


A COMPARISON OF SHIPBORNE RADAR ARCHITECTURES
FOR SMALL-UAV DETECTION

COMPARAISON DES ARCHITECTURES RADAR NAVALES
EMBARQUÉES POUR LA DÉTECTION DE PETITS DRONES



A Thesis Submitted
to the Department of Electrical and Computer Engineering
by

Nick Sanga, CD, PEng
Lieutenant-Commander

In Partial Fulfillment of the Requirements for the Degree of
Master of Applied Science in Electrical Engineering

April 15 2026

© This thesis may be used within the Department of National
Defence but copyright for open publication remains the property of the author.

Acknowledgement

I would like to express my sincere gratitude to Dr. Yahia Antar and Dr. Joey Bray for their support, guidance, and belief in me throughout this thesis. Their encouragement and confidence in my abilities have meant a great deal to me during this journey.

I would also like to extend my deepest thanks to my mother, Dr. Dinna M. Mussa, for the many sacrifices she has made for me and for nurturing my love of science from a young age. Her unwavering support, strength, and example have been a constant source of inspiration.

Abstract

Detecting small Uncrewed Aerial Vehicles (UAVs) from a naval radar platform is challenging because their weak, aspect-dependent radar cross section (RCS) and low-altitude flight occur in the presence of strong maritime interference mechanisms, including grazing-angle sea clutter, precipitation effects, and near-surface multipath. This thesis develops a simulation-based comparison framework for a ship-installable counter-UAV radar, evaluating pulsed Active Electronically Scanned Array (AESA) and Frequency Modulated Continuous Wave (FMCW) AESA candidates under a common set of naval constraints (frequency band, aperture/gain, average and peak power limits, range resolution, search volume, and refresh requirements). A MATLAB-based virtual maritime environment is used to generate repeatable scenarios and to model target RCS fluctuation alongside sea state and rain conditions, enabling controlled experiments with identical resource allocation. Detection performance is assessed using a consistent processing chain (range–Doppler formation, receiver dynamic-range control, Constant False Alarm Rate (CFAR) thresholding at fixed false-alarm probability, and probability-of-detection estimation) and is summarized using decision-oriented figures of merit such as high-confidence detection range, detection latency, and time-to-confirm. Monte Carlo evaluation across representative UAV mission profiles provides statistically defensible performance curves and sensitivity trends, which are then weighted to reflect risk-averse shipboard priorities, robustness in heavy sea states and precipitation. This thesis makes three main contributions: first, it develops a common-constraint simulation framework for fair comparison of ship-installable pulsed AESA and FMCW AESA counter-UAV radars; second, it provides a Monte Carlo evaluation method that links waveform performance to operationally meaningful detection and confirmation metrics in maritime environments; and third, it shows that FMCW is the preferred architecture within the tested constraint set, achieving 11 of 18 case wins and the strongest worst-case robustness, while also identifying the modelling assumptions that bound this recommendation.

Résumé

La détection de petits véhicules aériens sans équipage (UAV) à partir d'une plateforme radar navale constitue un défi important, car leur faible surface équivalente radar (SER), fortement dépendante de l'aspect, et leur vol à basse altitude s'inscrivent dans un environnement maritime dominé par de forts mécanismes d'interférence, notamment le fouillis de mer à faible angle de site, les effets des précipitations et le multitrajet près de la surface. Cette thèse propose un cadre de comparaison fondé sur la simulation pour l'évaluation d'un radar anti-UAV installable à bord d'un navire, en comparant des architectures AESA impulsives et AESA à onde continue modulée en fréquence (FMCW) sous un ensemble commun de contraintes navales, incluant la bande de fréquences, l'ouverture et le gain, les limites de puissance moyenne et de puissance de crête, la résolution en portée, le volume de recherche et les exigences de rafraîchissement. Un environnement maritime virtuel développé sous MATLAB est utilisé afin de générer des scénarios répétables et de modéliser à la fois les fluctuations de la SER des cibles, l'état de la mer et les conditions de pluie, de manière à permettre des expériences contrôlées avec une allocation identique des ressources. La performance de détection est évaluée au moyen d'une chaîne de traitement cohérente comprenant la formation distance-Doppler, le contrôle de la dynamique du récepteur, le seuillage à taux constant de fausses alarmes (CFAR) à probabilité de fausse alarme fixée, ainsi que l'estimation de la probabilité de détection. Les résultats sont ensuite résumés à l'aide de figures de mérite orientées vers la décision, telles que la portée de détection à haute confiance, la latence de détection et le temps de confirmation. Une évaluation de type Monte Carlo, appliquée à des profils de mission UAV représentatifs, permet d'obtenir des courbes de performance et des tendances de sensibilité statistiquement défendables. Ces résultats sont ensuite pondérés de façon à refléter des priorités navales prudentes, axées sur la robustesse en conditions de mer forte et de précipitations importantes. Cette thèse apporte trois contributions principales. Premièrement, elle développe un cadre de simulation à contraintes communes permettant une comparaison équitable entre des radars anti-UAV navals AESA impulsifs et AESA FMCW installables à bord d'un navire. Deuxièmement, elle propose une méthode d'évaluation de type Monte Carlo qui relie les performances des formes d'onde à des métriques de détection et de confirmation ayant une signification opérationnelle en environnement maritime. Troisièmement, elle montre que l'architecture FMCW constitue l'option privilégiée dans l'ensemble des contraintes étudiées, avec 11 gains sur 18 cas et la meilleure robustesse dans le pire cas, tout en mettant en évidence les hypothèses de modélisation qui encadrent cette recommandation.

Table of Contents

Acknowledgement	ii
Abstract.....	iii
Résumé.....	iv
Table of Contents.....	v
List of Figures.....	viii
List of Tables	xi
List of Abbreviations	xii
List of Symbols.....	xiv
Chapter 1 Introduction	1
1.1 Background	1
1.2 Problem Statement.....	2
1.3 Thesis Statement.....	2
1.4 Methodology	2
1.5 Scope	3
1.6 Thesis Organization.....	4
Chapter 2 Literature Review	6
2.1 Radar System.....	6
2.2 AESA Radar Beamsteering	7
2.3 AESA implementation scale and modelling assumptions	13
2.4 Pulsed Radar System	16
2.5 Pulsed Radar Systems in Literature.....	21
2.6 FMCW Radar System.....	24
2.7 FMCW Radar System in Literature.....	28
2.8 Radar Cross Section	32
2.9 Sea Clutter Modelling.....	36

2.10	Detection Processing Chain.....	39
2.11	Advanced Signal Processing	43
2.12	UAV Mission Profiles and Representative Flight-Path Scenarios	44
Chapter 3 Naval Radar Systems and Design Constraints from Literature		46
3.1	Simulation-Based Comparison Framework.....	46
3.2	Common Constraint Set and Baseline Assumptions	47
3.2.1	Shared radar constraints	48
3.2.2	Shared aperture and beamsteering assumptions.....	52
3.2.3	Pulse Radar Candidate	56
3.2.4	FMCW Radar Candidate.....	59
3.2.5	Fairness Basis for the Comparison.....	60
3.3	Threat-Informed UAV Scenarios and Flight-Path Definitions.....	62
3.4	Search Strategy and Dwell / Refresh Budget	66
3.5	Single-Dwell Intermediate Results and Monte Carlo Evaluation.....	68
Chapter 4 Results		78
4.1	Attack Scenarios (40 km Start Range)	78
4.1.1	Best-Case Weather (Sea State 1, 0 mm/h rain)	79
4.1.2	Medium-Case Weather (Sea State 3, 10 mm/h rain).....	79
4.1.3	Worst-Case Weather (Sea State 5, 50 mm/h rain)	80
4.1.4	Interpretation of the Attack 40 km Results	83
4.2	Remaining Engagement Time Results	84
4.3	Micro-Doppler Results	85
4.4	Overall Performance Comparisons.....	89
Chapter 5 Conclusion.....		93
5.1	Summary	93
5.2	Conclusions and Discussion	93
5.3	Future Work	95
5.4	Contributions to the field.....	96

References.....	97
Appendix A Supplementary Scenario Results.....	100

List of Figures

Figure 1-1 Thesis scope and simulation-based comparison framework	4
Figure 2-1 (a) radar Basic diagram and (b) Coherent receiver [8,7].....	7
Figure 2-2 AESA radar diagram [7]	8
Figure 2-3 Linear series of N isotropic antenna elements separated by a constant distance s [2]	9
Figure 2-4 Uniform rectangular array geometry [2]	10
Figure 2-5 Array factor power patterns for a 10-element broadside binomial array with $N = 10$ and $d = \lambda/4, \lambda/2, 3\lambda/4,$ and λ . [7].....	12
Figure 2-6 AMSAR demonstrator antenna array [9]	14
Figure 2-7 TRS-4D AESA radar [10].....	15
Figure 2-8 Determination of target range [7].....	17
Figure 2-9 LFM Tx Signal.....	19
Figure 2-10 Matched-filter response of the transmitted LFM pulse defined in (2.18). (a) Full normalized matched-filter magnitude versus delay, showing the compressed echo peak at the target delay and the surrounding sidelobe structure. (b) Zoomed view of the compressed mainlobe, illustrating the measured 3 dB width and its correspondence to the expected order of $1B$	20
Figure 2-11 Range coverage for X band waveform [11]	22
Figure 2-12 Minimum RCS detectable at a given range with 80% Pd 10 m^2 RCS target is detected at 18.5 km and a 5000 m^2 RCS target is detected at 90 km [11].	23
Figure 2-13 Smart-S MK2 [13].....	24
Figure 2-14 Basic FMCW diagram [2]	25
Figure 2-15 FMCW Tx Signal.....	26
Figure 2-16 (a) The figure shows the transmitted FMCW signal (blue) and its reflected signal (red). The round trip rt seen as τ in the diagram between the two signals maps to the frequency shift Δf . (b) Mixing the transmitted and reflected signals yields the beat signal fb whose frequency is Δf . (c) The FFT peak corresponds to Δf . A subtle change in the distance over the chirps shifts the phase of the FFT peak [14]	27
Figure 2-17 Adaptive CFAR threshold [15]	28
Figure 2-18 Experimental UAV detection on the Range-Doppler map with $Pfa = 10^{-5}$ and $N = 20$. The experiment is carried out using 25 GHz Ancortek radar [15].....	29
Figure 2-19 Alarm radar [16].....	30

Figure 2-20 A comparison between the probability of detection of the scan1 radar and stare1 radars with CPIs (T) set to different fractions of U (surveillance update interval) [16]	31
Figure 2-21 Small drones considered: (a) Trimble ZX5, (b) DJI Inspire 1 Pro, (c) DJI Phantom 4 Pro (DJI P4 Pro) [12].....	33
Figure 2-22 The measured RCS (dBsm) versus azimuth angles for the small drones: (a) Trimble ZX5, (b) DJI Inspire 1, (c) DJI Phantom 4 Pro [12].....	34
Figure 2-23 Example of grazing angle and impact of grazing angle to clutter scattering coefficient [8]	36
Figure 2-24 Example of sea clutter rough and smooth path [8].....	37
Figure 2-25 CFAR Logic Diagram [2]	42
Figure 2-26 Doppler radar echo of a two-blade rotor [21]	43
Figure 3-1 Simulation Framework.....	47
Figure 3-2 ITU Rain model vs Frequency	48
Figure 3-3 Radar Horizon adapted from [7]	50
Figure 3-4 Diagram of Radar location on a ship.....	52
Figure 3-5 60 x 40 URA azimuth vs taper cut	53
Figure 3-6 3-D beampattern for URA 60 x 40.....	54
Figure 3-7 Directivity versus steering angle.....	55
Figure 3-8 (a) Pulsed Radar candidate and (b) FMCW Radar candidate	56
Figure 3-9 (a) LFM Tx pulse waveform and (b) matched filter result.....	58
Figure 3-10 FMCW (a) Tx waveform and (b) matched filter result	60
Figure 3-11 UAV Attack scenario	63
Figure 3-12 UAV Surveillance scenario.....	64
Figure 3-13 UAV Intelligence scenario	65
Figure 3-14 Scenarios' radar horizon distances	66
Figure 3-15 Serpentine Search pattern.....	67
Figure 3-16 (a) FMCW range response, (b) FMCW Doppler response, (c) FMCW Range-Doppler response.....	71
Figure 3-17 (a) LFM range response, (b) LFM doppler response, (c) LFM range-Doppler response	73
Figure 3-18 UAV fluctuation with χ^2 model	75
Figure 4-1 (a) Classification Stats and (b) Detection Statistics for LFM radar during Attack Scenario under Worst Case Scenario	82
Figure 4-2 RPM sensitivity sweep for the LFM micro-Doppler classifier in the worst-case Attack 40 km condition: (a) later high-confidence classification snapshot and (b) earlier P_d , 50 snapshot. Each row corresponds to a different true rotor speed,	

with the left-hand panel showing the observed spectrogram and the right-hand panel showing the GLRT/NCC score across the candidate RPM templates. 88

Figure 4-3 Cross-case synthesis results of all cases..... 91

Figure 4-4 Overall summary across all cases 92

List of Tables

TABLE 2-1 WAVEFORM PULSE PARAMETERS: PEAK POWER, PULSE WIDTHS[15]	21
TABLE 2-2 PARAMETERS OF THREE SMALL UAVS [12].....	33
TABLE 2-3 FITTING ERRORS OF THE THREE DISTRIBUTION MODELS [12]	35
TABLE 2-4 DJI PHANTOM 4 PRO KEY PARAMETERS [24].....	44
TABLE 3-1 SHARED RADAR PARAMETERS	51
TABLE 4-1 ATTACK BEST-CASE WEATHER (SEA STATE 1, 0 MM/H RAIN)	79
TABLE 4-2 ATTACK MEDIUM-CASE WEATHER (SEA STATE 3, 10 MM/H RAIN)	80
TABLE 4-3 ATTACK WORST-CASE WEATHER (SEA STATE 5, 50 MM/H RAIN)	81
TABLE 4-4 REMAINING ENGAGEMENT TIME - SUMMARY	84
TABLE A-1 ATTACK BEST-CASE WEATHER (SEA STATE 1, 0 MM/H RAIN)	100
TABLE A-2 ATTACK MEDIUM-CASE WEATHER (SEA STATE 3, 10 MM/H RAIN)	100
TABLE A-3 ATTACK WORST-CASE WEATHER (SEA STATE 5, 50 MM/H RAIN)	100
TABLE A-4 SURVEILLANCE BEST-CASE WEATHER (SEA STATE 1, 0 MM/H RAIN)	101
TABLE A-5 SURVEILLANCE MEDIUM-CASE WEATHER (SEA STATE 3, 10 MM/H RAIN)	101
TABLE A-6 SURVEILLANCE WORST-CASE WEATHER (SEA STATE 5, 50 MM/H RAIN)	102
TABLE A-7 INTELLIGENCE BEST-CASE WEATHER (SEA STATE 1, 0 MM/H RAIN)	102
TABLE A-8 INTELLIGENCE MEDIUM-CASE WEATHER (SEA STATE 3, 10 MM/H RAIN)	102
TABLE A-9 INTELLIGENCE WORST-CASE WEATHER (SEA STATE 5, 50 MM/H RAIN) — SUMMARY.....	103

List of Abbreviations

Abbreviation	Meaning
ADC / DAC	Analog-to-Digital Converter / Digital-to-Analog Converter
AESA	Active Electronically Scanned Array
AN/SPY	USN shipborne phased-array radar designation (Aegis family)
AOPV	Arctic Offshore Patrol Vessel
C-UAV	Counter-UAV
CDS	Contract Design Specification
CFAR	Constant False Alarm Rate
COHO	Coherent Oscillator
CPF	Canadian Patrol Frigate
CPI	Coherent Processing Interval
CRR / RRR	Clutter Range Retention / Rain Range Retention
CW	Continuous Wave
DLT	Detection Latency
DNR	Director of Naval Requirements
EIRP	Effective Isotropic Radiated Power
FAD	False-Alarm Density
FFT	Fast Fourier Transform
FMCW	Frequency-Modulated Continuous Wave
FoM	Figure of Merit
GLRT	Generalized Likelihood Ratio Test
IF	Intermediate Frequency
ISTAR	Intelligence, Surveillance, Target Acquisition, and Reconnaissance
LFM	Linear Frequency Modulation
LNA	Low-Noise Amplifier

LO / STALO	Local Oscillator / Stable Local Oscillator
MATLAB	Matrix Laboratory
MD	Micro-Doppler
MIMO	Multiple-Input Multiple-Output
MTI / MTD	Moving Target Indication / Moving Target Detection
NCC	Normalized Cross-Correlation
PRF / PRI	Pulse Repetition Frequency / Pulse Repetition Interval
PSC	(P_d) -weighted Search Coverage
RCD	River Class Destroyer
RCS	Radar Cross Section
RCN	Royal Canadian Navy
RD/RDM	Range-Doppler/ Range-Doppler Map
RF	Radio Frequency
RPM	Revolution per Minute
SCR	Signal-to-Clutter Ratio
SDR	Software-Defined Radio
SIR	Signal-to-Interference Ratio
SNR	Signal-to-Noise Ratio
STC	Sensitivity Time Control
SWaP	Size, Weight and Power
T/R	Transmit / Receive
UAV	Uncrewed Aerial Vehicle
ULA / URA	Uniform Linear Array / Uniform Rectangular Array
USN	United States Navy

List of Symbols

Abbreviation	Meaning
A	Physical aperture area (m ²)
A_c	Illuminated clutter area (m ²)
A_e	Effective aperture area (m ²)
AF	Array factor
B	Signal bandwidth (Hz)
c	Speed of light (m/s)
C_s	Chirp slope (Hz/s)
d	Generic element spacing (m)
d_x, d_y	Element spacing in (x) and (y) (m)
$D(\theta, \phi)$	Directivity
$d\Omega$	Differential solid angle
Δf	Frequency shift (Hz)
Δf_d	Doppler-bin spacing (Hz)
ΔR	Range resolution (m)
Δv	Radial-velocity resolution (m/s)
E_0	Incident electric field at the target
E_e	Element factor
E_s	Scattered electric field
$F(\theta, \phi)$	Total radiation pattern
F_n	Noise factor
f_0	Starting frequency offset (Hz)
f_b	Beat frequency (Hz)
$f_{b,max}$	Maximum beat frequency (Hz)
f_c	Carrier frequency (Hz)
f_d	Doppler frequency shift (Hz)
G	Antenna gain (linear)
k	Boltzmann constant (J/K)

L	System losses (linear)
λ	Wavelength (m)
μ	Chirp rate (Hz/s)
N_c	Number of chirps in CPI
N_d	Number of Doppler samples
N_p	Number of pulses in CPI
N_r	Number of fast-time samples
Ω	Solid angle (sr)
$P[k, l]$	Range–Doppler power
P_d	Probability of detection
P_{fa}	Probability of false alarm
ϕ	Azimuth angle
π	Pi
P_t	Transmit power (W)
PRI	Pulse repetition interval (s)
PRF	Pulse repetition frequency (Hz)
R	Range (m)
R_0	Peak sidelobe ratio
R_{max}	Maximum range (m)
R_{min}	Minimum range (m)
R_{unamb}	Unambiguous range (m)
$s(t)$	Baseband waveform
S_{min}	Minimum detectable signal
σ	Radar cross section (m ²)
$\bar{\sigma}$	Mean radar cross section
σ_0	Sea clutter coefficient
σ_c	Sea clutter RCS
σ_w	Weather clutter RCS
T	Chirp duration (s)
T_0	Noise temperature (K)
T_{coh}	Coherent processing interval
T_{max}	Max GLRT/NCC score

T_{rep}	Repetition interval
τ	Pulse width / delay (s)
θ	Elevation angle
θ_{3dB}	Half-power beamwidth
θ_a	Azimuth beamwidth
θ_e	Elevation beamwidth
$u = \sin\theta \cos\phi$	Direction cosine x
u_0, v_0	Steering direction cosines
$v = \sin\theta \sin\phi$	Direction cosine y
v_{amb}	Max unambiguous velocity
v_r	Radial velocity
V_w	Weather clutter volume
$w_r[n]$	Range window
$w_d[m]$	Doppler window
$x[n, m]$	Raw data sample
$X[k, m]$	Range FFT output
$Z[k, l]$	Range-Doppler map
α	CFAR scaling factor
ψ_g	Grazing angle

Chapter 1

Introduction

1.1 Background

Modern naval vessels are increasingly exposed to small uncrewed aerial vehicles (UAVs) operating at low altitude and over short to medium ranges. These inexpensive, off-the-shelf systems can deliver surveillance and can engage and destroy targets at a fraction of the cost of traditional platforms [1]. They can be fielded in numbers to overwhelm defenses. They also provide actionable intelligence, such as force size, locations, and vulnerabilities. Therefore, radar technology has been forced to adapt in order to provide situational awareness to local commanders [2]. While examining the ongoing Ukraine-Russia war, drones allowed Ukrainian forces to stop the Russian advance to Kyiv. Therefore, non-state actors can exploit the same low-cost advantages for asymmetric warfare and terrorism, using drones for reconnaissance, swarming, or one-way attacks against civilian targets and military platforms alike [3,4].

