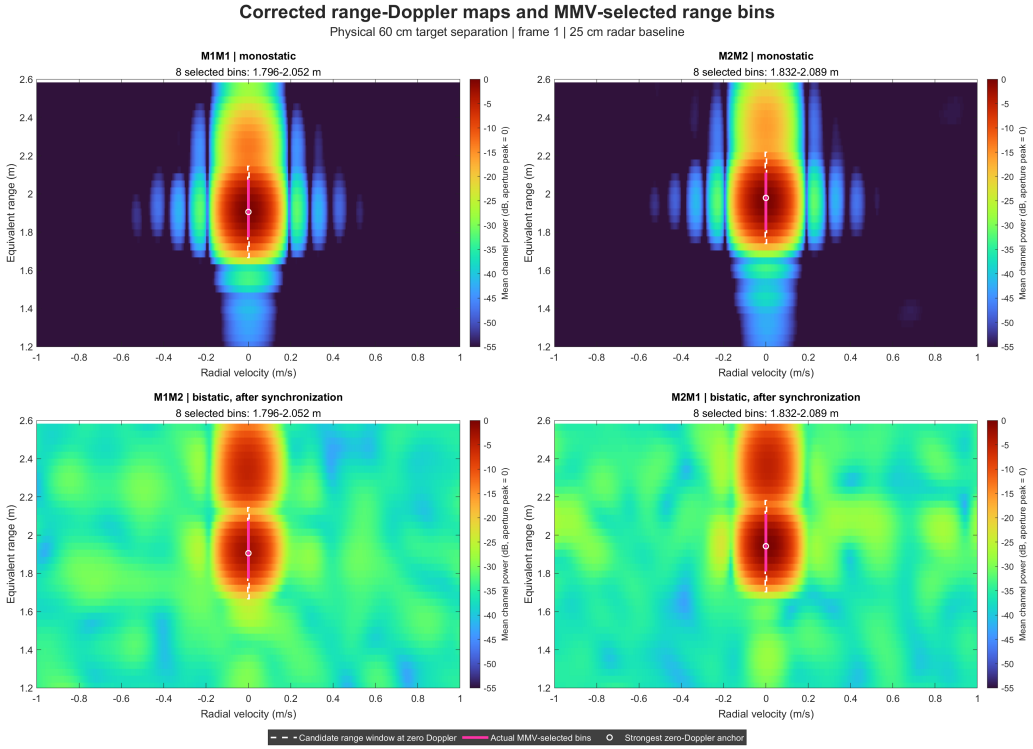


Figure 4.1: The white dashed stripes mark the region the selector searched for useful snapshots. The pink region marks the bins actually retained, while the white circle marks the anchor.

4.3.2 Range-Resolved Static Snapshots

A single detected range–Doppler cell is sufficient when the target response is well confined to one range bin. In the physical close-target measurements, however, two static reflectors can share zero Doppler while their energy is spread over neighbouring range bins. Reducing such a compact static response to only the constant-false-alarm-rate (CFAR) anchor cell can remove spatial information that is still present in the range–angle data. The implementation therefore preserves a compact set of zero-Doppler range-bin snapshots around the detected target region. In this multiple-measurement-vector (MMV) representation, a measurement task is one complex array snapshot from a particular aperture and retained range bin at the selected Doppler cell. Tasks can share candidate target angles while retaining their own complex coefficients. Figure 4.1 visualizes the retained bins.

For aperture p, let $\mathcal{L}_p$ denote the retained range-bin indices. The preserved complex data are written as

$$\mathbf{Y}_p = \begin{bmatrix} \mathbf{y}_{p,\ell_1} & \mathbf{y}_{p,\ell_2} & \cdots & \mathbf{y}_{p,\ell_{L_p}} \end{bmatrix}, \qquad \ell_j \in \mathcal{L}_p, \tag{4.6}$$

where each column is the complex spatial snapshot of one retained zero-Doppler range bin and has its own measured equivalent half-path range $R_{p,\ell}$. Each column is treated as a separate measurement task. The range bins are therefore not coherently summed before DoA estimation.

When range-resolved processing is used, normalization is applied once to the complete aperture matrix,

$$\widetilde{\mathbf{Y}}_p = \frac{\mathbf{Y}_p}{\max\left(\|\mathbf{Y}_p\|_{\mathrm{F}}, \epsilon\right)}, \tag{4.7}$$

rather than independently to every range bin. The task weight used by the non-coherent power-fusion branches is consequently

$$q_{p,\ell} = \|\widetilde{\mathbf{y}}_{p,\ell}\|_2^2 . \tag{4.8}$$

This retains the relative strength of the range bins and prevents weak or noisy bins from receiving the same weight as bins containing the dominant target response.

For the angle-domain relative dictionaries, every range task uses its own measured range. For candidate global bearing θ_k, the task-specific candidate position is

$$\mathbf{p}_{p,\ell,k} = \rho_{p,\ell,k} \begin{bmatrix} \sin\theta_k \\ \cos\theta_k \end{bmatrix}, \tag{4.9}$$

where $\rho_{p,\ell,k}$ is selected such that

$$\left\|\mathbf{p}_{p,\ell,k} - \bar{\mathbf{t}}_p\right\|_2 + \left\|\mathbf{p}_{p,\ell,k} - \bar{\mathbf{r}}_p\right\|_2 = 2R_{p,\ell}. \tag{4.10}$$

The relative steering vector for that task is then evaluated from the exact TX–target–RX element paths at $\mathbf{p}_{p,\ell,k}$. Thus, all tasks share the same angular support index k, but they are not forced to share one range. This is important when close static reflectors occupy neighbouring range bins.

For Bartlett, independent FOCUSS, and the displayed Block-FOCUSS power, the range-task outputs are reduced with the implemented `maxWithinAperture` rule. If $P_{p,\ell}(k)$ denotes a non-negative task power after task weighting, then

$$P_p^{\mathrm{MMV}}(k) = \max_{\ell \in \mathcal{L}_p} P_{p,\ell}(k), \tag{4.11}$$

followed by aperture normalization and non-coherent summation,

$$P^{\mathrm{MMV}}(k) = \sum_{p \in \mathcal{P}_c} \frac{P_p^{\mathrm{MMV}}(k)}{\max_{k'} P_p^{\mathrm{MMV}}(k')}. \tag{4.12}$$

The MMV angular score $P^{\mathrm{MMV}}(k)$ in Eq. (4.12) is normalized after this summation. Hence, at each candidate angle, an aperture contributes the range bin that gives its strongest support; responses from different range bins are not added coherently or directly summed into one broad angular response. This operation does not prescribe the number of targets. Model order is determined only from the estimator output using the truth-blind rules in Section 4.4.6.

Block-FOCUSS and MT-BCS additionally use the retained range bins as multiple measurement tasks during their shared-support estimation. Coherent FOCUSS also re-

tains the individual range tasks, but uses the corresponding absolute-path dictionaries and phase-aligned complex measurements in the coherent stack. The absolute coherent dictionary keeps one common Cartesian candidate for a given angular atom across the distributed system; the response of that candidate in each retained fast Fourier transform (FFT) range bin is modelled separately. This avoids assigning a different physical target position to the same coherent atom in different apertures.

4.3.3 Geometric Parameterization

The coordinate system and antenna-element positions are those defined in Chapter 3. A candidate target position is expressed in the global Cartesian coordinate system as

$$\mathbf{p}_k = \begin{bmatrix} x_k \\ y_k \end{bmatrix}, \tag{4.13}$$

where the x-axis follows the radar baseline and the y-axis points along the common broadside direction. The corresponding global bearing is

$$\theta(\mathbf{p}_k) = \operatorname{atan2}(x_k, y_k). \tag{4.14}$$

The local angles at radar nodes M_1 and M_2 remain denoted by θ_1 and θ_2, consistent with the system model. These local angles are generally different and are not assumed to be equal.

For aperture p, transmit element i, receive element q, and candidate position $\mathbf{p}_k$, the exact transmitter–target–receiver path is

$$L_{p,iq}(\mathbf{p}_k) = \|\mathbf{p}_k - \mathbf{t}_{p,i}\|_2 + \|\mathbf{p}_k - \mathbf{r}_{p,q}\|_2 . \tag{4.15}$$

This is the exact-path formulation introduced in Chapter 3, evaluated here for a candidate target position. The near-field geometry is therefore retained throughout the localization procedure.

The equivalent half-path range predicted for aperture p is

$$\widehat{R}_p(\mathbf{p}_k) = \frac{1}{2}\left(\|\mathbf{p}_k - \bar{\mathbf{t}}_p\|_2 + \|\mathbf{p}_k - \bar{\mathbf{r}}_p\|_2\right), \tag{4.16}$$

where $\bar{\mathbf{t}}_p$ and $\bar{\mathbf{r}}_p$ denote the transmit- and receive-array centre positions of aperture p.

For a monostatic aperture, the transmit and receive centres coincide and $\widehat{R}_p(\mathbf{p}_k)$ reduces to the distance between the radar node and the candidate target. For a bistatic aperture, it represents half the total transmitter–target–receiver path. This common convention allows monostatic and bistatic range measurements to be used in the same Cartesian localization framework.

4.3.4 Aperture-Local Localization and Target Association

For every selected range–Doppler cell, an aperture-local angular search is performed over

$$-80^\circ \le \theta \le 80^\circ, \tag{4.17}$$

using a spacing of 0.25°. For each candidate angle, the corresponding position is obtained from the intersection between the receiving-array bearing ray and the measured monostatic or bistatic range contour.

The spatial response is evaluated using a relative exact-path steering vector. For candidate position $\mathbf{p}_k$, its (i,q)-th entry is

$$a^{\mathrm{rel}}_{p,iq}(\mathbf{p}_k) = \exp\left[-\mathrm{j}\beta\left(L_{p,iq}(\mathbf{p}_k) - L_{p,11}(\mathbf{p}_k)\right)\right], \qquad \beta = \frac{2\pi}{\lambda}. \tag{4.18}$$

The complete relative steering vector is normalized to unit Euclidean norm,

$$\mathbf{a}^{\mathrm{rel}}_p(\mathbf{p}_k) \leftarrow \frac{\mathbf{a}^{\mathrm{rel}}_p(\mathbf{p}_k)}{\left\|\mathbf{a}^{\mathrm{rel}}_p(\mathbf{p}_k)\right\|_2}. \tag{4.19}$$

Referencing all element paths to $L_{p,11}(\mathbf{p}_k)$ removes one common phase from aperture p while preserving the spatial phase progression across its transmit–receive channels.

The aperture-local position estimate is

$$\widehat{\mathbf{p}}_p = \arg\max_{\mathbf{p}_k} \left|\mathbf{a}^{\mathrm{rel},\mathsf{H}}_p(\mathbf{p}_k)\mathbf{y}_p\right|^2. \tag{4.20}$$

The resulting aperture-local positions are associated using Cartesian proximity. Each aperture may contribute at most one detection to a cluster. Figure 4.2 shows an example.

For target cluster c, let $\mathcal{P}_c$ denote the set of associated apertures. The cluster centroid is defined as

$$\mathbf{p}_{\mathrm{cent},c} = \frac{1}{|\mathcal{P}_c|}\sum_{p\in\mathcal{P}_c}\widehat{\mathbf{p}}_p. \tag{4.21}$$

The centroid is the common geometric initialization for the next processing stages. It centres the target-specific angular search grid and the non-coherent Cartesian backpropagation grid. The non-coherent Cartesian position estimate obtained from that search subsequently centres the fine coherent Cartesian grid.

The centroid should therefore be distinguished from the final Cartesian position estimates returned by the non-coherent and coherent backpropagation methods.

4.3.5 Angular Search Grid and Steering-Dictionary Construction

The global bearing of target cluster c is obtained from its centroid,

$$\widehat{\theta}_{\mathrm{cent},c} = \theta(\mathbf{p}_{\mathrm{cent},c}). \tag{4.22}$$

For ordinary localization runs, the default target-specific angular search interval is

$$\Theta_c = \left[\widehat{\theta}_{\mathrm{cent},c} - 15^\circ, \widehat{\theta}_{\mathrm{cent},c} + 15^\circ\right] \cap [-90^\circ, 90^\circ], \tag{4.23}$$

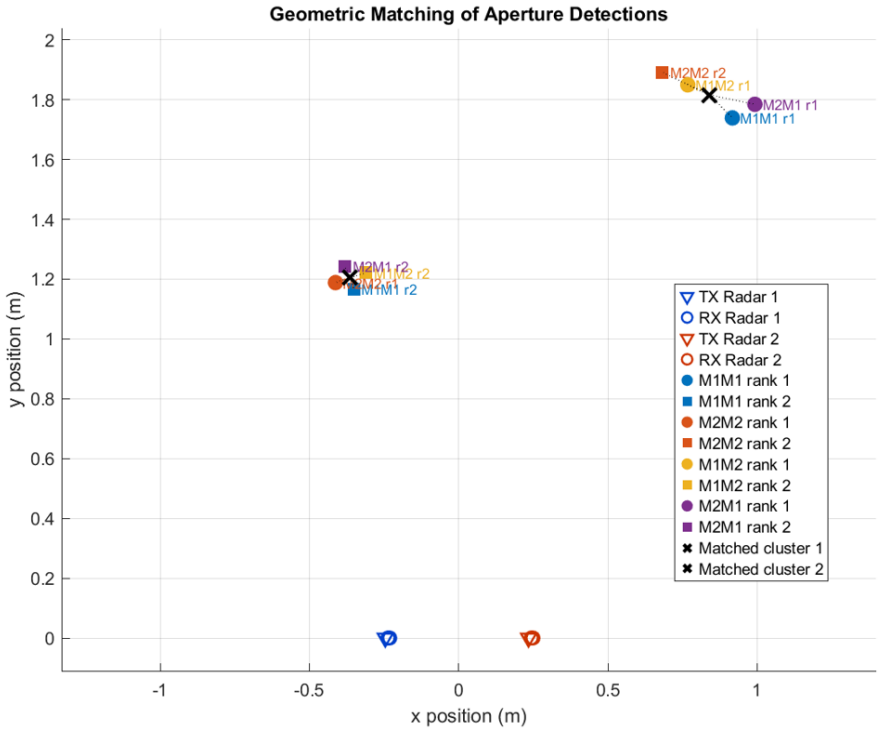


Figure 4.2: Aperture-local position estimates and cluster centroids for a two-target scenario.

sampled with a spacing of 0.1°. The evaluation runners override this half-width where required: the physical and measurement-matched pair studies use ±35°, while the controlled phase studies use the study-specific grids listed in Appendix A.2.

For angular-grid point θ_k, the corresponding candidate target lies on the global bearing ray

$$\mathbf{p}_k(\rho) = \rho \begin{bmatrix} \sin\theta_k \\ \cos\theta_k \end{bmatrix}. \tag{4.24}$$

The radial coordinate is selected by minimizing the mean squared mismatch between the predicted and measured equivalent half-path ranges,

$$\widehat{\rho}_k = \arg\min_{\rho} \frac{1}{|\mathcal{P}_c|} \sum_{p \in \mathcal{P}_c} \left[\widehat{R}_p\left(\mathbf{p}_k(\rho)\right) - R_p \right]^2. \tag{4.25}$$

The candidate position associated with angular atom k is therefore

$$\mathbf{p}_k = \mathbf{p}_k(\widehat{\rho}_k). \tag{4.26}$$

For single-cell processing and for the absolute coherent dictionary, $\mathbf{p}_k$ represents one common physical candidate position shared by the distributed system. In the range-resolved relative branch of Section 4.3.2, the angular index k remains common but the radial coordinate is replaced by the task-specific range intersection $\mathbf{p}_{p,\ell,k}$. The spatial

phase response is always evaluated using the corresponding transmit- and receive-element positions.

The consistency between candidate position $\mathbf{p}_k$ and the measured equivalent range of aperture p is represented by

$$w_{R,p}(\mathbf{p}_k) = \exp\left[-\frac{1}{2}\left(\frac{\widehat{R}_p(\mathbf{p}_k) - R_p}{\sigma_R}\right)^2\right], \tag{4.27}$$

where

$$\sigma_R = \max\left(2\Delta R,\, 0.04\,\mathrm{m}\right), \tag{4.28}$$

and ΔR is the spacing of the processed range axis. The range weight acts as a soft consistency constraint and is used by both the angular estimators and the Cartesian backpropagation methods.

Two steering dictionaries are constructed. The relative dictionary contains

$$A_p^{\mathrm{rel}}[(i,q),k] = a_{p,iq}^{\mathrm{rel}}(\mathbf{p}_k), \tag{4.29}$$

and is insensitive to one unknown common phase per aperture. It is used by methods that process the apertures independently or impose only a common angular support.

The absolute dictionary retains the complete propagation phase,

$$A_p^{\mathrm{abs}}[(i,q),k] = \exp\left[-\mathrm{j}\beta L_{p,iq}(\mathbf{p}_k) + \mathrm{j}\phi_{\mathrm{FFT},p,iq}(\mathbf{p}_k)\right], \tag{4.30}$$

where $\phi_{\mathrm{FFT},p,iq}(\mathbf{p}_k)$ represents the deterministic phase contribution associated with the selected Hann-windowed range-FFT bin.

The absolute dictionary is used for exact-path Cartesian backpropagation and for coherent multi-aperture processing. In non-coherent backpropagation, it is applied independently within each aperture, so an unknown aperture-wide phase does not affect the resulting power. For coherent processing, the aperture snapshots must additionally be phase aligned.

Every dictionary column is normalized to unit Euclidean norm,

$$\mathbf{A}_p(:,k) \leftarrow \frac{\mathbf{A}_p(:,k)}{\|\mathbf{A}_p(:,k)\|_2}. \tag{4.31}$$

For single-cell multi-aperture fusion, the normalized snapshot is

$$\widetilde{\mathbf{y}}_p = \frac{\mathbf{y}_p}{\max\left(\|\mathbf{y}_p\|_2, \epsilon\right)}. \tag{4.32}$$

For range-resolved processing, Eq. (4.7) is used instead so that all retained bins of one physical aperture share one normalization factor. In both cases, normalization limits differences caused solely by aperture-level received power while preserving the relative phase and amplitude variation between channels. The original measured snapshots remain

unchanged.

4.4 Direction-of-Arrival Estimators

The theoretical principles of the considered DoA estimators were introduced in Section 2.2. This section describes their implementation using the common angular grid and steering dictionaries defined in Section 4.3.5.

All estimators operate on the same target-specific angular support Θ_c. In single-cell processing they use the candidate positions $\mathbf{p}_k$ obtained from the range-conditioned radial fit. In the range-resolved relative branch, each aperture–range task instead uses the corresponding position $\mathbf{p}_{p,\ell,k}$ from Eq. (4.10), while retaining the same angular index k.

The methods differ primarily in how the aperture–range-bin measurement tasks defined in Section 4.3.2 are combined. Bartlett beamforming and independent FOCUSS process each task separately before combining powers. Block-FOCUSS and Bayesian compressive sensing impose a common angular support while retaining task-specific complex coefficients. Coherent fine-grid FOCUSS instead stacks phase-aligned complex measurements with their absolute-path dictionaries and estimates one complex coefficient vector for the complete coherent stack.

Unless stated otherwise, the normalized snapshot $\widetilde{\mathbf{y}}_p$ of Eq. (4.32) is used for single-cell processing, whereas the range-resolved measurements use the matrix normalization in Eq. (4.7).

4.4.1 Bartlett Beamforming

Bartlett beamforming is used as the conventional non-coherent reference method. For aperture p, retained range task ℓ, and angular-grid point k, the task-local Bartlett power is

$$P_{\mathrm{B},p,\ell}(k) = \left|\mathbf{A}_{p,\ell}^{\mathrm{rel},\mathsf{H}}(:,k)\widetilde{\mathbf{y}}_{p,\ell}\right|^2. \tag{4.33}$$

For a single-cell measurement, $\mathcal{L}_p$ contains one task and this reduces to the ordinary aperture-local Bartlett spectrum. For range-resolved processing, each task spectrum is first normalized and weighted by its retained snapshot energy,

$$Q_{\mathrm{B},p,\ell}(k) = q_{p,\ell}\frac{P_{\mathrm{B},p,\ell}(k)}{\max_{k'} P_{\mathrm{B},p,\ell}(k')}. \tag{4.34}$$

The implemented `maxWithinAperture` fusion then gives

$$P_{\mathrm{B},p}^{\mathrm{MMV}}(k) = \max_{\ell\in\mathcal{L}_p} Q_{\mathrm{B},p,\ell}(k), \tag{4.35}$$

and each physical aperture contribution is normalized before the four aperture powers are summed. The resulting multi-aperture spectrum can be written as

$$P_{\mathrm{B}}(k) = \overline{w}_R^2(k) \sum_{p \in \mathcal{P}_c} \frac{P_{\mathrm{B},p}^{\mathrm{MMV}}(k)}{\max_{k'} P_{\mathrm{B},p}^{\mathrm{MMV}}(k')}, \tag{4.36}$$

where, for the general case, the mean squared range-consistency weight is

$$\overline{w}_R^2(k) = \frac{1}{M_c} \sum_{m=1}^{M_c} q_m^{\mathrm{n}} w_{R,m}^2(k), \tag{4.37}$$

with m indexing the available measurement tasks and $q_m^{\mathrm{n}} = q_m / \sum_j q_j$. For the task-specific range-intersection dictionaries of Eq. (4.10), the range consistency is satisfied by construction.

The method uses the spatial phase progression within an individual task, but only powers are combined across range bins and physical apertures. It therefore does not require carrier-phase coherence between apertures. Reported DoA peaks are selected from the final spectrum using the truth-blind rule in Section 4.4.6.

4.4.2 FOCUSS

FOCUSS is applied independently to each available measurement task. For task m, corresponding to either one aperture snapshot or one retained aperture–range-bin snapshot, the measurement model is

$$\widetilde{\mathbf{y}}_m = \mathbf{A}_m^{\mathrm{rel}} \mathbf{x}_m + \mathbf{n}_m. \tag{4.38}$$

The implementation is initialized with

$$\mathbf{W}_m^{(0)} = \mathbf{I}. \tag{4.39}$$

At iteration s, the weighted dictionary is

$$\mathbf{A}_{w,m}^{(s)} = \mathbf{A}_m^{\mathrm{rel}} \mathbf{W}_m^{(s)}. \tag{4.40}$$

The coefficient estimate is updated according to

$$\mathbf{x}_m^{(s+1)} = \mathbf{W}_m^{(s)} \mathbf{A}_{w,m}^{(s),\mathsf{H}} \left[\mathbf{A}_{w,m}^{(s)} \mathbf{A}_{w,m}^{(s),\mathsf{H}} + \xi_m^{(s)} \mathbf{I} \right]^{-1} \widetilde{\mathbf{y}}_m, \tag{4.41}$$

where

$$\xi_m^{(s)} = \sigma_n^2 + \gamma_m^{(s)} \tag{4.42}$$

contains the assumed noise variance and an adaptive diagonal-loading term. The loading term is increased when the weighted dictionary becomes poorly conditioned.

The diagonal weights are updated from the current coefficient magnitudes,

$$\mathbf{W}_m^{(s+1)} = \mathrm{diag} \left[\left(|x_{m,k}^{(s+1)}| + \epsilon \right)^{1-p_{\mathrm{F}}} \right]_{k=1}^{K}, \tag{4.43}$$

where p_{F} is the sparsity parameter and ϵ prevents singular weights.

The matrix inversion in Eq. (4.41) is evaluated using an economy-size singular-value decomposition. Iteration terminates when

$$\frac{\left\|\mathbf{x}_m^{(s+1)} - \mathbf{x}_m^{(s)}\right\|_2}{\left\|\mathbf{x}_m^{(s)}\right\|_2 + \epsilon} < \tau_{\mathrm{F}}, \tag{4.44}$$

or when the maximum number of iterations is reached.

The task-local FOCUSS power is

$$P_{\mathrm{F},m}(k) = |x_{m,k}|^2. \tag{4.45}$$

For range-resolved measurements, the task powers are range weighted, normalized, multiplied by $q_{p,\ell}$, and fused with the same `maxWithinAperture` rule as Bartlett. Thus,

$$P_{\mathrm{F},p}^{\mathrm{MMV}}(k) = \max_{\ell \in \mathcal{L}_p} \left[q_{p,\ell} \frac{|x_{p,\ell,k}|^2 w_{R,p,\ell}^2(k)}{\max_{k'} |x_{p,\ell,k'}|^2 w_{R,p,\ell}^2(k')} \right]. \tag{4.46}$$

The physical-aperture spectra are then normalized and summed non-coherently. Each task therefore retains its own complex coefficient vector; agreement between range bins and apertures enters only through the final power score. The reported peaks are selected using the truth-blind rule in Section 4.4.6.

4.4.3 Block-FOCUSS

Block-FOCUSS uses the same relative dictionaries as independent FOCUSS but estimates all available measurement tasks jointly. Let M_c be the number of tasks in target cluster c. With range-resolved processing, a task is one retained range bin of one physical aperture. The coefficient matrix is

$$\mathbf{X} = \begin{bmatrix} \mathbf{x}_1 & \mathbf{x}_2 & \cdots & \mathbf{x}_{M_c} \end{bmatrix} \in \mathbb{C}^{K \times M_c}. \tag{4.47}$$

Each column corresponds to one task and each row to one global angular-grid point. The task measurement models are

$$\widetilde{\mathbf{y}}_m = \mathbf{A}_m^{\mathrm{rel}} \mathbf{x}_m + \mathbf{n}_m, \qquad m = 1, \ldots, M_c. \tag{4.48}$$

At iteration s, each task coefficient vector is updated using the same weighting matrix,

$$\mathbf{x}_m^{(s+1)} = \mathbf{W}^{(s)} \mathbf{A}_{w,m}^{(s),\mathsf{H}} \left[\mathbf{A}_{w,m}^{(s)} \mathbf{A}_{w,m}^{(s),\mathsf{H}} + \xi_m^{(s)} \mathbf{I} \right]^{-1} \widetilde{\mathbf{y}}_m, \tag{4.49}$$

with

$$\mathbf{A}_{w,m}^{(s)} = \mathbf{A}_m^{\mathrm{rel}} \mathbf{W}^{(s)}. \tag{4.50}$$

After updating all task coefficient vectors, the group magnitude associated with angular atom k is

$$r_k^{(s+1)} = \left(\sum_{m=1}^{M_c} \left| X_{k,m}^{(s+1)} \right|^2 \right)^{1/2}. \tag{4.51}$$

The common weighting matrix is then updated as

$$\mathbf{W}^{(s+1)} = \operatorname{diag}\left[\left(r_k^{(s+1)} + \epsilon \right)^{1-p_{\mathrm{B}}} \right]_{k=1}^{K}. \tag{4.52}$$

For a single snapshot per aperture, the Block-FOCUSS spectrum is simply r_k^2. When several range tasks are present, the implementation retains the taskwise coefficient powers and applies `maxWithinAperture` to obtain the displayed range-resolved score,

$$P_{\mathrm{BFOC},p}^{\mathrm{MMV}}(k) = \max_{\ell \in \mathcal{L}_p} |X_{k,p,\ell}|^2, \tag{4.53}$$

followed by normalization and summation over the physical apertures as in Eq. (4.12).