The Royal Canadian Navy (RCN) fleet is recapitalizing: Canadian Patrol Frigate (CPF) are nearing end of life and, under the National Shipbuilding Strategy, will be replaced by the River Class Destroyers (RCD) with more sensitive sensors, including the U.S. Navy's (USN) Aegis combat system and AN/SPY active electronically scanned array (AESA) radar [5]. AESA radars use multiple Transmit and Receive (T/R) elements that scan more quickly than conventional mechanical antennas, thereby providing faster refresh/update rates. The performance of an AESA also degrades gracefully if some elements fail. By contrast, the RCN's newest class of ships, the Arctic Offshore Patrol Vessels (AOPVs), currently lack an effective method for detecting small UAVs, potentially leaving them vulnerable in high UAV threat environments. Therefore, the RCN may consider installing a Counter-UAV (C-UAV) radar via a minor project as a stopgap solution [6].

Given this context, the RCN will likely prefer an AESA-based solution for commonality with the RCD. The open question is which radar architecture best detects small UAVs in the maritime environment: a monostatic pulsed radar that transmits short RF bursts and measures round-trip delay from the same aperture, or an FMCW radar that transmits a continuous, linearly frequency-modulated (chirp) [2,7,8]. This thesis compares these options under shared constraints to identify the most effective near-term C-UAV path for the fleet.

1.2 Problem Statement

Open literature rarely documents controlled comparisons between pulsed AESA and FMCW AESA architectures for detecting small UAV targets in maritime environments. When comparisons are reported, they often rely on differing baseline assumptions. As a result, observed performance differences may reflect unequal resource allocations rather than differences attributable to waveform family or architecture. For shipboard C-UAV applications, the design problem is therefore to quantify detection performance under a consistent specification set, while accounting for sea clutter, rain attenuation, multipath variability, and UAV RCS fluctuations, and to determine how parameter selection and search strategy constrain achievable detection performance and refresh behavior.

1.3 Thesis Statement

A simulation-based comparison framework can be developed for a ship-installable C-UAV radar to compare pulsed and FMCW performance under common naval constraints and to support a recommendation on a preferred radar type for potential RCN consideration as an AOPV-relevant minor project capability. Within this framework, pulsed AESA and FMCW AESA radar candidates are specified under consistent constraints, including operating frequency band, aperture size and gain, average transmit power, range resolution, scan volume, and refresh requirements.

1.4 Methodology

The methodology is structured to compare the two candidate radar architectures against a common naval operational problem. The first step is to define the AOPV-relevant detection requirement by specifying the platform geometry, the surveillance sector, and the range envelope over which a small UAV threat must be detected and classified. This establishes the operational baseline against which both radar types are assessed.

The second step is to define a set of representative UAV scenarios that impose different detection demands on the radar system. These scenarios are intended to capture credible maritime threat profiles by varying target range, altitude, speed, and approach conditions. Environmental severity is then introduced through combinations of rain and sea state so that system performance is evaluated under conditions representative of real shipboard operation rather than under ideal free-space assumptions.

Based on the threat scenarios, two radar candidates are formulated: a pulsed AESA radar and an FMCW AESA radar. Both radar types are designed under the same baseline shipboard constraints and with timing parameters chosen to remain operationally plausible.

This ensures that the comparison is conducted on an equivalent basis in terms of available search resources, rather than favouring one architecture through unrealistic parameter selection.

A simulation environment is then constructed to implement the radar processing chain for each candidate. This environment models the maritime operating conditions, the UAV target, and the relevant radar signal-processing steps required for detection and classification. The processing chain includes waveform dependent range processing, confirm-based detection logic, and micro-Doppler-based classification, thereby allowing both architectures to be assessed not only on whether a target is detected, but also on whether it can be classified as a UAV.

Finally, a Monte Carlo analysis is conducted across the full set of UAV scenarios and environmental conditions. The resulting performance measures are weighted according to operational priorities, with emphasis placed on detection and classification effectiveness within the defined shipboard mission context. These weighted results are then used to determine which radar architecture provides the most suitable overall performance for the naval counter-UAV application. The confirm logic, scenario-level milestone definitions, and Monte Carlo evaluation procedure are defined in detail later in Sections 2.10 and 3.5.

1.5 Scope

This thesis is limited to the assessment of two candidate ship-installable counter-UAV radar architectures for AOPV application: a pulsed AESA radar and an FMCW AESA radar. The study is focused on the detection and classification of small UAV threats in the maritime environment, with particular emphasis on the warning and local-area surveillance requirements of a naval platform operating in clutter-limited conditions.

The analysis considers the shipboard surveillance geometry, the sector of regard to be searched, and the operational detection ranges relevant to an AOPV mission set. Representative UAV threat scenarios are defined to reflect credible approaches to a surface combatant or patrol vessel, including variations in target range, altitude, velocity, and environmental severity. The simulated environment further includes the principal maritime propagation and interference effects expected to degrade radar performance, specifically rain attenuation and sea clutter.

Within this scope, the radar models include waveform generation, waveform-specific range processing, clutter suppression, confirm-based detection processing, and a target-classification stage based on micro-Doppler response. Fig. 1-1 summarizes this processing chain and the boundaries of the thesis scope.

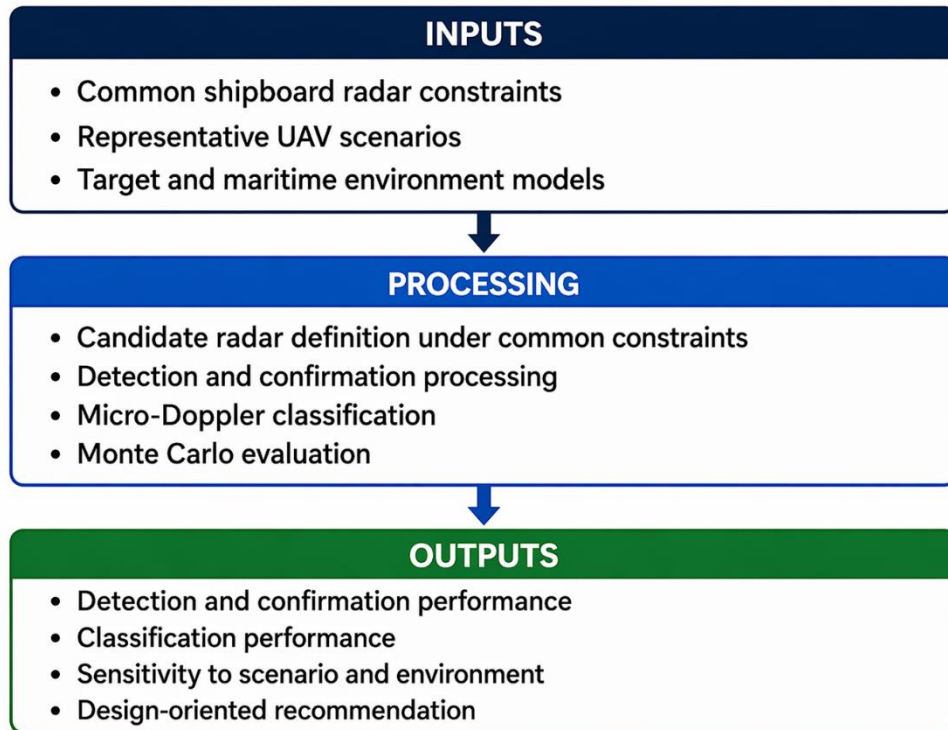


Figure 1-1 Thesis scope and simulation-based comparison framework

The study is restricted to simulation-based comparison and does not extend to hardware prototyping, detailed antenna manufacture, combat-system integration, or fleet-level operational analysis. The aim is not to produce a final radar design, but rather to determine which candidate architecture is better suited to the shipboard counter-UAV role when both systems are assessed under a common and operationally credible set of constraints.

1.6 Thesis Organization

This thesis is organized as follows.

Chapter 2 reviews the radar concepts relevant to both pulsed and FMCW systems. It also defines the operational problem and constraints for a ship-installable C-UAV radar on an AOPV, including the maritime environment, the target of interest, and the radar designs selected to support a fair comparison.

Chapter 3 describes and validates the simulation-based framework, including the environmental models and the representative UAV scenarios.

Chapter 4 presents and analyzes the results of the pulsed versus FMCW comparison, with the Attack 40 km case used as the primary deep-dive because it is the most operationally demanding and decision-relevant scenario. The remaining scenario results are summarized in appendices and referenced where needed to support the main trends.

Chapter 5 concludes the thesis by summarizing the principal findings, contributions, limitations, and future extensions of the comparison framework.

Chapter 2

Literature Review

This chapter presents the background needed for the comparison. It reviews basic radar components and contrasts pulsed with FMCW architectures. It then explains target RCS and sea clutter. Furthermore, it surveys practical mitigations for maritime conditions and concludes with an analysis of figures of merit.

2.1 Radar System

Radars transmit RF electromagnetic waves and detect echoes to estimate how far and how fast objects are. Distance (range) is estimated by using the round-trip time delay of the echo's return; speed (radial velocity) is estimated from changes in the echo's phase as the target moves [2,7,8]. Modern systems share a few core blocks: a waveform generator and power amplifier to create and boost the signal, an antenna to transmit and later collect echoes, a protected receiver to amplify weak returns above the noise, and a processor/display to turn data into range and speed for the operator. Unlike legacy systems with video amplifiers, contemporary radars digitize the signal early in the receiver's processing chain with analog-to-digital converters (ADCs), whereas the transmitter's analog waveforms are synthesized with digital-to-analog converters (DACs), enabling coherent processing as standard practice (Fig. 2-1 (a)) [2,7,8]. In coherent reception, the echo is boosted by a Low Noise Amplifier (LNA) and mixed to Intermediate Frequency (IF) against a stable local oscillator (STALO) so that phase is maintained; a coherent oscillator (COHO) supplies quadrature reference signals that allow the receiver to separate the echo into orthogonal in-phase (I) and quadrature (Q) components (Fig 2-1 (b)). These I/Q channels retain amplitude and phase information that can be used for subsequent processing [7]. Radars can be configured for different target sets by adjusting design settings so that performance aligns with the mission; the specific shipboard constraints, representative UAV scenarios, and candidate parameter choices are defined later in Sections 3.2 to 3.5.

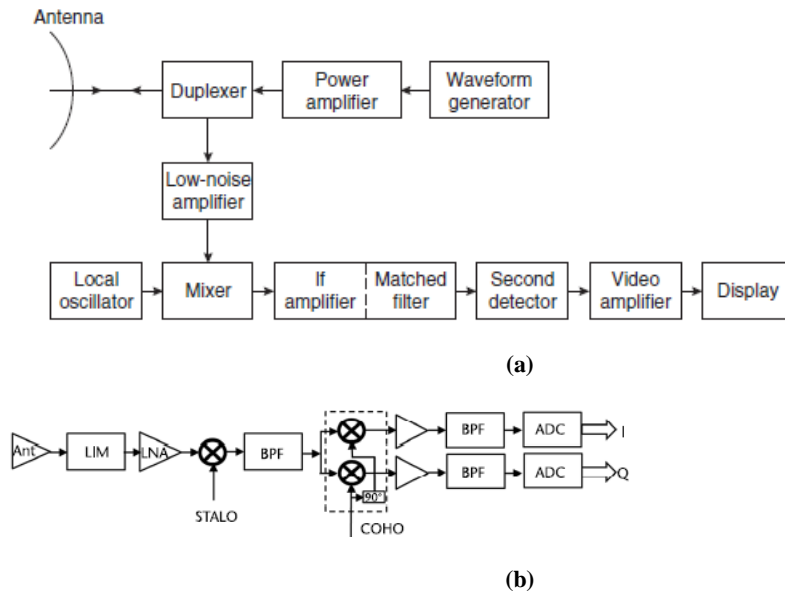


Figure 2-1 (a) radar Basic diagram and (b) Coherent receiver [8,7]

In a monostatic arrangement, one antenna handles T/R. In a bistatic arrangement, different antennas are used. The preceding subsection summarizes the radar-related parameters.

2.2 AESA Radar Beamsteering

Modern naval vessels increasingly employ AESA radars, which form and steer their beam by applying programmable complex weights (amplitude and phase) across many T/R modules, thereby removing any reliance on mechanical rotation and enabling millisecond-scale revisits of priority threat sectors. In the architecture shown in Fig. 2-2, each T/R module typically integrates a solid-state power amplifier for transmit, a low-noise amplifier for receive, a T/R switch or circulator, and digitally controlled phase and amplitude control [7]. These modules are synchronized and calibrated by a central exciter/Local Oscillator (LO), then digitized to support real-time beamforming and adaptive processing. Compared to mechanically steered antennas, AESA radars provide rapid, inertia-free steering, flexible multi-beam and multi-mission scheduling, and graceful degradation: if individual elements fail, overall performance degrades gradually rather than

through a single mechanical point of failure, reducing dependence on moving parts and the associated maintenance burden.




Figure 2-2 AESA radar diagram [7]

Fundamentally, an AESA radar is a phased array. The total radiated field pattern $F(\theta, \phi)$ is the product of the element factor $E_e(\theta, \phi)$ (the single-element pattern) and the array factor $AF(\theta, \phi)$, the latter determined solely by the element locations and their complex weights




Figure 2-3 Linear series of N isotropic antenna elements separated by a constant distance s [2]

Fig 2-3 shows a uniform linear array (ULA) made of N isotropic radiating elements with a constant inter-element spacing labeled s in the figure. The array is a straight line of elements; the reference direction $\theta = 0$ is taken as broadside, meaning normal to the axis of the array [2]. If the array is electronically steered to an angle θ_0 (measured from broadside) using uniform excitation, the phase progression between adjacent elements is chosen to align the main beam in the desired look direction.

For this ULA with spacing s , the array factor can be written as [2]

$$AF(\theta) = \frac{1}{N} \frac{\sin\left(\frac{N}{2}ks[\sin\theta - \sin\theta_0]\right)}{\sin\left(\frac{1}{2}ks[\sin\theta - \sin\theta_0]\right)} \quad (2.1)$$




Figure 2-4 Uniform rectangular array geometry [2]

For a uniform rectangular array (URA), seen in Fig 2-4, consisting of $M \times N$ elements with spacings d_x , and d_y along the x and y axes, respectively, steered to (θ_0, ϕ_0) , it is convenient to use direction cosines $u = \sin\theta\cos\phi$, and $v = \sin\theta\sin\phi$. With weights $w_{m,n} = |A_{m,n}| e^{j\phi_{m,n}}$ applied to element (m,n) at location (md_x, nd_y) , the array factor is:

$$AF(u, v) = \frac{1}{MN} \sum_{m=0}^{M-1} \sum_{n=0}^{N-1} |A_{m,n}| e^{jk[md_x(u-u_0)+nd_y(v-v_0)]} \quad (2.2)$$

Beam steering has practical limits that directly affect detection performance. First, as scan angle increases, the main-beam directivity falls, and the beamwidth broadens, reducing gain in the commanded look direction and increasing susceptibility to interference and

clutter returns outside the main lobe [7]. The directivity of an AESA system is derived from its total radiation pattern:

$$D(\theta, \phi) = \frac{4\pi|F(\theta, \phi)|^2}{\int_{4\pi}|F(\theta, \phi)|^2 d\Omega} \quad (2.3)$$

where Ω denotes solid angle and $d\Omega$ is the differential solid-angle element. In spherical coordinates, $d\Omega = \sin\theta d\theta d\phi$, and the integral over 4π indicates integration over all directions. Gain is related to directivity by $G(\theta, \phi) = \eta D(\theta, \phi)$, where η is the total efficiency accounting for aperture and ohmic losses.

Grating lobes are spatial-aliasing lobes that can appear when the inter-element spacing exceeds $\frac{\lambda}{2}$ in the scan plane. When the beam is steered, these lobes can approach the main-beam level and redirect transmit or receive gain into ambiguous directions, which is unacceptable for wide-angle search [7]. Sidelobes, in contrast, are the secondary maxima that exist even when grating lobes are avoided, arising from the finite, discretized aperture and the chosen element weighting. This distinction is illustrated in Fig. 2-5: as spacing increases from $d = \frac{\lambda}{4}$ toward $d = \lambda$, additional strong lobes emerge and angular selectivity degrades, demonstrating why spacing must be treated as a baseline constraint for practical scanning arrays [7].

In a monostatic phased-array radar, the two-way spatial response is governed by the product of the transmit and receive array factors, so unwanted lobes can admit clutter and interference on both transmit and receive paths and thereby reduce detection margin [1]. For this thesis, grating-lobe avoidance is treated as a baseline array-design constraint by assuming dense element spacing at or below $\lambda/2$ across the intended steering sector. With grating lobes controlled by design, the remaining practical lever for reducing sidelobe levels is amplitude tapering, which suppresses sidelobes at the cost of increased mainlobe width and reduced aperture efficiency [1], [7].




Figure 2-5 Array factor power patterns for a 10-element broadside binomial array with $N = 10$ and $d = \lambda/4, \lambda/2, 3\lambda/4,$ and λ . [7]

One classical tapering approach is the binomial (Pascal) distribution, in which the element amplitudes follow coefficients of $(1 + x)^{m-1}$. For an element spacing $d \leq \frac{\lambda}{2}$ binomial arrays can achieve no sidelobes, at the cost of a wider beamwidth. The half-power beamwidth is commonly approximated by [7]:

$$HPBW_{binominal} = 1.06 \left(\frac{2L}{\lambda} \right)^{0.5} \quad (2.7)$$

where L is the total aperture length.

A widely used alternative is the Dolph–Tschebyscheff (Chebyshev) taper, which enforces an equal-ripple sidelobe level by choosing excitations associated with Chebyshev polynomials. In continuous-aperture form, the corresponding space factor is [7]:

$$SF(\theta) = \frac{\cosh[(\pi A)^2 - u^2]^{0.5}}{\cosh(\pi A)}, u = \frac{\pi L}{\lambda} \cos\theta \quad (2.8)$$

A is selected such that $\cosh(\pi A) = R_0$. Here, R_0 is the peak sidelobe ratio, defined as the ratio of the mainlobe peak field magnitude to the maximum sidelobe field magnitude, i.e., $R_0 = |F_{main,peak}/F_{side,peak}|$ which corresponds to a sidelobe level of $20\log_{10}\left(\frac{1}{R_0}\right)$ dB relative to the mainlobe peak [7]

The Taylor distribution is often preferred when the design goal is not only low peak sidelobes, but also rapidly decaying off-set sidelobes (useful for suppressing distant clutter and spurious returns). In the Taylor design, the first $n - 1$ inner sidelobes are constrained to a constant field ratio $\frac{1}{R_0}$, whereas higher-order sidelobes decay monotonically with angle. For a continuous line source of length L , the space factor is [7]:

$$SF(\theta) = \begin{cases} (L/((\pi B)^2 - u^2)^{0.5} \sinh[(\pi B)^2 - u^2], u^2 < (\pi B)^2 \\ (L/(-(\pi B)^2 + u^2)^{0.5} \sin[-(\pi B)^2 + u^2], u^2 > (\pi B)^2 \end{cases}, u = \frac{\pi l}{\lambda} \cos\theta \quad (2.9)$$

where $\frac{\sinh(\pi B)}{\pi B} = R_0$ to enforce the selected peak sidelobe level [3].

2.3 AESA implementation scale and modelling assumptions

Operational AESA radars are implemented as large collections of repeatable transmit-receive (T/R) modules rather than a single monolithic RF chain [9]. As illustrated conceptually in Fig. 2-2, each module contains the front-end transmit and receive functions needed for electronic beamforming, and the full antenna face is formed by combining many such modules into a common aperture. the AMSAR X-band AESA demonstrator is a representative example of this implementation approach. In that system, each T/R module is less than 4 cm^3 in volume and has a mass of 15 g, and these modules are combined to form a circular array face of 0.6 m diameter, as shown in Fig. 2-6 [9]. This is important for the present work because it shows that practical AESA faces are built from many small, standardized elements, which supports modelling the shipboard radar face as a modular aperture composed of a large number of repeatable T/R elements rather than as a single transmit and receive chain.

The electronic field of view of a single AESA face is typically limited to $\pm 60^\circ$ in azimuth and elevation, with a wider field of obtained through the use of multiple faces or mechanical repositioning in some applications [9]. This motivates the steering-sector limits

adopted in the shipboard AESA model and explains why scan loss and grating lobe risk at large steering angles must be explicitly accounted for in the performance comparison.



Figure 2-6 AMSAR demonstrator antenna array [9]

Naval integration examples further support these modelling assumptions. The TRS-4D on the German F125 frigate is a naval AESA system that uses four fixed faces to provide full 360° coverage and redundancy [10]. Literature also shows that multiple digitally generated receive beams can be processed in parallel and electronically stabilized to compensate for ship motion [10]. These features are directly consistent with modelling the radar as an electronically steered face capable of multi-beam receive processing within a bounded field of view, and an electronically scheduled search dwell budget.



Figure 2-7 TRS-4D AESA radar [10]

Finally, aperture size to beamwidth considerations show that achieving narrow surveillance beams at X band requires meter-class apertures and correspondingly high array directivity. For that reason, the simulated AESA face is constrained to an area of 1 m^2 so that it remains consistent with a ship-installable radar aperture while still supporting a surveillance-relevant beamwidth. At this stage, the important modelling assumption is the beamwidth associated with a meter-class X-band AESA face over the intended steering sector, rather than a unique array layout or element count. The specific array instantiation used for the candidate radars is introduced later in Chapter 3 when the common constraint set is converted into a concrete implementation. Once the aperture limit is imposed and dense element spacing is adopted as a baseline design constraint for grating-lobe control, the corresponding element count follows directly from geometry. In this way, the assumed face size and implementation scale remain consistent with literature examples such as the AMSAR demonstrator, namely a large aperture formed from many small and repeatable T/R elements [9].

The discussion above concerns spatial response and array implementation constraints. With AESA implementation scale, scan limits, face size, and element count established, the next subsection defines the pulsed waveform and timing parameters used as the first candidate in the ship installable counter UAV comparison framework.

2.4 Pulsed Radar System

A pulsed radar transmits short bursts of RF energy and then listens for the returning echo during the interval between pulses. Range is determined from the round-trip travel time of the echo and, with a coherent receiver, radial velocity is inferred from the phase change measured across successive pulses, as illustrated in Fig. 2-8, where R is the target range, τ_{out} is the outward travel time, τ_{back} is the return travel time, and c is the speed of light [2,7,8]. In the monostatic arrangement considered, the same antenna is used for both transmission and reception. This architecture concentrates energy in time and is well suited to coherent detection, but it also introduces practical timing constraints. In particular, while the radar is transmitting, and during the short recovery time immediately afterward, the receiver cannot observe echoes, which creates a finite near-range blind interval for very close targets. Pulsed radar is examined here because it remains widely used in surveillance and naval radar systems [2,7,8]. The maximum range of a pulsed radar is determined by the radar range equation in terms of transmitted power P_t , antenna gain G , target radar cross section σ , minimum detectable signal S_{min} , system losses L , and wavelength λ , where $\lambda = \frac{c}{f}$. For a single unmodulated pulse radar, the radar range equation is given by [7]:

$$R_{max} = \left(\frac{P_t G A_e \sigma}{(4\pi)^2 S_{min} L} \right)^{0.25} \quad (2.10)$$

where

$$S_{min} = k T_0 B F_n SNR \quad (2.11)$$

and

$$A_e = \frac{G \lambda^2}{4\pi} \quad (2.12)$$

where k is the Boltzmann's constant, T_0 is the standard equipment temperature of 290 K, B is the receiver bandwidth in Hz, F_n is the noise factor (unitless), SNR is the signal-to-noise ratio required to detect a signal, and finally, R_{max} is the maximum detection range (m). For a target to be detected by a receiver, the signal strength must be above the radar noise. Analysis indicates that an SNR of 14 dB is necessary to reduce clutter-induced false alarms [2]. By using (2.7), (2.8), (2.9), for example, with a 10 GHz radar transmitting at 1 MW, assuming 40 dB gain, a target RCS of 0.01 m², a minimum SNR of 14 dB, system losses of 12 dB, and a 5 dB noise figure with a receiver bandwidth of 20 kHz, the maximum detection range is estimated to be about 46 km. Enhancing detection range can be achieved

by improving antenna gain or by applying advanced processing, most notably pulse integration, which combines multiple echoes to increase the SNR even further (processing gain) [7].




Figure 2-8 Determination of target range [7]

Pulse Repetition Interval (PRI) is the delay between pulses. The Pulse Repetition Frequency (PRF) is how often a pulse is transmitted and is $1/PRI$. Due to these factors, a system has a maximum Unambiguous Range R_{unamb} , which is defined as the minimum range that the system can distinguish between returns from two separately transmitted pulses:

$$R_{unamb} = \frac{c}{2PRF} \quad (2.13)$$

The system transmits a pulsed RF signal for a duration equal to its pulse width τ ; therefore, a system's minimum range R_{min} and the radar range resolution ΔR are:

$$R_{min} = \frac{c\tau}{2} = \Delta R \quad (2.14)$$

Blind speed is a specific radial velocity at which a moving target becomes undetectable by a moving target indication (MTI) radar system lacking Doppler processing. This occurs when the radial velocity V_n in (m/s) is a multiple n of a factor:

$$V_n = \frac{n\lambda}{2PRI} = \frac{n\lambda PRF}{2} \quad n = 1, 2, 3 \dots \quad (2.15)$$

Pulsed radar performance is also limited by the maximum ambiguous velocity [7]:

$$V_{amb} = \frac{\lambda PRF}{4} \quad (2.16)$$

As discussed in Section 2.1, modern radar systems use a coherent receiver: the received signal is down-converted into I and Q components which are two baseband signals 90° apart that together retain amplitude and phase. Preserving the phase enables Doppler estimation by examining the phase progression across pulses. Doppler is estimated by taking a Fast Fourier Transform (FFT) across a block of N_p pulses. The coherent processing interval or dwell time is $T_{coh} = N_p PRI$. The Doppler bin spacing is $1/T_{coh}$, and since two-way Doppler is $f_d = 2v/\lambda$, the radial-velocity resolution is [7]

$$\Delta v_{pulse} = \frac{\lambda}{2N_p PRI} = \frac{\lambda}{2T_{coh}} \quad (2.17)$$

Modern pulsed surveillance radars transmit a longer pulse to increase transmitted energy $E_t = P_t \tau$ and apply pulse compression at the receiver to recover fine range resolution. Linear frequency modulation (LFM) is the most common pulse-compression approach because it is simple to generate, robust in hardware, and provides a predictable matched-filter response [2,7,8]. The central idea is that the radar transmits a pulse of duration τ with its instantaneous frequency swept over a wide bandwidth B so that the pulse energy increases with τ while the range resolution is set primarily by B , not by τ .

A baseband LFM pulse can be written as:

$$s(t) = \text{rect}\left(\frac{t}{\tau}\right) e^{j2\pi\left(f_0 t + \frac{\mu}{2} t^2\right)}, 0 \leq t \leq \tau \quad (2.18)$$

where f_0 is the starting frequency offset and $\mu = B/\tau$ is the chirp rate in Hz/s.

Fig. 2-9 shows an LFM transmit pulse.

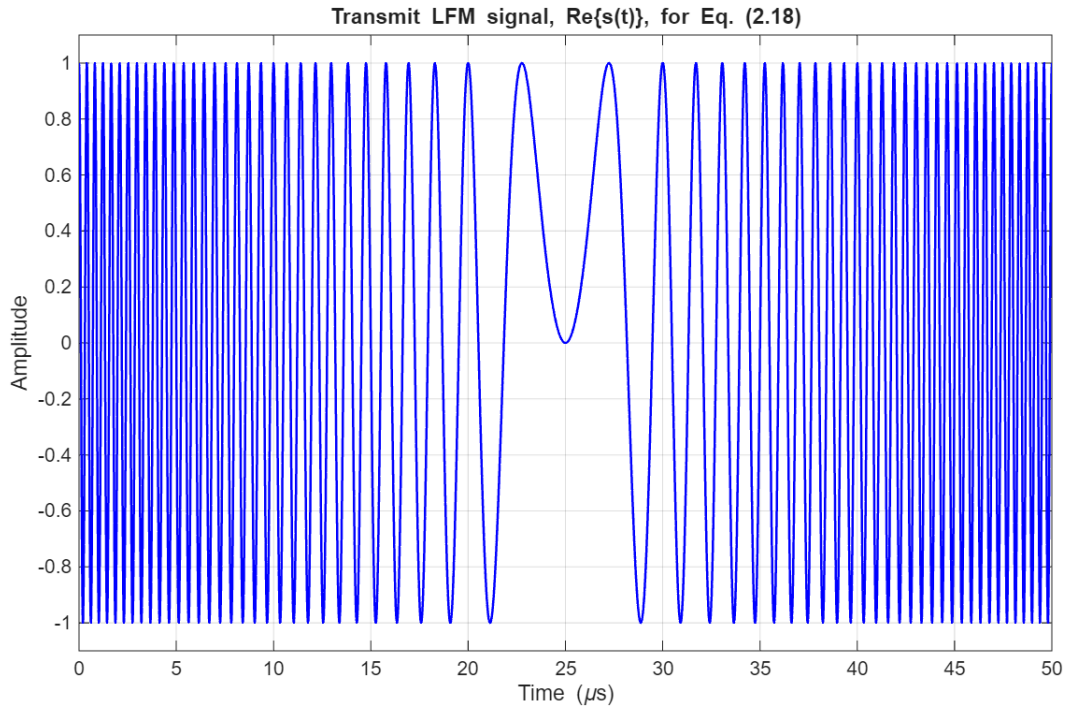


Figure 2-9 LFM Tx Signal

Fig. 2-10 shows the matched-filter output for the LFM pulse defined in (2.18). In Fig. 2-10(a), the full matched-filter response is plotted as a function of delay. The dominant peak occurs at the echo delay, indicating that the matched filter compresses the long-transmitted pulse into a localized response at the target return time. Away from this delay, the response falls rapidly, while the smaller oscillations around the main peak are the sidelobes associated with pulse compression. Fig. 2-10(b) shows a zoomed view of the same response near the peak. This panel is included to show that the compressed mainlobe width is proportional to $1/B$, consistent with the standard LFM pulse-compression result in (2.19). Therefore, although the transmitted pulse duration is τ , the receiver resolves the target using the much narrower compressed response determined primarily by waveform bandwidth [2,7,8].

$$\Delta R_{LFM} \approx \frac{c}{2B} \quad (2.19)$$




Figure 2-10 Matched-filter response of the transmitted LFM pulse defined in (2.18). (a) Full normalized matched-filter magnitude versus delay, showing the compressed echo peak at the target delay and the surrounding sidelobe structure. (b) Zoomed view of the compressed mainlobe, illustrating the measured 3 dB width and its correspondence to the expected order of $\frac{1}{B}$.

Another advantage of LFM is the processing gain that is associated with its matched filter [2,7,8]:

$$G_{pg} \approx \tau B \quad (2.20)$$

The discussion of pulsed radar shows that detection performance is governed primarily by the integrated energy on target and the resulting post-processing SNR, rather than by peak transmit power alone. In practice, pulsed systems trade short on-target illumination for higher peak power and coherent integration across multiple pulses to accumulate sufficient energy within each revisit. LFM pulse compression provides the practical shipboard baseline for this approach: it increases energy-on-target by using longer pulses while preserving fine range resolution through waveform bandwidth.

2.5 Pulsed Radar Systems in Literature

The preceding pulsed-radar discussion can be grounded with two representative examples. These examples are not adopted directly as design templates; rather, they are used to show how practical shipboard radars balance dwell time, waveform selection, clutter performance, and detection range under real operating constraints.

As a first example, a low-power solid-state X-band marine radar uses a composite three-pulse waveform to improve small-target detection in sea and coastal clutter through coherent pulse-Doppler processing [11]. The study explicitly links scan mechanics to available dwell time: with a 48-rpm antenna rotation rate and a 1.8° horizontal beamwidth, the radar has an available dwell time of 6.25 ms per azimuth cell and a full refresh time of 1.25 s [11]. This is directly relevant because it shows that, even before detailed signal processing is considered, the available coherent integration time is already constrained by the surveillance geometry and scan schedule.

Table 2-1 summarizes the pulse parameters shown in [11], while Fig. 2-11 shows the corresponding operational range regions of the short, medium, and long-pulse modes. The discontinuities in Fig. 2-11 are a direct consequence of the composite three-pulse design. They do not indicate an error or a physical discontinuity in propagation; rather, they reflect the fact that the radar switches between three different pulse widths, each optimized for a different range regime. The short pulse is used at near range because its small pulse width reduces the minimum blind range, but it carries less transmitted energy. The medium and long pulses provide more energy on target and therefore better long-range detection performance, but they also increase the blind range and change the usable operating region. As a result, the overall coverage is piecewise rather than continuous, with transition boundaries where one pulse becomes preferable to another [11]. This is an important practical point because it shows that pulsed radars often require multiple waveform modes to span the full operational range interval.

TABLE 2-1 WAVEFORM PULSE PARAMETERS: PEAK POWER, PULSE WIDTHS[15]

	X Band	S Band
Peak Power	71W	46 W
Short pulse	233 ns	233 ns
Medium pulse	19.44 μ s	18.46 μ s
Long pulse	72.02 μ s	71.29 μ s
PRF	1329	1329

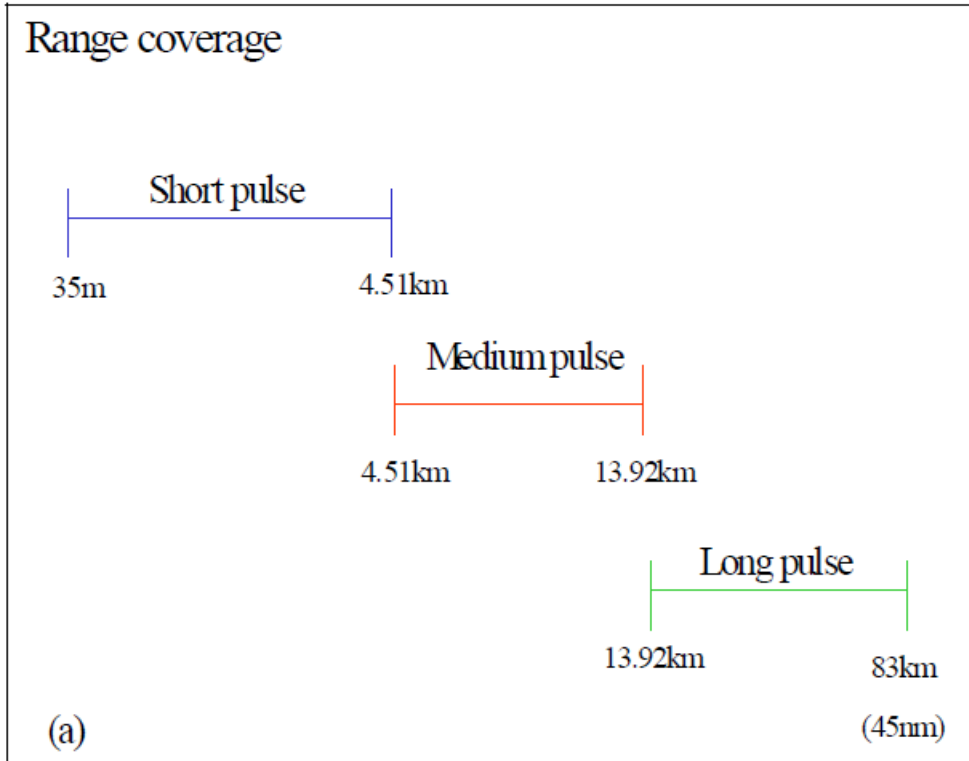


Figure 2-11 Range coverage for X band waveform [11]

Fig. 2-12 further illustrates the performance limits of this radar at $P_d = 0.8$ [11]. The results show that a 10 m^2 target can be detected at 18.5 km, whereas a 5000 m^2 target can be detected at 90 km [11]. The key implication is not that these values transfer directly to a shipboard C-UAV radar, but that they demonstrate how strongly detection range contracts as target RCS decreases. A target with an RCS representative of a small UAV, -20dBsm , is 30 dB below a 10 m^2 target. Under the radar-range scaling $R_{\max} \propto \sigma^{0.25}$, this corresponds to an approximate 5- to 6-fold reduction in detection range. This supports the conclusion that a marine radar in the roughly 70 W peak-power class is not, by itself, an adequate baseline for tens-of-kilometre UAV warning unless it is paired with substantially greater effective isotropic radiated power (EIRP), where EIRP is the transmit power multiplied by antenna gain, and or additional coherent processing gain [11,12].