Block-FOCUSS therefore promotes one common angular support across all tasks while allowing every task to retain an independent complex coefficient. Unknown inter-aperture block phases are absorbed into these task-specific coefficients and are not used directly in coherent summation. The final number of reported DoA peaks is not fixed by the block-sparse solver and is selected from the resulting spectrum using Section 4.4.6.

4.4.4 Coherent Fine-Grid FOCUSS

Coherent FOCUSS uses the phase-aligned measurements produced by the fine synchronization procedure in Section 4.2.2. Only physical apertures for which a valid block-phase estimate is available are included. If range-resolved processing is active, every retained aperture–range pair remains an individual complex measurement task in the coherent stack.

Let $\widetilde{\mathbf{y}}_m^{\mathrm{aligned}}$ and $\mathbf{A}_m^{\mathrm{abs}}$ denote the aligned snapshot and absolute-path dictionary of coherent task m. Tasks belonging to the same physical aperture receive the same estimated inter-aperture block-phase correction. The stacked measurement vector and dictionary are

$$\widetilde{\mathbf{y}}_{\mathrm{stack}} = \begin{bmatrix} \widetilde{\mathbf{y}}_1^{\mathrm{aligned}} \\ \widetilde{\mathbf{y}}_2^{\mathrm{aligned}} \\ \vdots \\ \widetilde{\mathbf{y}}_{M_{\mathrm{C}}}^{\mathrm{aligned}} \end{bmatrix}, \qquad \mathbf{A}_{\mathrm{stack}}^{\mathrm{abs}} = \begin{bmatrix} \mathbf{A}_1^{\mathrm{abs}} \\ \mathbf{A}_2^{\mathrm{abs}} \\ \vdots \\ \mathbf{A}_{M_{\mathrm{C}}}^{\mathrm{abs}} \end{bmatrix}, \tag{4.54}$$

where M_{C} is the number of valid coherent tasks. In contrast to the relative range-task dictionaries, the absolute dictionaries preserve one common Cartesian candidate for each angular atom across all tasks. The candidate response in each retained FFT range bin is evaluated with the corresponding absolute-path and deterministic range-FFT phase model.

The coherent sparse model is

$$\widetilde{\mathbf{y}}_{\text{stack}} = \mathbf{A}^{\text{abs}}_{\text{stack}}\mathbf{x}_{\text{C}} + \mathbf{n}_{\text{stack}}, \tag{4.55}$$

where one complex coefficient vector $\mathbf{x}_{\text{C}}$ is shared by all coherent channels and tasks. The iterative update follows the same weighted formulation as independent FOCUSS,

$$\mathbf{x}^{(s+1)}_{\text{C}} = \mathbf{W}^{(s)}_{\text{C}}\mathbf{A}^{(s),\mathsf{H}}_{w,\text{C}}\left[\mathbf{A}^{(s)}_{w,\text{C}}\mathbf{A}^{(s),\mathsf{H}}_{w,\text{C}} + \xi^{(s)}_{\text{C}}\mathbf{I}\right]^{-1}\widetilde{\mathbf{y}}_{\text{stack}}, \tag{4.56}$$

where

$$\mathbf{A}^{(s)}_{w,\text{C}} = \mathbf{A}^{\text{abs}}_{\text{stack}}\mathbf{W}^{(s)}_{\text{C}}. \tag{4.57}$$

The weights are updated as

$$\mathbf{W}^{(s+1)}_{\text{C}} = \text{diag}\left[\left(|x^{(s+1)}_{\text{C},k}| + \epsilon\right)^{1-p_{\text{C}}}\right]_{k=1}^{K}. \tag{4.58}$$

The resulting coherent spectrum is

$$P_{\text{CF}}(k) = |x_{\text{C},k}|^2. \tag{4.59}$$

In contrast to Block-FOCUSS, one angular atom has one common complex coefficient for the complete distributed stack. The method therefore requires both the absolute path model and the inter-aperture block-phase correction to be consistent.

The main coherent branch optionally restricts the solve to ambiguity-safe intervals around phase-robust Block-FOCUSS support. This restriction only limits the candidate search region; it does not transfer a target count from Block-FOCUSS. After the coherent solve, the number of reported sources is selected independently from the coherent result by the complex-data model-order test in Section 4.4.6. An additional ungated coherent branch evaluates the complete coarse angular sector as a diagnostic output. It is excluded from the reported five-method comparison, which uses the gated coherent FOCUSS branch.

4.4.5 Bayesian Compressive Sensing

The Bayesian estimator follows the Bayesian compressive-sensing approach to DoA estimation in [32] and is implemented as a proper-complex multi-task sparse model. With range-resolved processing, every retained aperture–range snapshot is a Bayesian task. Each task retains an individual coefficient vector $\mathbf{x}_m$, while all tasks share one relevance precision α_k for angular atom k.

For task m, the likelihood is

$$\widetilde{\mathbf{y}}_m \mid \mathbf{x}_m, \sigma^2_m \sim \mathcal{CN}\left(\mathbf{A}^{\text{rel}}_m\mathbf{x}_m, \sigma^2_m\mathbf{I}\right). \tag{4.60}$$

The coefficient prior is

$$x_{m,k} \mid \alpha_k \sim \mathcal{CN}\left(0, \alpha_k^{-1}\right). \tag{4.61}$$

Define

$$\mathbf{\Gamma} = \operatorname{diag}\left(\alpha_1^{-1}, \ldots, \alpha_K^{-1}\right). \tag{4.62}$$

For fixed hyperparameters, the posterior mean of the coefficient vector for task m is

$$\boldsymbol{\mu}_m = \mathbf{\Gamma}\mathbf{A}_m^{\mathrm{rel},\mathsf{H}}\left[\mathbf{A}_m^{\mathrm{rel}}\mathbf{\Gamma}\mathbf{A}_m^{\mathrm{rel},\mathsf{H}} + \sigma_m^2\mathbf{I}\right]^{-1}\widetilde{\mathbf{y}}_m. \tag{4.63}$$

The relevance precisions are updated jointly from the posterior statistics of all tasks. The shared hyperparameters promote a common angular support, whereas the individual posterior means allow task-dependent complex amplitudes and phases.

The implementation uses Bartlett-based initialization, an independent noise-variance update for each task, logarithmic hyperparameter damping, and iterative relevance updates. Iteration terminates when the relative hyperparameter change falls below the convergence threshold or when the maximum number of iterations is reached.

The multi-task Bayesian power associated with angular atom k is

$$P_{\mathrm{BCS}}(k) = \sum_{m=1}^{M_c} |\mu_{m,k}|^2. \tag{4.64}$$

Unlike the `maxWithinAperture` display fusion used for Bartlett, independent FOCUSS, and Block-FOCUSS, MT-BCS uses the retained range bins as joint Bayesian tasks and obtains its support from the posterior energy. The truth-blind conversion of this energy into a reported source count and DoA set is given in Section 4.4.6.

As with Block-FOCUSS, MT-BCS estimates a common angular support without requiring the task coefficient vectors to share the same complex phase.

4.4.6 Truth-Blind Support and Model-Order Selection

The estimator output is converted into reported target directions without using the known target positions or the true number of targets. Ground truth is used only afterwards when simulation results are scored in Chapter 6. Consequently, an inferred $K = 2$ means that the estimator itself selected two source directions; whether those two directions correspond to the true pair is a separate evaluation criterion.

Visible-Peak Selection for Bartlett, FOCUSS, and Block-FOCUSS

Bartlett, independent FOCUSS, and Block-FOCUSS determine their reported source count directly from visible local maxima of their own final normalized spectrum. Search-boundary samples are not treated as resolved local maxima. Candidate maxima are ordered from strongest to weakest and a maximum is considered only when its level exceeds the method-specific floor in Table 4.1.

The strongest accepted maximum defines the first source. A second source is reported only when its angular separation from every stronger selected maximum satisfies

$$|\theta_i - \theta_j| \geq 0.20^\circ, \tag{4.65}$$

and the spectrum contains a sufficiently deep valley between the two maxima. For maxima at indices i and j, the implemented valley drop is

$$D_{ij} = \min\left[S(i), S(j)\right] - \min_{k \in [i,j]} S(k), \tag{4.66}$$

where $S(k)$ is the normalized spectrum in dB. The second maximum is retained only if D_{ij} exceeds the method-specific threshold. At most two peaks are reported by these methods in the present two-target evaluation.

Table 4.1: Truth-blind visible-peak selection settings for the non-coherent spectra.

Method	**Peak floor**	**Minimum separation**	**Minimum valley drop**
Bartlett	$-6\,\mathrm{dB}$	0.20°	$0.35\,\mathrm{dB}$
FOCUSS	$-8\,\mathrm{dB}$	0.20°	$1.0\,\mathrm{dB}$
Block-FOCUSS	$-8\,\mathrm{dB}$	0.20°	$3.0\,\mathrm{dB}$

If no local maximum exceeds the normal floor, the implementation retains the strongest interior local maximum as a single-source estimate. If the spectrum is degenerate and contains no interior local maximum, its strongest sample is used as a $K = 1$ fallback. These fallbacks prevent the post-processing rule itself from inventing a second source.

Posterior-Energy Support for MT-BCS

MT-BCS does not use the visible-peak rule above. Let $e_k = P_{\mathrm{BCS}}(k)$ denote the posterior support energy of atom k, and let $e_{(1)} \geq e_{(2)} \geq \cdots$ be the energies sorted in descending order. The smallest number J is retained for which

$$\frac{\sum_{j=1}^{J} e_{(j)}}{\sum_k e_k} \geq 0.95. \tag{4.67}$$

This cumulative-energy rule determines which angular atoms survive. Because an off-grid source can distribute posterior energy over neighbouring atoms, retained atoms separated by no more than 0.25° are merged into one support group. The reported angle of a group $\mathcal{G}_g$ is the posterior-energy centroid

$$\widehat{\theta}_g = \frac{\sum_{k \in \mathcal{G}_g} e_k \theta_k}{\sum_{k \in \mathcal{G}_g} e_k}. \tag{4.68}$$

The inferred MT-BCS model order is the number of support groups after this merge. No known source count enters either the energy-retention or merging step.

Complex eBIC Selection for Coherent FOCUSS

A large coherent distributed aperture can produce several narrow ambiguity lobes. Therefore, selecting the two largest coherent FOCUSS coefficients would not provide a reliable source-count rule. Coherent FOCUSS first uses only local maxima that are visibly present in its final coefficient spectrum at or above $-8\,\mathrm{dB}$ as candidate atoms. The best one-atom and two-atom supports are then compared using the original complex coherent measurement stack rather than the plotted spectrum alone.

For candidate support $\mathcal{S}_K$, the residual sum of squares is

$$\mathrm{RSS}_K = \left\| \widetilde{\mathbf{y}}_{\mathrm{stack}} - \mathbf{A}_{\mathcal{S}_K} \mathbf{A}_{\mathcal{S}_K}^{\dagger} \widetilde{\mathbf{y}}_{\mathrm{stack}} \right\|_2^2 . \tag{4.69}$$

Before evaluating a two-atom model, the normalized dictionary columns must satisfy

$$\left| \mathbf{a}_i^{\mathsf{H}} \mathbf{a}_j \right| \leq 0.90. \tag{4.70}$$

Pairs with higher coherence are excluded because they do not form a sufficiently identifiable two-source dictionary model and can otherwise represent one off-grid source with two nearly identical atoms.

For the remaining candidates, the implemented extended Bayesian information criterion is

$$\mathrm{eBIC}_K = N_r \log\left(\frac{\mathrm{RSS}_K}{N_r}\right) + 3K \log N_r + 2\gamma \log \binom{G_v}{K}, \tag{4.71}$$

where $N_r = 2N$ is the number of real observations corresponding to the N complex samples, three real parameters are charged per source (one angle and one complex amplitude), G_v is the number of visible coherent FOCUSS candidate maxima, and $\gamma = 1$. The two-source model is selected only when its eBIC is not larger than that of the best one-source model. Otherwise, $K = 1$ is retained.

This test uses neither synthetic truth nor the source count selected by Block-FOCUSS. An optional Block-FOCUSS gate may restrict where coherent FOCUSS is allowed to search, but the coherent source count is inferred independently from the final coherent candidates and their complex-data fit.

4.4.7 Numerical Settings

The principal numerical settings of the implemented FOCUSS-based estimators are listed in Table 4.2. The independent and block variants use separate sparsity exponents and assumed noise variances; this suppresses the broad coefficient pedestal while retaining the shared-support behavior of Block-FOCUSS.

The BCS implementation uses a maximum of 250 iterations, a convergence tolerance of 10^{-4}, and logarithmic hyperparameter damping of 0.35. The same estimator settings are used throughout the simulation and measurement evaluations unless stated otherwise.

Table 4.2: Numerical settings of the implemented FOCUSS-based DoA estimators.

Setting	**FOCUSS**	**Block-FOCUSS**	**Coherent FOCUSS**
Sparsity parameter	$p_{\mathrm{F}} = 0.55$	$p_{\mathrm{B}} = 0.50$	$p_{\mathrm{C}} = 0.50$
Base loading	10^{-4}	10^{-4}	10^{-4}
Assumed noise variance	0.3	0.3	10^{-3}
Maximum iterations	120	120	120
Convergence tolerance	10^{-6}	10^{-6}	10^{-6}
Weight stabilization ϵ	10^{-6}	10^{-6}	10^{-6}

The physical-data sparsity and peak-visibility settings were fixed after a preliminary diagnostic sweep of representative measured spectra and then held fixed across the reported cases. They should therefore be interpreted as empirically configured rather than as independently validated hyperparameters.

4.5 Cartesian Backpropagation Localization

In addition to the one-dimensional DoA estimators of Section 4.4, two Cartesian backpropagation methods are used to estimate the target position directly. The first combines the aperture responses non-coherently, whereas the second combines their complex responses after inter-aperture phase alignment.

Following aperture-local angle estimation and geometric association, target cluster c contains a set of aperture-local position estimates. Let $\mathcal{P}_c$ denote the set of apertures associated with that cluster. Its initial Cartesian position is the centroid

$$\mathbf{p}_{\mathrm{cent},c} = \frac{1}{|\mathcal{P}_c|} \sum_{p \in \mathcal{P}_c} \widehat{\mathbf{p}}_p, \tag{4.72}$$

where $\widehat{\mathbf{p}}_p$ is the provisional position obtained from the range and aperture-local angle estimate of aperture p. The non-coherent backpropagation search is centred on the centroid, $\mathbf{p}_{\mathrm{cent},c}$. Local maxima are subsequently detected in the non-coherent heatmap using a level floor, a minimum spatial separation, and a minimum valley-depth criterion. Up to two retained maxima are used as truth-blind seeds for separate fine coherent search windows.

4.5.1 Non-Coherent Backpropagation

The non-coherent search is performed over a $1.2\,\mathrm{m} \times 1.2\,\mathrm{m}$ Cartesian region centred on $\mathbf{p}_{\mathrm{cent},c}$, using a grid spacing of 1 cm.

For candidate position $\mathbf{p}_k$, the consistency between its predicted equivalent half-path range and the measured range R_p is represented by

$$w_{R,p}(\mathbf{p}_k) = \exp\left[-\frac{1}{2}\left(\frac{\widehat{R}_p(\mathbf{p}_k) - R_p}{\sigma_R}\right)^2\right], \tag{4.73}$$

with

$$\sigma_R = \max\left(2\Delta R,\, 0.04\,\mathrm{m}\right), \tag{4.74}$$

where ΔR is the spacing of the processed range axis.

The aperture-local projection is

$$z_p(\mathbf{p}_k) = \mathbf{a}_p^{\mathrm{abs},H}(\mathbf{p}_k)\widetilde{\mathbf{y}}_p, \tag{4.75}$$

and the non-coherent score is

$$S_{\mathrm{NC}}(\mathbf{p}_k) = \frac{1}{|\mathcal{P}_c|}\sum_{p\in\mathcal{P}_c} |z_p(\mathbf{p}_k)|^2\, w_{R,p}(\mathbf{p}_k). \tag{4.76}$$

Only the aperture powers are combined. The score therefore uses the exact spatial phase progression within each aperture but does not require a known carrier-phase relation between apertures.

The strongest non-coherent estimate is

$$\widehat{\mathbf{p}}_{\mathrm{NC},c} = \arg\max_{\mathbf{p}_k\in\mathcal{G}_{\mathrm{NC},c}} S_{\mathrm{NC}}(\mathbf{p}_k), \tag{4.77}$$

where $\mathcal{G}_{\mathrm{NC},c}$ is the Cartesian grid centred on $\mathbf{p}_{\mathrm{cent},c}$.

4.5.2 Coherent Backpropagation

Let $\widehat{\mathbf{p}}_{\mathrm{NC},c,s}$, $s = 1,\ldots,K_{\mathrm{NC}}$, denote the retained coarse maxima, with $K_{\mathrm{NC}} \le 2$. For each retained maximum, the coherent backpropagation search is performed on a separate fine Cartesian grid centred on its seed $\widehat{\mathbf{p}}_{\mathrm{NC},c,s}$. This restricts the computationally demanding coherent search to the target region identified by the preceding non-coherent stage.

In the default configuration, the coherent grid covers a $0.10\,\mathrm{m} \times 0.10\,\mathrm{m}$ region with spacing

$$\Delta_{\mathrm{C}} = 0.5\,\mathrm{mm}. \tag{4.78}$$

The resulting grid for seed s is denoted by $\mathcal{G}_{\mathrm{C},c,s}(\widehat{\mathbf{p}}_{\mathrm{NC},c,s})$. Before coherent fusion, the aperture snapshots are corrected using the inter-aperture block-phase estimates obtained in Section 4.2.2. For aperture p, the normalized and phase-aligned snapshot is

$$\widetilde{\mathbf{y}}_p^{\mathrm{aligned}} = \widetilde{\mathbf{y}}_p \exp\left(-\mathrm{j}\Delta\widehat{\beta}_p\right), \tag{4.79}$$

where $\Delta\widehat{\beta}_p$ is the estimated phase offset relative to the selected reference aperture.

For candidate position $\mathbf{p}_k$, the range-weighted absolute steering vectors are stacked over the apertures associated with target cluster c,

$$\mathbf{a}_{\mathrm{stack},w}(\mathbf{p}_k) = \begin{bmatrix} \sqrt{w_{R,p_1}(\mathbf{p}_k)}\mathbf{a}^{\mathrm{abs}}_{p_1}(\mathbf{p}_k) \\ \sqrt{w_{R,p_2}(\mathbf{p}_k)}\mathbf{a}^{\mathrm{abs}}_{p_2}(\mathbf{p}_k) \\ \vdots \end{bmatrix}, \qquad p_1, p_2, \ldots \in \mathcal{P}_c. \tag{4.80}$$

The corresponding weighted measurements are stacked in the same aperture order,

$$\widetilde{\mathbf{y}}_{\mathrm{stack},w}(\mathbf{p}_k) = \begin{bmatrix} \sqrt{w_{R,p_1}(\mathbf{p}_k)}\widetilde{\mathbf{y}}^{\mathrm{aligned}}_{p_1} \\ \sqrt{w_{R,p_2}(\mathbf{p}_k)}\widetilde{\mathbf{y}}^{\mathrm{aligned}}_{p_2} \\ \vdots \end{bmatrix}. \tag{4.81}$$

The dependence of $\widetilde{\mathbf{y}}_{\mathrm{stack},w}(\mathbf{p}_k)$ on the candidate position is caused only by the range-consistency weights; the underlying measured snapshots remain unchanged. Applying the square root of the range weight to both the steering and measurement vectors gives one factor $w_{R,p}(\mathbf{p}_k)$ to the coherent contribution of aperture p.

The coherent Cartesian score is

$$S_{\mathrm{C}}(\mathbf{p}_k) = \frac{\left|\mathbf{a}^{\mathsf{H}}_{\mathrm{stack},w}(\mathbf{p}_k)\widetilde{\mathbf{y}}_{\mathrm{stack},w}(\mathbf{p}_k)\right|^2}{\|\mathbf{a}_{\mathrm{stack},w}(\mathbf{p}_k)\|_2^2 + \epsilon}. \tag{4.82}$$

In contrast to non-coherent backpropagation, the complex aperture projections are summed before taking the magnitude squared. The score therefore favours candidate positions for which the exact-path steering model and the estimated inter-aperture phase corrections cause the aperture responses to add constructively.

The coherent position estimate is

$$\widehat{\mathbf{p}}_{\mathrm{C},c} = \arg \max_{\mathbf{p}_k \in \mathcal{G}_{\mathrm{C},c}(\widehat{\mathbf{p}}_{\mathrm{NC},c})} S_{\mathrm{C}}(\mathbf{p}_k). \tag{4.83}$$

The individual fine maps are retained as diagnostics. When two coarse seeds are available, their refined positions are not selected independently, because a rank-one score would treat the coherent contribution of the other target as model mismatch. Instead, a short list of owned local maxima is extracted from each fine window. For every candidate pair $(\mathbf{p}_i, \mathbf{p}_j)$, the two-target steering matrix is formed as

$$\mathbf{A}_{ij} = \begin{bmatrix} \mathbf{a}_{\mathrm{stack},w}(\mathbf{p}_i) & \mathbf{a}_{\mathrm{stack},w}(\mathbf{p}_j) \end{bmatrix}, \tag{4.84}$$

and the unknown complex target amplitudes are estimated by least squares,

$$\widehat{\boldsymbol{\alpha}}_{ij} = \mathbf{A}^{\dagger}_{ij}\widetilde{\mathbf{y}}_{\mathrm{stack},w}. \tag{4.85}$$

The selected pair maximizes the normalized energy explained by the joint projection,

$$(\widehat{\mathbf{p}}_1, \widehat{\mathbf{p}}_2) = \arg \max_{(\mathbf{p}_i, \mathbf{p}_j)} \frac{\|\mathbf{A}_{ij}\widehat{\boldsymbol{\alpha}}_{ij}\|_2^2}{\|\widetilde{\mathbf{y}}_{\mathrm{stack},w}\|_2^2 + \epsilon}. \tag{4.86}$$

The search is restricted to cross-window pairs, so the two refinements remain associated with distinct coarse supports. If the coarse map retains only one reliable maximum, no second target is forced and the single-target coherent estimate is reported.

Because the aperture responses are combined coherently, the estimate is sensitive to errors in the radar geometry, antenna-element positions, reference-target position, absolute steering model, and residual inter-aperture phase. The influence of these errors on localization performance is evaluated in Chapter 6.

Chapter 5

Simulation and Experimental Evaluation Methodology

This chapter describes how the processing and estimation methodology of Chapter 4 is assessed using both generated IF-level data and physical measurements. The purpose is to distinguish algorithmic behavior under controlled assumptions from achievable performance with independently clocked radar nodes. The simulator provides known target geometry and independently adjustable impairments; the physical setup tests whether its predicted phase and localization behavior transfers to measured data. Evaluation geometries, baseline conditions, performance metrics, and statistical protocols are defined here before any outcome is presented in Chapter 6.

The methodology separates target-count recovery, relative target separation, and correct absolute localization. It also distinguishes scene-assisted calibration from independent reference transfer, and simulation frame realizations from repetitions obtained in one physical measurement campaign. In support of the thesis contribution, the controlled studies are organized by where the residual phase disagreement enters the processing chain: within an aperture, between aperture blocks, or over the chirps of a CPI. This provides a common basis for testing estimator architectures with different coherence assumptions rather than assigning one universal phase-error tolerance.

5.1 Simulation Architecture

Two radar nodes, denoted by M_1 and M_2, are positioned symmetrically about the global origin. The simulator reproduces the four-aperture configuration and TDMA channel ordering specified in Chapter 4. In 2TX mode, consecutive chirps alternate between TX1 and TX2, resulting in the virtual-channel ordering given in Eq. (4.1).

The raw data of radar 1 contain the local M_1M_1 response and the remote M_2M_1 response. Similarly, radar 2 contains the M_2M_2 and M_1M_2 responses. The bistatic contributions are separated from the monostatic scene by the same nominal 5 MHz frequency offset used in the measurement setup.

For each radar, the generated data cube has dimensions

$$N_{\mathrm{RX}} \times N_{\mathrm{ADC}} \times N_{\mathrm{chirp}} \times N_{\mathrm{frame}}. \tag{5.1}$$

Before saving, the cube is packed into the same layout returned by the DCA1000 reader. As a result, synthetic and measured data use the same reshaping, TDMA separation, range processing, synchronization, detection, and localization pipeline.

The simulation parameters are listed in Table 5.1.

Table 5.1: Parameters of the synthetic radar generator.

Parameter	**Value**
Center frequency f_c	77 GHz
Wavelength $\lambda = c/f_c$	3.90 mm
Chirp slope μ	29.982 THz/s
Analogue-to-digital converter (ADC) sampling frequency f_s	30 MHz
Samples per chirp N_{ADC}	1024
Sampled chirp bandwidth B_{ADC}	1.0234 GHz
Raw chirps per frame	128
Frames per statistical capture	4
Independent captures per geometry	2
Pulse repetition interval T_c	160 μs
Receivers per radar	4
Transmitters per radar	1 or 2 TDMA
Radar baseline	b
Nominal bistatic frequency offset	5 MHz
Measurement-matched nominal signal-to-noise ratio (SNR)	6 dB
Measurement-matched common-noise fraction η	0.35
Measurement-matched slow-time correlation ρ	0.68
RX-chain phase standard deviation	0.30°
TDMA TX-chain phase standard deviation	0.07°
Reference phase-centre standard deviation	85°
Image-path copies per primary target	2
Structured-clutter range interval	3.05–9.50 m

5.2 Target Geometry and IF Signal Model

Targets are represented as point reflectors in a two-dimensional Cartesian plane. Target k has initial position $\mathbf{p}_{k,0}$, velocity $\mathbf{v}_k$, and complex amplitude A_k. At raw chirp index n, its position is

$$\mathbf{p}_k[n] = \mathbf{p}_{k,0} + \mathbf{v}_k n T_c. \tag{5.2}$$

The target is assumed stationary during one ADC interval and is updated between chirps. This stop-and-hop approximation is accurate because the displacement during one chirp is negligible for the velocities considered.

For transmitter element i and receiver element j of aperture p, the exact propagation path is

$$L_{p,ij,k}[n] = \|\mathbf{p}_k[n] - \mathbf{t}_{p,i}\| + \|\mathbf{p}_k[n] - \mathbf{r}_{p,j}\|, \tag{5.3}$$

with corresponding delay

$$\tau_{p,ij,k}[n] = \frac{L_{p,ij,k}[n]}{c}. \tag{5.4}$$

The exact path is used for all monostatic and bistatic channels. Therefore, the generated data include near-field curvature, baseline-dependent phase differences, and the correct local transmit and receive directions.

Fast time is denoted by

$$t'_\ell = \frac{\ell}{f_s}, \qquad \ell = 0, \ldots, N_{\mathrm{ADC}} - 1. \tag{5.5}$$

The FMCW frequency ramp restarts at every chirp. The target-dependent beat frequency is modeled as

$$f_{b,p,ij,k}[n] = \mu\tau_{p,ij,k}[n] + s_p\Delta f[n], \tag{5.6}$$

where $\Delta f[n]$ is the inter-radar frequency mismatch and

$$s_p = \begin{cases} +1, & p = M_2M_1, \\ -1, & p = M_1M_2, \\ 0, & p \in \{M_1M_1, M_2M_2\}. \end{cases} \tag{5.7}$$

The opposite signs reproduce the reciprocal bistatic frequency shifts observed in the measurement setup.