Figure 2-12 Minimum RCS detectable at a given range with 80% P_d 10 m² RCS target is detected at 18.5 km and a 5000 m² RCS target is detected at 90 km [11]

Next, consider an RCN specific example, consisting of the Thales SMART-S Mk2, which is a 3D surveillance radar installed on the Canadian Patrol Frigate (CPF) class of ships and is shown in Fig 2-13. The SMART-S Mk2 data list shows that the system transmits a horizontal beamwidth of 2° and selectable illumination patterns, which reinforces that beam shaping is part of practical clutter management in naval radar design, not just a theoretical antenna concept [13]. The important takeaway is that naval pulsed radars manage clutter-limited performance through beam shaping, coherent pulse-Doppler processing, and waveform or mode selection for different operating conditions.



Figure 2-13 Smart-S MK2 [13]

Viewed together, these two examples sharpen the pulsed-radar design narrative. The available coherent dwell is constrained by scan rate and beamwidth, composite pulse sets are often needed to balance minimum range and long-range detection performance, and operational systems trade refresh rate against surveillance range. These points motivate the pulsed waveform assumptions used and provide a practical basis for comparison with FMCW.

2.6 FMCW Radar System

Frequency Modulated Continuous Wave (FMCW) radar provides an alternative approach that accumulates energy through continuous illumination rather than discrete bursts. For a ship-installable C-UAV radar, FMCW is therefore evaluated as the second candidate architecture under the same overall design constraints used for the pulsed comparison. In this configuration, the transmitter emits a continuous linear chirp while a separate receiver captures the returned echo. Figure 2-14 shows a simplified coherent radar receiver concept. A continuous-wave oscillator generates the transmitted signal at frequency f_t and also provides the reference signal to the receiver. The transmitted wave propagates toward the target, and the returned echo is received with a frequency shift of $f_t \pm f_d$, where f_d is the Doppler frequency caused by target radial motion. Inside the receiver, the returned signal is compared with the reference signal from the oscillator, and the resulting difference frequency f_d is passed to the indicator for display. The figure therefore illustrates the basic principle of coherent Doppler extraction: the radar measures target motion by comparing the received echo with a stable reference derived from the transmitter.

For the FMCW case, the same comparison principle is extended to a linearly swept waveform. The FMCW candidate is treated as a bistatic shipboard AESA implementation in which the transmit and receive antennas are separated by one metre. By mixing the received echo with a reference derived from the transmitted chirp, the radar forms a beat frequency f_b that maps to target range, while Doppler is estimated from phase progression across successive chirps [2,7,8]. The main implementation constraints are transmitter-to-receiver leakage, sufficient beat-frequency headroom in the analog and digital processing chain, and the need to separate range and Doppler contributions during signal processing [2,7,8]. For a linear sweep of the form $f(t) = F_0 + C_s t$, where $C_s = B/T_c$ is the chirp slope, B is the total frequency excursion of one chirp (the sweep bandwidth), and T_c is the time required to complete that sweep (the chirp duration), a target at range R produces a beat frequency given by [2,7,8]:

$$f_b = \frac{2C_s R}{c}, \text{ if no doppler effect} \quad (2.21)$$

Therefore, range can be determined from the beat frequency:

$$R = \frac{c f_b}{2C_s} \quad (2.22)$$

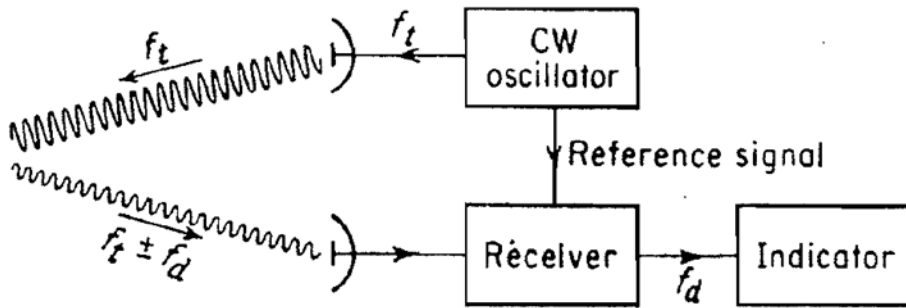


Figure 2-14 Basic FMCW diagram [2]

A single-ramp FMCW waveform may be written in complex baseband form as

$$s_{FMCW}(t) = \exp \left\{ j 2\pi \left[f^0 t + \left(\frac{\mu}{2} \right) t^2 \right] \right\}, 0 \leq t \leq T_c$$

where f_0 is the sweep start frequency, T_c is the chirp duration, B is the sweep bandwidth, and $\mu = B/T_c$ is the chirp slope. The instantaneous frequency is therefore $f(t) = f_0 + \mu t$, which varies linearly over the chirp interval. This shows that the basic FMCW transmit signal has the same quadratic-phase chirp form as an LFM signal, but in FMCW radar the chirp is transmitted continuously and processed by mixing the received echo with a reference copy of the transmitted waveform as seen in Fig. 2-15.

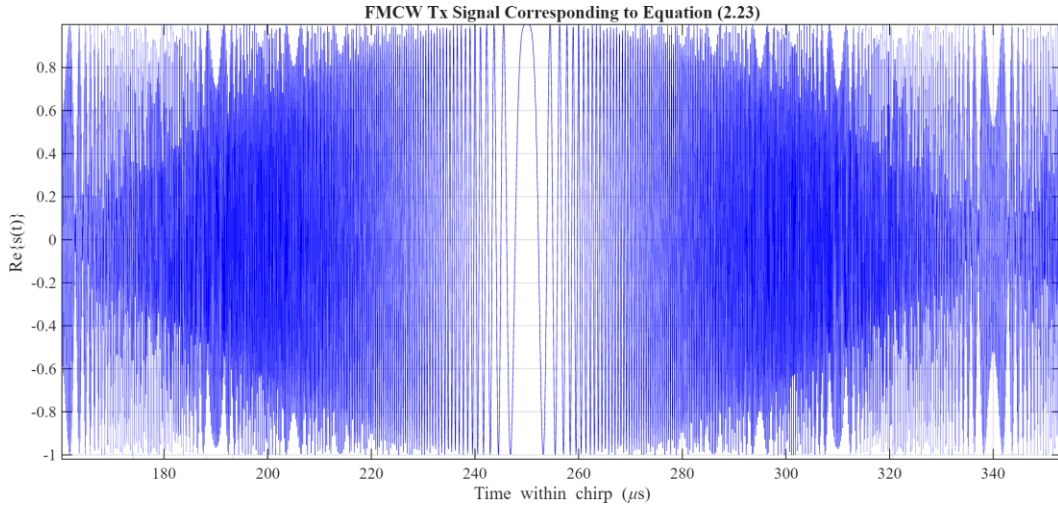


Figure 2-15 FMCW Tx Signal

Range resolution is set by sweep bandwidth

$$\Delta R = \frac{c}{2B} \quad (2.24)$$

The maximum unambiguous range is limited by the maximum beat frequency that the ADC can sample

$$R_{max} = \frac{c f_{bmax}}{2C_s} \quad (2.25)$$

Doppler processing in FMCW radar is performed across a sequence of repeated chirps. The time between adjacent chirps is the chirp repetition interval, denoted T_{rep} , and a block of N_c chirps forms the coherent processing interval used for Doppler estimation. In contemporary FMCW processing, range is first estimated from each individual chirp using fast-time processing, and Doppler is then estimated from the phase progression across the

repeated chirps using a slow-time FFT, as shown in Fig. 2-16. The coherent processing interval is therefore $T_{coh} = N_c T_c$. Which gives the radial-velocity resolution [14]:

$$\Delta v_{FMCW} = \frac{\lambda}{2N_c T_c} = \frac{\lambda}{2T_{coh}} \quad (2.25)$$

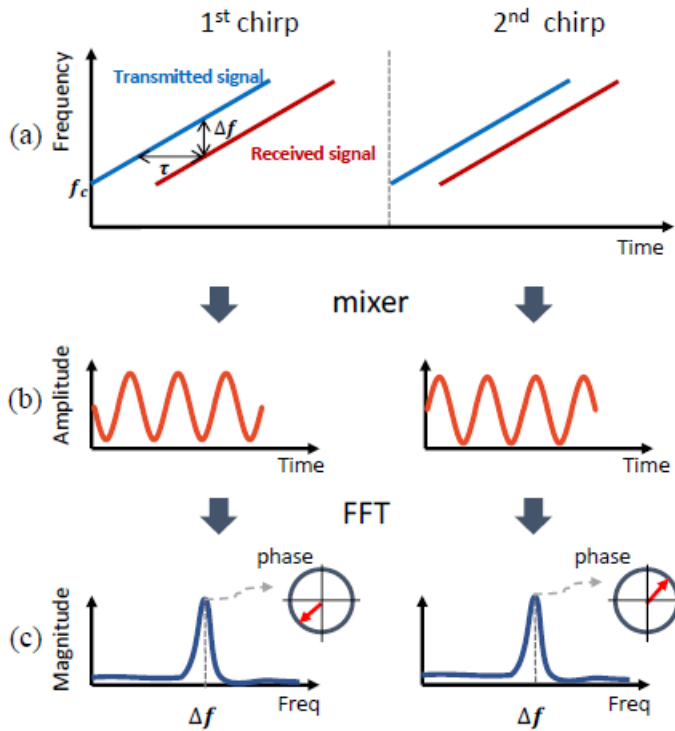


Figure 2-16 (a) The figure shows the transmitted FMCW signal (blue) and its reflected signal (red). The round trip r_t seen as τ in the diagram between the two signals maps to the frequency shift Δf . **(b)** Mixing the transmitted and reflected signals yields the beat signal f_b whose frequency is Δf . **(c)** The FFT peak corresponds to Δf . A subtle change in the distance over the chirps shifts the phase of the FFT peak [14]

The coherent interval T_{coh} also limits the amount of coherent integration, and thereby processing gain available in the FMCW processing chain [8,14].

$$G_{coh} \approx BT_{coh} \quad (2.26)$$

The same concept applies to pulsed radar: coherent integration over a CPI improves the SNR, but the signal energy is only present during the pulse, whereas FMCW maintains continuous illumination during the ramp [14]. As such, FMCW allows the ADC to provide far more samples of the target than could be achieved by a pulsed radar. As such, the peak power required for an FMCW radar is generally lower than that of a pulsed radar.

2.7 FMCW Radar System in Literature

A dual-polarized FMCW radar for the detection of small-RCS UAVs operating close to the sea surface in a cluttered maritime background models sea clutter using non-Rayleigh statistics and applies an adaptive CFAR detector, shown in Fig. 2-17, to maintain detection performance in a non-homogeneous clutter background [15]. Rotor-induced micro-Doppler (MD) signatures are exploited in the time-frequency domain to improve discrimination between UAV echoes and clutter returns [15]. Here, micro-Doppler refers to the additional Doppler modulation caused by moving target substructures, such as rotating blades, superimposed on the bulk translational Doppler of the UAV.

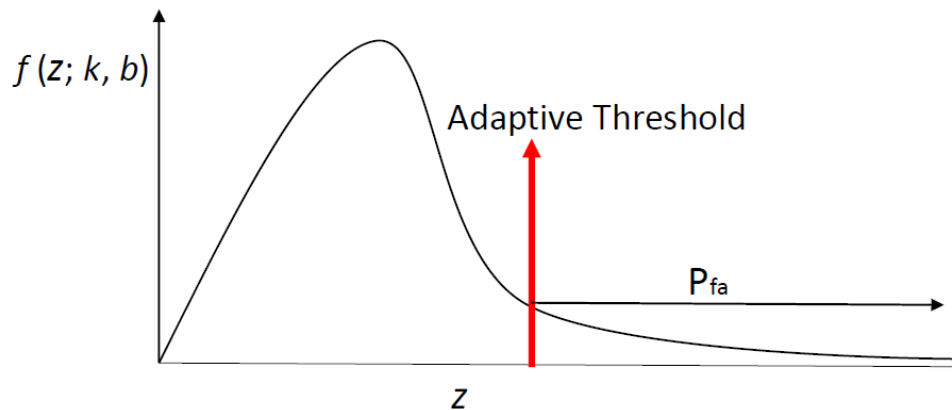


Figure 2-17 Adaptive CFAR threshold [15]

The example system operates at 24 GHz with a sweep bandwidth of 0.5 GHz, a sweep time of 1 ms, a chirp repetition frequency of 1 kHz, a transmit power of 12 dBm, and an antenna gain of 20 dBi, which corresponds to a range resolution of 0.3 m [15].




Figure 2-18 Experimental UAV detection on the Range-Doppler map with $P_{fa} = 10^{-5}$ and $N = 20$. The experiment is carried out using 25 GHz Ancortek radar [15].

This example further supports the requirement that low-grazing-angle UAV detection be evaluated in a clutter-limited framework, with clutter statistics and CFAR processing included explicitly in the performance assessment rather than by considering signal-to-noise ratio alone.

A second relevant FMCW example is the Alarm™ staring radar, which operates between 4 to 6 GHz, with 100 W of transmitted power, a field of view of 120° in azimuth by 30° in elevation, and multiple narrow receive beams, as shown in Fig. 2-19 [16]. In its scanning mode, the transmit beam is mechanically steered through a sequence of discrete look directions, so each direction is illuminated only briefly during a surveillance update. In its staring mode, the transmit beam remains fixed over the coverage sector while multiple static receive beams continuously observe that illuminated region [16]. The significance of this distinction is that persistent observation in staring mode provides a longer coherent processing interval (CPI), thereby improving the resolution to extract MD or blade-flash

features for discrimination between UAVs and birds when conventional detection variables alone are insufficient [16].




Figure 2-19 Alarm radar [16]

Very long coherent processing intervals do not always improve performance, because target decorrelation and migration through range or Doppler cells can offset the gain from longer coherent integration [16]. FMCW performance must be assessed together with coherent processing design and target fluctuation effects, which supports the use of coherent integration and MD processing [16].

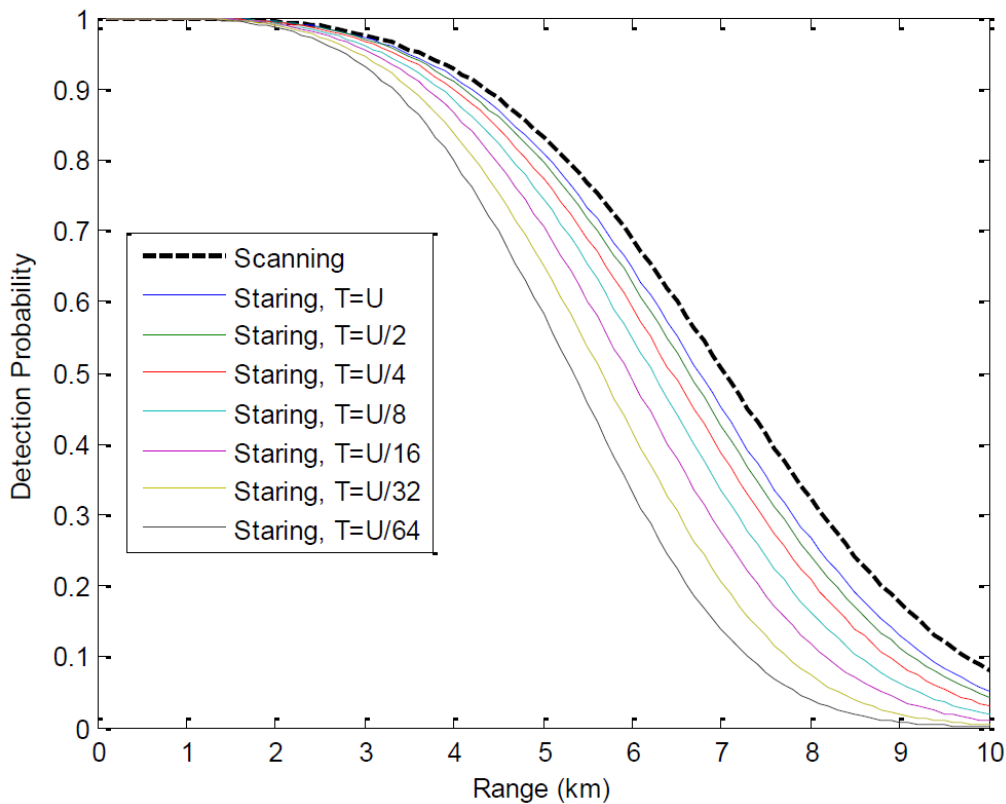


Figure 2-20 A comparison between the probability of detection of the scan1 radar and stare1 radars with CPIs (T) set to different fractions of U (surveillance update interval) [16]

As shown in Fig. 2-20, The Alarm™ radar is primarily a short-range UAV sensor, given that the probability of detection decreases rapidly beyond only a few kilometres, and the scanning benchmark outperforms the staring configuration over the plotted range [16], [17]. As such, a floodlight FMCW radar with a transmit power on the order of 100 W does not, by itself, provide the tens-of-kilometre warning horizon that is desired for shipboard counter-UAV surveillance.

This section has demonstrated that FMCW remains an attractive candidate for shipboard counter-UAV sensing because it maintains continuous illumination during each dwell and achieves fine range resolution through sweep bandwidth rather than very short pulses. Its practical limitations are more naturally stated in receiver and implementation terms. First, leakage from the transmitter into the receiver can elevate the near-range interference floor and reduce sensitivity to weak close-in targets. Second, the intermediate-frequency and ADC bandwidth must be large enough to accommodate the expected beat frequencies associated with the required detection range. These constraints do not invalidate FMCW

operation, but they do show that detection performance depends not only on waveform theory, but also on isolation, receiver dynamic range, and signal-processing design.

2.8 Radar Cross Section

RCS refers to a target's effective area as the radar perceives it. When an RF signal is transmitted from an antenna and strikes a target, the target absorbs some of the energy while the rest is scattered. The ratio of the scattered power to the reflected power back to the radar is proportional to the RCS [2,7]:

$$\sigma = \lim_{R \rightarrow \infty} \left(4\pi R^2 \frac{|E_s|^2}{|E_o|^2} \right) \quad (2.28)$$

where E_s is the scattered electric field observed in the far field in the direction of observation, and E_o is the incident electric field at the target location. In the monostatic case, the scattered field is evaluated in the backscatter direction, that is, back toward the radar. Simple objects allow relatively straightforward RCS calculation. However, complex objects, characterized by irregular geometries, heterogeneous material properties, and intricate internal structures, pose significant challenges for accurate electromagnetic modelling. One approximate approach is to subdivide the object into a series of simpler shapes and evaluate each at different orientations. However, this does not fully capture the electromagnetic interactions among the constituent parts, which can significantly influence the total scattering behavior [7]. In contrast, computational electromagnetics methods solve the scattering problem numerically and are therefore better suited to complex targets. For small UAVs, the RCS is not constant; it varies with aspect angle, frequency, polarization, and aircraft attitude in flight. Rotor blades, attitude changes, and material composition can therefore cause the echo to fluctuate with time. To represent this behavior in simulation, statistical models such as chi-square or lognormal distributions may be fitted to simulated or measured data, providing a practical and computationally efficient description of fluctuating target reflectivity [12].

Physical size does not uniquely determine RCS as it varies with shape, materials, aspect, frequency, and polarization, so small airframes can appear large at some aspects and weak at others. To obtain repeatable, angle- and frequency-dependent data, RCS measurements were taken in a compact-range anechoic chamber of three small UAVs, Fig. 2-21, listed in Table 2-2 [12]. Figure 2-22 shows the monostatic RCS of each target over a full 360° in azimuth at 15 GHz (blue) and 25 GHz (red). The curves peak near 0° and 180° (nose-on and tail-on), where structural features align with the radar line of sight and produce RCS flashes [2,12]. Away from those headings the RCS falls and varies with aspect. The 25 GHz trace and variance are generally higher because smaller features reflect more strongly

at the shorter wavelength. The drone RCS peaks are on the order of only 0.01 m², whereas an anti-ship missile has a typical RCS of 1 m² = 0 dBsm. [12].

TABLE 2-2 PARAMETERS OF THREE SMALL UAVS [12]

Drones	Trimble ZX5	DJI Inspire 1 Pro	DJI Phantom 4 Pro
Width in mm	850	438	289.5
Length in mm	490	451	289.5
Volume in m ³	0.15	0.059	0.016
15 GHz RCS	-14.39 dBsm ≈ 0.036 m ²	-14.24 dBsm ≈ 0.0378 m ²	-15.03 dBsm ≈ 0.031 m ²
25 GHz RCS	-9.64 dBsm ≈ 0.109 m ²	-11.09 dBsm ≈ 0.0778 m ² ;	-12.40 dBsm ≈ 0.0576 m ²




Figure 2-21 Small drones considered: (a) Trimble ZX5, (b) DJI Inspire 1 Pro, (c) DJI Phantom 4 Pro (DJI P4 Pro) [12]

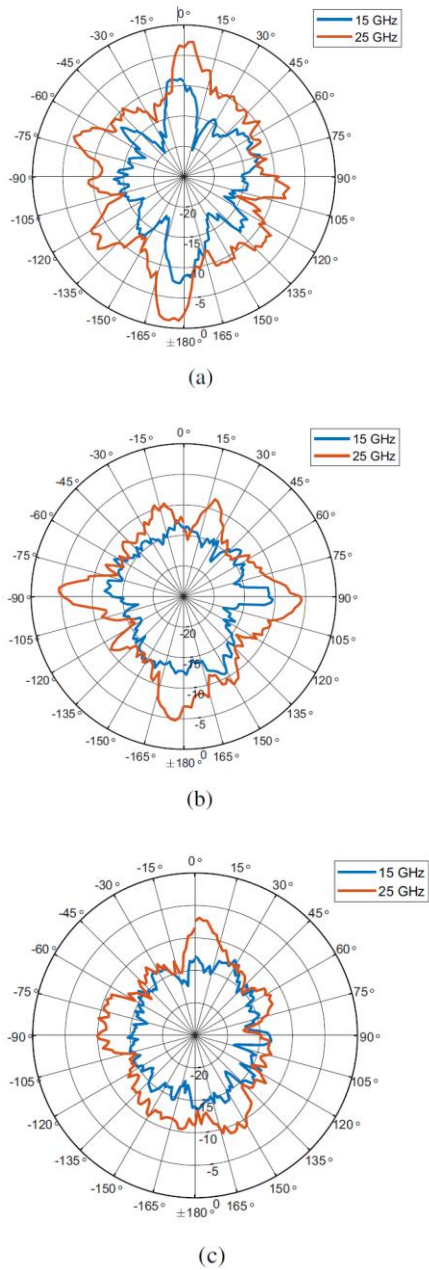


Figure 2-22 The measured RCS (dBsm) versus azimuth angles for the small drones: (a) Trimble ZX5, (b) DJI Inspire 1, (c) DJI Phantom 4 Pro [12]

The variability of the RCS becomes even more severe in real flight because the UAV attitude is not perfectly steady, rather, it exhibits wave-like behavior due to attitude changes and jitter, such that dynamic measurements cannot be replaced by static measurements

when the goal is to represent operational conditions; however, a statistical model may be a suitable substitute [12].

For example, the χ^2 (gamma-family) probability density function, given by

$$p(\sigma) = \frac{1}{\Gamma(k)} \frac{k}{\sigma_{mean}} \left(\frac{k\sigma}{\sigma_{mean}} \right)^{k-1} e^{-\frac{k\sigma}{\sigma_{mean}}} \quad (2.29)$$

where σ_{mean} is the mean RCS value, and k controls the fluctuation severity, i.e higher values produces smoother behavior and smaller values procures larger undulations. Other PDFs used in [12] include lognormal and Rice.

The goodness of fit of the models was evaluated at four aspects (head, tail, left, right) according to [12]:

$$ef = \sum_{i=1}^m |p_i - \hat{p}_i| \times 100\% \quad (2.30)$$

where p_i is the measured probability and \hat{p}_i is the fitted-model probability. The results shown in Table 2-3 indicate that χ^2 provides the most consistent fit across all aspects, whereas lognormal and Rice can fail to match the more irregular distributions in some aspects (particularly the tail and left)

TABLE 2-3 FITTING ERRORS OF THE THREE DISTRIBUTION MODELS [12]

Three model Fitting errors	Head(%)	Tail(%)	Left(%)	Right(%)
χ^2	6.3	9.4	5.1	3.6
Lognormal	6.3	11.6	6.7	13.3
Rice	10.0	27.3	27.6	7.6

Given these results, a χ^2 fluctuation model appears to be the preferred numerical model for small UAVs. However, target fluctuations are only one part of the detection problem in the maritime setting: even a favorable $\sigma(t)$ can be masked if the background returns are

dominated by sea backscatter or precipitation clutter. The next subsection introduces sea clutter modelling and shows how grazing-angle-dependent backscatter and weather volume clutter determine the signal-to-clutter conditions that ultimately bound shipboard UAV detection performance.

2.9 Sea Clutter Modelling

Radar clutter is defined as “any object that may generate unwanted radar returns that may interfere with normal radar operations” [8]. As with targets, clutter can be characterized by an effective RCS σ_c :

$$\sigma_c = \sigma^o A_c \quad (2.31)$$

where σ^o is the sea-clutter scattering coefficient, and A_c is the illuminated clutter area within the radar resolution cell. Fig. 2-23 shows how the sea’s backscatter per unit area σ^o varies with grazing angle ψ_g , the angle between the sea surface and the radar beam, in a non-linear manner. At very low grazing angles, sea clutter is weak; it then rises to a plateau, and beyond about 60° it increases sharply [8].




Figure 2-23 Example of grazing angle and impact of grazing angle to clutter scattering coefficient [8]

Fig. 2-24 highlights that sea-surface roughness introduces a transition between smooth-like and rough-like behavior: at shallow angles the surface reflects in a specular manner, whereas increasing roughness and incidence angle introduces multiple reflection paths.

This multipath behavior produces phase differences that distort the received echo and can degrade detection performance in cluttered maritime environments [8].




Figure 2-24 Example of sea clutter rough and smooth path [8]

For modelling, the practical requirement is therefore not just a qualitative understanding of sea clutter, but an explicit model for σ^o as a function of frequency, grazing angle, sea state, and polarization. In this thesis, σ^o is obtained using the NRL low-grazing empirical sea clutter model (as implemented in MATLAB's Sea clutter/reflectivity functions). This choice is motivated by published comparisons showing that older widely used models (e.g., GTI) can substantially underestimate σ^o at low grazing angles for moderate sea states. The NRL model is derived from a large compiled database and provides an empirically validated, conservative baseline for low-grazing X-band operation across sea states, with separate parameter sets for horizontal and vertical polarization [17,18]. This makes it appropriate for the sea-skimming C-UAV geometries considered in this work.

Radars in clutter situations must now be designed to overcome the Signal to Clutter Ratio (SCR), which is defined as [8]:

$$SCR_{seaclutter} = \frac{2\sigma_{target}\cos\psi_g}{\sigma^o\theta_{3db}Rc\tau} \quad (2.32)$$

where θ_{3db} is the antenna beamwidth. (2.32) highlights two important implications for shipboard radar design. First, reducing beamwidth reduces the illuminated clutter area and therefore improves the signal-to-clutter ratio. Second, sea clutter is more problematic at low grazing angles, which makes low-altitude threats near the horizon inherently more

difficult to detect. Sea-skimming missiles exploit this effect by maintaining very low altitude over the sea surface.

In addition to sea-surface clutter, weather introduces volume clutter and propagation loss. Weather clutter arises from distributed scatterers such as rain droplets, hail, and spray. Rain droplets may be approximated as small spheres, and when the droplet diameter is small relative to wavelength, their scattering is considered to be in the Rayleigh regime. The individual droplet RCS can be expressed as [8]:

$$\sigma_{rain} = 9\pi r^2 (br)^4 \quad (2.33)$$

where $b = 2\pi/\lambda$ and r is the radius of the rain droplet. The RCS per unit resolution volume, which is defined as “the three-dimensional “cell” that one pulse of the radar interrogates,” can be defined as the following [8]:

$$\sigma_w(rain) = \sum_{i=1}^n \sigma_{rain_i} V_w \quad (2.34)$$

where $V_w = (\pi/8)\theta_a\theta_e R^2 c\tau$ is the weather-clutter resolution volume, θ_a and θ_e are the antenna azimuth and elevation beamwidths in radians, respectively, and R is range. Since weather clutter is distributed throughout a volume rather than confined to the sea surface, the relevant clutter cross section is the equivalent radar cross section of the rain-filled resolution volume, denoted $\sigma_w(rain)$. The corresponding weather-clutter signal-to-clutter ratio is therefore defined as [8]:

$$SCR_{(weatherclutter)} = \sigma_{target}/\sigma_w(rain) \quad (2.35)$$

Rain is treated primarily as a propagation impairment in addition to any volume-clutter contribution. Specifically, rain attenuation is modelled using the ITU-R method implemented in MATLAB’s `rainpl` function. This choice follows established practice: ITU-R P.618, together with P.837/P.838, are widely used to estimate rain attenuation for microwave and millimetre-wave links, and published studies demonstrate the expected behavior of increasing attenuation with rain rate and frequency, supporting their use for link-budget and detection-range degradation analysis in heavy-rain conditions [18, 19].

Having defined the target reflectivity model and the dominant maritime background and propagation impairments, the next subsection formalizes the detection framework.

2.10 Detection Processing Chain

The received radar waveform is not directly suitable for decision-making; it must be transformed through a sequence of signal-processing stages to produce meaningful observables and a statistically controlled detection decision. This subsection defines the detection processing chain, including range processing and Doppler processing to form range–Doppler cells, and CFAR thresholding at a specified P_{fa} [2,7,8].

Given that the small UAV will usually be closing on the ship, a range–Doppler map allows it to be distinguished from near-zero-Doppler interference such as sea clutter. Let $x[n, m]$ denote the complex baseband (I/Q) samples, where n indexes fast time and m indexes slow time. The digital receiver samples the complex baseband return in two dimensions: fast time within each PRI/chirp, and slow time across pulses/chirps. Let the (Analog to Digital) ADC sampling rate be f_s , giving N_{fast} samples per PRI/chirp. The sampling rate must be chosen to prevent aliasing of the highest frequency content in the processed waveform; in general, $f_s \geq 2f_{\text{max}}$ (Nyquist), with practical systems often sampling several times higher (5 to 10 times) to preserve fidelity and margin [20]. In the pulsed LFM case, fast-time sampling supports matched filtering and the resulting range-bin formation; in the FMCW case, fast-time sampling must accommodate the maximum beat-frequency, which limits unambiguous range via ADC/receiver bandwidth. In slow time, the effective sampling frequency is the PRF (or chirp repetition rate), which sets unambiguous Doppler and Doppler-bin spacing over the coherent processing interval. The range–Doppler map (RDM) is formed by applying a range transform in fast time followed by a Doppler transform in slow time [2,7,8]:

$$X[k, m] = \sum_{n=0}^{N_r-1} x[n, m] w_r[n] e^{-\frac{j2\pi kn}{N_r}} \quad (2.36)$$

$$Z[k, l] = \sum_{n=0}^{N_d-1} X[k, l] w_d[l] e^{-\frac{j2\pi kn}{N_d}} \quad (2.37)$$

where $w_r[n]$ and $w_d[l]$ are range and Doppler smoothing window functions that reduce sidelobes, N_r is the number of fast-time samples used for range processing, and N_d is the number of pulses/chirps in the coherent processing interval. The detector typically operates on cell power (square law detector)

$$P[k, l] = |Z[k, l]|^2 \quad (2.38)$$

For pulsed radar, fast-time processing consists of matched filtering, or pulse compression, followed by sampling into range bins. Doppler processing is then performed in slow time by collecting a sequence of consecutive pulses from the same range bin and applying a Doppler FFT across that sequence. For FMCW radar, fast-time processing consists of dechirping, followed by an FFT to form beat-frequency and therefore range bins. Doppler processing is then performed by collecting a sequence of consecutive chirps and applying a Doppler FFT across the slow-time sequence, as discussed in Section 2.4. In both radar types, the output is a Range-Doppler map (RDM) in which the horizontal axis represents range bins and the vertical axis represents Doppler bins, or equivalently radial-velocity bins. This is the standard fast-time range and slow-time Doppler processing structure used in Range-Doppler (RD) analysis.

The Doppler-bin spacing is given by

$$\Delta f_d = \frac{1}{T_{coh}} = \frac{1}{N_d T_{rep}} \quad (2.39)$$

where N_d is the number of slow-time samples used in the Doppler FFT, T_{coh} is the coherent processing interval, and T_{rep} is the repetition interval between adjacent slow-time samples. For pulsed radar, T_{rep} is the PRI. For FMCW radar, T_{rep} is the chirp repetition interval. Using the two-way monostatic Doppler relationship $f_d = 2v_r/\lambda$ where v_r is the target radial velocity and f_d is the Doppler frequency shift in hertz, the Doppler-bin spacing maps directly to radial-velocity resolution as

$$\Delta v = \frac{\lambda}{2} \Delta f_d \quad (2.40)$$

Thus, the spacing of the vertical axis of the RDM is set by (2.40), while the maximum unambiguous velocity remains limited by (2.16). The spacing of the horizontal axis is determined by the radar range resolution, given by (2.19) for pulsed LFM radar and by (2.23) for FMCW radar.

In a shipboard environment, short-range returns can be dominated by strong clutter, leakage, and platform scattering. To prevent receiver saturation and ADC clipping, the receiver applies a Sensitivity Time Control (STC) gain variation, which is a range-dependent attenuation applied early in the processing chain. In its general form [7,8]:

$$x_{stc}(t) = G_{stc}(R(t))x(t) \quad (2.41)$$

where $G_{stc}(R)$ increases with range to remove the implicit R^{-4} variation. This ensures that the subsequent range–Doppler processing and thresholding operate on signals within the receiver’s linear dynamic range.

Moving Target Indication (MTI) is then used as an interference-rejection step to suppress residual near-zero-Doppler energy (sea clutter, platform scattering, and leakage terms) that can persist even after STC and windowed range–Doppler processing. MTI is implemented as a slow-time high-pass filter applied pulse-to-pulse (or chirp-to-chirp) for each range bin. Let $x_r[n]$ denote the complex slow-time sequence at range bin r and pulse index n . These standard MTI cancellers are obtained by applying first- and second-difference operators to the slow-time sequence, giving the 2-pulse form [2,7,8]

$$y_r[n] = x_r[n] - x_r[n - 1] \quad (2.42)$$

and a stronger 3-pulse form is

$$y_r[n] = x_r[n] - 2x_r[n - 1] + x_r[n - 2] \quad (2.43)$$

These filters significantly attenuate stationary or very-slow returns while preserving the Doppler content of moving targets. This is well aligned with the shipboard small-UAV problem because the UAV radial velocity is typically separated from the near-zero-Doppler clutter ridge, so MTI improves contrast, reduces clutter-induced sidelobes, and reduces false alarms with only a modest penalty for very slow targets. After MTI, detections are declared using CA-CFAR in the range–Doppler domain to maintain the specified P_{fa} under changing sea state and weather conditions [2,7,8].

A fixed detection threshold is not robust in maritime conditions because the interference floor varies with range, grazing angle, and sea state. To maintain a controlled false-alarm probability, Constant False-Alarm Rate (CFAR) is used, as shown in Fig. 2-25. Let P_0 be the cell under test (CUT) and $\{P_i\}_{i=1}^N$ be N training cells selected around the CUT (excluding guard cells). For cell-averaging CFAR (CA-CFAR), the interference estimate is [2,7,8]

$$\hat{P} = \frac{1}{N} \sum_{i=1}^N P_i \quad (2.44)$$

with the following decision rule:

$$P_0 > T \rightarrow H_1, P_0 \leq T \rightarrow H_0, T = \alpha \hat{P} \quad (2.45)$$

If noise and clutter power are exponentially distributed in the training cells, such as during heavy sea and rain rates, the CA-CFAR scaling factor that will maintain P_{fa} is:

$$\alpha = N \left(P_{fa}^{-\frac{1}{N}} - 1 \right) \quad (2.46)$$

This links the threshold directly to a user-specified P_{fa} and allows consistent comparison between radar candidates under varying background conditions.

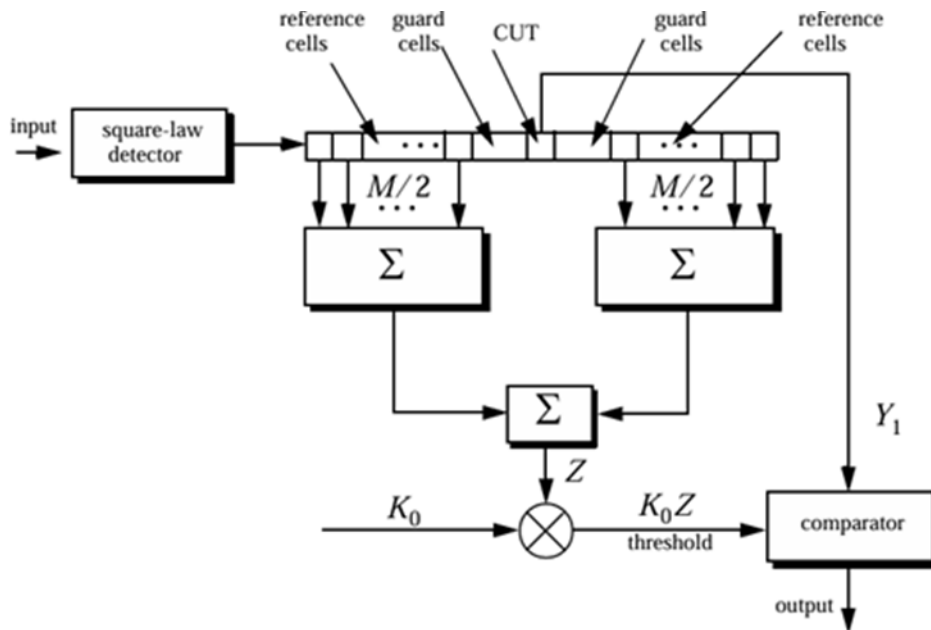


Figure 2-25 CFAR Logic Diagram [2]

Finally, because detections may be intermittent in heavy clutter and multipath, the simulated detection sequence is optionally post-processed using an M -of- N binary integration rule. Here, a scan denotes one radar revisit of the relevant search direction, that is, one update opportunity that produces a binary detection or non-detection result. A confirmed detection is declared when at least M detections occur within a sliding window of N scans. This rule is applied to the Monte Carlo scan outcomes to produce an operational confirmation metric, rather than to derive an analytical expression for P_d . The resulting measures are later reported as detection latency and confirmed-detection range.