The complex dechirped IF sample is

$$\begin{aligned} x_{p,ij,k}[\ell, n] = {} & g_{p,j}\alpha_{p,ij,k}[n] \exp\big(j2\pi f_{b,p,ij,k}[n]t'_\ell\big) \\ & \times \exp(-j\beta L_{p,ij,k}[n]) \exp(j\phi_p[n]), \end{aligned} \tag{5.8}$$

where $\beta = 2\pi/\lambda$, $g_{p,j}$ is the complex receive-channel gain, $\alpha_{p,ij,k}[n]$ is the propagation amplitude, and $\phi_p[n]$ contains the aperture-level oscillator and hardware phase.

The first exponential places the target at its correct fast-time beat frequency. The second represents the exact carrier propagation phase required for spatial processing. Since $L_{p,ij,k}[n]$ is updated for every chirp, target motion also produces the physical Doppler phase

$$f_{D,p,ij,k} = -\frac{1}{\lambda}\frac{\mathrm{d}L_{p,ij,k}}{\mathrm{d}t}. \tag{5.9}$$

Target Doppler is therefore introduced only through the changing propagation path

and is not added as a separate phase term.

A simplified path-loss model is used,

$$\alpha_{p,ij,k}[n] = A_k \left(\frac{4}{\max(L_{p,ij,k}[n], 1)} \right)^2. \tag{5.10}$$

This retains a controlled distance dependence without attempting to model the complete radar cross section of each target.

5.3 Simulation Impairment Model

5.3.1 Inter-Radar Frequency and Phase Errors

Whereas Chapter 4 describes how the frequency mismatch is estimated and corrected from received data, this section defines the underlying frequency and phase errors used to generate the synthetic measurements. The inter-radar frequency mismatch is modeled as

$$\Delta f[n] = f_{\text{off}} + n f_{\text{drift}} + \eta_f[n], \tag{5.11}$$

where f_{off} contains the nominal 5 MHz separation, f_{drift} produces a linear chirp-to-chirp frequency drift, and $\eta_f[n]$ represents residual jitter.

The jitter is generated as a first-order autoregressive (AR(1)) process,

$$\eta_f[n] = \rho \eta_f[n-1] + \sigma_f \sqrt{1-\rho^2}\, w_f[n], \qquad w_f[n] \sim \mathcal{N}(0,1). \tag{5.12}$$

The oscillator phase is obtained by integrating the same frequency history,

$$\psi[n+1] = \psi[n] + 2\pi \Delta f[n] T_c. \tag{5.13}$$

This link is important because the frequency error that shifts the bistatic range response must also generate the corresponding slow-time phase progression. The reciprocal bistatic aperture phases are

$$\phi_{M_2M_1}[n] = +\psi[n] + \phi_{M_2M_1,0} + a_{M_2M_1} n + q_\phi[n], \tag{5.14}$$

$$\phi_{M_1M_2}[n] = -\psi[n] + \phi_{M_1M_2,0} + a_{M_1M_2} n - q_\phi[n]. \tag{5.15}$$

Here, $\phi_{p,0}$ is an initial block-phase mismatch, a_p is an optional residual linear phase slope, and $q_\phi[n]$ is an optional random-walk disturbance.

After fast-time frequency correction, the common oscillator phase remains on the bistatic returns. Removing the phase of a static reference therefore removes the common link error. For another static target, the corrected phase should be constant over slow time. A moving target, however, must retain the phase variation caused by its changing propagation path. The intended outcome of the synchronization procedure is therefore:

- clock-induced range migration is removed;
- clock-induced slow-time phase is removed;
- physical target Doppler is preserved.

5.3.2 Additional Nonidealities

The target responses can be supplemented with receive-channel errors, direct coupling, structured clutter, and receiver noise.

Each physical receive channel is assigned a constant complex gain

$$g_j = 10^{a_j/20} \exp(j\gamma_j), \tag{5.16}$$

where a_j and γ_j represent amplitude and phase errors. In 2TX operation, the same physical RX gain is applied to both transmitter blocks. Each TDMA transmitter additionally receives a constant chain factor

$$h_t = 10^{b_t/20} \exp(\mathrm{j}\xi_t), \tag{5.17}$$

which is shared by the monostatic and bistatic paths originating from that physical transmitter. The measurement-matched preset uses standard deviations of 0.25 dB and 0.30° for the RX-chain amplitude and phase, and 0.10 dB and 0.07° for the TX-chain amplitude and phase. These factors are drawn once per capture and remain fixed over its four frames.

The wall reflector used for block-phase calibration is not an ideal point scatterer in the physical experiment. Its effective phase centre can therefore differ between aperture paths. This nuisance is represented by one constant reference-only phase term per aperture,

$$\beta_{\mathrm{ref},p} \sim \mathcal{N}(0, \sigma_{\mathrm{ref}}^2), \qquad \sigma_{\mathrm{ref}} = 85^\circ, \tag{5.18}$$

in the measurement-matched preset. The term is not applied to the scene targets, whose phases continue to follow their exact geometric paths. It should therefore be interpreted as an effective calibration-reference uncertainty and not as a range-independent scene-target phase error.

Direct radar coupling is modeled as a bistatic return with a propagation path approximately equal to the radar baseline. On the equivalent half-path range axis, this contribution appears near

$$R_{\mathrm{coupling}} \approx \frac{b}{2}. \tag{5.19}$$

Measurement-like clutter is represented by stationary targets with a mixture of common and channel-dependent complex amplitudes. A correlated slow-time phase texture broadens these target responses around zero Doppler without changing their physical ranges. To avoid introducing artificial reflectors in the close-range localization region,

background clutter is generated only between approximately 3.05 and 9.50 m. The static synchronization reflector is modeled separately.

Weak specular multipath is included through image scatterers. For primary target position $\mathbf{p}_k$, an image position is generated as

$$\widetilde{\mathbf{p}}_{k,q} = \mathbf{p}_k + \delta_{R,q}\mathbf{e}_{R,k} + \delta_{\perp,q}\mathbf{e}_{\perp,k}, \tag{5.20}$$

where $\mathbf{e}_{R,k}$ and $\mathbf{e}_{\perp,k}$ denote the radial and lateral unit vectors. In the measurement-matched preset, $\delta_{R,q} \in [1.20, 2.80]\,\mathrm{m}$ and $\delta_{\perp,q} \sim \mathcal{N}(0, 0.25^2\,\mathrm{m}^2)$. Image positions are constrained to ranges above 3.05 m and are propagated with the same exact transmitter–image–receiver path model as the primary target. Their complex amplitude is a weak random fraction of the primary amplitude,

$$\widetilde{\alpha}_{k,q} = 0.035\,\xi_q\alpha_k, \qquad \xi_q \sim \mathcal{U}(0.65, 1.35). \tag{5.21}$$

Receiver noise combines an independent term for every RX chain and a common term shared by the four RX chains of one radar,

$$w_{r,j}[\ell, n] = \sqrt{1-\eta}\,u_{r,j}[\ell, n] + \sqrt{\eta}\,c_r[\ell, n], \tag{5.22}$$

where η is the common-mode power fraction. Both components follow an AR(1) evolution over chirps,

$$u[n] = \rho u[n-1] + \sqrt{1-\rho^2}\,\varepsilon[n], \qquad \varepsilon[n] \sim \mathcal{CN}(0, 1), \tag{5.23}$$

and the same recursion is used for $c_r[n]$. The resulting noise cube is scaled such that

$$\sigma_w^2 = \frac{P_x}{10^{\mathrm{SNR_{dB}}/10}}. \tag{5.24}$$

The measurement-matched preset uses $\eta = 0.35$ and $\rho = 0.68$. A fixed random seed makes an individual capture reproducible, while different captures use different deterministic seeds and therefore independent hardware and background realizations.

5.4 Complete Receiver Data and Simulation Validation

For receiver node r, the complete simulated sample is

$$\begin{aligned} x_r[\ell, n, j] = &\sum_{p\in\mathcal{P}_r}\sum_{i\in\mathcal{I}_p(n)}\sum_{k} x_{p,ij,k}[\ell, n] + x_{r,j}^{\mathrm{coupling}}[\ell, n] \\ &+ \sum_{u} x_{r,j,u}^{\mathrm{clutter}}[\ell, n] + w_r[\ell, n, j], \end{aligned} \tag{5.25}$$

where $\mathcal{I}_p(n)$ denotes the active transmitter index set for aperture p during chirp n. In 1TX mode this set contains one transmitter, while in 2TX TDMA mode it depends on the

active transmitter slot. The target index k contains both primary scatterers and, when enabled, their weak image-path copies. The receiver-side aperture sets are

$$\mathcal{P}_1 = \{M_1M_1, M_2M_1\}, \qquad \mathcal{P}_2 = \{M_2M_2, M_1M_2\}. \tag{5.26}$$

All components are combined in the complex IF domain before range or Doppler processing. Consequently, the influence of synchronization errors and other nonidealities propagates naturally through the complete processing chain.

Several scene configurations are used, including isolated static targets, static and moving targets in the same frame, and closely spaced target pairs. Each synthetic capture also contains a static synchronization reflector. The perfect and nonideal versions of a scene use identical geometry and target parameters; only the selected error components differ.

The statistical measurement-matched experiment mirrors the repeated-frame structure of the physical captures. Four frames are generated under one fixed hardware realization, after which a second capture is generated with an independent hardware and background draw. The resulting eight observations per geometry therefore contain two levels of variation: within-capture frame variation and between-capture realization variation.

Model calibration and validation use disjoint target-separation sets. Physical captures with separations of 7, 20, 40, and 60 cm determine the measurement-matched nuisance preset. Captures at 10, 30, 50, 70, and 100 cm are held out and used only after the preset has been fixed. Agreement is assessed from corrected bistatic range–Doppler background statistics, cross-channel coherence, intra-TX spatial-phase residuals, and inter-aperture phase spread. Localization and source-count results on the held-out set are consequently validation outcomes rather than tuning targets.

The simulator produces diagnostic plots of the geometry, range profiles, range–slow-time maps, range–Doppler maps, frequency history, and integrated oscillator phase. Examples are shown in Figs. 5.1 and 5.2.

The frequency and phase models are checked internally by verifying that

$$\angle\left(e^{j(\psi[n+1]-\psi[n])}\right) = \angle\left(e^{j2\pi\Delta f[n]T_c}\right) \tag{5.27}$$

within numerical precision.

The generated files are processed using the synchronization procedure described in Chapter 4. Validation verifies that the corrected bistatic returns appear at their physical equivalent ranges, that stationary targets are concentrated near zero Doppler, and that the physical Doppler of moving targets is preserved.

5.5 Physical Measurement Setup and Data Acquisition

The physical measurements provide a test of synchronization and localization under errors not fully represented in the synthetic signal model. The same four monostatic and reciprocal bistatic apertures are available as in the simulator, but node phase coherence

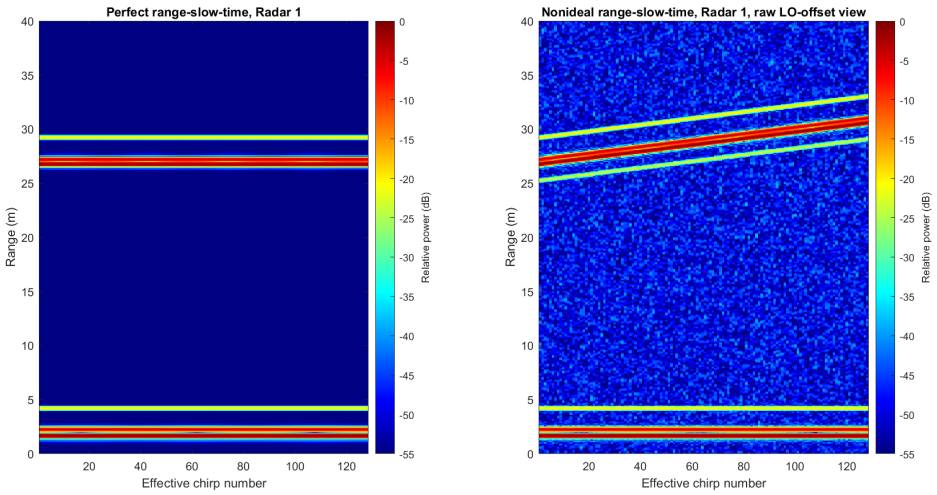


Figure 5.1: Perfect and nonideal synthetic range–slow-time responses. The nonideal bistatic return contains the imposed frequency migration and oscillator phase variation.

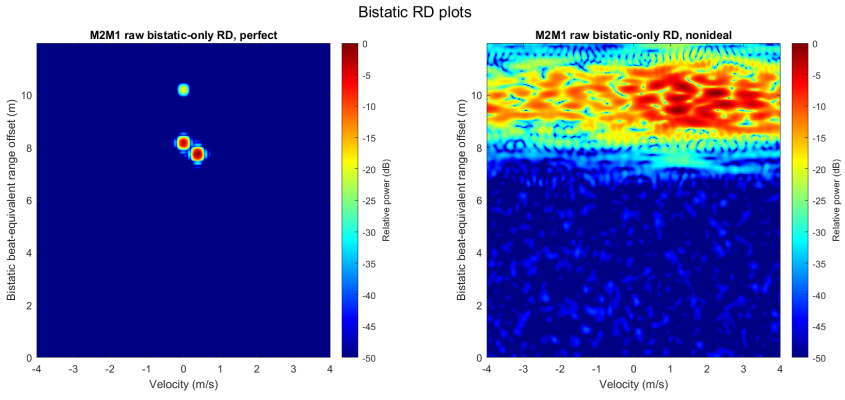


Figure 5.2: Perfect and nonideal synthetic range–Doppler responses before receive-side synchronization.

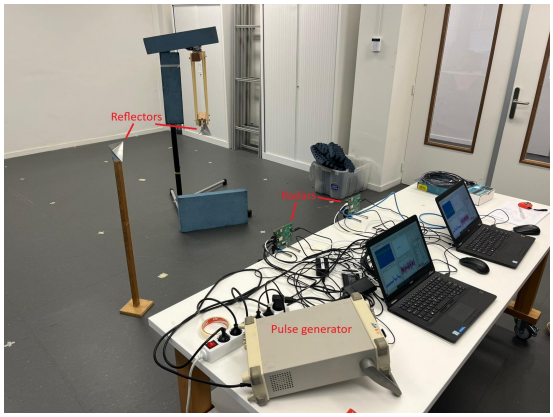
and the effective measurement geometry must be established from observed data rather than assumed.

The measurement system consists of two Texas Instruments IWR1443BOOST evaluation modules, each based on an IWR1443 radar sensor. Each sensor provides three transmit and four receive antennas. In the azimuth processing used in this thesis, two azimuth transmit antennas are combined with the four receive antennas to form an eight-element virtual azimuth array with approximately $\lambda/2$ element spacing. For this configuration, the conventional single-module azimuth resolution benchmark is approximately 15°. This value describes the ability to resolve two nearby targets and is distinct from single-target angular accuracy and from the angular field of view.

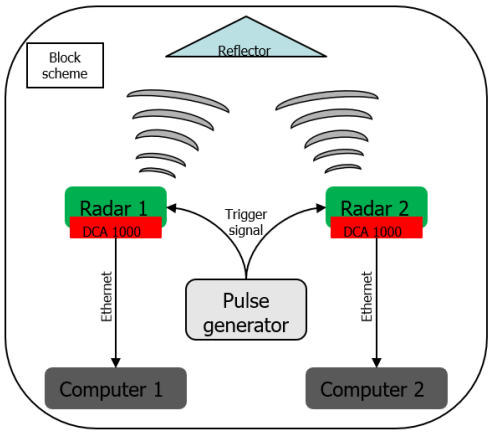
Each radar is connected to a DCA1000 evaluation module, which captures the raw digital ADC samples and transfers them to a separate computer through Ethernet. The radar units are mounted on a measurement table using fixed clamps, as shown in Fig. 5.3.

Both radar nodes are triggered by an external pulse generator. The generator provides a 3.3 V square pulse every 50 ms. Since the radar frame duration is set to 40 ms, each frame is triggered separately. This provides a common coarse acquisition time for both nodes, although it does not create frequency or carrier-phase coherence between their independent oscillators.

The recorded data cubes are transferred to two separate laptops for offline processing. After import, the raw samples are organized into fast-time samples, slow-time chirps, and spatial channels. The two-transmitter TDMA virtual-channel ordering is specified in Eq. (4.1) in Chapter 4. Practical acquisition issues, including missing frames and missing chirps, are discussed in Section A.1.



(a) Photograph of the experimental setup, including the radar units, reference reflectors, pulse generator, and data-capture hardware.



(b) Schematic representation of the radar nodes, trigger source, DCA1000 capture boards, and processing computers.

Figure 5.3: Experimental measurement setup. The photograph in (a) shows the physical arrangement of the hardware, while the schematic in (b) summarizes the acquisition and trigger connections.

For the later physical baseline measurements, a separate spherical reflector provides

a reference for estimating relative aperture-block phase. This setup is distinct from the scene-assisted, in-sample target-cluster reference used in the original 25 cm baseline study. These two calibration conditions are evaluated separately and are not interchangeable.

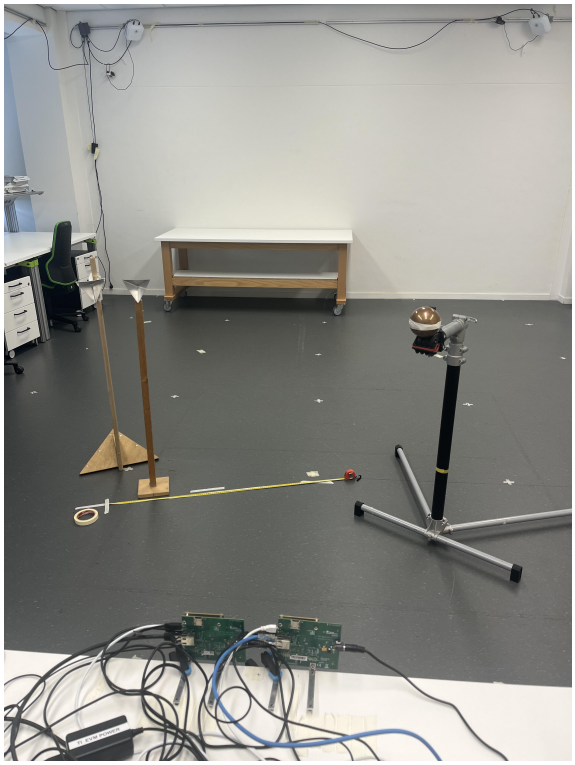


Figure 5.4: Measurement setup containing the spherical reflector used for inter-aperture phase compensation. The radar apertures are configured with a 14.03 cm baseline.

5.6 Evaluation Geometries and Performance Metrics

All tests use two transmit channels and four receive channels per radar, and the exact near-field monostatic and bistatic path model described in Chapter 4. The controlled baseline comparisons use 6.23 cm, 14.03 cm, and 25 cm; the first is an idealized configuration that cannot be realized mechanically with the equipment used here. Physical measurements are reported for baseline distances of 14.03 cm and 25 cm, with an additional 19.5 cm acquisition represented in the supplied physical-result figures. Its results are shown separately from the three controlled baseline comparisons. At 77 GHz, the three controlled baselines correspond to approximately 16λ, 36λ, and 64.2λ, respectively.

5.6.1 Single-Target Scene

The single-target phase-sensitivity sweeps use a static target at $+15^\circ$ and 2.0 m radial range and a reference reflector at $[0, 4.2]$ m. Figure 5.5 shows the single target schematically at the same bearing and 1.5 m radial range.

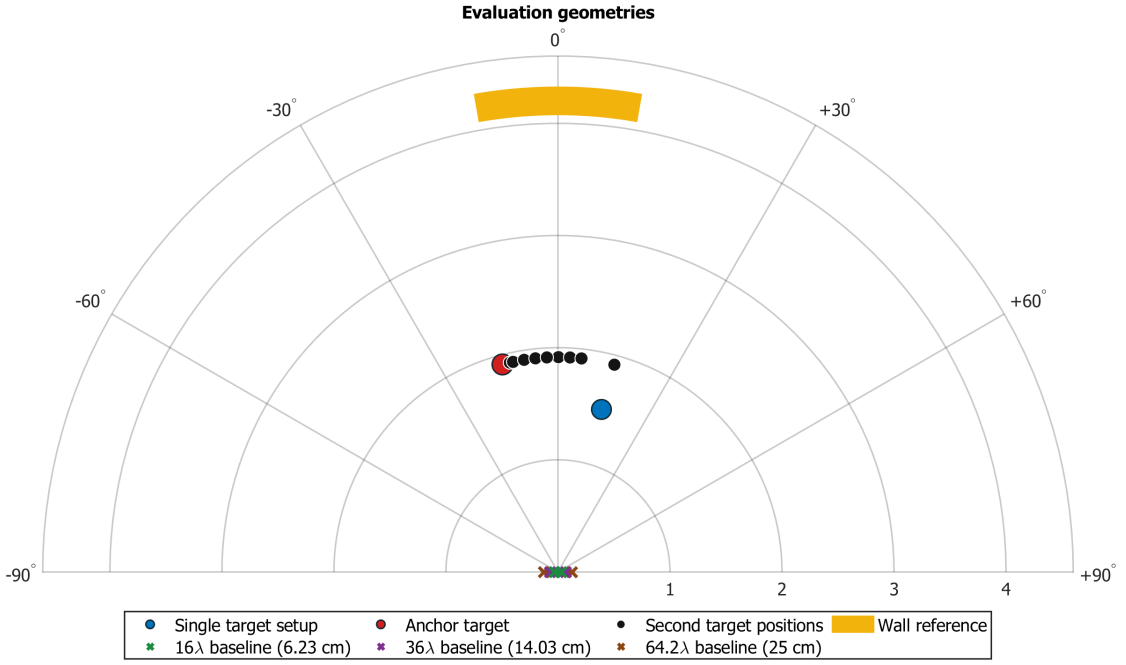


Figure 5.5: Evaluation geometries in range–angle coordinates. The single-target setup, fixed anchor target, translated second-target positions, and extended wall reference are shown together with the 16λ, 36λ, and 64.2λ distributed baselines at the origin.

5.6.2 Same-Cell Two-Target Scene

The same-range-cell study contains two static targets whose aperture-dependent path differences remain below one range-resolution cell. This means that range–Doppler processing is not intended to resolve the pair into two distinct detections; separation must instead arise from spatial information used by the DoA estimators. In the physical and simulated experiments, one reflector remains fixed at approximately $-15°$ and 1.85 m down-range, while a second reflector is translated laterally by 7, 10, 20, 30, 40, 50, 60, 70, and 100 cm, with approximately unchanged down-range position. This configured lateral physical separation is denoted by Δx; it identifies the repeatable geometry, while the relevant DoA-resolution coordinate is the corresponding expected angular separation $\Delta\theta_{\mathrm{nom}}$. Because range remains approximately fixed, the mapping is summarized in Table 5.2. Each standard physical geometry contains eight processed frames. The frame-eligibility conditions are applied before performance scoring.

Pairs below approximately 15° are therefore below the practical resolution benchmark of the local eight-element virtual azimuth array described in Section 4.1. Their use here tests whether distributed spatial information can separate a same-range pair beyond that local benchmark; centimetre spacing alone does not determine resolvability.

For a two-target scene, let the ordered nominal bearings be $\theta_1 < \theta_2$ and the ordered estimated bearings be $\widehat{\theta}_1 < \widehat{\theta}_2$. The nominal and estimated angular separations are

$$\Delta\theta_{\mathrm{nom}} = |\theta_2 - \theta_1|, \qquad \widehat{\Delta\theta} = |\widehat{\theta}_2 - \widehat{\theta}_1|, \tag{5.28}$$

Table 5.2: Configured lateral physical separation Δx and expected angular separation $\Delta\theta_{\text{nom}}$ for the same-range pair geometry at approximately 1.85 m down-range. The 15° value is the practical local-radar resolution benchmark.

Δx (cm)	$\Delta\theta_{\text{nom}}$ (°)
7	2.04
10	2.93
20	5.92
30	8.96
40	12.04
50	15.13
60	18.23
70	21.30
100	30.25

and the relative-separation error is

$$e_{\Delta\theta} = \left|\widehat{\Delta\theta} - \Delta\theta_{\text{nom}}\right|. \tag{5.29}$$

The same relative-resolution score is used for the synthetic and measured target-separation tests. A case or frame is successful when

$$K_{\text{est}} = 2 \quad \text{and} \quad e_{\Delta\theta} \leq \max(0.75^\circ,\, 0.20\,\Delta\theta_{\text{nom}}). \tag{5.30}$$

This relative-separation score is the primary resolution metric for both the synthetic and measured close-pair experiments. It is distinct from merely selecting two supports, since the spacing criterion must also be satisfied.

The separate absolute-angle success criterion requires $K_{\text{est}} = 2$ and a one-to-one assignment of the two selected bearings to the two nominal bearings, with both absolute angular errors not exceeding 2°. This criterion is reported separately and is not substituted for the relative-separation metric.

A correct estimate of $K_{\text{est}} = 2$ and a correct pair separation are necessary but do not by themselves establish correct absolute target angles. Absolute localization is assessed separately in the results; success under Eq. (5.30) must not be reported as proof of correct absolute DoA or Cartesian position.

5.7 Evaluation Protocols

5.7.1 Ideal Controls and Estimator Baselines

The ideal-control simulator removes additive noise and imposed synchronization errors while retaining finite sampling, the implemented FFT processing, exact path geometry, and the truth-blind model-order rules defined in Chapter 4. A single-target sanity control checks that the close-pair evaluation does not rely on a hard-coded target count. All five implemented DoA estimators are evaluated on the same scenes; theoretical MVDR,

MUSIC, ESPRIT, and maximum-likelihood approaches discussed in Chapter 2 are not additional experimental baselines.

5.7.2 Baseline Ambiguity Comparison

The ideal ambiguity comparison uses the exact-path steering responses at the controlled baselines. Non-coherent power fusion and coherent stacking are evaluated using comparable normalized metrics. The aim is to isolate the angular discrimination and periodic ambiguity introduced by the baseline, without treating a narrow peak as proof of correct absolute localization.