2.11 Advanced Signal Processing

Many techniques can be used to enhance detection or discrimination in difficult interference, such as MD to sense UAV rotor motion. MD arises when parts of a target undergo additional motion (rotation/vibration) relative to the target's center of mass, producing time-varying Doppler frequency components in the received signal [21]. For a multicopter UAV, the dominant MD components are driven by rotor blade rotation and appear as harmonic structures in the Doppler spectrum.

Figure 2-26 provides an example of such a rotor Doppler echo obtained using a CW radar at 35 GHz, for a two-blade rotor with blade length 10 cm rotating at 20 revolutions per second [21]. The plot shows that the return is not a single Doppler line; instead, it contains a dominant periodic component associated with blade rotation and a set of harmonic structures tied to the blade passage process. These harmonics arise because scattering from the rotating blades repeats periodically, producing Doppler components at multiples of the blade rate (and therefore persistent “lines” in the time–frequency view).




Figure 2-26 Doppler radar echo of a two-blade rotor [21]

What changes with environment are contrast: sea-clutter spikes and multipath increase the background variance, and short coherent intervals reduce Doppler resolution, making those harmonics harder to separate.

2.12 UAV Mission Profiles and Representative Flight-Path Scenarios

Recent conflict literature shows that small UAVs are used with repeatable, tactically distinct movement patterns, and those patterns map directly to the shipboard detection problem. In the Ukraine context, UAVs are at the center of air operations conducting Intelligence, Surveillance, Target Acquisition, and Reconnaissance (ISTAR) missions on a daily basis, with tactical value driven by sustained observation and rapid tasking rather than large payloads or high speed [1].

At the small end of the spectrum, consumer quadcopters, similar to the DJI UAV described in Section 2.5, have become the standard for small-unit reconnaissance and artillery spotting, and are also adapted into “drop drones” for close-in attack, implying flight behavior that frequently includes slow loiter/hover near the objective, short repositioning moves, and repeated short-range passes [22]. Canadian commentary on modern combat similarly emphasizes that small drones are dropping single bombs, and that drone activity occurs in high daily volume, reinforcing that the small-UAV threat is persistent and can be present in large numbers rather than appearing as a single, isolated track [23]. These sources provide defensible bounds for “small UAV” kinematics and support selecting representative scenario parameters rather than arbitrary speeds and altitudes. As a concrete baseline, the DJI Phantom 4 Pro will be used as a representative small quadcopter and its key manufacturer-style specifications are summarized in Table 2-4 [24]. These values anchor the nominal small-UAV motion envelope used in the scenario definitions.

TABLE 2-4 DJI PHANTOM 4 PRO KEY PARAMETERS [24]

Parameter	Value
Max Speed	72 km/h
Max altitude	6000 m

Overall, the literature supports three representative shipboard threat scenarios that stress the radar in distinct ways: (i) low-grazing inbound motion where sea clutter and short warning time dominate, (ii) crossing motion where search burden and revisit rate determine detection opportunity, and (iii) slow hover/loiter motion where longer dwells and persistent observables are feasible [1,22,24].

In summary, this chapter established the technical basis required to conduct a simulation-based comparison of ship-installable C-UAV radar architectures. It defined the two-candidate sensing approach, the AESA implementation constraints that govern beam

steering and revisit, the small-UAV target model, and the dominant maritime impairments that drive detection performance (sea clutter, rain attenuation, and multipath). It also formalized the detection processing chain that will be applied consistently to both radar types R-D mapping, STC, and CFAR thresholding at a specified false-alarm rate. The next section presents the methodology: it specifies the common constraint set and radar parameterization used for the pulsed and FMCW candidates, defines the representative UAV flight scenarios and environmental conditions, and details the Monte-Carlo simulation workflow and validation checks used to generate comparable detection performance results.

Naval Radar Systems and Design Constraints from Literature

Chapter 3 defines the simulation-based comparison framework used to evaluate the performance of ship-installable counter-UAV radars. It converts the modelling assumptions established in Chapter 2 into a practical experimental design consisting of a common constraint set for both radar candidates, representative UAV flight scenarios, maritime environmental conditions, and an end-to-end detection processing chain executed in Monte Carlo. Parameter values are selected to remain consistent with published shipboard radar practice and then held fixed so that the pulsed and FMCW candidates can be compared on an equal technical basis.

3.1 Simulation-Based Comparison Framework

The methodology implements the candidate radar configurations in a virtual environment in MATLAB R2025a [18]. Fig. 3-1 presents the simulation-based comparison framework as an input-process-output chain. The inputs comprise: (i) a common naval constraint set held fixed for both radar candidates, including single-face AESA architecture, carrier frequency, aperture, steering sector, and transmit-power limits; (ii) representative UAV operational scenarios, including attack, surveillance, and intelligence flight paths; (iii) a target reflectivity model defined by mean RCS and statistical fluctuation; and (iv) maritime environment models, including sea-state clutter, rain attenuation, and multipath. These inputs are propagated through a common detection-processing chain that forms RD cells, applies receiver dynamic-range management where required through sensitivity time control (STC), and declares detections using CFAR thresholding at a specified P_{fa} . Random effects, such as RCS fluctuation and time-varying propagation and clutter, are captured through Monte Carlo trials. The outputs are detection-oriented performance measures, including $P_d(R)$ at fixed P_{fa} , detection range at selected operating points, and time to detect or confirm under the assumed scan schedule. Collectively, these measures provide a consistent basis for comparing pulsed and FMCW radar candidates under identical constraints.

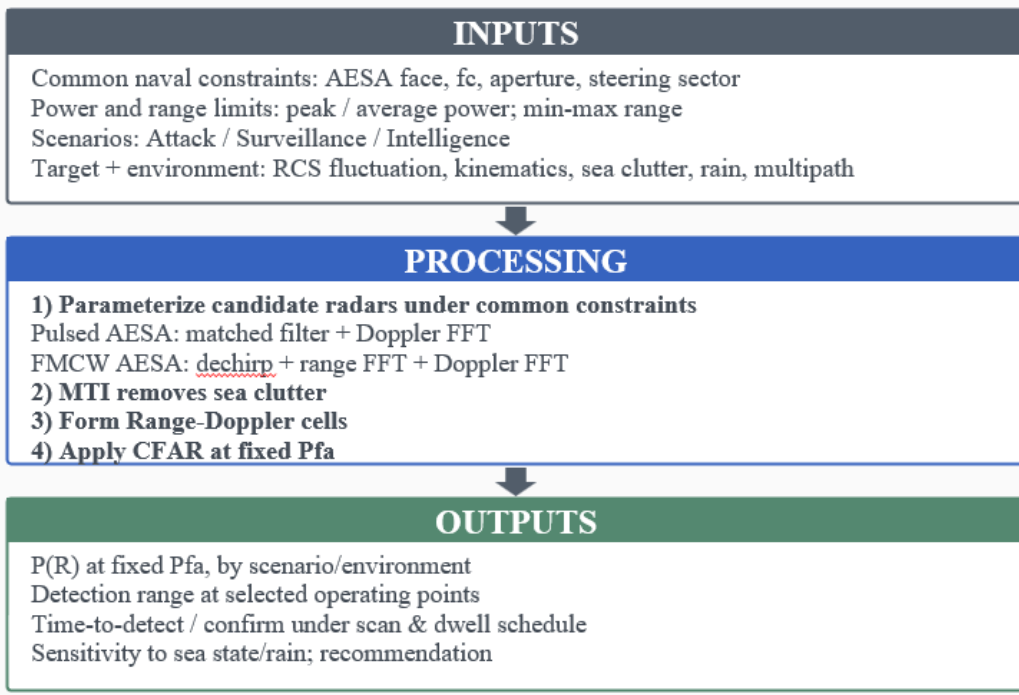


Figure 3-1 Simulation Framework

The remainder of Chapter 3, therefore, instantiates each block of Fig. 3-1 in turn by defining the common constraint set, specifying the threat-driven scenarios and target/environment models, and then locking the two radar candidates and the detection processing parameters used to generate comparable P_d results.

3.2 Common Constraint Set and Baseline Assumptions

The comparison is only meaningful if both radar candidates share the same first-order shipboard constraints; this section therefore defines the common constraint set first and then instantiates each waveform family within that same bounded design space.

3.2.1 Shared radar constraints

The first step in implementing the framework is to define the common constraint set that represents a ship-installable radar face, and that is applied identically to both candidate radar waveforms. A single AESA face was chosen with a fixed physical aperture and element spacing, a single carrier frequency and common polarization assumptions, a bounded electronic steering sector in azimuth and elevation, and a transmitter power constraint expressed through peak power and duty factor (and therefore average power). The evaluation range window is also fixed to represent the detection problem of interest, with a minimum range driven by near-range interference management and a maximum range set to exceed the engagement range of shipboard effectors, allowing assessment of detection margin and warning time. However, frequency selection is not driven by rain loss alone. In a ship-installable AESA, the operating band must simultaneously satisfy: (i) maritime propagation robustness, including rain and atmospheric loss; (ii) achievable antenna performance within a constrained face area; and (iii) sufficient Doppler sensitivity to separate slow UAV motion and rotor MD from sea clutter. Fig 3-2 shows that attenuation increases with frequency; therefore, S-band offers the strongest rain robustness, with losses remaining comparatively low even under heavy rain. Nevertheless, S-band imposes a different penalty for a size-limited AESA. For a fixed physical aperture area, antenna gain scales as $G \propto A/\lambda^2$. As wavelength increases at lower frequency, the same 1 m²-class face produces lower gain and a broader beamwidth. The broader beam does reduce the number of beam positions required to cover a given sector, but it also reduces angular resolution and lowers the antenna gain available in each look direction. The result is weaker detection performance against small UAV targets and less effective spatial discrimination of clutter within the searched sector.

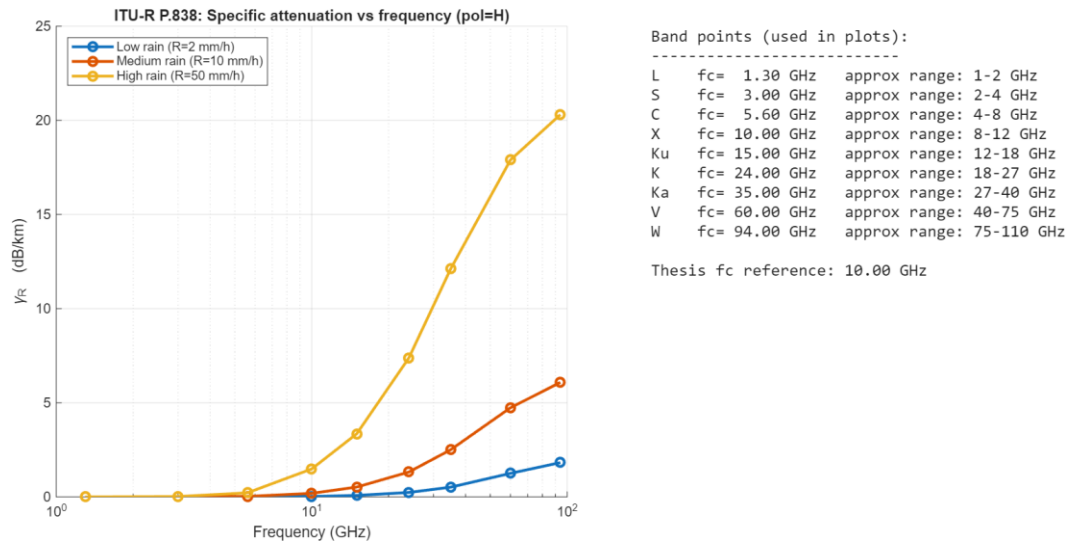


Figure 3-2 ITU Rain model vs Frequency

At the other extreme, K/Ka-band provides excellent range and Doppler resolution for a given waveform bandwidth and increases sensitivity to fine micro-motion, but it incurs substantially larger rain attenuation, which is undesirable for a shipboard radar intended to remain effective across sea states and weather. X-band therefore represents a balanced operating point for the constraint set.

To maintain a ship-installable and defensible transmit-power assumption, the pulsed radar is limited to 10 kW peak power for the single AESA face considered in the model. With a duty factor of 10%, this yields an average transmit power of 1 kW. This assumption is consistent with a modular active-antenna architecture and remains broadly aligned with the radar examples discussed in Chapter 2 [18].

The number of T/R elements is set directly by the selected aperture size and half-wavelength element spacing. For the chosen AESA face dimensions, this yields 1000 T/R elements, consistent with the scale of the AMSAR X-band demonstrator [9]. Each element is assigned a 6-dBi cosine field pattern, representing a broad single-element response suitable for electronic scanning, while the full-array aperture provides the narrow main beam and overall array directivity. Cosine element patterns are a standard way to represent broad antenna element field responses in phased-array modelling.

The evaluation range window is selected to reflect both shipboard geometry constraints at short range and tactically meaningful warning distance at long range. The maximum range is set to 40 km because the ship's hard-kill engagement range for the counter-UAV problem, using typical RCN weapon systems, is on the order of 20 km, and a detection range that is double provides defensible margin for classification time, decision time, and weapon cueing [27]. The minimum range is set to 1 km, as at shorter distances a member of a ship's company would be able to visually detect the UAV using binoculars. Therefore, $R_{\min} = 1$ km provides a realistic lower bound at which the receiver and CFAR processing are operating in their intended linear regime, while $R_{\max} = 40$ km provides a meaningful upper bound for assessing whether either architecture can deliver adequate detection probability at ranges that support shipboard reaction timelines under adverse sea state and rain conditions.

A separate geometric consideration on the long-range boundary is the radar horizon as shown by Fig. 3-3 [7,8]. For a radar at height h_r above the sea surface, the horizon distance d_r is

$$d_r = \sqrt{2R_e h_r + h_r^2} \quad (3.1)$$

where R_e is the effective Earth radius. For a target at height h_t , the corresponding line-of-sight limit is obtained by summing the radar and target horizon distances:

$$d_{\text{LOS}} = \sqrt{2R_e h_r + h_r^2} + \sqrt{2R_e h_t + h_t^2} \quad (3.2)$$

When $h_r, h_t \ll R_e$, this reduces to the familiar approximation

$$d_{\text{LOS}} \approx \sqrt{2R_e h_r} + \sqrt{2R_e h_t} \quad (3.3)$$




Figure 3-3 Radar Horizon adapted from [7]

This relationship is used as a geometric consistency check on the selected evaluation window rather than as a scenario definition. Because the admissible line-of-sight range depends on target altitude, a single radar-horizon distance is not quoted in the common-constraint subsection; instead, the scenario-specific target heights and their implications for geometric visibility are introduced in Section 3.3 [7,8].

Table 3-1 highlights the design decisions, whereas Fig 3-3 shows the location of the radar relative to the ship:

TABLE 3-1 SHARED RADAR PARAMETERS

Parameters	Value
Radar faces modelled	1
Radar height above sea	20 m
Carrier frequency	10 GHz
Face physical area	$\approx 1 \text{ m}^2$
Array Geometry	URA (60 x 40)
Element Spacing	$\lambda/2$
Element Gain	6 dBi
AESA Tapering	Chebyshev
Steering Limits in az and el	+/- 60 degrees
Peak Power	10 kW
Average Power	1 kW
Duty factor	10 %
Waveform Bandwidth	5 MHz
Fast time Sampling Rate	40 MHz
PRF	2 kHz
PRI	0.5 ms
Samples per PRI	20000




Figure 3-4 Diagram of Radar location on a ship

3.2.2 Shared aperture and beamsteering assumptions

With the common operating band and installation limits established, the next step is to instantiate one representative AESA face that is shared by both radar candidates. The aperture, element spacing, taper, and steering limits are held fixed so that both waveforms operate through the same beamsteering geometry and scan-loss environment. In this way, the comparison isolates waveform effects without confounding them with a different antenna realization.

The AESA face used for both candidates is the single 60×40 URA at 10 GHz with spacing and a Chebyshev 30 dB taper. The 60×40 URA is an illustrative instantiation consistent with the assumed aperture size, steering sector, and beamwidth requirements; it is not claimed to be the only valid array realization. The key modelling assumption is the beamwidth consistent with a ship-installable metre-class AESA aperture, while the exact element count is one representative implementation under dense-spacing assumptions. The azimuth cut versus angle, 3-D beam pattern and directivity versus steering angle are provided in Fig. 3-5, 3-6 and 3-7. In the implemented comparison, scan loss is applied through the steering-angle-dependent directivity reduction shown in Fig. 3-7, and this effect is held common across both waveform candidates as part of the shared antenna model

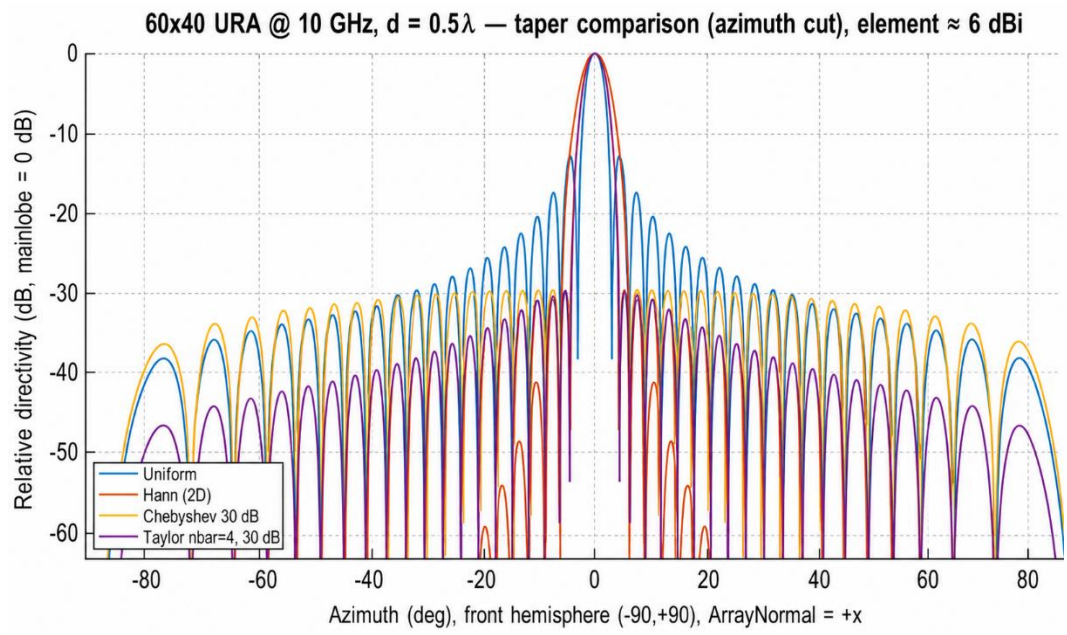


Figure 3-5 60 x 40 URA azimuth vs taper cut

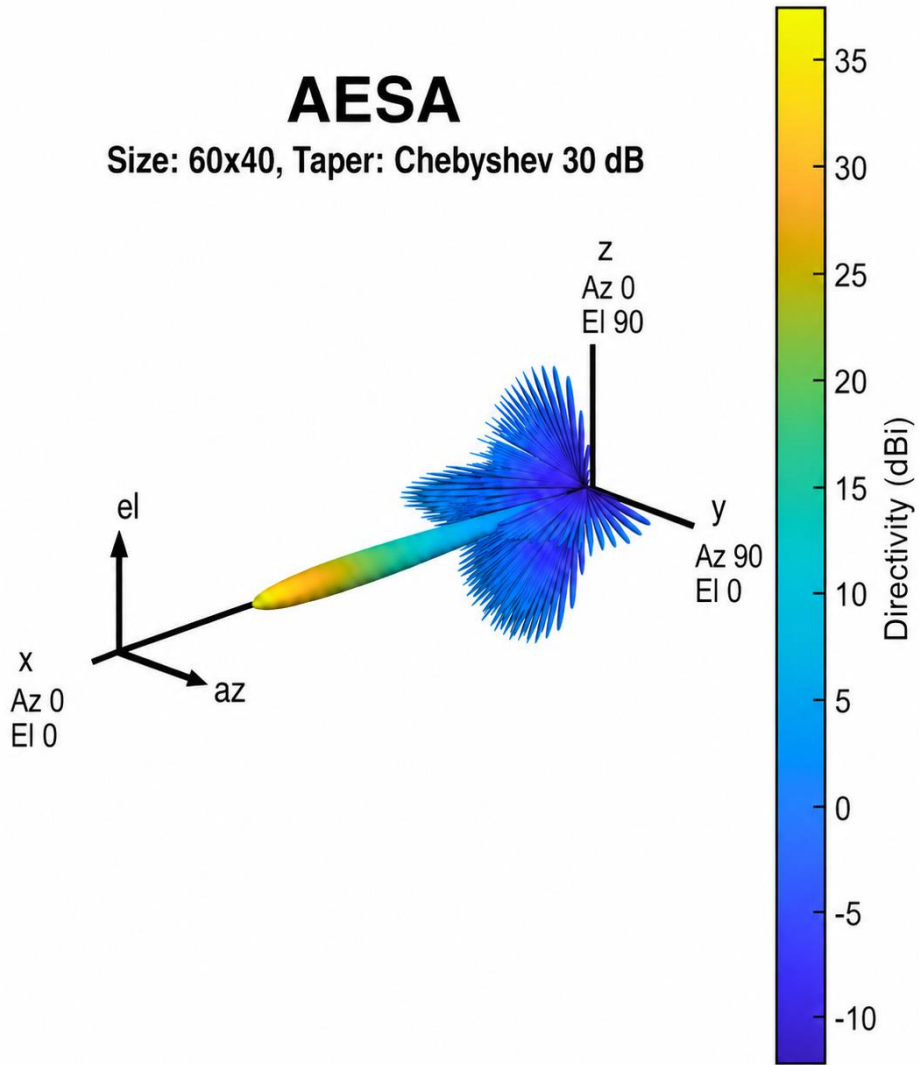


Figure 3-6 3-D beam pattern for URA 60 x 40

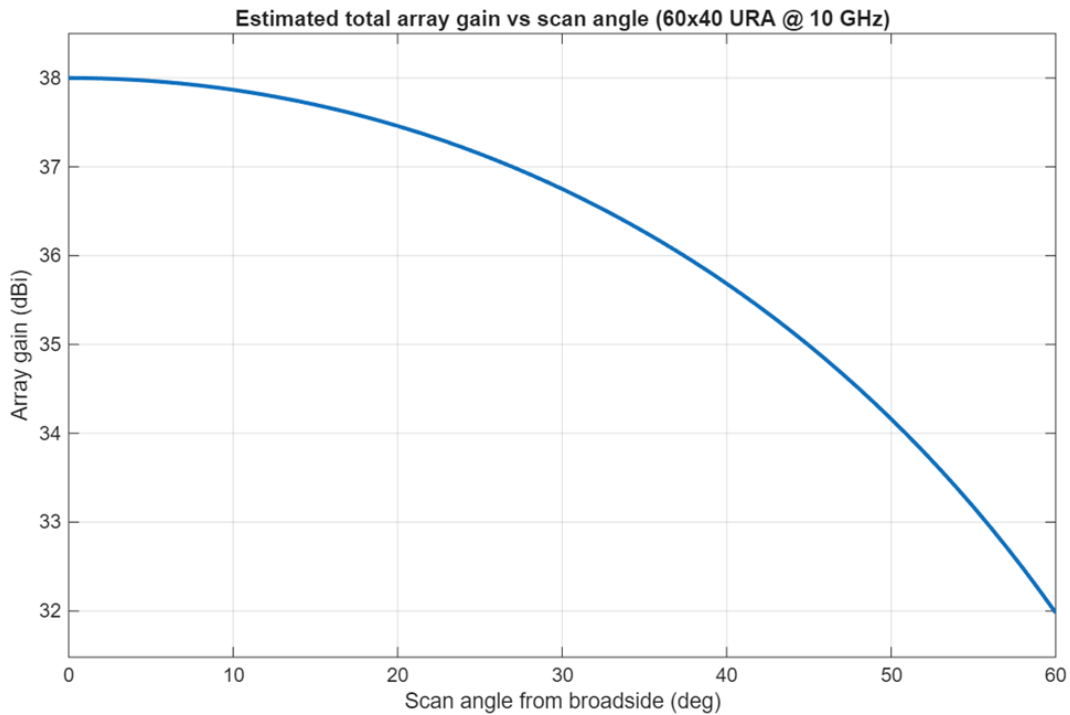


Figure 3-7 Directivity versus steering angle

The AESA panel produces a half-power beamwidth of approximately 3.2° in azimuth and 2.1° in elevation. The important point is that both radar candidates share the same beam shape, steering sector, and scan-loss environment, so these aperture effects are treated as common antenna assumptions rather than waveform-dependent effects. Fig. 3-8 summarizes the two candidate radar architectures as they are instantiated under this shared AESA face and common shipboard constraint set. In Fig. 3-8(a), the pulsed candidate is implemented as a monostatic AESA in which the same face is used for both transmission and reception. Detection is based on short LFM bursts followed by matched filtering and slow-time Doppler processing. In Fig. 3-8(b), the FMCW candidate is implemented as a bistatic AESA in which separate transmit and receive faces are used, with continuous linear chirps processed through dechirp, fast-time FFT range estimation, and slow-time Doppler FFT. Thus, the figure highlights the architectural distinction between the candidates at a system level, while the following subsections define the waveform timing and processing details used in the comparison.

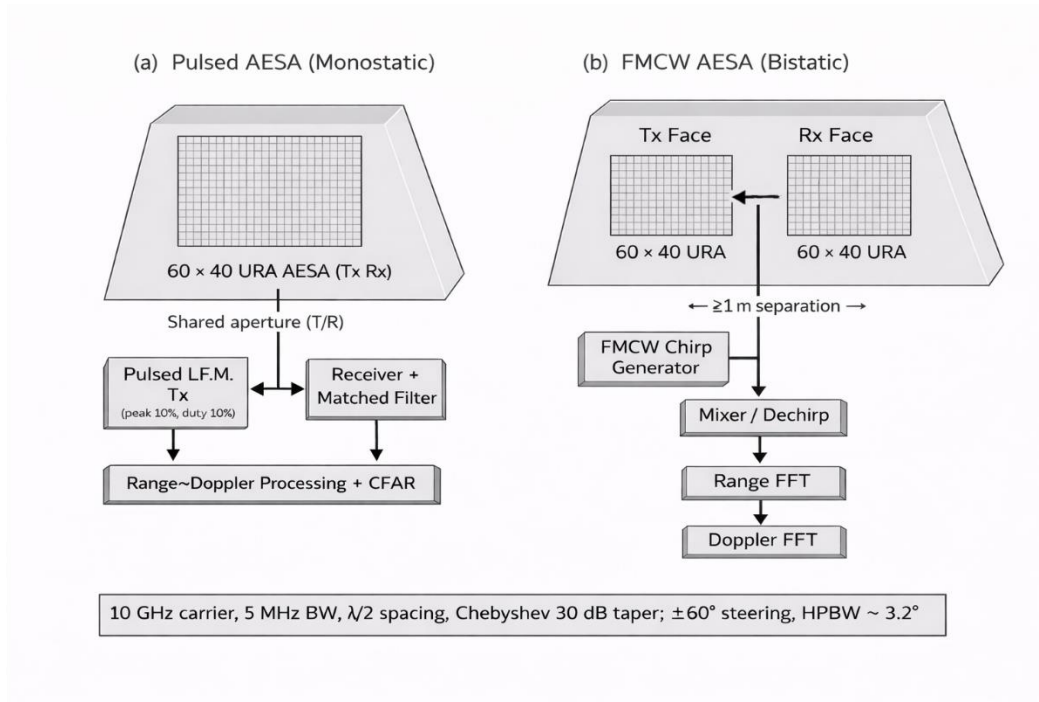


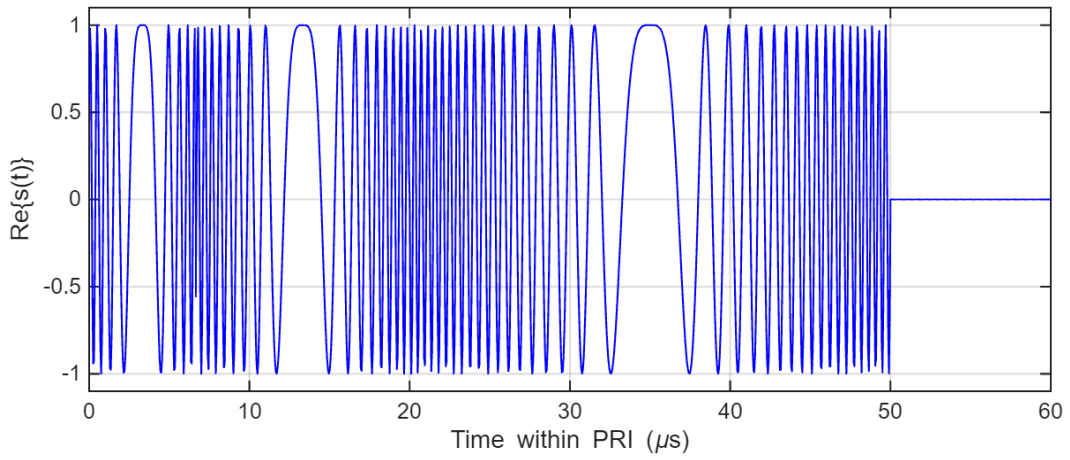
Figure 3-8 (a) Pulsed Radar candidate and (b) FMCW Radar candidate

3.2.3 Pulse Radar Candidate



The pulsed candidate is defined as a monostatic AESA transmitting an LFM waveform with pulse-compression processing.

The panel transmits a LFM waveform with a total on-time of $\tau_{total} = 50 \mu\text{s}$ and bandwidth $B = 5 \text{ MHz}$. Matched filtering is then used to achieve pulse-compressed range resolution while preserving energy on target through the longer effective transmit duration. To support close-range coverage without reducing τ_{total} , the $50 \mu\text{s}$ transmit energy is implemented as a three-subpulse stagger train packed contiguously at the start of each PRI, with pulse widths $[6.65, 13.35, 30.00] \mu\text{s}$ and start times $[0, 6.65, 20.00] \mu\text{s}$. Each subpulse is matched-filtered using its own replica to produce a compressed fast-time response, after which the three responses are aligned to a common delay reference and coherently summed to form a single composite range profile. In this pulsed mode, pulses are transmitted at $\text{PRF} = 2 \text{ kHz}$ ($\text{PRI} = 0.5 \text{ ms}$) corresponding to an unambiguous range of $R_u \approx 75 \text{ km}$; performance is evaluated over the 0–40 km region of interest (ROI). Fig. 3-9 illustrates this pulsed implementation. In Fig. 3-9(a), the transmit waveform is shown within one PRI, where the three contiguous LFM subpulses can be seen occupying the first $50 \mu\text{s}$ of the repetition interval, followed by receive-only time. This construction shows how the waveform preserves the required total transmit energy while still maintaining near-range

access within the PRI. Fig. 3-9(b) then shows the corresponding matched-filter response. The full response and the zoomed view around the truth range confirm that pulse compression recovers a narrow mainlobe at the target location, with a measured -3dB width of approximately 30 m, which is consistent with the expected range resolution for a 5 MHz LFM waveform.



(a)



(b)


Figure 3-9 (a) LFM Tx pulse waveform and (b) matched filter result

Taken together, these choices define a practical pulsed shipboard baseline under the common constraint set. The LFM candidate preserves fine range resolution through pulse compression, but it must distribute its transmit energy within short, staggered segments inside each PRI so that close-range access, long-range energy on target, and ambiguity control are all maintained within the same waveform structure. In this sense, the pulsed case provides the reference architecture against which FMCW is compared under the same aperture, bandwidth, steering burden, and coherent slow-time framework.


3.2.4 FMCW Radar Candidate

Within the same shared aperture, frequency, steering, and coherent slow-time framework, the FMCW candidate is defined using a train of identical linear up-chirps and dechirp-based range processing. The purpose of this candidate is to evaluate the continuous-illumination alternative under the same first-order shipboard resource constraints used for the pulsed comparison. In contrast to the pulsed case, the FMCW waveform occupies the full repetition interval, so the transmitted energy is distributed continuously over the PRI rather than in short staggered transmit segments.

For the FMCW candidate, each dwell is formed from a train of identical linear up-chirps. Range is estimated by dechirping the received signal and applying a fast-time FFT to the resulting beat signal, with the beat-frequency-to-range mapping governed by the chirp slope $C_s = B/T_c$, where $T_c = 0.5$ ms. Doppler is then estimated by an FFT across the 16 chirps in slow time to form a Range–Doppler map. For the common coherent processing interval of $T_{coh} = 8$ ms, the velocity bin spacing is $\Delta v \approx \lambda/2T_{coh} \approx 1.9$ m/s, while the unambiguous velocity is $v_{amb} = \lambda PRF/4 \approx 15$ m/s. The FMCW waveform therefore provides continuous transmission over the PRI while remaining within the same overall design space and operational search burden as the pulsed candidate. Fig. 3-10 shows (a) the FMCW transmit waveform and (b) the corresponding FFT-based range response.



(a)



(b)

Figure 3-10 FMCW (a) Tx waveform and (b) matched filter result

The FMCW candidate represents the continuous-illumination alternative under the same ship-installable aperture, steering, bandwidth, and coherent slow-time framework used for the pulsed case. The next subsection states explicitly the fairness basis used to ensure that the comparison reflects waveform-dependent differences rather than unequal resource allocation.

3.2.5 Fairness Basis for the Comparison

A controlled comparison between pulsed and FMCW radar candidates is only meaningful if both architectures are evaluated under the same first-order shipboard resource constraints. The carrier frequency, AESA face aperture, element spacing and taper, steering sector, waveform bandwidth, search volume, false-alarm requirement, and evaluation range window are therefore held common. The same target, sea-clutter, rain, and radar-horizon models are also applied to both candidates so that environmental severity enters the comparison on an equal basis.

In addition, both radar candidates are evaluated using the same coherent slow-time observation structure. Each dwell contains 16 coherent slow-time samples separated by $T_{rep} = 0.5$ ms. In the pulsed case these slow-time samples are pulses, whereas in the

FMCW case they are chirps. The coherent processing interval is therefore $T_{coh} = 16 \times 0.5\text{ms} = 8\text{ms}$ for both modes. And the Doppler parameters are the same as in Section 3.2.4.

The fairness basis is therefore not that the two waveforms are made physically identical, but that they are given the same ship-installable aperture, the same search burden, the same coherent slow-time observation interval, and the same environmental and detection framework. What is allowed to differ is the waveform family itself and the associated processing consequences. The pulsed candidate concentrates its transmit energy into staggered LFM subpulses within each PRI and uses matched filtering for range processing, whereas the FMCW candidate occupies the full repetition interval with continuous chirps and uses dechirp and FFT-based beat-frequency estimation. Under this basis, any resulting differences in detection, confirmation, and classification performance can be interpreted as waveform-dependent architectural differences rather than artifacts of unequal aperture, steering allocation, or coherent integration time.

Although both candidates are constrained to the same bandwidth and the same coherent slow-time observation interval, their per-transmission time-bandwidth product differs because their time structures differ. For the pulsed candidate, the 10% duty factor at $PRF = 2\text{ kHz}$ implies $\tau = 50\ \mu\text{s}$, so the pulse-compression time-bandwidth product is:

$$(BT)_{pulse} = B\tau = (5 \times 10^6)(50 \times 10^{-6}) = 250.$$

For the FMCW candidate, each chirp uses $T_c = 0.5\text{ ms}$ with the same bandwidth, giving

$$(BT)_{chirp} = BT_c = (5 \times 10^6)(0.5 \times 10^{-3}) = 2500,$$

which is ten times higher, corresponding to a nominal processing-gain advantage when the waveform is processed. This higher BT is an inherent FMCW waveform feature because the same bandwidth is maintained while the target is illuminated continuously over a longer interval per range measurement.

Importantly, the comparison remains fair because the energy placed on the channel per measurement and per dwell is held equal by the common average-power constraint. With $P_t = 1\text{ kW}$ per panel and $T_c = 0.5\text{ ms}$, the energy per FMCW chirp is

$$E_{chirp} = P_t T_c = 1000 \times 0.0005 = 0.5\text{ J}.$$

For the pulsed radar, $P_{peak} = 10\text{ kW}$ and $\tau = 50\ \mu\text{s}$ gives the same energy per pulse:

$$E_{pulse} = P_{peak} \tau = 10,000 \times 50 \times 10^{-6} = 0.5\text{ J}.$$

Over the common dwell of 16 transmissions, both therefore deliver the same energy:

$$E_{CPI} = 16 \times 0.5 = 8 \text{ J},$$

equivalently,

$$E_{CPI} = P_t T_{coh} = 1000 \times 0.008 = 8 \text{ J}.$$

Thus, FMCW's larger BT is preserved as a true waveform feature, while energy delivery and average-power feasibility are controlled so that the comparison does not give one waveform an artificial advantage through greater energy on target.

In addition to the standalone pulsed LFM and FMCW candidates, a Hybrid operating mode is evaluated under the same common constraint set. In this mode, LFM provides the primary detection chain, while FMCW is reserved for classification through the micro-Doppler processing path once a prior detection has been established and a classification-enable condition is satisfied. This condition may be imposed either as a fixed range limit or as an adaptive gate based on estimated classification SNR, so that classification is invoked only when observability is sufficient. The Hybrid case is therefore included as a task-partitioned operating concept to test whether LFM-based detection combined with FMCW-based classification provides a more robust shipboard solution than either standalone waveform under the same aperture, power, and scan constraints.