For aperture q, the normalized power ambiguity relative to reference angle θ_0 is

$$\mathrm{AF}_q(\theta;\theta_0) = \frac{|\mathbf{a}_q^H(\theta_0)\mathbf{a}_q(\theta)|^2}{\|\mathbf{a}_q(\theta_0)\|_2^2\,\|\mathbf{a}_q(\theta)\|_2^2}, \tag{5.31}$$

where $\mathbf{a}_q(\theta)$ denotes the exact-path steering vector of aperture q. The non-coherent comparison averages the aperture powers,

$$\mathrm{AF}_{\mathrm{NC}}(\theta;\theta_0) = \frac{1}{4}\sum_{q=1}^{4}\mathrm{AF}_q(\theta;\theta_0), \tag{5.32}$$

whereas the coherent ambiguity is computed after stacking the four absolute-phase steering blocks into $\mathbf{a}_{\mathrm{C}}(\theta)$,

$$\mathrm{AF}_{\mathrm{C}}(\theta;\theta_0) = \frac{|\mathbf{a}_{\mathrm{C}}^H(\theta_0)\mathbf{a}_{\mathrm{C}}(\theta)|^2}{\|\mathbf{a}_{\mathrm{C}}(\theta_0)\|_2^2\,\|\mathbf{a}_{\mathrm{C}}(\theta)\|_2^2}. \tag{5.33}$$

The curves use the same 1.85 m down-range geometry and a reference direction of $-15°$.

5.7.3 Residual Phase-Sensitivity Protocol

Residual phase error is decomposed into intra-aperture, inter-aperture, and within-CPI components because the three perturbations enter the processing chain differently, as modeled in Section 3.3. The intra-aperture metric is the circular standard deviation of the geometry-removed element residual after one common phase is removed independently from each transmitter group. The inter-aperture metric is the circular spread of the geometry-removed block projections after the simulated external reference-target compensation used in these synthetic ablations. Within-CPI error is described by the residual phase history over chirp index after the nominal target motion and synchronization model have been removed.

For each baseline, the ideal IF cubes are processed through the complete synchronization and detection chain before zero-mean residual phase perturbations are imposed at root-mean-square (RMS) levels of 0°, 2°, 5°, 10°, 20°, 30°, 45°, 60°, and 90°. Each nonzero level contains $N_{\mathrm{MC}} = 20$ Monte Carlo realizations. For the intra-aperture sweep, zero-mean random receive-channel phase errors are drawn and shared between the two transmit groups that use the same physical receive chain. For the inter-aperture sweep,

one random common phase rotation is applied to each aperture block relative to the reference block. Each realization is scaled to the requested RMS value. The same normalized random draw is reused at successive RMS levels, so performance changes can be attributed primarily to the perturbation magnitude rather than to unrelated sampling variation. The unperturbed control is evaluated once because it is deterministic.

The same perturbation schedule is applied to the single-target scene and to the selected close-pair scenes. The two-target results are evaluated using Eq. (5.30); a separate reported target-count outcome is retained even when relative separation fails. Results of these controlled ablations are presented in Section 6.3.

5.7.4 Within-CPI Phase-Evolution Study

The within-CPI study isolates phase evolution remaining after synchronization. The generated data first pass through the receive-side frequency and slow-time phase corrections. An additional residual is then applied to the corrected, range-resolved bistatic data before the Doppler FFT. The synchronization stages are not rerun after this injection. This distinguishes the experiment from the raw oscillator-error model in Eqs. (5.11)–(5.15), whose impairments are present before correction.

For same-transmitter slow-time index $n = 0, \ldots, N-1$, the imposed relative phase is

$$\psi[n] = 2\pi \Delta f_{\text{res}} n \Delta t + W[n], \qquad W[0] = 0, \tag{5.34}$$

where Δf_{res} is a constant residual frequency offset and Δt is the interval between successive chirps from the same transmitter. The random walk has independent Gaussian increments,

$$W[n+1] - W[n] \sim \mathcal{N}(0, q\Delta t), \tag{5.35}$$

with phase-diffusion coefficient q in rad^2/s. The frequency-offset and random-walk components are evaluated separately. A constant frequency offset produces a linear phase progression; it is not an increasing frequency drift or a static block-phase offset.

Writing the corrected vector of virtual-channel samples at range bin r as $\mathbf{z}_p[r, n]$, the perturbed data are

$$\widetilde{\mathbf{z}}_p[r, n] = \exp(\mathrm{j} s_p \psi[n])\, \mathbf{z}_p[r, n], \tag{5.36}$$

where $s_{M_1M_2} = +1$, $s_{M_2M_1} = -1$, and $s_p = 0$ for the two monostatic apertures. Thus the reciprocal bistatic paths receive opposite phase histories, common to their spatial channels, while the monostatic paths remain unchanged in this slow-node-phase approximation. The same factor is applied at every range bin and affects both the scene target and the reference response before Doppler processing.

The additional phase is constant within each chirp. Fast-time phase evolution and changes between the two interleaved transmitter slots are not modeled. The waveform and CPI timing are held fixed: the experiment varies the imposed post-correction residuals rather than bandwidth, chirp duration, or CPI length. It is a controlled slow-time

sensitivity test, not a complete oscillator or phase-locked-loop model. A single-target scene is used to distinguish temporal coherence effects from two-target separation, and the five DoA methods are assessed through the reported peak angle and inferred source count.

The two-node baseline is $16\lambda \approx 6.23\,\mathrm{cm}$ at $77\,\mathrm{GHz}$. The scene target has bearing -15° and down-range coordinate $y = 1.85\,\mathrm{m}$, corresponding to approximately $1.92\,\mathrm{m}$ radial distance from the origin. A separate static wall response at approximately $4.2\,\mathrm{m}$ supplies the phase reference. The nominal acquisition SNR is $30\,\mathrm{dB}$; additional clutter and preset phase perturbations are disabled. With two interleaved transmitters, 128 raw chirps provide $N = 64$ samples per transmitter at $\Delta t = 320\,\mu\mathrm{s}$. The first-to-last sample span is $T_{\mathrm{span}} = (N-1)\Delta t = 20.16\,\mathrm{ms}$, while the nominal Doppler observation period is $N\Delta t = 20.48\,\mathrm{ms}$.

The frequency offsets are $\Delta f_{\mathrm{res}} \in \{0, 5, 10, 25, 50, 100\}\,\mathrm{Hz}$. The separate random-walk sweep uses terminal phase standard deviations $\sigma_W \in \{0, 30, 60, 120, 180, 360\}^\circ$ at T_{span}, with $q = \sigma_W^2 / T_{\mathrm{span}}$ after conversion to radians. The shared zero-error control gives eleven distinct conditions, each with ten Monte Carlo realizations. Conditions within a realization reuse the same simulated capture. An additional moving-target control uses radial velocity $1\,\mathrm{m/s}$ away from the origin, with three realizations for zero error, $25\,\mathrm{Hz}$, $100\,\mathrm{Hz}$, and a 180° random walk. There are therefore 110 static and 12 moving cases. The large walk values are stress conditions, not residual levels measured in the hardware.

Each case is evaluated with the four paired protocols in Table 5.3. Fixed-cell readout retains the range and Doppler cells selected in the corresponding unperturbed control; it is a diagnostic readout, not a claim of renewed detection. Peak-tracked readout uses the current detected peak in each aperture. Detection selects one strongest CFAR anchor within the 1.5–2.2 m target range gate without a static-velocity restriction. The coarse angular candidate sector has a 40° half-width around the centroid in every condition. Coherent refinement retains the existing Block-FOCUSS gate, with settings 6° and $0.12\,\mathrm{m}$, rather than searching the entire coarse sector.

Frozen calibration retains the wall-reference snapshot from the paired unperturbed capture; it does not represent an independently measured physical calibration. Same-CPI calibration instead uses the affected wall snapshot from the current CPI, extracted at zero Doppler. Both supply one block-phase correction per aperture rather than a new chirp-by-chirp synchronization pass.

The strongest-peak angular RMSE for each of the five DoA methods uses every available reported angle, including cases with an incorrect source count, without conditioning on a successful truth association. The fraction reporting exactly one target is evaluated separately over all trials. Missing angles are excluded from the RMSE; no-estimate cases remain in the denominator of the source-count outcome. The four protocols give 2440 evaluations for the five DoA methods across the 122 cases. Target truth is used for scoring

Table 5.3: Paired RD-readout and reference-calibration protocols for the residual-CPI study.

Protocol	Target readout	Reference calibration
1	Fixed control cells	Frozen control reference
2	Fixed control cells	Same-CPI reference
3	Current detected peaks	Frozen control reference
4	Current detected peaks	Same-CPI reference

only; these single-target results do not establish close-pair resolution.

5.7.5 Physical Frame Evaluation and Phase Calibration

The physical close-pair study processes each eligible frame separately, retaining the observed source count and selected angular supports. Pair separation and absolute-angle accuracy are assessed as distinct outputs. The range-resolved multiple-measurement-vector (MMV) representation introduced in Chapter 4 preserves angular information across neighboring range bins within the same associated geometric cluster. The relevant frame-selection and acquisition settings are documented in Section A.1 and the subsequent appendix configuration tables.

The original 25 cm-baseline target-cluster phase adjustment uses the evaluated scene itself as reference and is therefore in-sample. Later calibration with a separate spherical reflector instead probes reference transfer to the scene. The two calibration modes must be identified separately when interpreting physical phase residuals and localization performance.

5.7.6 Backpropagation and Measurement-Matched Validation

Cartesian backpropagation is assessed through ideal and physical localization examples after the coarse non-coherent hypotheses are available. Coherent refinement and joint two-target fitting use the implementation of Section 4.5; their output is not substituted for the truth-blind same-cell DoA model-order benchmark.

The measurement-matched simulator is also compared with physical static and moving scenes after applying the same range–Doppler processing. The comparisons assess the transfer of background statistics and localization behavior separately; matching a range–Doppler map does not imply matching source-count or absolute-localization performance.

5.8 Ground Truth, Model Scope, and Evaluation Limits

For each generated capture, the simulator stores the exact target positions, velocities, radar geometry, channel gains, inter-radar frequency history, and aperture-level phase disturbances. These quantities provide complete ground truth for the synchronization and localization experiments presented in the following chapters.

For an aperture centered at $\mathbf{c}_p$, the true local angle of target k is defined as

$$\theta_{p,k,\mathrm{true}} = \operatorname{atan2}\left(x_{k,\mathrm{true}} - c_{p,x}, y_{k,\mathrm{true}} - c_{p,y}\right). \tag{5.37}$$

The stored target positions and local angles are later used to evaluate Cartesian and angular estimation performance using the metrics defined in Section 5.6. The simulator is a controlled signal-level model rather than a complete electromagnetic environment. Targets are represented as point scatterers, propagation is generated at dechirped IF level, and target motion follows the stop-and-hop approximation. Nevertheless, the model preserves the mechanisms most relevant to this thesis: exact monostatic and bistatic paths, physical Doppler, linked frequency and phase errors, TDMA timing, channel imbalance, coupling, exact-path image multipath, structured clutter, and correlated receiver noise.

Because the resulting files use the same format and processing chain as the measurements, the simulator provides a reproducible basis for evaluating the synchronization and localization performance under controlled circumstances.

5.9 Evaluation Methodology Summary

This chapter defined the synthetic IF-level radar model and the physical evaluation protocol used to assess the synchronization and localization methods. The simulator models two radar nodes and all four monostatic and bistatic apertures using exact transmitter–target–receiver path lengths. These paths determine the fast-time beat frequency, carrier propagation phase, near-field spatial response, and physical Doppler of moving targets.

The inter-radar frequency mismatch is represented by a nominal offset, linear drift, and correlated jitter. The corresponding oscillator phase is obtained by integrating the same frequency history, ensuring that range migration and slow-time phase distortion remain physically consistent. Additional nonidealities, including initial block-phase offsets, receive-channel errors, direct coupling, exact-path image multipath, structured clutter, and correlated complex receiver noise, can be enabled independently.

Together, these controls support the mechanism-resolved comparison central to the thesis: the evaluation can distinguish local-manifold distortion, static cross-aperture disagreement, and phase evolution over coherent integration time before their effects are compared across estimator assumptions.

The generated radar cubes use the same format as the measured DCA1000 data and are processed by the same synchronization, detection, and localization pipeline. This enables controlled validation of whether clock-induced range and phase errors are removed while physical target Doppler and spatial phase are preserved. Since all target and error parameters are stored as ground truth, the simulator provides a reproducible basis for evaluating angular and Cartesian estimation performance under progressively more realistic conditions. A calibration/held-out split prevents the measurement-matched nuisance preset from being evaluated only on the physical geometries used to tune it, while the multi-frame capture structure separates within-capture repeatability from between-capture variation.

The physical setup, same-cell test geometries, model-order and separation criteria, and phase-sensitivity protocol are specified in the same chapter, so that Chapter 6 can focus on observed estimator behavior and its interpretation rather than introducing the experiments.

Chapter 6

Results and Discussion

This chapter presents the controlled-simulation and physical-measurement findings of the two-node FMCW radar study. The evaluation geometry, performance criteria, calibration protocols, and truth-blind support-selection rules are defined in Chapters 5 and 4. Here, the emphasis is on what the implemented estimators actually recover and on whether the assumptions required for coherent localization are supported by the measured phase quality.

The results progress from ideal same-cell angular resolution and baseline-dependent ambiguity to residual phase-error sensitivity, physical close-pair measurements, Cartesian backpropagation, and measurement-matched simulation. Throughout the chapter, selecting two sources, recovering their relative angular separation, and identifying their correct absolute angles or positions are treated as different outcomes. Target positions and the true source count are used to score estimates retrospectively; they are not supplied to the estimators when selecting support or model order.

6.1 Ideal Angular Resolution of the Distributed Aperture

The ideal controls use the processing and scoring rules described in Section 5.7.1: additive noise and imposed synchronization errors are absent, while finite sampling, FFT processing, exact-path geometry, and truth-blind model-order selection remain. The single-source control confirms that two-target results are not obtained by supplying a fixed source count. Figures 6.1–6.3 show the ideal same-cell outcomes separately for the three tested radar baselines: 16λ (0.0623 m), 36λ (0.1403 m), and 64.2λ (0.25 m). In each cell, the plotted annotation gives the estimated angular separation in degrees when the estimator returns $K_{\mathrm{est}} = 2$. Blue crosses indicate $K_{\mathrm{est}} \neq 2$, orange triangles indicate $K_{\mathrm{est}} = 2$ without satisfying Eq. (5.30), and green circles indicate successful relative-separation recovery.

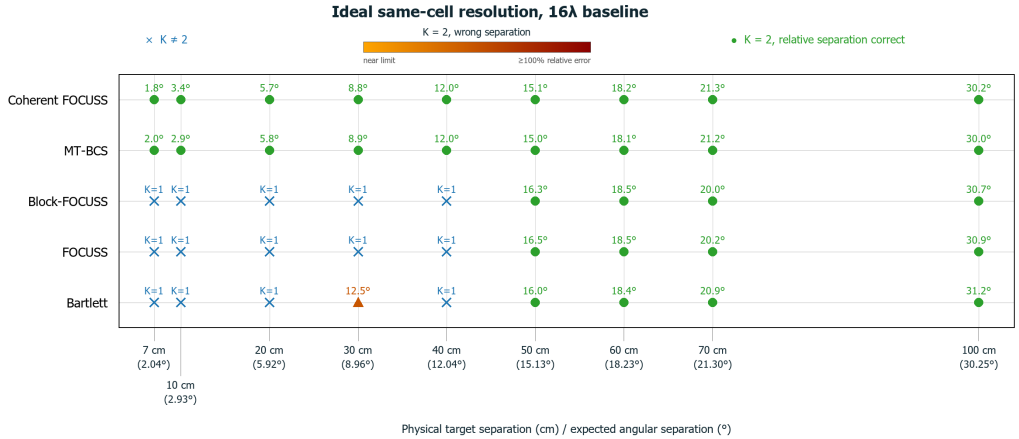


Figure 6.1: Ideal same-cell model-order and estimated angular separation for the 16λ (0.0623 m) radar baseline. Numeric annotations show $\widehat{\Delta\theta}$ in degrees when $K_{\text{est}} = 2$.

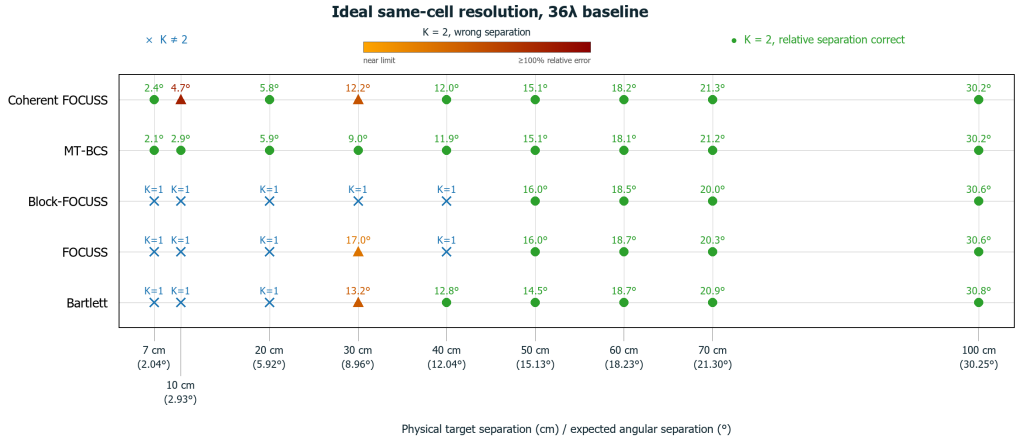


Figure 6.2: Ideal same-cell model-order and estimated angular separation for the 36λ (0.1403 m) radar baseline. Numeric annotations show $\widehat{\Delta\theta}$ in degrees when $K_{\text{est}} = 2$.

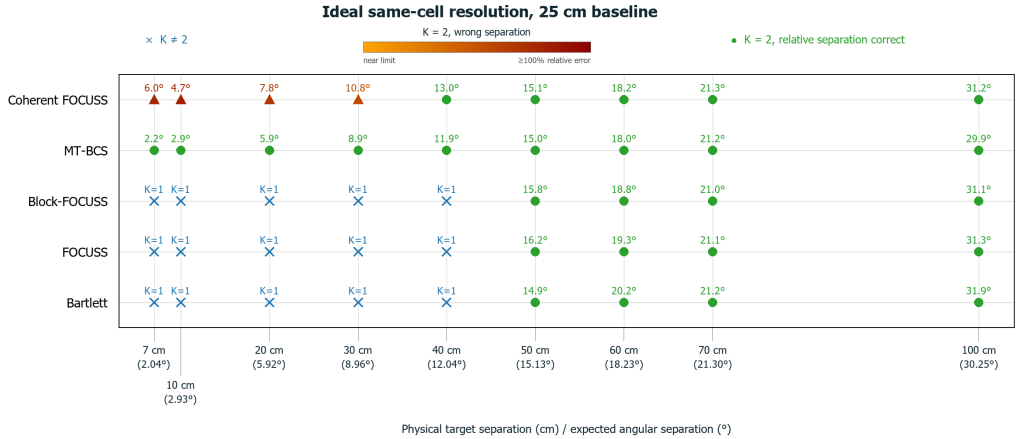


Figure 6.3: Ideal same-cell model-order and estimated angular separation for the 64.2λ (0.25 m) radar baseline. Numeric annotations show $\widehat{\Delta\theta}$ in degrees when $K_{\text{est}} = 2$.

Across the three baselines, MT-BCS returns $K_{\text{est}} = 2$ for every ideal pair, with estimated angular separations close to the nominal values. Coherent FOCUSS also reports two supports throughout the ideal close-pair set, although its estimated separations deviate more at some small target spacings and longer baselines. Bartlett, independent FOCUSS, and Block-FOCUSS begin returning two supports only at larger separations. This shows that correct source-count inference and correct pair spacing are not interchangeable, even without imposed phase errors. The baseline-dependent ambiguity behind some coherent deviations is examined next.

6.2 Baseline-Dependent Coherent and Non-Coherent Ambiguity

Figures 6.4 and 6.5 compare coherent and non-coherent exact-path ambiguity responses. The first contrasts fusion modes at the 25 cm baseline; the second illustrates how the coherent fringe pattern varies with baseline.

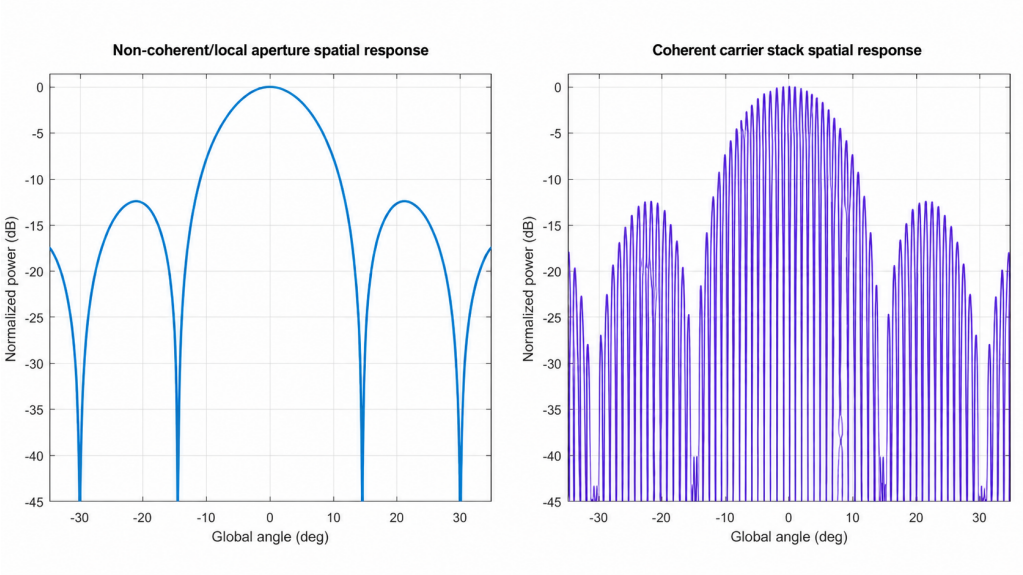


Figure 6.4: Coherent and non-coherent exact-path spatial responses for the 25 cm baseline.

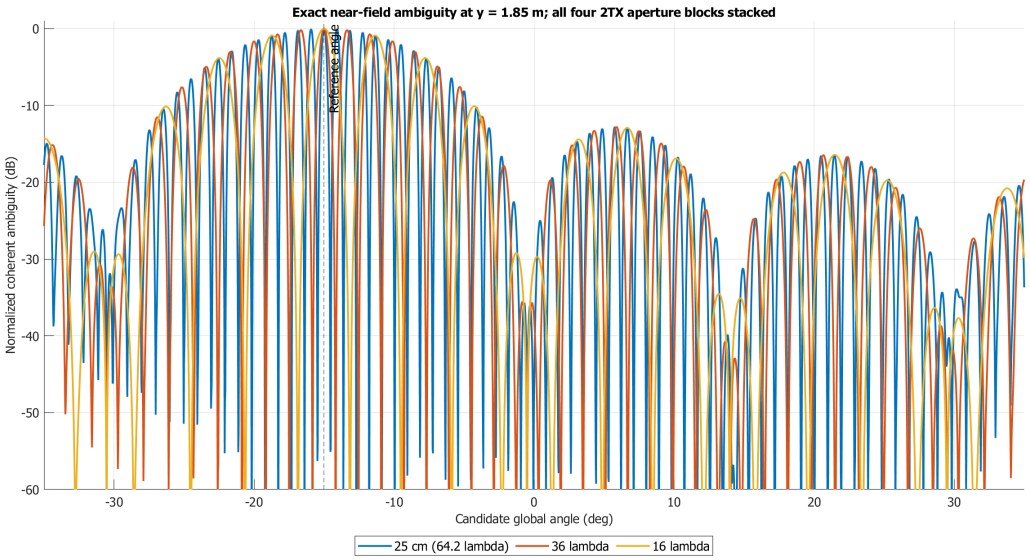


Figure 6.5: Coherent ambiguity functions for the three evaluated baselines. Increasing baseline narrows the local response while producing denser spatial ambiguity fringes.

The exact-path, normalized coherent and non-coherent ambiguity functions are defined in Eqs. (5.31)–(5.33) in Section 5.7.2. The non-coherent response remains broad in these comparisons, whereas coherent stacking produces narrower local maxima accompanied by increasingly dense periodic fringes as the baseline grows. This is a conditional resolution gain: a narrow peak is not proof that the global maximum identifies the correct absolute target position.

6.3 Sensitivity to Chirp-Invariant Residual Phase Errors

The isolated intra-aperture and inter-aperture phase perturbations follow the Monte Carlo protocol specified in Section 5.7.3; their error models are defined in Section 3.3. The separate within-CPI branch applies post-correction chirp-dependent phase evolution after range processing and before the Doppler FFT. The static-sweep results distinguish angular error conditional on a matched estimate from the success of source-count or pair-separation recovery. In particular, an apparently small conditional angle error can coexist with a low probability of producing a valid estimate.

6.3.1 Single-Target Phase Error Sensitivity

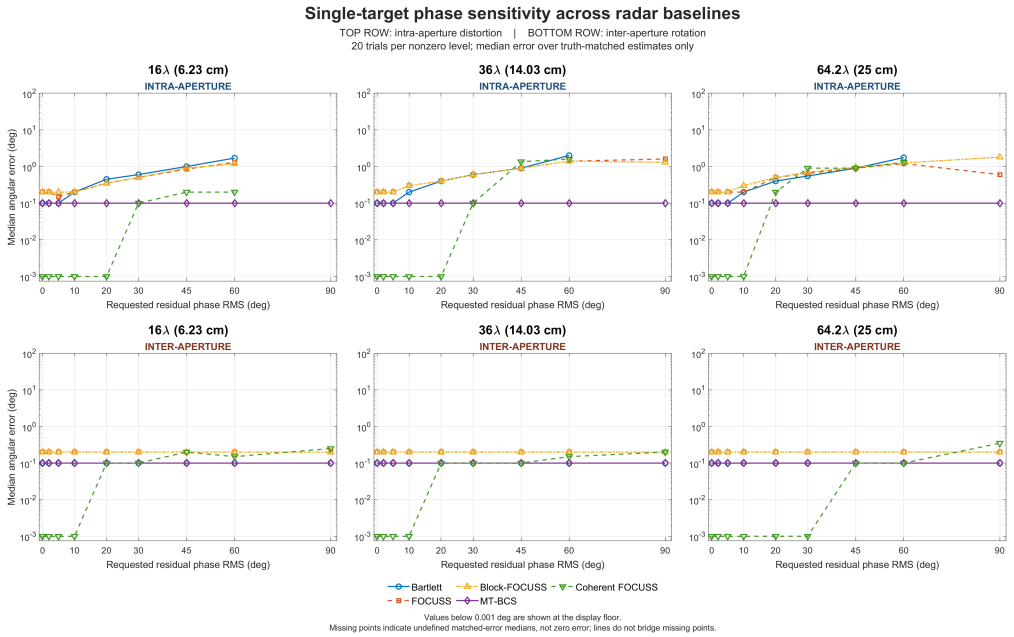


Figure 6.6: Median single-target angular error of the five DoA estimators for the 16λ (0.0623 m), 36λ (0.1403 m), and 64.2λ (0.25 m) radar baselines under isolated residual intra-aperture and inter-aperture phase perturbations.

Across all three baselines, MT-BCS retains its 0.1° median matched-support error and a 100% correct-K rate through 90° of inter-aperture rotation. Coherent FOCUSS still matches the target but retains the correct one-source model in only 35%, 40%, and 40% of the corresponding trials. Figure A.4 shows the source-count rates at each baseline, complementing the matched-support angular errors in Figure 6.6. Under 60° intra-aperture distortion, correct matched one-target recovery is at most 10% for any plotted method and baseline, and it is zero throughout at 90°. The apparently low MT-BCS angular error at severe intra-aperture levels is therefore a surviving-match statistic that coexists with misses and false splits, rather than evidence of robustness.

6.3.2 Two-Target Pair Phase Error Sensitivity

Figures 6.7 and 6.8 apply the same relative-separation success criterion from Equation (5.30) to the 20 cm and 50 cm pairs, respectively.

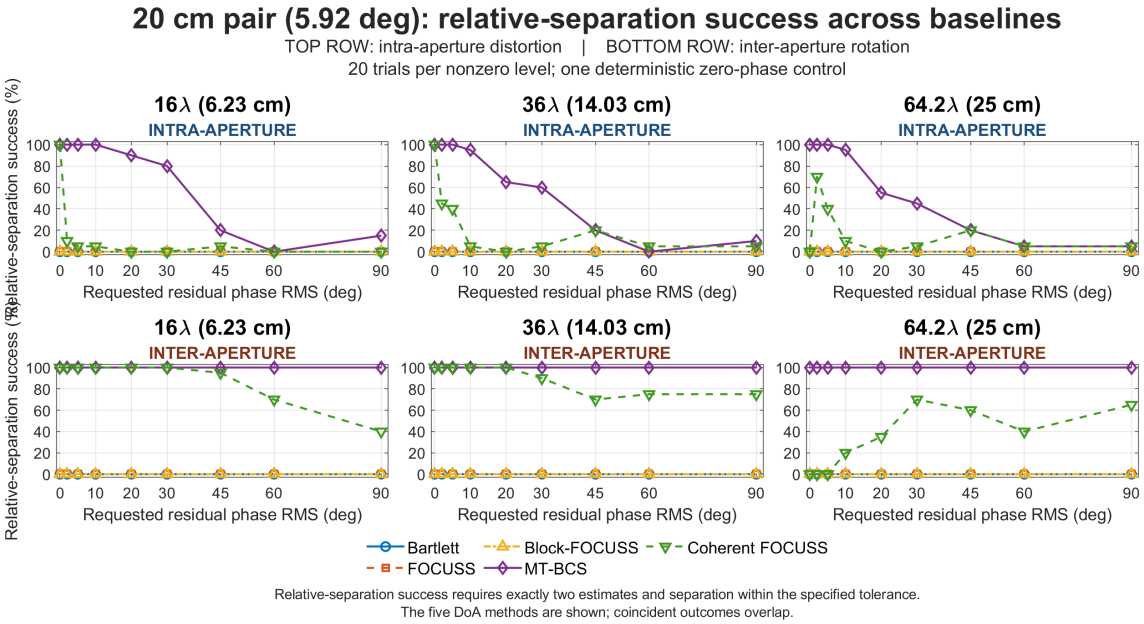


Figure 6.7: Relative-separation success rates of Bartlett, independent FOCUSS, Block-FOCUSS, MT-BCS, and coherent FOCUSS for the 20 cm (5.92°) pair at the three tested baselines under isolated residual intra-aperture and inter-aperture phase perturbations. Success requires two estimates and an estimated separation within the tolerance of Equation (5.30).

Block-FOCUSS does not resolve the 5.92° pair in the deterministic control, so its zero rate throughout the sweep reflects insufficient ideal resolving power rather than demonstrated phase robustness. MT-BCS retains at least 80% relative-separation success under intra-aperture perturbation through 30°, 10°, and 10° for the 16λ, 36λ, and 0.25 m baselines, respectively. Under inter-aperture rotation, MT-BCS remains at 100% through 90° for all three baselines. Coherent FOCUSS reaches 70%, 75%, and 40% at 60° and 40%, 75%, and 65% at 90° for the same baseline order. Its non-monotone baseline-dependent rates reflect phase-wrapped coherent ambiguity rather than one transferable scalar threshold.

At 50 cm, Bartlett and MT-BCS retain high intra-aperture relative-separation success over a substantially wider phase range than independent FOCUSS, Block-FOCUSS, and coherent FOCUSS. Inter-aperture rotation leaves Bartlett, independent FOCUSS, Block-FOCUSS, and MT-BCS at 100% throughout the sweep, whereas coherent FOCUSS decreases slightly only at the largest rotations.

6.4 Sensitivity to Phase Evolution Within the CPI

The post-correction residual model of Section 5.7.4 extends the static phase sweeps by introducing a constant residual frequency offset or a random phase walk before the Doppler FFT. The results below compare Bartlett, independent FOCUSS, Block-FOCUSS, MT-BCS, and coherent FOCUSS for the stationary single target at the 16λ baseline, with ten

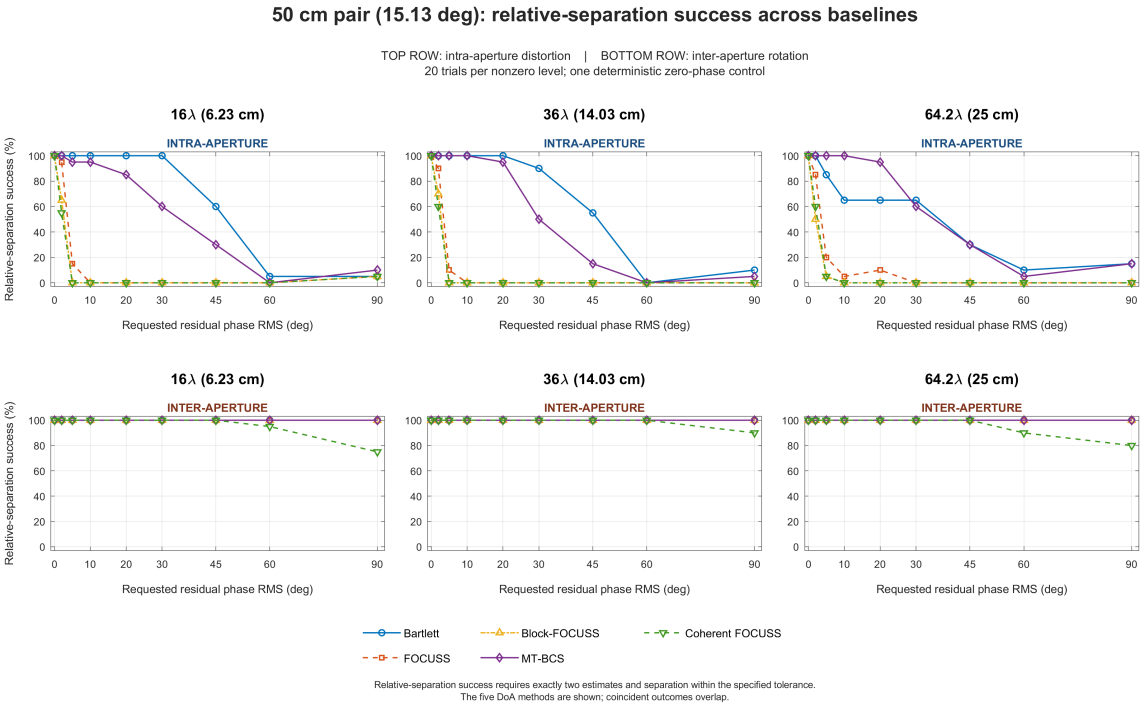


Figure 6.8: Relative-separation success rates for the 50 cm (15.13°) pair across the three baselines under isolated intra-aperture and inter-aperture phase perturbations. The success criterion is identical to that used for the 20 cm pair.

realizations per condition and fixed waveform settings. Angular RMSE of the strongest reported peak and the fraction reporting exactly one target are evaluated separately. The paired RD-readout and reference-calibration protocols are defined in Table 5.3.

Figure 6.9 illustrates the two residual mechanisms for one representative realization. The numerical means reported below are evaluated over all ten realizations.

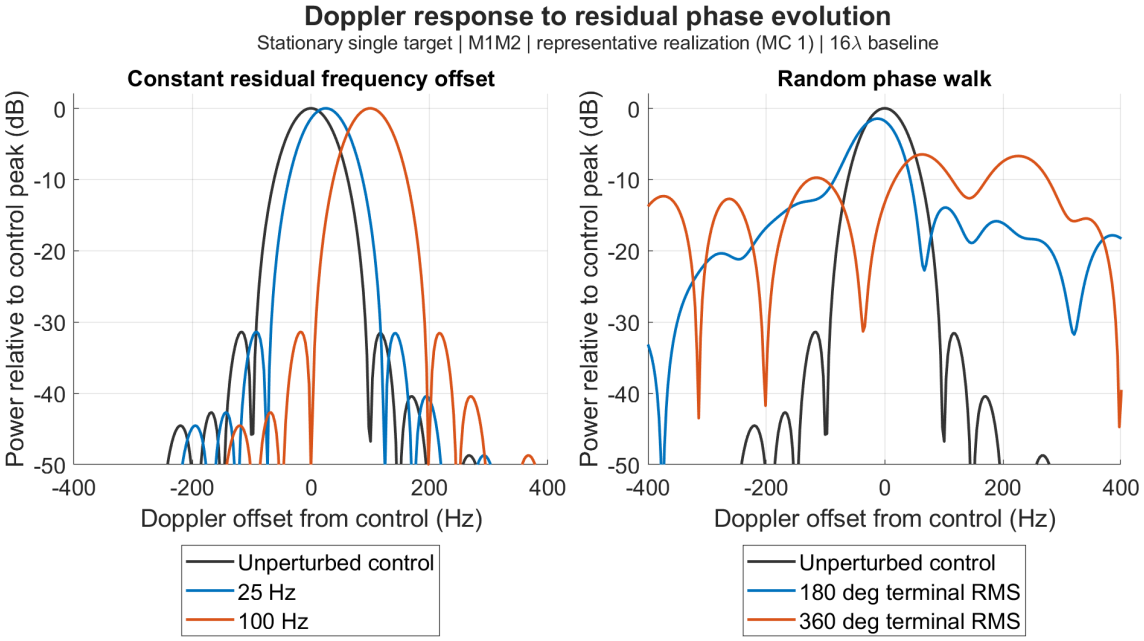


Figure 6.9: Representative stationary-target Doppler spectra in aperture M_1M_2 at the 16λ baseline. The left panel varies residual frequency offset; the right panel varies random-walk terminal RMS. Powers are averaged over virtual channels and normalized to the same unperturbed control peak, before block-phase calibration. Curves show one Monte Carlo realization rather than an average across realizations.

Figure 6.10 summarizes the corresponding angle and source-count outcomes of the five DoA methods.

With peak-tracked readout, Bartlett, independent FOCUSS, and Block-FOCUSS report one target in every stationary trial, with angular RMSE between approximately 0.11° and 0.15°. Their curves overlap in this single-target scene. MT-BCS retains small strongest-peak errors, but its one-target rate falls to 10% at 100 Hz and to 80% at the 360° phase-walk condition. These failures concern additional selected supports rather than missing angular estimates. For these four methods, frozen and same-CPI reference calibration give identical outcomes.

A constant frequency offset mainly displaces the Doppler response. At 100 Hz, the mean M_1M_2 peak displacement is 100.7 Hz and its mean peak-power change is approximately 0 dB relative to the unperturbed control. With peak-tracked readout and frozen reference calibration, coherent FOCUSS retains angular RMSE below 0.05° and reports one target in all ten trials at every tested offset. Thus, a residual linear phase progression does not necessarily cause a large angle error when the displaced Doppler peak is selected.

For coherent FOCUSS, the same-CPI reference gives a different outcome. With peak-tracked readout, angular RMSE reaches 1.85° at 50 Hz and 1.82° at 100 Hz. The reference snapshot is extracted at zero Doppler while the selected target peak can move to another Doppler cell; the two snapshots therefore need not retain a compatible integrated phase. At 25 Hz, the strongest-peak RMSE is only 0.010°, yet all ten trials report two targets. An accurate dominant peak does not ensure correct source count.

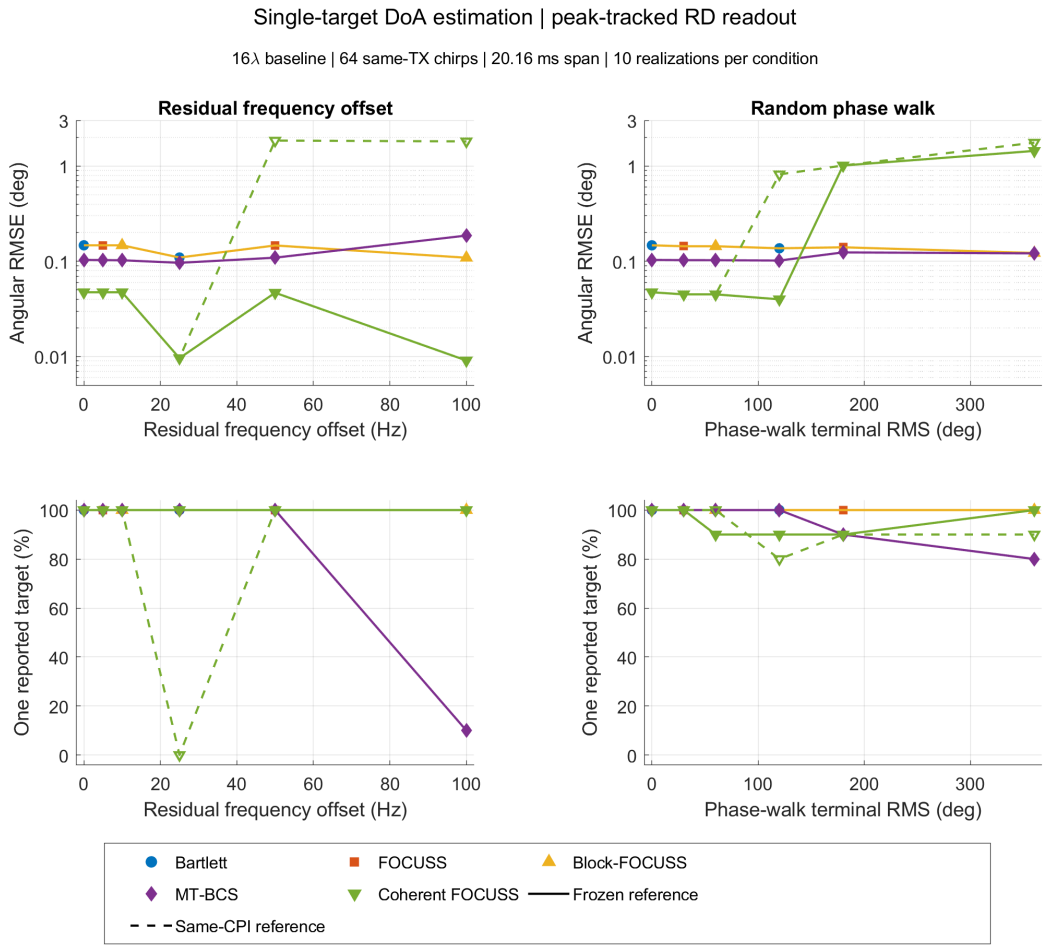


Figure 6.10: Peak-tracked RD readout for the stationary single target at the 16λ baseline. Upper panels show strongest-peak angular RMSE on a logarithmic scale; lower panels show the fraction reporting exactly one target. Colors identify the five DoA methods; solid and dashed lines distinguish frozen control-reference and same-CPI calibration. Each condition contains ten realizations. The reference protocols coincide for the four non-coherent or shared-support methods; Bartlett, independent FOCUSS, and Block-FOCUSS also have coincident outcomes in this scene. RMSE includes trials with incorrect source count.

Strong random-walk perturbations produce Doppler spreading and a reduction in peak power. For M_1M_2, the 360° terminal-RMS condition gives a mean peak-power change of $-5.5\,\mathrm{dB}$ relative to the control. With peak-tracked readout and frozen calibration, coherent FOCUSS angular RMSE increases to 1.01° at 180° and 1.44° at 360°. With same-CPI calibration, its angular RMSE is 1.76° at 360°. Conversely, all frozen-reference coherent-FOCUSS trials at 360° still report exactly one target despite the larger angle error. Source count and angular accuracy therefore describe different aspects of temporal-phase robustness. The non-monotone source-count rates should be interpreted as outcomes of this scene and finite Monte Carlo sample, rather than universal operating limits.

These imposed post-correction offsets should be distinguished from the coarse receive-side synchronization error. In the separate controlled synchronization ablations, the Vecchi-inspired coarse-correction stage estimates the imposed 5 MHz bistatic frequency

offset with a median absolute error between 3.7 and 4.9 kHz. The residual in the perfect condition remains of the order of one fast-time spectral-grid interval; those values are not the residual-frequency levels swept in the present CPI experiment.

Supplementary fixed-cell results and the moving-target control are presented in Appendix A.5 (Figures A.1 and A.2). The moving control uses three realizations per condition and serves as a consistency check.

For an isolated return, the imposed phase history changes the complex slow-time integration factor at the selected Doppler cell. Since the perturbation is common to the virtual channels within each aperture, it does not directly impose an element-dependent intra-aperture distortion. It can nevertheless change the relative phase and amplitude of the aperture snapshots. A linear progression can largely be followed as a Doppler shift; a random walk cannot generally be removed by selecting one displaced cell. In the fixed-cell stationary test, the target and reference signal components sample the same imposed temporal factor near zero Doppler, allowing its phase to cancel during same-CPI calibration. This cancellation is not assured when the target cell is tracked but the reference remains at zero Doppler. The calibration applies a block rotation rather than correcting the full phase history. Subsequent snapshot normalization likewise does not restore lost coherent gain or SNR.

This distinction is consistent with the spectral effects of oscillator phase noise discussed by Siddiq *et al.* [10] and the separation of systematic and stochastic phase distortions emphasized by Tschapek *et al.* [22]. Here, these effects are traced through RD-cell selection, reference calibration, and the five DoA estimators. This controlled post-correction study does not fit a hardware oscillator spectrum or reproduce the delay-dependent fast-time phase-noise mechanisms examined in those works.

6.5 Physical Close-Pair Measurements and Phase Quality

The close-pair measurement protocol and frame-scoring definitions are specified in Sections 5.6 and 5.7.5. The primary experimental question is whether the estimators distinguish two targets occupying one range–Doppler neighbourhood when the local array does not reliably resolve them. The observations are therefore reported separately as source-count recovery, relative pair spacing, and absolute-angle success with a 2° tolerance.

6.5.1 Target-Count and Relative-Separation Results

Figures 6.11–6.13 summarize the eligible frame-level model-order and relative-spacing outcomes at three physical baselines. The results obtained with radar baseline distances of 14.03 and 25 cm form the principal physical comparisons; the additional acquisition with a 19.5 cm radar baseline is displayed separately rather than treated as a third matched Monte Carlo baseline. Labels in the plots give the median estimated separation over eligible frames returning $K_{\mathrm{est}} = 2$. Colors and symbols distinguish failure to select two sup-

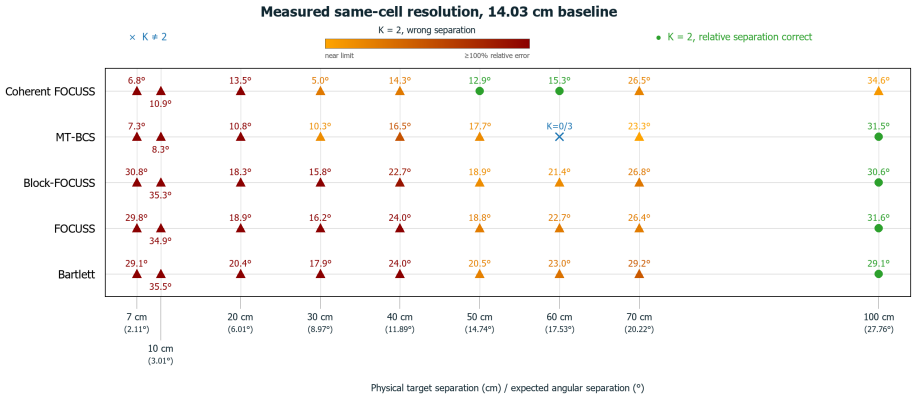


Figure 6.11: Measured same-cell model-order outcomes and annotated median estimated angular separation for the 14.03 cm baseline without inter-aperture phase compensation. Numeric annotations show the median $\widehat{\Delta\theta}$ over eligible frames with $K_{\text{est}} = 2$.

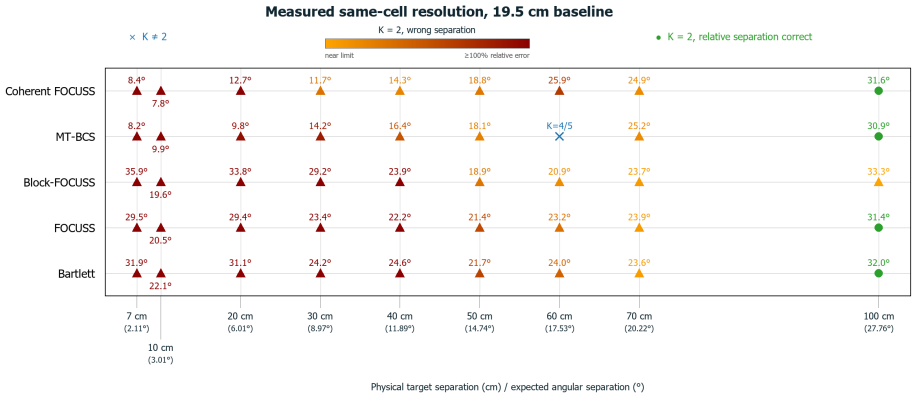


Figure 6.12: Measured same-cell model-order outcomes and annotated median estimated angular separation for the 19.5 cm baseline without inter-aperture phase compensation. Numeric annotations show the median $\widehat{\Delta\theta}$ over eligible frames with $K_{\text{est}} = 2$.

ports, selection of two supports with incorrect spacing, and successful relative-separation recovery under Eq. (5.30). This criterion does not establish correct absolute target angles.

Across the two later baselines, the frame-level outcomes again show that selecting $K_{\text{est}} = 2$ is more common than recovering the correct relative spacing. Aggregated relative-separation success across the nine tested separations remains at or below 15% for every method. These campaigns extend the qualitative observation across baselines without establishing absolute-angle accuracy.

In the 25 cm-baseline close-pair study, the distinction between these outcomes is consequential. For the 7 cm (2.04°) pair, MT-BCS and coherent FOCUSS select two supports in all eight frames but do not recover the required relative separation. At 30 cm (8.96°), several methods select two sources without satisfying the relative-spacing criterion. At 60 cm (18.23°), the relative-separation success rates are 100% for Bartlett, independent FOCUSS, and Block-FOCUSS, 62.5% for MT-BCS, and 50% for coherent FOCUSS. At 70 cm (21.30°) the corresponding rates are 100%, 100%, 100%, 0%, and 75%; at 100 cm

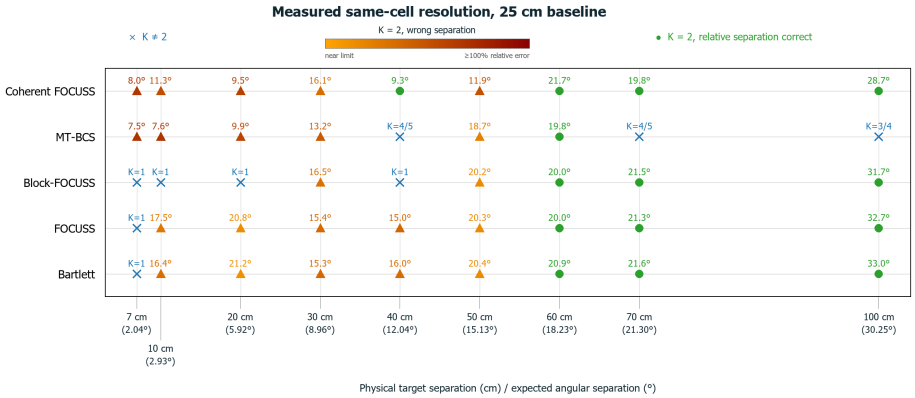


Figure 6.13: Measured same-cell model-order outcomes and annotated median estimated angular separation for the 25 cm baseline. Numeric annotations show the median $\widehat{\Delta\theta}$ over eligible frames with $K_{\text{est}} = 2$.

(30.25°) they are 100%, 87.5%, 87.5%, 0%, and 87.5%. These figures describe relative separation, not correct absolute positions. Across the 72 evaluated frames per method, none satisfies the separate criterion of two selected sources with both assigned bearing errors within 2° of the nominal bearings. Among frames reporting exactly two sources, angular RMSE relative to the nominal geometry ranges from 3.80° to 7.78° across the five DoA methods (Table A.4).

For the measured captures, Figures 6.14 and 6.15 show representative range-resolved DoA score curves at 7, 30, 60, and 100 cm. Splitting the four cases over two figures makes the individual spectra and selected supports readable while retaining the progression from unresolved close pairs to wider, but sometimes misidentified, supports.

6.5.2 Residual Phase Quality in the Measured Scenes

The physical phase diagnostics complement the model-order and pair-spacing results. Figure 6.16 presents the 14.03 cm baseline as an illustrative measurement condition; the residuals cannot be interpreted as isolated oscillator errors.

The figure shows the measured intra-aperture residuals and the inter-aperture phase spread before and after compensation. The residual errors remain relatively large, even after compensation. Possible contributors include uncertainty in the positions and effective phase centres of the radars, targets, and reference reflector. At the operating wavelength, millimetre-scale position errors already correspond to substantial phase errors. Multipath propagation, target superposition, channel-dependent hardware errors, and array-geometry mismatch can also contribute; these diagnostics do not separate their individual effects.

The spherical reflector provides a common reference from which a phase correction can be estimated for each aperture. This reduces the average inter-aperture spread, but does not remove it completely because the reflector position and phase centre are not known with sub-millimetre accuracy. Moreover, the applied correction consists of one phase ro-

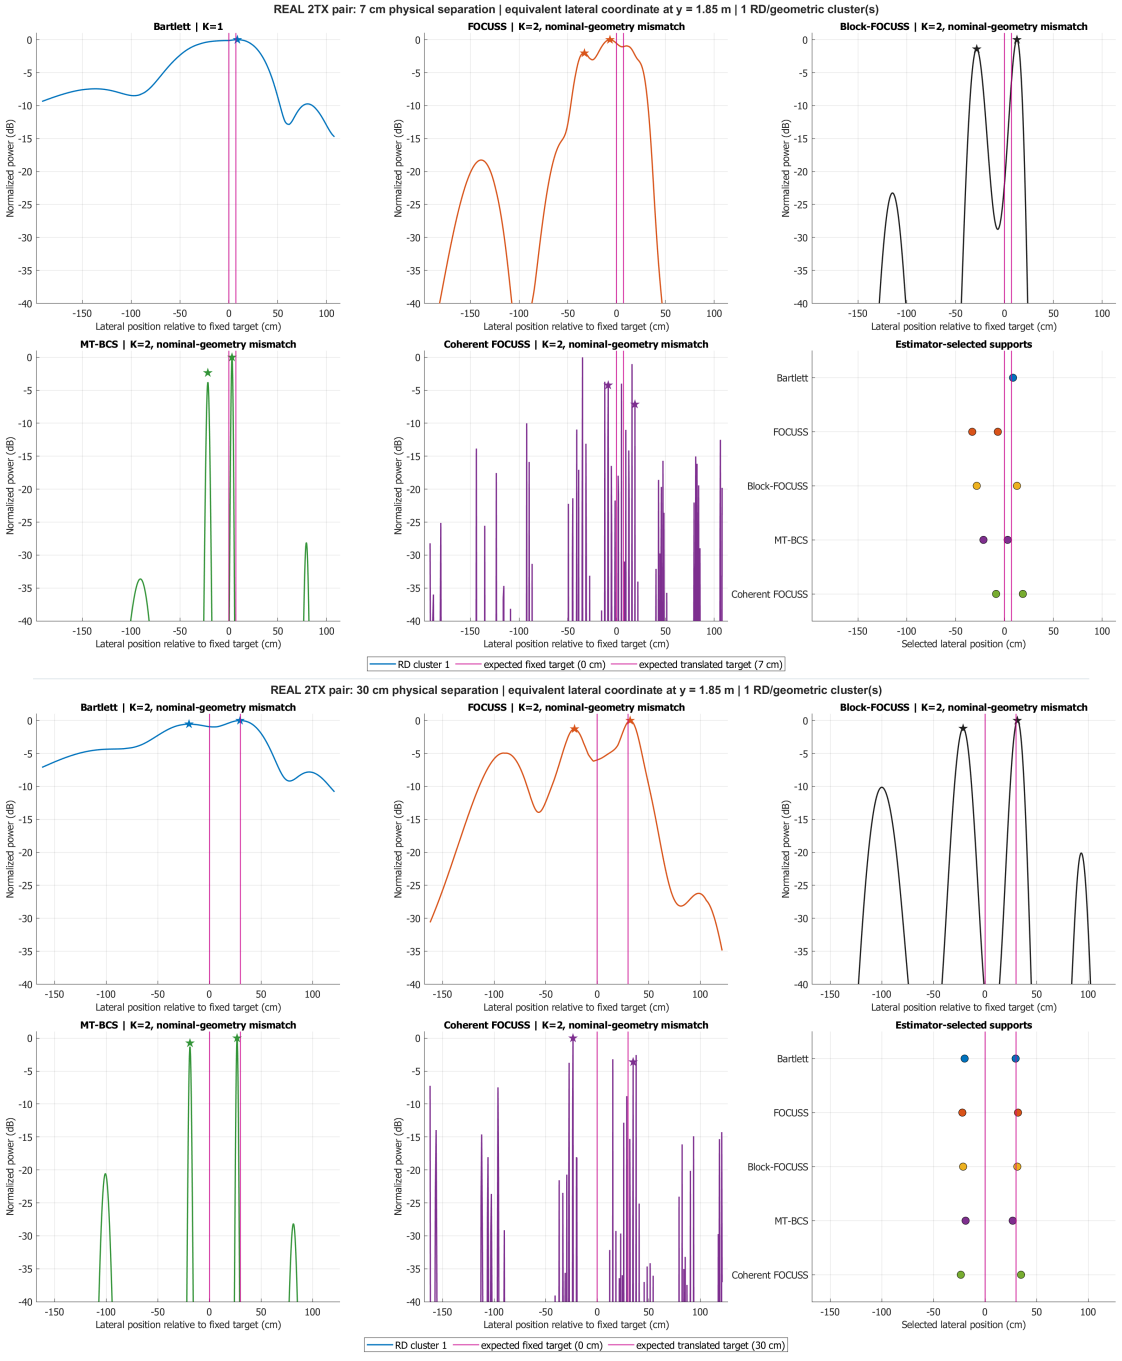


Figure 6.14: Representative range-resolved DoA score curves at 7 and 30 cm physical separation for the 25 cm baseline. The final panel in each row summarizes the supports selected by the five DoA estimators.

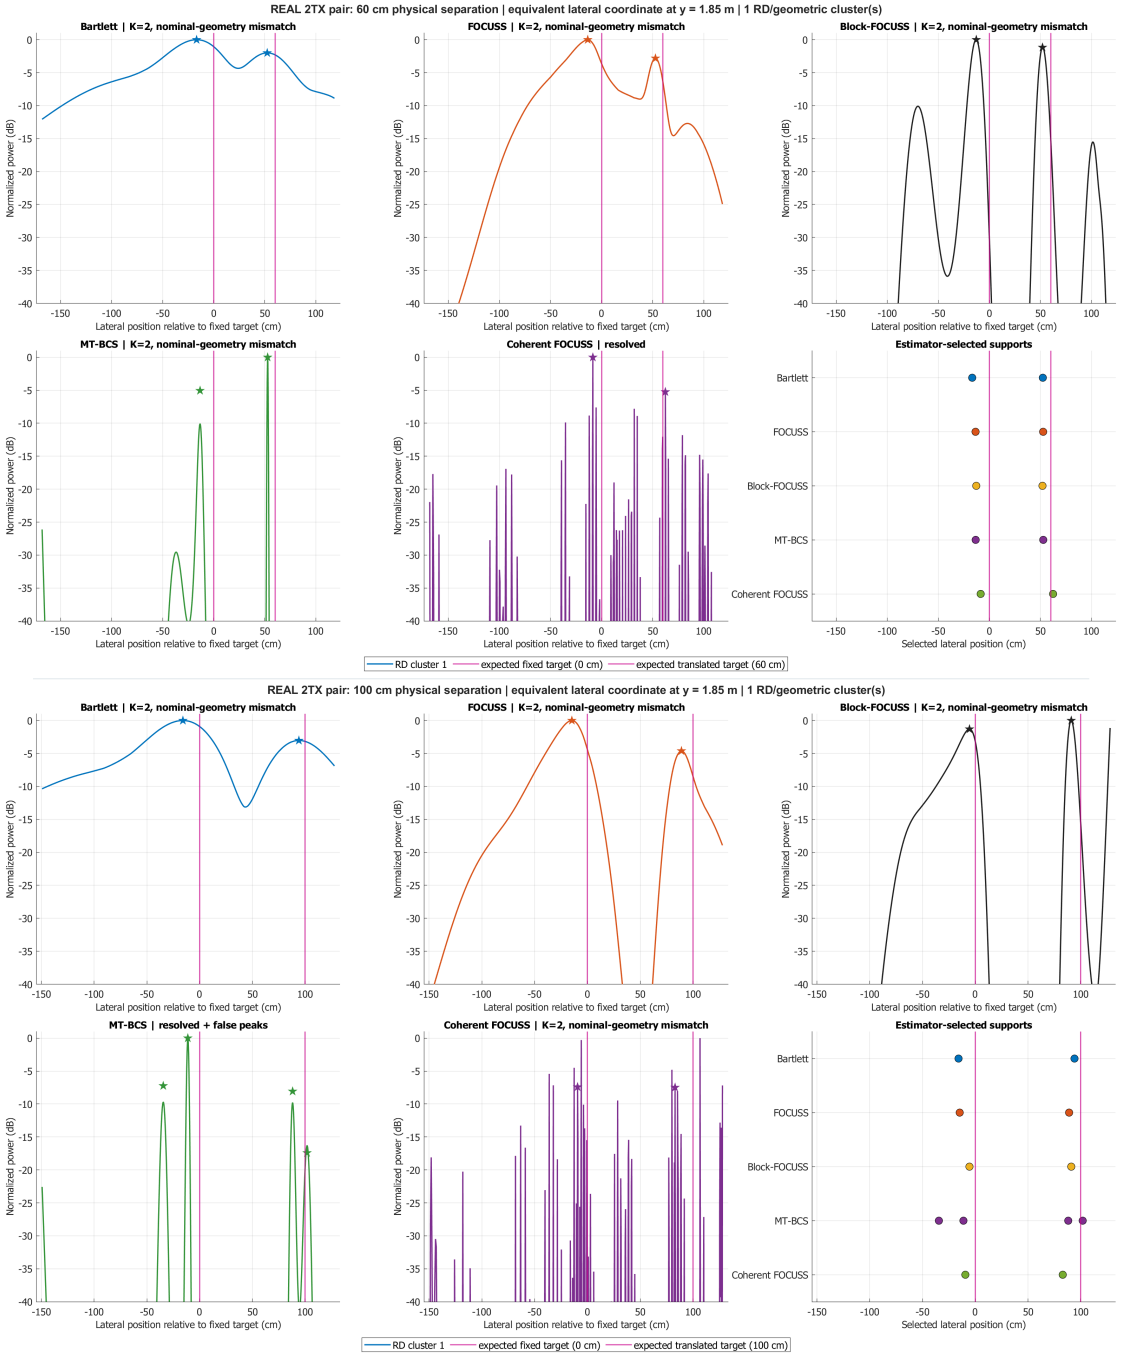


Figure 6.15: Representative range-resolved DoA score curves at 60 and 100 cm physical separation for the 25 cm baseline. The final panel in each row summarizes the supports selected by the five DoA estimators.

tation per aperture and therefore cannot correct channel-dependent intra-aperture errors. The values in Figure 6.16 should consequently be interpreted as effective residual errors containing both phase-calibration errors and scene-dependent model mismatch, rather than as pure oscillator phase errors.

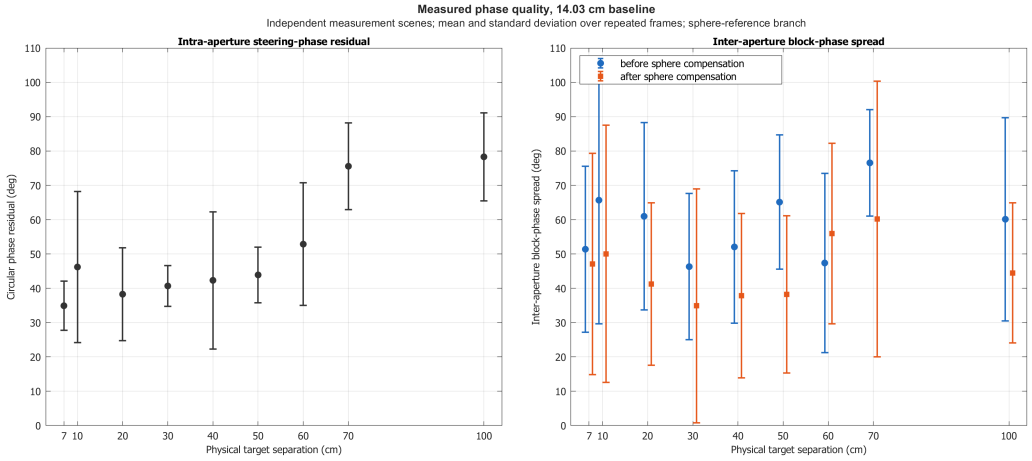


Figure 6.16: Mean phase-quality diagnostics for the 14.03 cm baseline across independent target-separation scenes. The left panel shows the circular intra-aperture steering-phase residual relative to the selected steering atom; the right panel shows the circular inter-aperture block-phase spread before and after spherical-reference compensation. Markers and error bars indicate the mean and standard deviation across repeated frames.

The spherical reflector used for independent phase-reference compensation is shown in Figure 5.4 in Chapter 5. Ideally, this procedure requires the effective phase centres of the reflector and the radar apertures to be known with sub-millimetre accuracy. Achieving this accuracy was not feasible in the experimental setup. The resulting compensation should therefore be regarded as a practical test of reference-based phase correction rather than an ideal calibration. Nevertheless, the reduction in inter-aperture spread indicates that the reflector captures part of the systematic phase mismatch between the apertures.

6.6 Cartesian Backpropagation: Ideal and Physical Localization

Cartesian backpropagation provides a complementary localization demonstration and is not counted as an additional same-cell model-order estimator. Under the procedures in Section 4.5, non-coherent combination produces a broad spatial response, while coherent combination can form narrower maxima if its inter-aperture phase and exact-path geometry are sufficiently accurate.

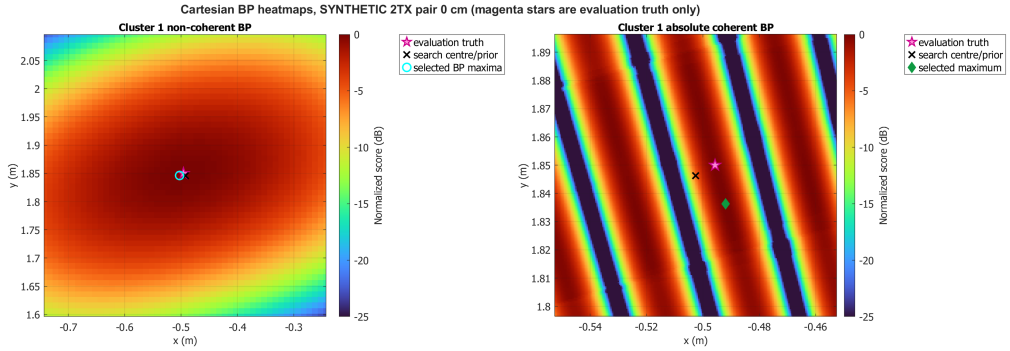


Figure 6.17: Ideal single-target Cartesian backpropagation at the 25 cm baseline. The non-coherent map produces a broad spatial maximum, whereas absolute coherent combination produces narrow periodic fringes and selects the fringe containing the target.

Figure 6.18 extends the ideal demonstration to two targets in the same range region. Two maxima are first retained from the coarse non-coherent map. A local coherent region is subsequently evaluated around each maximum, after which the two positions are selected jointly using a two-target coherent model.

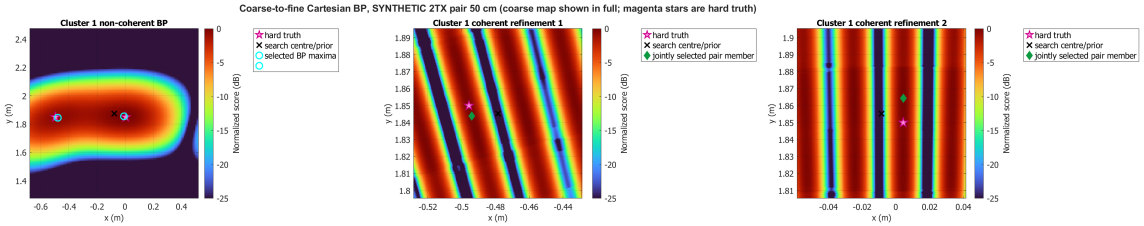


Figure 6.18: Coarse non-coherent backprojection and coherent refinement for two ideal synthetic targets separated by 50 cm, using a 25 cm baseline. The jointly selected coherent positions closely coincide with the ground-truth positions shown in magenta.

Under ideal synthetic conditions, the procedure localizes both targets accurately. The coarse map provides two suitable initialization regions, while the joint coherent refinement selects the correct fringe within each region. This demonstrates the spatial resolution that coherent backpropagation can provide when the signal model, array geometry, and inter-aperture phase correction are mutually consistent.

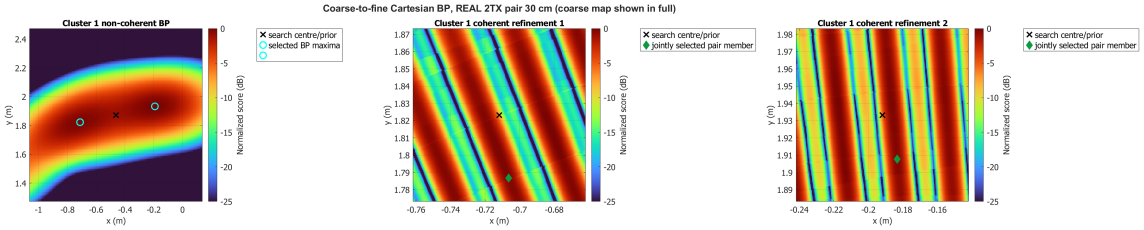


Figure 6.19: Coarse non-coherent backprojection and joint coherent refinement for the physical 30 cm target-separation experiment at the 25 cm baseline. Cyan markers denote the retained coarse maxima and green markers denote the jointly selected coherent positions.

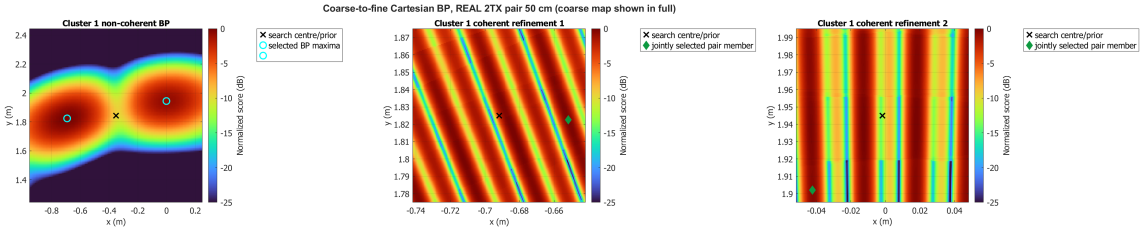


Figure 6.20: Coarse non-coherent backprojection and joint coherent refinement for the physical 50 cm target-separation experiment at the 25 cm baseline. Cyan markers denote the retained coarse maxima and green markers denote the jointly selected coherent positions.

Figures 6.19 and 6.20 show the corresponding physical measurements. The physical measurements retain the same general coarse-to-fine structure: the non-coherent maps identify two broad spatial regions, and the coherent maps contain a series of much narrower phase-dependent fringes. However, the separation between the two selected coherent positions does not reproduce the commanded physical separation exactly. The disagreement is already visible in the coarse estimates and can be further affected by residual inter-aperture phase error, uncertainty in the radar and target geometry, multipath, and the coherent mixture of both targets within the same range region. These effects can shift the strongest jointly fitted solution to a neighbouring fringe even when both target regions have been identified correctly. Consequently, the ideal result demonstrates the potential resolution of coherent backpropagation, whereas the physical results show its sensitivity to modeling and calibration accuracy; these limitations are examined further in the following section.

6.7 Simulation-to-Measurement Comparison

The simulator was configured from the measured room geometry and processed with the same range–Doppler chain as the measured signals. Figure 6.21 compares a scene containing two stationary targets, while Figure 6.22 replaces one target by a moving reflector. In both cases, the simulation reproduces the principal target locations, the zero-Doppler structure, and the bistatic Doppler displacement observed in the measurements.

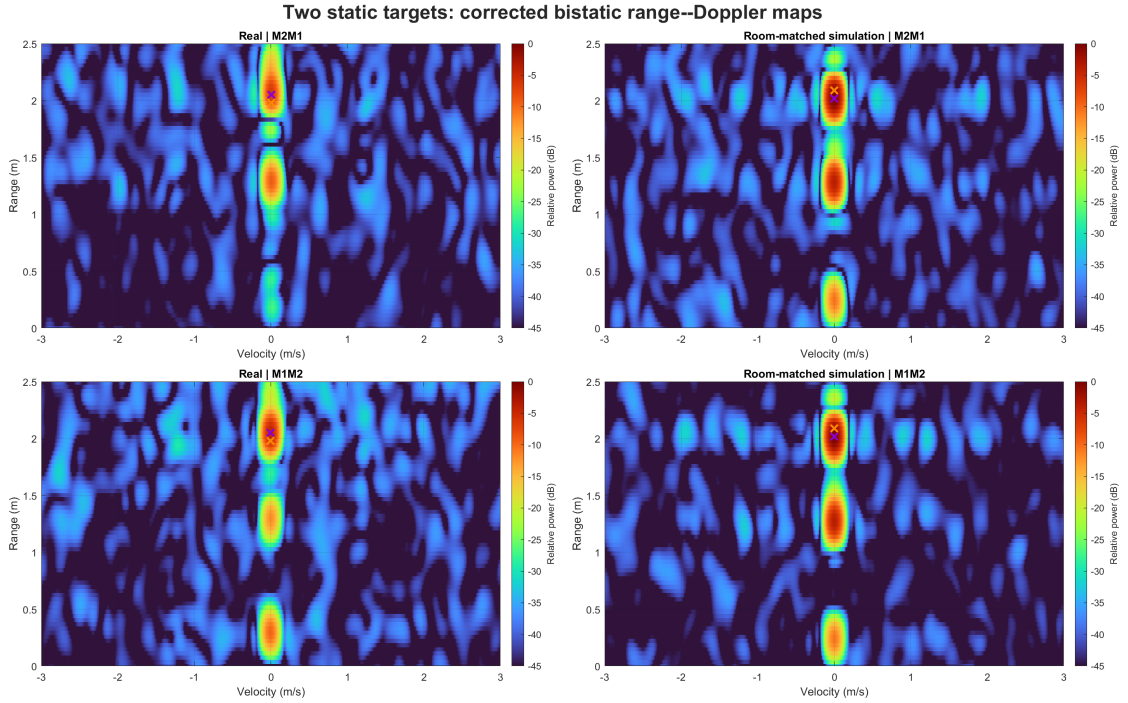


Figure 6.21: Measured and simulated range–Doppler maps for the scene containing two stationary targets. Both bistatic receive paths are shown over the range interval used for target analysis.

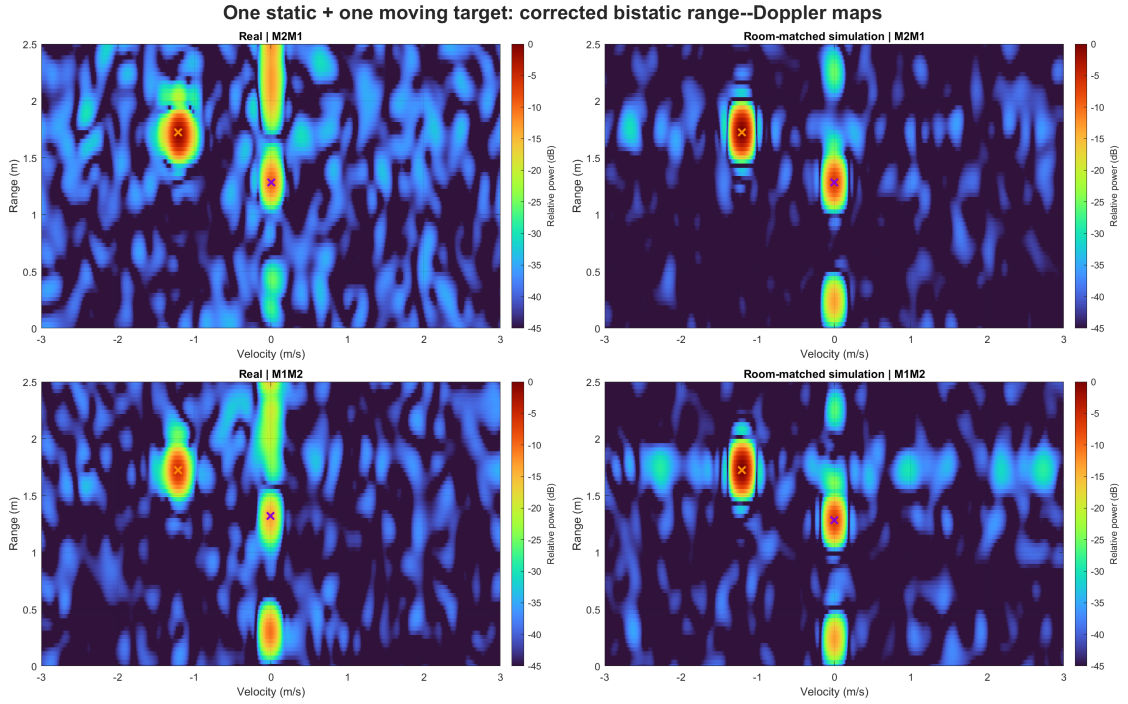


Figure 6.22: Measured and simulated range–Doppler maps for a scene containing one stationary and one moving target. The simulated target position and velocity are derived from the corresponding measurement.

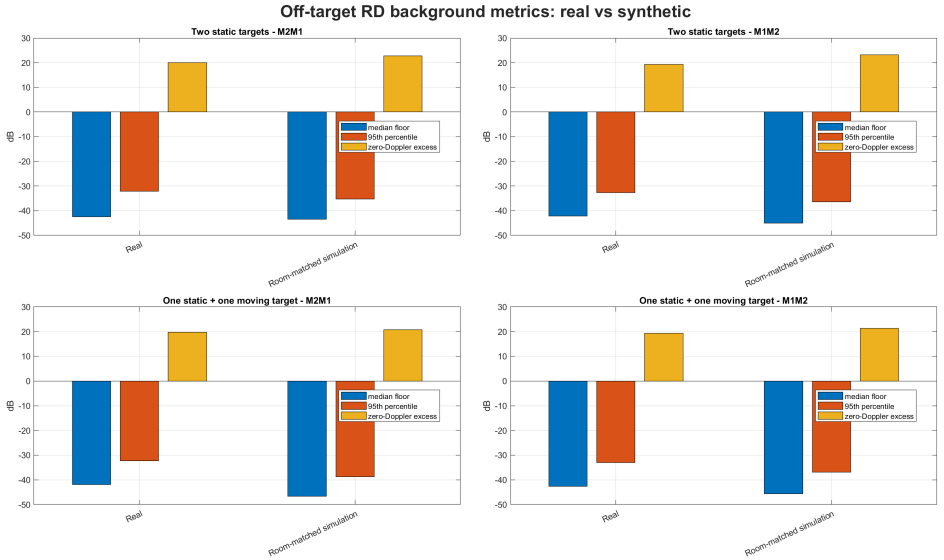


Figure 6.23: Comparison of the measured and simulated background statistics for the stationary and moving scenes.

The measurement-matched model reproduces selected range–Doppler structures but remains several decibels cleaner in the background statistics than the physical measurements. In particular, agreement in target position, Doppler displacement, or average clutter structure does not imply that the simulator predicts the observed truth-blind model-order outcomes or the physical estimator ranking. These simulations are therefore diagnostic approximations, not quantitative replicas of the measured localization task.

6.8 Discussion

The measurements show that the pipeline achieved coarse range–Doppler detection and geometric association, but did not demonstrate reliable absolute-angle success with a 2° tolerance. This distinction is consistent with the hierarchy of localization assumptions introduced in Chapter 2: distributed support can be estimated without a common carrier phase, while exploiting coherent baseline interference requires tighter calibration and geometric consistency [6, 9]. The range–Doppler processing and geometric association produced spatially plausible target regions, and multiple estimators recovered approximately correct relative target separations in the wider-separation experiments. This distinguishes successful upstream processing from failure of the absolute-angle success criterion. Since Bartlett and MT-BCS also miss the absolute-angle criterion without requiring a common carrier phase, the measurements do not isolate one dominant cause of the localization errors. Geometry uncertainty, local-manifold mismatch, and calibration errors can affect several estimator branches; coherent processing additionally requires the relative complex aperture phases to agree with the assumed propagation model.

At 77 GHz, a small error in the assumed propagation distance produces a substantial carrier-phase error. Consequently, coherent aperture combination requires sub-millimetre-

level knowledge of the radar phase centres, reference-reflector position, target geometry, and relevant propagation paths. Two compensation procedures must therefore be distinguished. In the original 25 cm-baseline study, the target cluster itself was used to estimate the inter-aperture block rotation. The reduction from approximately 73.3° to 22.9° belongs to this scene-assisted, in-sample procedure and not to the spherical-reflector calibration. Because the same mixed target response is used for calibration and evaluation, the correction may absorb target superposition, multipath, and scene-specific model error; it is best interpreted as an optimistic conditional diagnostic rather than independent validation. In the later baseline experiments, the spherical reflector supplies an external reference that is separate from the evaluated target cluster. This is methodologically stronger and tests phase transfer between spatial locations, but its performance is limited by uncertainty in the reflector position and effective phase centre. The outcomes of the two procedures are therefore not directly comparable: the cluster-based result demonstrates how much alignment is possible when the evaluated scene assists its own calibration, whereas the spherical-reference result tests the more realistic independent-calibration problem.

The remaining phase disagreement was not limited to one common phase rotation between apertures. The measured residuals also contained approximately 42.0° intra-aperture error, 36.1° TX2–TX1 mismatch, and 61.9° error over the complete 2TX virtual aperture. This distinction is important because block-phase compensation can only correct a common inter-aperture rotation; it cannot restore an element-to-element phase progression that is already distorted within an aperture. The measured phase residuals may furthermore contain contributions from target superposition, reflector orientation, multipath, imperfect phase-centre modeling, and range-bin selection. They should therefore be interpreted as effective scene-level coherence indicators rather than as isolated hardware-phase errors.

Relative to prior analyses that prescribe one random phase-offset variance or derive a corresponding performance bound [17, 18], the controlled results show why the location of the residual in the signal model matters for the tested estimators. A common block rotation can leave local and shared-support processing unchanged while invalidating coherent stacking, whereas channel-relative distortion degrades the steering information used by every phase-based branch. The added value is therefore not another scalar phase threshold, but an estimator-specific interpretation of which localization claim remains supported for each residual mechanism.

The within-CPI results in Section 6.4 add a distinction that the static spatial-error sweeps cannot establish: temporal phase error must be interpreted together with Doppler-cell selection and reference calibration. A linear residual frequency offset mainly shifts the response, whereas a strong random walk also spreads energy. This separation is consistent with the temporal and spectral phase-distortion mechanisms discussed in [10, 22]. The present comparison connects these effects to the reported angle and source count of five DoA methods. In particular, the coherent result changes substantially when a tracked target snapshot is calibrated using a reference retained at zero Doppler. This

is a processing-protocol interaction, rather than evidence that reference compensation is intrinsically ineffective. Small strongest-peak errors can also coexist with incorrect source counts, as MT-BCS and coherent FOCUSS illustrate. The controlled single-target experiment at fixed waveform settings therefore complements the static pair studies, without establishing hardware oscillator tolerances or new two-target resolution limits.

The difference between relative and absolute localization supports this interpretation. At the larger target separations, several DoA estimators recovered approximately correct pair spacing even though the absolute target-angle criterion was not satisfied. A common angular bias can preserve the separation between two estimates while translating both estimates away from their expected locations. For applications in which relative spacing or source multiplicity is more important than global coordinates, this behavior may still be useful. Absolute localization, however, requires a more accurate geometric reference and cannot be inferred from relative-separation success alone.

A similar effect is visible in the Cartesian backpropagation results. Under ideal synthetic conditions, two maxima in the coarse non-coherent map provide suitable initialization regions, and the joint coherent refinement accurately selects the corresponding target positions. In the physical measurements, the same procedure identifies plausible spatial regions, but the selected coherent maxima can move to neighbouring phase fringes. This behavior is consistent with the periodic ambiguity of a long coherent baseline: improved nominal resolution is accompanied by increased sensitivity to small phase and geometry errors. The quality of the final coherent localization therefore depends strongly on whether the applied phase correction places the measurements on the physically correct fringe.

The comparison with the geometry-matched simulator further indicates that waveform and range–Doppler similarity are not sufficient to predict estimator behavior in the measured scene. Although several background and signal characteristics could be reproduced, the simulator did not reproduce the complete physical estimator ranking. Unmodelled multipath, antenna patterns, coupling, reflector scattering, and calibration drift are plausible sources of this discrepancy. The ideal and measurement-matched simulations should therefore be regarded as controlled diagnostic cases rather than complete representations of the laboratory environment.

These observations suggest, but do not yet validate, a phase-quality-dependent choice between coherent and non-coherent processing. When the intra-aperture manifold remains reliable but the relative block phases are uncertain, Bartlett, MT-BCS, or another shared-support method can retain useful angular information without requiring a common carrier-phase origin. Fully coherent processing becomes attractive when both the intra-aperture manifold quality and the inter-aperture phase consistency pass an independently defined quality criterion. Such a rule would require independent calibration-quality thresholds and out-of-sample validation; the present experiments do not establish an operational switching threshold.

Future physical validation would benefit from accurately surveyed radar and reflector coordinates, mechanically fixed phase-centre positions, and repeated calibration across

independently configured scenes. Better-known geometry would help distinguish estimator limitations from positioning errors and test how closely the practical system approaches the ideal coherent performance.

Chapter 7

Conclusions and Future Work

This thesis developed and evaluated a common synchronization and localization framework that connects residual phase quality to estimator-specific localization performance. Controlled simulations and physical measurements of two independently clocked FMCW radar nodes quantify the conditions under which their two monostatic and two reciprocal bistatic apertures support angular and Cartesian localization. The main contribution is the separation of intra-aperture, inter-aperture, and within-CPI phase effects and their consequences for target-count inference, relative pair separation, and angular or Cartesian error. This assessment distinguishes range–Doppler alignment from the carrier-phase coherence needed to exploit the distributed aperture reliably.

7.1 Summary of Main Findings

The receive-side synchronization and processing framework supported range–Doppler detection, multi-aperture snapshot extraction, and geometric target association. The Vecchi-inspired coarse correction reduced the dominant bistatic frequency offset and slow-time drift, but that alignment alone did not establish a phase-consistent distributed aperture. The evaluation distinguishes channel-relative errors within an aperture, common rotations between aperture blocks, and phase evolution over a CPI. The completed static-error studies show that these error structures cannot be represented by one estimator-independent tolerance.

The fixed-waveform within-CPI study further showed that a residual frequency offset need not produce a large angular error when the displaced Doppler peak is followed. With peak-tracked readout and frozen-reference calibration, coherent FOCUSS retained RMSE below 0.05° throughout the tested 0–100 Hz range, but reached 1.44° under the 360° terminal-RMS phase walk. Bartlett, independent FOCUSS, and Block-FOCUSS retained correct single-target count throughout the stationary tests. Temporal robustness consequently depends on the phase history and processing protocol, as well as the estimator.

In the controlled static-phase studies, shared-support processing, especially MT-BCS,

provided consistent close-pair estimates without imposing a common carrier phase across apertures. Coherent FOCUSS and coherent Cartesian backpropagation could exploit the distributed baseline for sharper ideal responses, but were sensitive to calibration, geometry, and baseline-dependent phase ambiguities. In the physical close-pair measurements, several estimators recovered plausible target regions and, at larger separations, approximately correct relative spacing. For frames reporting exactly two sources in the main 25 cm-baseline study, angular RMSE relative to the nominal target bearings ranged from 3.80° to 7.78° across the five DoA methods (Table A.4). Across the 72 evaluated frames per method, none met the separate criterion requiring two sources and both assigned bearing errors within 2°. Thus, ideal distributed-aperture resolution was not equivalent to experimentally established absolute localization. The later campaigns at radar baseline distances of 14.03 and 19.5 cm showed the same distinction: aggregate relative-separation success across the nine separations did not exceed 15% for any method. Spherical-reference compensation changed coherent FOCUSS from 8% to 12% at 14.03 cm and from 10% to 7% at 19.5 cm, so it did not provide a consistent cross-baseline improvement in the present setup.

7.2 Answers to the Research Questions

The four questions introduced in Section 1.2 are answered below using the controlled and physical results of Chapter 6.

7.2.1 RQ1: Synchronization and Residual Errors

To what extent can receive-side correction align monostatic and bistatic measurements from independent FMCW nodes, and what phase residuals remain?

The receive-side procedure reduced the dominant bistatic frequency displacement, chirp-dependent migration, and slow-time phase drift sufficiently to support common range–Doppler processing and associated target measurements. In the controlled synchronization efforts, the median absolute residual in estimating the imposed 5 MHz bistatic frequency offset was 3.7–4.9 kHz, corresponding to approximately 19–24 mm of equivalent-range error, below the 146.6 mm waveform range-resolution interval. This demonstrates useful coarse alignment, not carrier-level synchronization.

The physical measurements retained substantial effective phase disagreement. In the original 0.25 m study, median residuals were 42.0° within a TX group, 61.9° across the full 2TX virtual aperture, and 36.1° for the TX2–TX1 mismatch. Across the 68 frames with finite phase diagnostics, a target-cluster-derived block correction reduced the median inter-aperture spread from 73.3° to 22.9°. This reduction was obtained *in-sample*, using the same close-pair cluster for phase adjustment and evaluation, and does not establish transferable calibration. Later experiments used a separate spherical reflector for a more independent phase-reference test; their remaining errors were affected by uncertain reference geometry and effective phase centres. Accordingly, the supported conclusion

is that range–Doppler alignment was achieved while reliable inter-aperture carrier-phase coherence remained unverified in the physical setup.

7.2.2 RQ2: Estimator Sensitivity to Residual Phase Errors

How do the evaluated local, sparse, shared-support, and coherent methods respond to residual intra-aperture, inter-aperture, and within-CPI phase perturbations, particularly for close target pairs?

The controlled experiments demonstrate that the *structure* of residual phase error matters as much as its magnitude. A common inter-aperture block rotation preserves the local steering progression; methods operating independently on aperture snapshots or enforcing geometrically aligned shared support can therefore remain insensitive to that rotation. At the 0.25 m baseline, Bartlett, independent FOCUSS, Block-FOCUSS, and MT-BCS retained 100% correct single-target model order through 90° RMS of isolated block rotation. Coherent FOCUSS, which requires the blocks to agree in complex phase, retained 60% at 60° and 40% at 90° in that experiment.

Intra-aperture errors instead distort the channel-to-channel phase progression used by every evaluated phase-based estimator. At 60° intra-aperture RMS in the 0.25 m single-target control, the four non-coherent or shared-support DoA methods had 0% correct one-source model order and coherent FOCUSS reached only 5%. Low angular error calculated only on surviving matched estimates must therefore not be interpreted as reliable recovery when support or model order fails.

The close-pair experiments reinforce this distinction. The 20 cm separation corresponds to a nominal 5.92°, substantially below the approximately 15° local-aperture benchmark, whereas 50 cm corresponds to 15.13°. For the 5.92° pair, MT-BCS retained at least 80% relative-separation success under intra-aperture perturbations through 30°, 10°, and 10° RMS at the 16λ (0.0623 m), 36λ (0.1403 m), and 64.2λ (0.25 m) radar baselines, respectively. Under isolated inter-aperture rotations it remained at 100% through 90° in these controls. Block-FOCUSS did not resolve this pair even in its deterministic ideal control; its insensitivity to block rotation must not be confused with successful close-pair resolution. The results do not support one universal phase-error tolerance: the outcome depends on the estimator, baseline, target scene, and whether success means source count, relative spacing, or correct absolute angle.

The within-CPI comparison in Section 6.4 extends this finding to temporal residuals at the 16λ baseline. With peak-tracked readout, Bartlett, independent FOCUSS, and Block-FOCUSS maintained 100% one-target recovery and approximately 0.11°–0.15° angular RMSE. MT-BCS retained small strongest-peak errors, but its one-target rate fell to 10% at 100 Hz. Coherent FOCUSS was more dependent on reference treatment: peak tracking with a frozen reference tolerated the tested frequency offsets, whereas same-CPI calibration at zero Doppler gave approximately 1.8° RMSE at 50–100 Hz. Strong random walks also degraded its angle estimates. These are controlled single-target outcomes at fixed waveform settings, not measured operating limits or close-pair resolution thresholds.

7.2.3 RQ3: Conditions for Coherent Localization

Under which tested conditions can reference-based calibration and geometric accuracy support coherent localization, and how does this compare with less phase-dependent fusion?

Coherent processing is advantageous when the channel-relative manifolds, exact propagation geometry, and relative aperture-block phases all agree sufficiently closely with the assumed model. Under ideal synthetic conditions, coherent FOCUSS returned two supports throughout the evaluated pair scenes. Its relative-separation successes were 9/9 at 16λ, 7/9 at 36λ, and 5/9 at the 0.25 m baseline, illustrating how a longer coherent baseline can introduce additional ambiguity despite a narrower local response. The 16λ configuration was a simulation-only baseline that could not be physically realized with the available radar units.

The same trade-off appeared in Cartesian backpropagation. Non-coherent processing produced broader but less phase-sensitive spatial regions. Ideal coherent backpropagation formed narrow periodic fringes; a coarse-to-fine procedure using non-coherent initialization and joint coherent two-target fitting recovered the ideal synthetic pair. In physical measurements, however, uncertain phase and geometry could displace the selected coherent maxima to neighbouring fringes. Reference-target compensation with the separate spherical reflector was a more independent test than the in-sample cluster correction described under RQ1, but the available reflector and antenna phase-centre positions were not known accurately enough to demonstrate reliable absolute coherent localization. Consequently, coherent gain is supported as a conditional ideal-model result, not as a repeatable physical super-resolution result for this setup.

7.2.4 RQ4: Simulation-to-Measurement Transfer

How well do controlled and measurement-matched simulations predict physical target-count, relative-separation, and absolute-localization outcomes?

The controlled simulations identified which performance limitations arose from the assumed aperture geometry, source-selection procedure, and imposed phase errors. In particular, the ideal same-cell tests showed that the tested configuration contains sufficient modeled spatial information to recover close pairs under favorable conditions. The physical 25 cm-baseline study examined nine target separations. Several estimators achieved correct relative spacing at wider separations; for example, at 60 cm (18.23°) the relative-separation success rates were 100% for Bartlett, independent FOCUSS, and Block-FOCUSS, 62.5% for MT-BCS, and 50% for coherent FOCUSS. The corresponding angular RMSE relative to the nominal geometry, evaluated only in frames returning exactly two sources, ranged from 3.80° to 7.78° across the five methods (Table A.4). Relative-spacing recovery and absolute-bearing error therefore provide distinct assessments of the measured localization performance.

The measurement-matched simulator reproduced selected features of the measured range–Doppler maps, including principal target locations and Doppler displacement, but

Table 7.1: Consolidated answers to the four research questions and their supporting evidence.

RQ	Evidence	Conclusion
RQ1	Coarse-frequency correction; physical intra-/inter-aperture diagnostics; two reference strategies.	Range–Doppler alignment is supported, but stable transferable carrier-phase coherence is not established.
RQ2	Static single- and two-target phase sweeps at three baselines; fixed-waveform single-target CPI sweeps at 16λ.	Local and shared-support processing tolerate common block rotations better. Temporal residuals interact with Doppler selection and reference calibration; a small dominant-peak error does not guarantee correct source count. Tolerance remains mechanism-, scene-, and estimator-dependent.
RQ3	Ideal coherent FOCUSS and backpropagation; physical reference-based calibration and spatial maps.	Coherent gain is demonstrable under calibrated ideal models, but phase/geometry mismatch and ambiguity prevent a comparable physical claim.
RQ4	Truth-blind physical close-pair scoring; measurement-matched range–Doppler comparisons.	Some relative spacing is recovered physically; absolute-angle success within 2° is not. The simulator does not predict all physical failures.

its background remained several decibels cleaner and it did not reproduce the physical estimator ranking or truth-blind model-order outcomes. A close match in the range–Doppler domain is therefore insufficient to validate the spatial phase and scattering assumptions required for localization. The controlled model is useful for mechanism isolation and conditional comparisons, while the present room-matched model is not a quantitative predictor of measured absolute-localization performance.

Table 7.1 summarizes the answers without collapsing the different success criteria into one resolution claim.

7.3 Research Contributions

The principal contribution is an estimator-specific residual-phase assessment, rather than the invention of a new synchronization algorithm or a new sparse DoA estimator. Building on established receive-side correction and multi-aperture estimation concepts [5, 9], the work connects the phase quality remaining after synchronization and calibration to the localization capabilities that can still be supported. The common framework spans measured and synthetic dechirped-IF data, synchronization, detection, aperture association, truth-blind DoA estimation, and coherent and non-coherent Cartesian localization.

The framework implements exact monostatic and bistatic propagation paths and compares Bartlett, independent and Block-FOCUSS, MT-BCS, coherent FOCUSS, and Cartesian backpropagation under explicitly different phase assumptions. Its mechanism-resolved evaluation distinguishes intra-aperture manifold distortion, inter-aperture block rotation, and phase evolution over coherent integration time. The completed static-error studies make clear why an estimator that is insensitive to one residual component may still fail under another. The physical experiments add a necessary validation boundary by separat-

ing source-count inference, relative pair separation, and correct absolute localization, and by distinguishing scene-assisted calibration from independent reference transfer. Finally, a configurable IF-level simulator and measurement-matched room model expose where successful signal-domain matching does not translate into matching estimator behavior.

The within-CPI extension links controlled post-correction frequency offsets and random phase walks to Doppler-response changes and five-method DoA outcomes. Four paired readout and reference-calibration protocols distinguish sensitivity to a fixed Doppler cell from sensitivity to the transferred block-phase correction. This provides an additional explanation of estimator failure beyond a single static phase-error variance.

7.4 Limitations

The dominant limitation is the absence of independently verified phase and geometry accuracy adequate for reliable coherent aperture combination. Uncertain antenna and reflector phase centres, approximate positions, multipath, target-dependent scattering, and residual receive-channel errors can all distort the measured steering manifold. At approximately 77 GHz, millimetre-scale path uncertainty is already a substantial fraction of the carrier wavelength. A correction consisting of one common rotation per aperture block cannot compensate a channel-dependent phase distortion. The residual metrics reported here are effective scene-and-model disagreements and should not be interpreted as pure oscillator noise.

The shortest (16λ) simulation baseline could not be built with the available modules. The conclusions must therefore be confined to the reported baselines, geometries, processing settings, and experimental data rather than generalized to all distributed radar configurations.

The temporal-phase experiment uses one single-target geometry, fixed waveform and CPI timing, and an imposed phase that is common to the virtual channels within each aperture. It does not model intra-chirp evolution or fit measured oscillator spectra. Its numerical outcomes therefore describe sensitivity within this controlled model rather than validated hardware residual limits.

The measurement-matched simulator includes only an approximate representation of the laboratory environment. Its noise, clutter, scattering, coupling, and reference-phase models do not reproduce all physical estimator failures. Matching selected range–Doppler statistics is not proof that the spatial manifold or calibration transfer is correct. Table 7.2 collects the distinction between required conditions and observed evidence.

7.5 Future Work

The most immediate priority is a mechanically rigid and independently surveyed measurement setup. Future measurements should mount the radar modules, targets, and phase reference on a stable rail or optical bench with repeatable fixtures, and use laser or optical

Table 7.2: Requirements for reliable coherent localization compared with the evidence obtained in the present evaluation.

Condition	**Required for interpretation**	**Observed evidence and boundary**
Range–Doppler alignment	Corresponding echoes must be associated after correcting dominant clock-induced distortions.	Controlled 3.7–4.9 kHz residual frequency error supports coarse alignment but does not establish carrier coherence.
Local steering quality	Relative channel phases must match the assumed aperture manifold to within a scene- and estimator-dependent tolerance.	Physical medians of 42.0° intra-TX and 61.9° full-2TX are substantial effective scene-and-model residuals.
Cross-aperture phase	Coherent stacking needs calibrated block phases and sufficiently accurate reference geometry.	The 73.3° to 22.9° reduction is in-sample; independent spherical-reference transfer remains geometry-limited.
Absolute target localization	Correct target count and spacing must be accompanied by correct absolute angles or positions.	Some physical relative-separation successes were observed, but no absolute-angle success within 2° in the 72 principal frame-level cases.

metrology to determine their relative positions as accurately as possible. At the operating wavelength, sub-millimetre knowledge of the relevant propagation paths is desirable for coherent processing. The setup should also use a stable, point-like reference target or validated redundant reference channels, with explicit characterization of the radar-element and reflector effective phase centres. In addition, the experiment should ideally be moved to a more controlled, low-scattering environment or use absorbing material to suppress unwanted room reflections. These improvements would make it easier to separate hardware phase offsets from geometry error and multipath, and to determine whether a phase correction transfers beyond the reference used to estimate it.

A second priority is to improve simulation-to-measurement transfer without fitting directly to the cases used for evaluation. The simulator should incorporate experimentally calibrated distributions of target radar-cross-section imbalance, near-target multipath, antenna-pattern and coupling effects, geometry uncertainty, and range-bin or target-association errors. Parameters should be estimated on a calibration subset and evaluated on independent captures and geometries. Future comparisons should continue to report detection eligibility, truth-blind source count, relative spacing, absolute localization, and conditional estimator success separately.

Finally, the measured intra-aperture steering residual and compensated inter-aperture block-phase spread suggest a phase-quality-aware selection strategy. Coherent processing would be enabled only when both types of diagnostic satisfy thresholds validated for the relevant scene and baseline; otherwise shared-support or non-coherent processing could be used, or a frame flagged for recalibration if the local manifold is also unreliable. This is a proposed extension rather than a validated contribution of the current thesis: no operational switching rule or transferable numerical threshold has yet been established.

The overall conclusion is that distributed radar provides valuable additional spatial

information, but the ability to exploit it coherently must be demonstrated through independent phase and geometric validation. For the evaluated two-node system, synchronization supported coarse and relative localization, while physical absolute super-resolution remained unverified. The framework identifies which assumptions require improvement before the theoretical advantage of a larger distributed aperture can be claimed in practice.

Bibliography

[1] S. M. Patole, M. Torlak, D. Wang, and M. Ali, "Automotive Radars: A Review of Signal Processing Techniques," *IEEE Signal Processing Magazine*, vol. 34, no. 2, pp. 22–35, 2017.

[2] I. Bilik, "Comparative Analysis of Radar and Lidar Technologies for Automotive Applications," *IEEE Intelligent Transportation Systems Magazine*, vol. 15, no. 1, pp. 244–269, 2023.

[3] H. D. Griffiths, "From a Different Perspective: Principles, Practice and Potential of Bistatic Radar," in *Proceedings of the International Conference on Radar (IEEE Cat. No. 03EX695)*, pp. 1–7, 2003.

[4] Y. Yang and R. S. Blum, "Phase Synchronization for Coherent MIMO Radar: Algorithms and Their Analysis," *IEEE Transactions on Signal Processing*, vol. 59, no. 11, pp. 5538–5557, 2011.

[5] C. Vecchi, R. Güneri, and A. Yarovoy, "Model-Based Correction Algorithm for Synchronization of Two Cooperative MIMO Radars," in *2023 IEEE Conference on Antenna Measurements and Applications (CAMA)*, pp. 898–902, 2023.

[6] V. Janoudi, P. Schoeder, T. Grebner, N. Appenrodt, J. Dickmann, and C. Waldschmidt, "Signal Model for Coherent Processing of Uncoupled and Low Frequency Coupled MIMO Radar Networks," *IEEE Journal of Microwaves*, vol. 4, no. 1, pp. 69–85, 2024.

[7] D. Griffiths, G. Donlan, J. Kannanthara, M. Jahangir, C. Baker, M. Antoniou, and Y. Singh, "Phase Synchronization of Bistatic Radar by Exploiting the Urban Scene," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 61, no. 3, pp. 6083–6097, 2025.

[8] A. Dürr, D. Böhm, D. Schwarz, S. Häfner, R. Thomä, and C. Waldschmidt, "Coherent Measurements of a Multistatic MIMO Radar Network With Phase Noise Optimized Non-Coherent Signal Synthesis," *IEEE Journal of Microwaves*, vol. 2, no. 2, pp. 239–252, 2022.

[9] M. Mo, "Two-Stage DOA Estimation via Multi-Aperture FOCUSS in Distributed Radar Systems," msc thesis, Delft University of Technology, 2025.

[10] K. Siddiq, M. K. Hobden, S. R. Pennock, and R. J. Watson, "Phase Noise in FMCW Radar Systems," *IEEE Transactions on Aerospace and Electronic Systems*, vol. 55, no. 1, pp. 70–81, 2019.

[11] D. Werbunat, L. Paulus, J. Aguilar, V. Janoudi, J. Kanz, M. Basler, M. Vossiek, and C. Waldschmidt, "Coherent Radar Networks: Terminology, Synchronization, and Applications," *IEEE Journal of Microwaves*, vol. 6, no. 4, pp. 852–873, 2026.

[12] M. Gottinger, M. Hoffmann, M. Christmann, M. Schütz, F. Kirsch, P. Gulden, and M. Vossiek, "Coherent Automotive Radar Networks: The Next Generation of Radar-Based Imaging and Mapping," *IEEE Journal of Microwaves*, vol. 1, no. 1, pp. 149–163, 2021.

[13] A. Dürr, R. Kramer, D. Schwarz, M. Geiger, and C. Waldschmidt, "Calibration-Based Phase Coherence of Incoherent and Quasi-Coherent 160-GHz MIMO Radars," *IEEE Transactions on Microwave Theory and Techniques*, vol. 68, no. 7, pp. 2768–2778, 2020.

[14] B. Meinecke, D. Werbunat, Q. Haidari, M. Linder, and C. Waldschmidt, "Near-Field Compensation for Coherent Radar Networks," *IEEE Microwave and Wireless Components Letters*, vol. 32, no. 10, pp. 1251–1254, 2022.

[15] P. Fenske, T. Koegel, R. Ghasemi, and M. Vossiek, "Constellation Estimation, Coherent Signal Processing, and Multiperspective Imaging in an Uncoupled Bistatic Cooperative Radar Network," *IEEE Journal of Microwaves*, vol. 4, no. 3, pp. 486–500, 2024.

[16] P. J. Beasley, N. Peters, C. Horne, and M. A. Ritchie, "Global Navigation Satellite Systems Disciplined Oscillator Synchronisation of Multistatic Radar," *IET Radar, Sonar & Navigation*, vol. 18, no. 1, pp. 23–40, 2024.

[17] H. Godrich and A. M. Haimovich, "Localization Performance of Coherent MIMO Radar Systems Subject to Phase Synchronization Errors," in *2010 4th International Symposium on Communications, Control and Signal Processing (ISCCSP)*, pp. 1–5, 2010.

[18] N. Rojhani, M. S. Greco, and F. Gini, "Phase Synchronization Sensitivity for Widely Separated MIMO Radar in CES Disturbance," *Signal Processing*, vol. 199, p. 108627, 2022.

[19] H. Li, F. Wang, C. Zeng, and M. A. Govoni, "Signal Detection in Distributed MIMO Radar With Non-Orthogonal Waveforms and Sync Errors," *IEEE Transactions on Signal Processing*, vol. 69, pp. 3671–3684, 2021.

[20] B. Sun, H. Chen, and H. Zou, "Sparsity-Aware Multitarget Localisation for Distributed MIMO Radar Against Phase Synchronisation Mismatch," *IET Communications*, vol. 10, no. 16, pp. 2269–2275, 2016.

[21] P. Chen, Z. Chen, Z. Cao, and X. Wang, "A New Atomic Norm for DOA Estimation With Gain-Phase Errors," *IEEE Transactions on Signal Processing*, vol. 68, pp. 4293–4306, 2020.

[22] P. Tschapek, G. Körner, C. Carlowitz, and M. Vossiek, "Detailed Analysis and Modeling of Phase Noise and Systematic Phase Distortions in FMCW Radar Systems," *IEEE Journal of Microwaves*, vol. 2, no. 4, pp. 648–659, 2022.

[23] V. Janoudi, P. Schoeder, T. Grebner, C. Waldschmidt, J. Dickmann, and N. Appenrodt, "Coherent Virtual Aperture Calibration in Uncoupled MIMO Radar Networks," in *ICMIM 2024; 7th IEEE MTT Conference on Microwaves for Intelligent Mobility*, pp. 61–64, 2024.

[24] S. Edstaller and D. Mueller, "A Cooperative Radar System With Active Reference Target Synchronization for Kinematic Target Analysis," *IEEE Transactions on Microwave Theory and Techniques*, vol. 69, no. 9, pp. 4118–4131, 2021.

[25] J. Capon, "High-Resolution Frequency-Wavenumber Spectrum Analysis," *Proceedings of the IEEE*, vol. 57, no. 8, pp. 1408–1418, 1969.

[26] R. O. Schmidt, "Multiple Emitter Location and Signal Parameter Estimation," *IEEE Transactions on Antennas and Propagation*, vol. 34, no. 3, pp. 276–280, 1986.

[27] R. Roy and T. Kailath, "ESPRIT—Estimation of Signal Parameters via Rotational Invariance Techniques," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. 37, no. 7, pp. 984–995, 1989.

[28] P. Stoica and R. L. Moses, *Spectral Analysis of Signals*. Upper Saddle River, NJ: Pearson Prentice Hall, 2005.

[29] D. Malioutov, M. Cetin, and A. S. Willsky, "A Sparse Signal Reconstruction Perspective for Source Localization With Sensor Arrays," *IEEE Transactions on Signal Processing*, vol. 53, no. 8, pp. 3010–3022, 2005.

[30] Z. Yang, L. Xie, and C. Zhang, "Off-Grid Direction of Arrival Estimation Using Sparse Bayesian Inference," *IEEE Transactions on Signal Processing*, vol. 61, no. 1, pp. 38–43, 2013.

[31] I. F. Gorodnitsky and B. D. Rao, "Sparse Signal Reconstruction from Limited Data Using FOCUSS: A Re-Weighted Minimum Norm Algorithm," *IEEE Transactions on Signal Processing*, vol. 45, no. 3, pp. 600–616, 1997.

[32] M. Carlin, P. Rocca, G. Oliveri, F. Viani, and A. Massa, "Directions-of-Arrival Estimation Through Bayesian Compressive Sensing Strategies," *IEEE Transactions on Antennas and Propagation*, vol. 61, no. 7, pp. 3828–3838, 2013.

Appendix A

Implementation and Reproducibility Details

A.1 Common Acquisition Parameters

All synthetic and measured studies use the same nominal FMCW waveform and the same two-transmitter azimuth processing configuration. Table A.1 lists only parameters that are common to the reported studies. Study-specific baselines, spatial search regions, and phase-reference choices are listed separately in Section A.2.

Table A.1: Common waveform and acquisition parameters used in the reported studies.

Parameter	Value	Qualification
Carrier frequency, f_c	77 GHz	Gives $\lambda \approx 3.90$ mm
ADC sampling rate, f_s	30 MHz	Complex baseband sampling
Chirp slope	29.982 THz/s	Shared by measured and synthetic data
ADC samples per chirp	1024	Sampled bandwidth 1.0234 GHz
Nominal range resolution	146.6 mm	From the sampled chirp bandwidth
Raw chirps per frame	128	Before TX de-interleaving
TDM–MIMO schedule	2 alternating TX	64 slow-time samples per TX group
Active azimuth channels	2 TX and 4 RX	Eight virtual azimuth channels per receiving aperture
Chirp repetition interval	160 μs	Applied to the raw interleaved chirp sequence
RX element spacing	$\lambda/2$	Uniform receive subarray
TX element spacing	$4\lambda/2$	Produces the eight-element virtual azimuth array
TX–RX phase-centre separation	12 mm	Included in the exact bistatic path model

For the physical acquisitions, the radar configuration used a 40 ms frame-period parameter and the external controller issued a 3.3 V trigger every 50 ms. These quantities describe different parts of the acquisition chain and should not be interpreted as two estimates of the same frame interval. Synthetic data use the same waveform sampling and

TDM–MIMO structure, but do not emulate the external trigger electronics.

A.2 Study-Specific Processing Configurations

The estimator formulations are given in Chapter 4. Table A.2 therefore records only the processing choices that differ between the final evaluation studies. Where range–Doppler detection is performed, the shared detector is OS–CFAR with 12/32 training cells, 3/10 guard cells, rank fraction 0.75, $P_{\mathrm{FA}} = 10^{-3}$, a $-15\,\mathrm{dB}$ offset, and a $-35\,\mathrm{dB}$ floor. Static-scene processing retains $|v| \leq 0.15\,\mathrm{m/s}$, applies non-maximum suppression over one range bin and 50 Doppler bins, and passes at most one range–Doppler cell per aperture to the association stage.

Table A.2: Study-specific processing settings used for the final evaluations. A dash denotes a stage that is not used in that study.

Setting	Synthetic phase sensitivity	Physical 25 cm	Physical 14.03/19.5 cm	Measurement-matched simulation
Baseline(s)	16λ (6.23 cm), 36λ (14.03 cm), 25 cm	25 cm	14.03 cm and 19.5 cm	25 cm
Range input / gate	Synchronized base cell; perturbation injected after synchronization	1.50–2.20 m	1.55–2.50 m	1.50–2.20 m
DoA search grid	Single target: $\pm 180°$; pair: $\pm 20°$; $0.1°$ step	Centroid $\pm 35°$; $0.1°$ step	Centroid $\pm 35°$; $0.1°$ step	Centroid $\pm 35°$; $0.1°$ step
Association / evaluation gate	0.15 m synthetic truth match, evaluation only	1.25 m geometric association	1.25 m geometric association	1.25 m geometric association
Coarse Cartesian BP region	–	1.20×1.20 m; 10 mm cells	0.50×0.50 m; 10 mm cells	1.20×1.20 m; 10 mm cells
Coherent BP refinement	–	0.20×0.20 m; 0.5 mm cells	0.20×0.20 m; 0.5 mm cells	0.20×0.20 m; 0.5 mm cells
Phase-reference treatment	Known synthetic reference; residual phase injected afterwards	Evaluated target cluster used in-sample for one block rotation per aperture	Two variants: no block compensation and estimated spherical reference (not surveyed)	Same target-cluster reference path as the reported 25 cm analysis

The 25 cm and 14.03/19.5 cm measurements use different inter-aperture phase-reference strategies. The reported 25 cm analysis estimates the block correction from the evaluated target cluster itself and is therefore an in-sample diagnostic. The later study compares an uncompensated branch with compensation based on a separate spherical reflector whose position is estimated from the radar data rather than surveyed with sub-wavelength accuracy.

Within-CPI sensitivity configuration. The study of Section 5.7.4 uses a 16λ baseline, a target at $-15°$ and $y = 1.85\,\mathrm{m}$, and a separate wall reference near 4.2 m, at 30 dB SNR. Additional clutter and preset phase perturbations are disabled. The waveform of Table A.1 is held fixed: 64 same-TX samples at $320\,\mu\mathrm{s}$ spacing give a 20.16 ms first-to-last span and 20.48 ms nominal Doppler observation period. Residual frequency offsets are 0,

5, 10, 25, 50, and 100 Hz; the separate phase-walk sweep has terminal RMS 0°, 30°, 60°, 120°, 180°, and 360°, with a shared zero-error control. Injection is after synchronization and range processing, before the Doppler FFT, with constant phase within each chirp.

Both stationary and moving CPI readouts use the 1.50–2.20 m target range gate without the static-velocity restriction above. The coarse angular sector is centroid ±40°; coherent refinement retains the Block-FOCUSS gate of 6° and 0.12 m. Fixed/tracked readout is paired with frozen/same-CPI reference calibration (Table 5.3); the same-CPI reference remains at zero Doppler. Five DoA methods are reported. Receiver-noise seeds are 97001–97010 for the static cases and 98001–98003 for the moving control. The phase-walk seed adds 1000 times the frequency in hertz and the terminal RMS in degrees to the receiver-noise seed. Conditions within each realization reuse the same capture. Dataset totals are given in Table A.3.

A.3 Dataset and Evaluation Inventory

Table A.3 summarizes the cases that directly support the reported results. A case denotes one scene realization or one processed frame before multiplication by the number of estimators. Target truth is used for scoring and for constructing the synthetic scene, but is not supplied to the estimators as a model-order decision.

A.4 Reproducibility and Data Provenance

The final analyses were executed with MATLAB R2024b Update 6. The August captures are stored under `Data&code(given)/data/2026_08_04`; the later baseline captures are under `Data&code(given)/data/2026_09_16/16_sep_data_msc`. Processed caches and final outputs are stored under `processedradardata/real_2026_08_04`, `presentation_outputs/chapter7_physical_0804_final`, and `presentation_outputs/physical_2026_09_16_v2`.

The principal entry points are `run_chapter7_validation_overnight.ps1`, `run_real_2026_08_04_resolution_study.m`, `run_real_2026_09_16_study.m`, `recover_physical_0916.ps1`, and `render_physical_0916_results.m`. Synthetic captures store deterministic random seeds; the calibrated August ensemble uses $8040000 + 10000s + c$, where s is the target separation in centimetres and $c \in \{1, 2\}$ is the capture index. Batch runs store logs, manifests, completion markers, and compact processed caches. Exact reproduction of the physical studies additionally requires the paired raw DCA1000 files and saved receive-channel calibration files. The CPI entry point is `run_cpi_single_target_experiment.m`, using `run_single_target_cpi_study.m`; its checkpoints and tables are stored under `presentation_outputs/single_target_cpi_20261104`.

Table A.3: Inventory of the datasets used in the final evaluation.

Study	Baseline(s)	Scene geometries	Realizations and total
Synthetic phase sensitivity	6.23, 14.03, and 25 cm	One single-target scene and the reported 20 and 50 cm pair scenes; residual levels 0°, 2°, 5°, 10°, 20°, 30°, 45°, 60°, and 90°	One deterministic zero-error control per scene. Each nonzero level has 20 Monte Carlo realizations for both intra- and inter-aperture perturbations: 963 cases per baseline and 2889 cases in total.
Within-CPI phase sensitivity	16λ (6.23 cm)	One target at $-15°$ and $y = 1.85$ m; stationary or moving radially at 1 m/s. Waveform and CPI timing fixed.	Static: $11 \times 10 = 110$ cases; moving: $4 \times 3 = 12$ cases. Four protocols and five reported DoAs yield 2440 case–protocol–method rows. The archive contains 4392 rows including additional estimator and diagnostic outputs.
Physical close-pair study	25 cm	Separations 7, 10, 20, 30, 40, 50, 60, 70, and 100 cm at approximately 1.85 m range	Eight processed frames per separation: 72 frame cases.
Physical baseline study	14.03 and 19.5 cm	The same nine relative target separations at approximately 1.9 m range	2 baselines $\times 9$ separations $\times 8$ frames $= 144$ frame cases. Each is processed without compensation and with the estimated spherical reference, giving 288 analysis variants and 1440 frame–estimator rows for five DoA methods.
Measurement-matched simulation	25 cm	Calibration separations 7, 20, 40, and 60 cm; held-out separations 10, 30, 50, 70, and 100 cm	Two independently generated captures with four frames each: eight cases per separation, comprising 32 calibration and 40 held-out cases, 72 cases in total.

A.5 Supplementary Within-CPI Results

These supplementary results use the single-target configuration and four processing protocols of Section 5.7.4. They complement the stationary peak-tracked results in Section 6.4; the same five DoA methods are shown.

A.5.1 Fixed Range–Doppler Cells

Figure A.1 holds each aperture's range and Doppler cells at their unperturbed control values. The complex snapshot is recomputed at those cells after the perturbation, so this is a diagnostic readout rather than a fresh detection test. For coherent FOCUSS, same-CPI reference calibration keeps the angular RMSE below 0.05° throughout these fixed-cell tests. This does not imply recovery of the coherent amplitude lost at the selected cell.

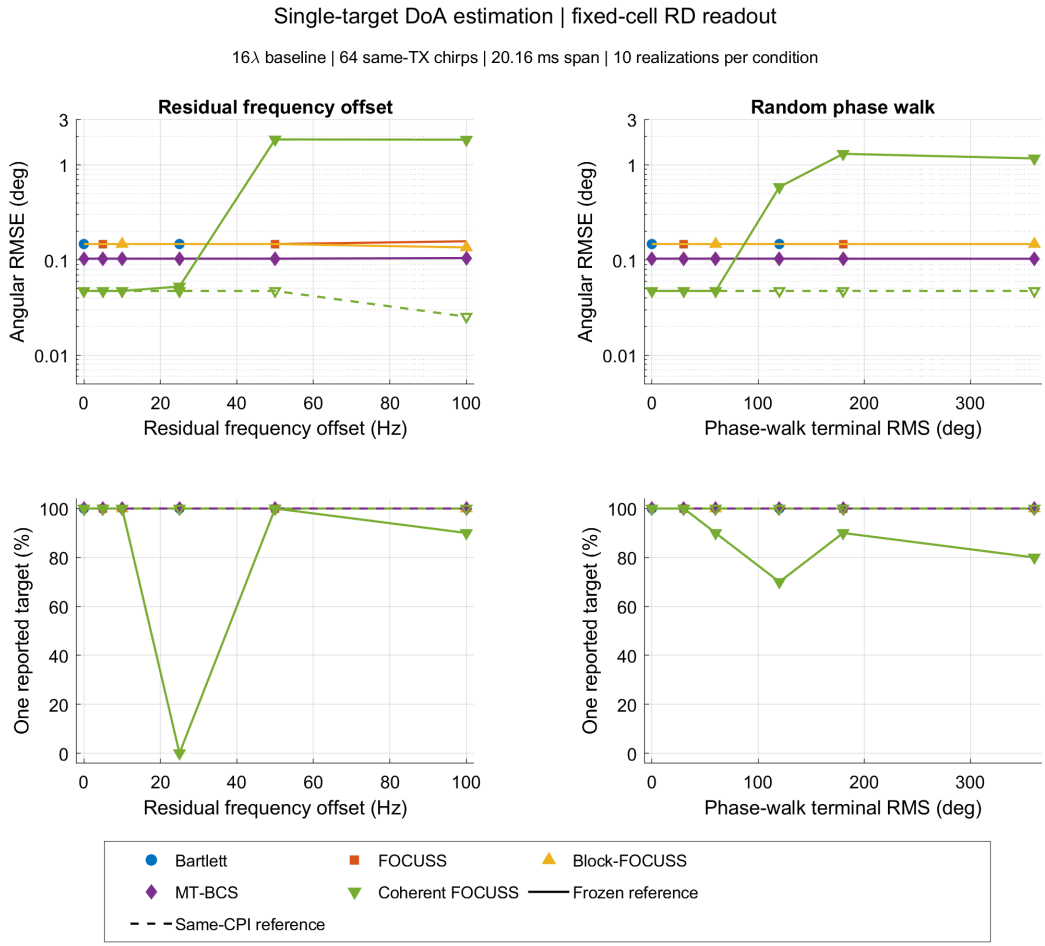


Figure A.1: Fixed range–Doppler cells for the stationary single target at the 16λ baseline. Upper panels show strongest-peak angular RMSE on a logarithmic scale, including trials with incorrect model order; lower panels show the fraction reporting exactly one target. Colors identify the five DoA methods, with solid and dashed lines for frozen and same-CPI reference calibration. Each condition contains ten realizations. The two reference protocols have coincident outcomes for the four non-coherent or shared-support methods.

A.5.2 Moving-Target Control

The supplementary control in Figure A.2 uses a target moving radially away at 1 m/s, with the same initial bearing and waveform. The four separate conditions are the unperturbed control, residual offsets of 25 and 100 Hz, and a phase walk with 180° terminal RMS. Each contains three realizations; these cases provide a consistency check rather than a densely sampled tolerance limit. An angle is reported in every trial, so a zero one-target rate here denotes incorrect model order rather than a missing angular estimate.

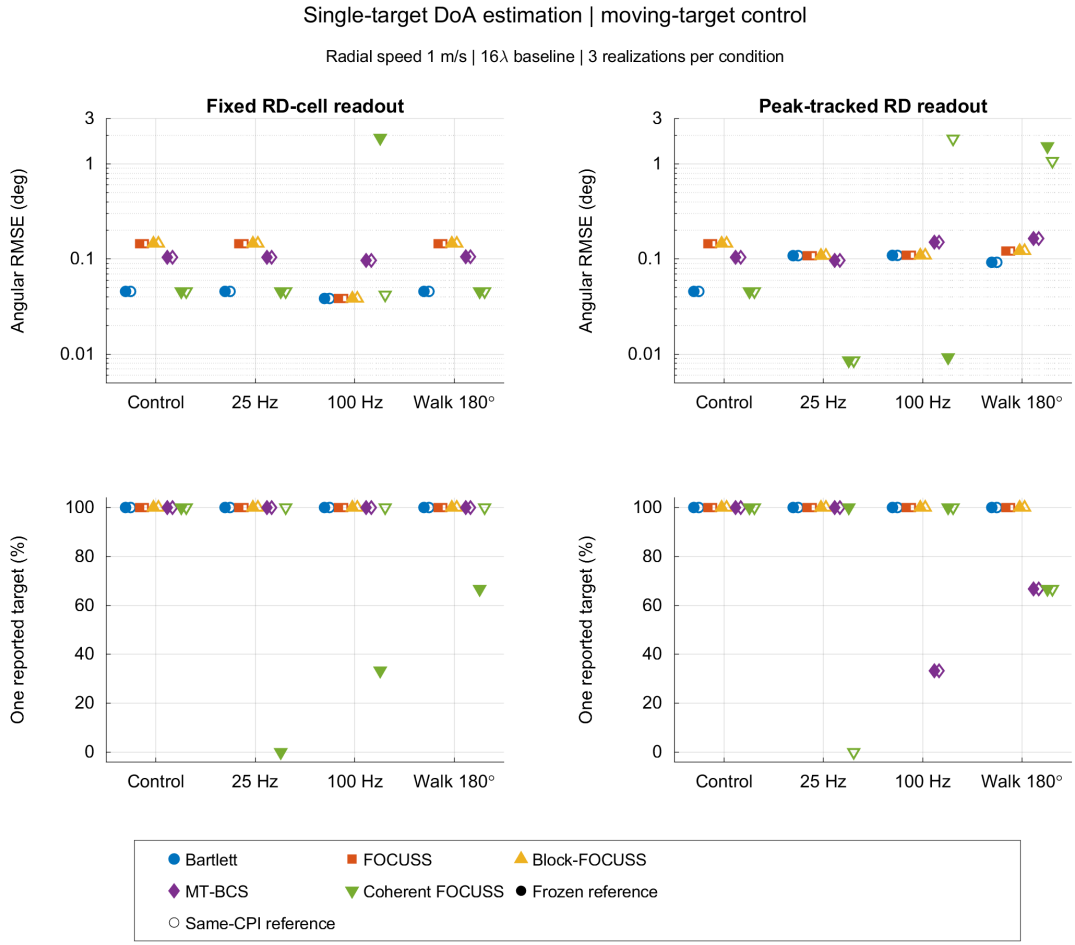


Figure A.2: Five-method DoA comparison for the 1 m/s moving-target control at the 16λ baseline. Columns compare fixed-cell and peak-tracked readout; rows show strongest-peak angular RMSE on a logarithmic scale and the fraction reporting exactly one target. Colors identify methods; filled and open markers distinguish frozen and same-CPI reference calibration. Four discrete conditions are evaluated with three realizations each. Coincident protocol outcomes overlap. RMSE includes trials with incorrect source count.

A.6 Additional Estimator Diagnostics

A.6.1 Measured and Synthetic Pair Comparison

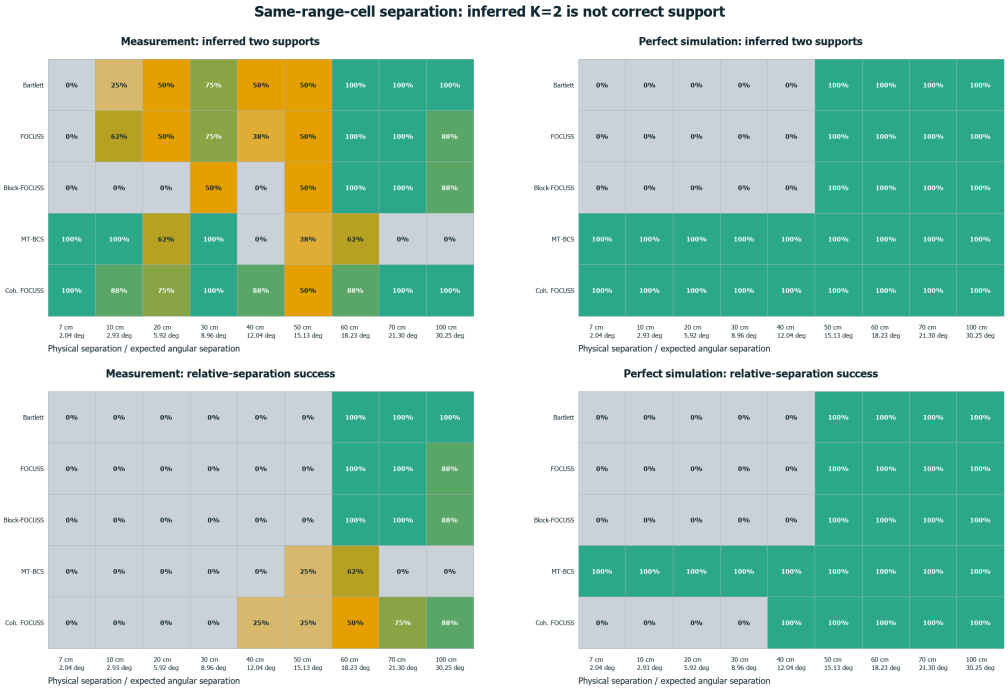


Figure A.3: Matched measured and synthetic target-separation tests for the five DoA estimators at the 0.25 m baseline. Upper panels show the fraction with exactly two selected supports; both lower panels show relative-separation success using Eq. (5.30). Relative-separation success does not establish correct absolute bearings.

Figure A.3 compares the same nominal target geometry and 0.25 m radar baseline under synthetic and measured conditions. Both datasets use identical model-order and relative-separation scoring definitions and tolerances. The upper panels report the two-support rate; both lower panels report relative-separation success. All percentages in this comparison use all evaluated cases or frames as the denominator, without conditioning on successful geometric association. The synthetic entries represent one deterministic control case per separation and method, whereas measured percentages summarize eight evaluated frames per separation. Conditional estimator performance is a separate diagnostic.

A.6.2 Single-Target Model-Order Recovery

Figure A.4 complements the matched-support angular errors in Figure 6.6. It reports the fraction of trials returning exactly one source, independently of whether the reported angle satisfies the matching criterion.

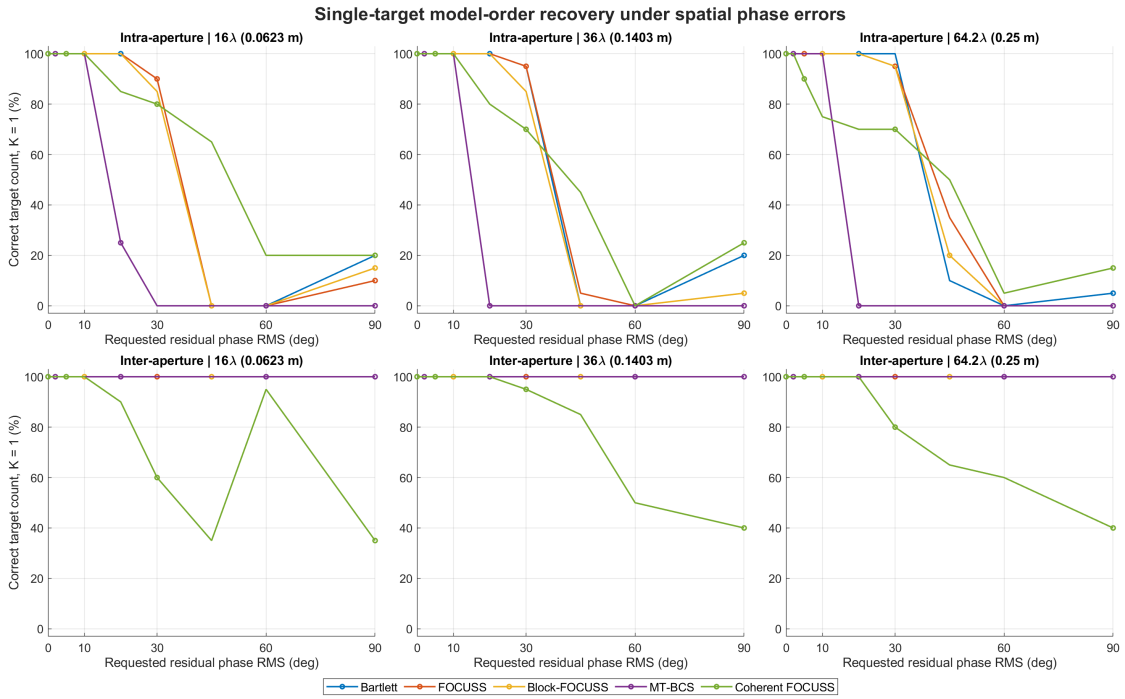


Figure A.4: Single-target model-order recovery under isolated intra-aperture (top) and inter-aperture (bottom) phase perturbations at the three controlled radar baselines. Each nonzero level contains 20 Monte Carlo realizations; zero error is a deterministic control. Coincident curves overlap. Correct target count alone does not establish a correct bearing.

A.6.3 Physical Angular Error Relative to the Nominal Geometry

Table A.4 reports angular error for the original 25 cm-baseline measurements, separately from source-count and relative-separation success. For each method, only frames returning exactly two selected bearings are included. Within each such frame, the sorted estimates are assigned to the sorted nominal bearings; both angular errors contribute to the RMSE. No truth-association gate is applied to discard poorly located pairs. Frames with a different source count remain excluded from this conditional RMSE and are accounted for by the reported N_2 and the separate model-order results.

The reference bearings follow the nominal anchor-and-translation geometry in Table 5.2. They are not independently surveyed exact target positions, so these errors include discrepancies in the assumed measurement geometry. The separate absolute-angle success criterion requires two reported sources and both assigned errors to be at most 2°; it does not require zero error.

Table A.4: Angular RMSE relative to the nominal target bearings for the physical 25 cm radar baseline. N_2 is the number of frames reporting exactly two sources out of 72 evaluated frames per method. RMSE includes both assigned bearings in these frames, without a truth-matching filter.

Method	$N_2/72$	Angular RMSE (°)
Bartlett	44/72	7.73
Independent FOCUSS	45/72	7.78
Block-FOCUSS	31/72	3.80
MT-BCS	37/72	4.30
Coherent FOCUSS	63/72	6.14