To conclude, this section has fixed the pulsed, FMCW, and Hybrid operating modes under a single ship-installable constraint set and explicit fairness basis. With the common radar designs now defined, the next step is to specify the representative UAV threat scenarios used in the comparison. The following section therefore introduces the threat-informed flight-path definitions, including target number, starting range, altitude, radial speed, and approach geometry, which together provide the operational context for the subsequent simulation and evaluation.

3.3 Threat-Informed UAV Scenarios and Flight-Path Definitions

To evaluate detection performance in a repeatable and operationally relevant manner, the simulation uses a small set of representative UAV scenarios that reflect distinct employment patterns and impose different stresses on the radar. The three scenarios are defined in a ship-fixed Cartesian frame with the ship at the origin and the effective antenna location taken to be 20 m above the waterline [5,27]. Each scenario is specified by an initial range R_0 , a set of azimuth bearings, a target altitude, a radial closing speed v_r , and the number of UAVs present. The ship attitude dynamics, including pitch, roll, and yaw, are considered out of scope and are therefore not modelled. Together, these scenario definitions are intended to span the principal detection stressors relevant to shipboard

counter-UAV sensing: low-altitude inbound motion with short warning time, moderate-altitude sector search, and slower intelligence-like approaches for which longer dwell is available. These cases are then used throughout the remainder of Chapter 3 to derive the required search sector, the number of beam positions, the dwell time available per look direction, and the resulting coherent processing interval for both the pulsed and FMCW candidates.

Scenario A — Attack (low altitude, multi-UAV inbound).

The attack case represents the short-warning, clutter-limited geometry. Four UAVs are initialized at a common starting range of $R_0 = 40$ km with low altitude $at = 25$ m and inbound radial speed $v_r = 13.5$ m/s. The UAV bearings are laterally separated in azimuth at $\{-40^\circ, -15^\circ, +15^\circ, +40^\circ\}$ to represent a deconflicted multi-contact raid inside a single forward sector. This geometry yields very small elevation angles, less than 0.1° , which is an intentionally challenging low-grazing condition for sea clutter and multipath. The resulting top-down geometry and the implied search sector are shown in Fig. 3-11 [1,22,23].

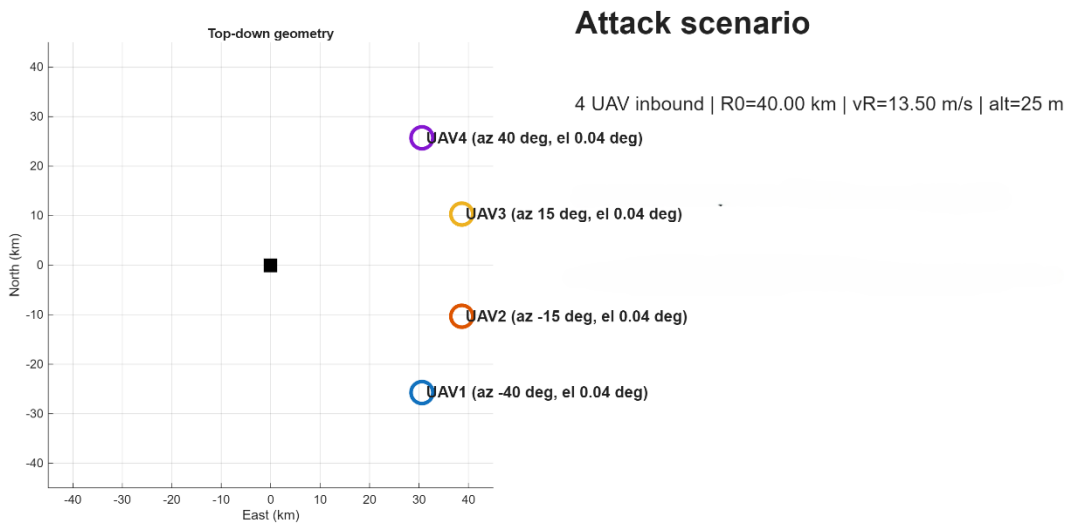


Figure 3-11 UAV Attack scenario

Scenario B — Surveillance (medium altitude, two-UAV sector coverage).

The surveillance case is defined to exercise azimuth/elevation sector search burden rather than only head-on closure. Two UAVs are initialized at $R_0 = 40$ km and $alt = 600$ m with inbound radial speed $v_r = 10$ m/s, placed symmetrically at azimuth bearings {

$-15^0, +15^0$ }. Relative to the attack case, the higher altitude increases the elevation look angle (order $\sim 0.9^0$ at 40 km), which changes both the clutter regime and the beam scheduling burden. The geometry and sector definition used for this scenario are shown in Fig. 3-12 [1,22,23].




Figure 3-12 UAV Surveillance scenario

Scenario C — Intelligence (slow inbound, single UAV).

The Intelligence case represents a slow-speed contact for which longer observation time is available. A single UAV is initialized at $R_0 = 40$ km, at an altitude of 300 m, with an inbound radial speed of $v_r = 3$ m/s and an azimuth of $+15^0$. This geometry yields a modest elevation angle, on the order of 0.4^0 at 40 km, while preserving a long observation interval before intercept. This makes the case suitable for examining how dwell allocation and coherent integration influence detection probability under the same environmental models. The geometry and implied search sector are shown in Fig. 3-13 [1,22,23].




Figure 3-13 UAV Intelligence scenario

Using the radar-horizon relation (3.3) introduced in Section 3.2, the three scenarios imply different geometric line-of-sight limits because the target altitude perceived by the radar changes while the radar height remains fixed at 20 m above sea level [7,8]. For the lowest-altitude attack case ($h_t = 25$ m), the combined radar horizon is 39 km, so the opening range of 40 km is effectively a near-horizon geometry. For the intelligence case ($h_t = 300$ m), the corresponding line-of-sight limit increases to over 89 km, and for the surveillance case ($h_t = 600$ m) it increases further to 119 km as shown in Fig. 3-14. Thus, only the attack scenario is close to the geometric horizon at initialization; the higher-altitude scenarios are well within line of sight, so any long-range detection degradation in those cases should be interpreted primarily as clutter, attenuation, or sensitivity-limited rather than horizon-limited.




Figure 3-14 Scenarios' radar horizon distances

Collectively, these three scenarios provide a consistent and repeatable operational test set for the remainder of Chapter 3. With the threat geometries now defined, the next section determines how the common AESA sector is searched in practice, including the number of beam positions required and the dwell time available per look direction for the imposed refresh requirement.

3.4 Search Strategy and Dwell / Refresh Budget

For the representative scenarios defined in Section 3.3, the search is executed over a single AESA sector spanning $\pm 60^\circ$ in azimuth and 0° to 60° in elevation. Accordingly, full 360° azimuth coverage would require at least three sector positions. The AESA panel used in this work produces boresight half-power beamwidths of approximately 3.2° in azimuth and 2.1° in elevation. To avoid coverage gaps under steering while preserving a simple common scan grid, a 2° step is adopted in both azimuth and elevation. The sector is searched using a serpentine, or raster, scan order in which adjacent elevation rows are traversed in alternating azimuth directions in order to minimize beam-to-beam repositioning overhead.

For the $\pm 60^\circ$ azimuth by 0° to 60° elevation sector, the number of beam positions is therefore

$$N_{az} = \left\lceil \frac{120^\circ}{2^\circ} \right\rceil + 1 = 61, N_{el} = \left\lceil \frac{60^\circ}{2^\circ} \right\rceil + 1 = 31,$$

so that

$$N_{beams} = N_{az}N_{el} = 61 \times 31 = 1891.$$

With the imposed surveillance refresh requirement of $T_{refresh} \leq 2s$, the per-beam integration time is set to

$$T_{dwell} = T_{int} = 1.0 \text{ ms},$$

which gives a full-sector refresh time of

$$T_{refresh} = N_{beams}T_{dwell} = 1891 \times 1.0 \text{ ms} = 1.891 \text{ s}.$$

This satisfies the 2 s refresh constraint while preserving a dense search grid consistent with the serpentine visitation pattern shown in Fig. 3-15.




Figure 3-15 Serpentine Search pattern

The dwell time then fixes the coherent burst length available for search-mode RD processing. With a common slow-time repetition interval of $T_{rep} = 0.5$ ms, corresponding to a pulse or chirp repetition rate of 2 kHz, the number of coherent transmissions that fit within one beam dwell is

$$N_d = \left\lfloor \frac{T_{dwell}}{T_{rep}} \right\rfloor = \left\lfloor \frac{1.0 \text{ ms}}{0.5 \text{ ms}} \right\rfloor = 2,$$

giving a search-mode coherent processing interval of

$$T_{coh} = N_d T_{rep} = 1.0 \text{ ms}.$$


These values, $N_{beams} = 1891$, $T_{dwell} = 1.0$ ms, $T_{refresh} = 1.891$ s, and $T_{coh} = 1.0$ ms, are used for the pulsed, FMCW, and Hybrid operating modes so that the comparison is conducted under a common scan schedule and dwell budget. The next section then uses representative processor-validation examples to confirm that the waveform-specific range processing, Doppler processing, and CFAR stages produce the expected intermediate observables under these common radar settings before the dwell-by-dwell outcomes are aggregated statistically through Monte Carlo evaluation.

3.5 Single-Dwell Intermediate Results and Monte Carlo Evaluation


This section presents representative processor-validation examples and the corresponding statistical evaluation framework used for the comparative results reported in Chapter 4. The purpose is not to re-derive the underlying equations, which were established in Chapter 2, but to verify that the implemented signal-processing steps produce the expected observables under the common radar settings defined earlier in Chapter 3. The examples shown here use longer dwells than the operational search-mode dwell in order to validate waveform-specific range processing, MTI behaviour, and CFAR detection under controlled conditions. They are therefore intended as implementation checks rather than literal beam-by-beam search dwells.

Fig. 3-16 shows the FMCW processor-validation example for a single moving target embedded in sea clutter with rain-dependent attenuation applied through the selected environment model, using $f_c = 10$ GHz, $B = 5$ MHz, $PRF = 2$ kHz, and $N = 128$ chirps ($T_{coh} = 64$ ms). Fig. 3-16(a) shows the fast-time range response after dechirping and fast-time FFT processing. The dominant peak is aligned with the injected truth range at $R = 10$ km, confirming correct range recovery. Fig. 3-16(b) shows the slow-time Doppler cut taken at the truth range bin. The pre- and post-MTI responses demonstrate suppression of


stationary and near-zero-Doppler interference while preserving the moving-target return near $v = 13.5$ m/s. Fig. 3-16(c) shows the zoomed Range-Doppler map with the CFAR overlay, confirming that the target is detected in the expected region of interest. Standard fast-time and slow-time windowing are applied to control sidelobes in the range and Doppler responses.



(a)




(b)




(c)

Figure 3-16 (a) FMCW range response, (b) FMCW Doppler response, (c) FMCW Range-Doppler response


Fig. 3-17 shows the corresponding pulsed LFM example under the same radar settings, but with a staggered PRI composed of three sequential subpulses within a single PRI. Fig. 3-17(a) shows the resulting range response after pulse compression and subpulse alignment. Although three subpulses are transmitted, the final range response remains well behaved because each subpulse is matched-filtered separately, shifted to remove its start-time offset, and aligned to a common delay reference before combination. As a result, the compressed returns coincide at the same true delay, and the final range profile is centered at the injected truth range of $R = 10$ km rather than being smeared across fast time by the staggered transmission structure. Fig. 3-17(b) shows the Doppler cut at the truth range bin. As in the FMCW case, the pre- and post-MTI responses show suppression of near-zero-Doppler clutter while preserving the moving-target return near $v = 13.5$ m/s. Fig. 3-17(c) shows the zoomed Range-Doppler map with the CFAR overlay, demonstrating detection at the expected target location. These results show that, despite the different front-end range processing, both waveform types ultimately produce comparable downstream observables in the common R-D domain.



(a)



(b)



(c)

Figure 3-17 (a) LFM range response, (b) LFM doppler response, (c) LFM range-Doppler response

Taken together, Fig. 3-16 and Fig. 3-17 establish three points before proceeding to the statistical evaluation. First, both waveforms recover a consistent range within the shared bandwidth, with the dominant response aligned with the injected truth range. Second, the coherent record provides a usable Doppler observable, and the pre- and post-MTI responses show that stationary or near-zero-Doppler interference can be suppressed without erasing the moving-target return. Third, the two-dimensional CA-CFAR stage provides an adaptive threshold across the Range-Doppler plane rather than a fixed cut, which is necessary for maritime scenes where clutter and weather losses vary with range and environment. The remainder of this section formalizes how these dwell-level CFAR outcomes are aggregated across Monte Carlo trials to obtain the reported $\hat{P}_d(R)$, time-to-detect, time-to-confirm, and classification metrics.

The single-dwell examples above illustrate the behaviour of the processor for one representative realization. The performance metrics reported in this work, however, require these per-dwell outcomes to be converted into statistical measures under variability in UAV RCS, sea clutter, rain attenuation, and multipath. In the present framework, each dwell produces a binary CFAR decision within the target region of interest, and Monte Carlo evaluation is then used to aggregate these dwell-by-dwell outcomes into $\hat{P}_d(R)$, time-

to-detect, and time-to-confirm results across the representative scenarios and environments [25,28].

To ensure that the adopted RCS fluctuation severity is defensible, the fluctuation statistics are anchored to measured small-UAV results. Here, fluctuation severity refers to the degree to which the instantaneous target RCS varies about its mean value. In the χ^2 model used here, this severity is controlled by the effective degree-of-freedom parameter ν : lower values of ν correspond to stronger, more intermittent fluctuations, while higher values correspond to smoother behaviour. Static and dynamic UAV RCS statistics reported in [29] were used to select the fluctuation model. The dynamic measurements consistently exhibit both a larger mean and a larger variance than the corresponding static measurements, which indicates that a time-varying fluctuation model is required if realistic UAV echo behaviour is to be represented. For the 8.5 GHz VV case, which is the closest reported case to the X-band frequency used here, the measured dynamic statistics are a mean of -15.83 dBsm and a variance of -29.16 dBsm [29]. A conservative mean small-UAV RCS of -20 dBsm is adopted in the present model, while the fluctuation severity reported in [29] is retained. This preserves the spiky target behaviour implied by the measured data while enforcing the lower mean RCS representative of the counter-UAV threat considered throughout the comparison.

To implement this behaviour in simulation, the RCS power fluctuation is modelled as a χ^2 , or equivalently gamma, distributed random process with an effective degree-of-freedom parameter. This choice is also consistent with dynamic-flight UAV RCS literature, which shows that a χ^2 distribution can provide a better fit to dynamic-flight RCS variation than lognormal or Rice models in the evaluated viewing sectors [12]. A convenient moment-matched χ^2 power model is

$$\sigma = \bar{\sigma} \left(\frac{X}{\nu} \right), X \sim \chi^2(\nu) \quad (3.4)$$

which gives

$$E[\sigma] = \bar{\sigma}, \text{Var}[\sigma] = \bar{\sigma}^2 \left(\frac{2}{\nu} \right) \quad (3.5)$$

so that matching a measured coefficient of variation, or equivalently, a measured mean and variance, implies

$$\nu = \frac{2\bar{\sigma}^2}{\text{Var}[\sigma]} \quad (3.6)$$

Using the 8.5 GHz VV dynamic statistics, the reported mean and variance imply an equivalent χ^2 fluctuation parameter of $\nu \approx 1.12$. This low value of ν indicates a strongly fluctuating target, meaning that the return power is expected to vary sharply from dwell to dwell rather than remain close to its mean. A mean value of -20 dBsm is then imposed with a floor of -50 dBsm, as a modelling choice to maintain stable Monte Carlo behavior under the strongly fluctuating χ^2 RCS model and to limit excessively deep lower-tail fades in the sampled dwell-to-dwell RCS sequence [12,29]. This yields a mean-consistent baseline target model that preserves the measured fluctuation severity while enforcing the more conservative mean RCS used in the Monte Carlo detection experiments. A representative realization of the resulting χ^2 fluctuation process is shown in Fig. 3-18.

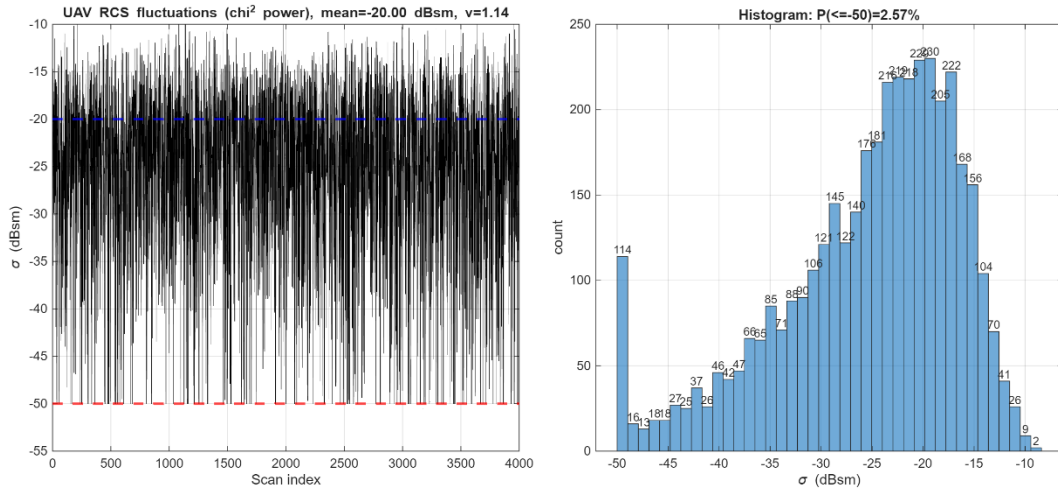


Figure 3-18 UAV fluctuation with χ^2 model

For a fixed range R (or for each scan index that maps to a range along a scenario trajectory), the simulation is run for N_{exp} independent trials. Each trial draws independent realizations of the random components (RCS draw, clutter/interference realization, multipath fade where applicable), generates the range–Doppler cell under test, applies CA-CFAR at the specified P_{fa} , and produces a binary outcome:

$$D_e(R) = \begin{cases} 1, & \text{CFAR declares detection in the ROI} \\ 0, & \text{otherwise} \end{cases} \quad e = 1, 2, \dots, N_{\text{exp}}. \quad (3.7)$$

The Monte Carlo estimate of the probability of detection is then:

$$\hat{P}_d(R) = \frac{1}{N_{\text{exp}}} \sum_{e=1}^{N_{\text{exp}}} D_e(R) \quad (3.8)$$

In the present model, each dwell also produces a binary CFAR outcome in time. Let $D_e[i]$ denote the detection event at dwell index i for Monte Carlo experiment e . Time-to-detect is defined as the first dwell index at which a detection occurs:

$$i_{\text{det},e} = \min\{i: D_e[i] = 1\} \quad (3.9)$$

Time-to-confirm uses an M -of- N sliding-window rule over dwells. Defining the confirmation event as

$$C_e[i] = \mathbf{1} \left\{ \sum_{k=i-N+1}^i D_e[k] \geq M \right\} \quad (3.10)$$

the corresponding confirmation time is

$$i_{\text{conf},e} = \min\{i: C_e[i] = 1\} \quad (3.11)$$

Aggregating $i_{\text{det},e}$ and $i_{\text{conf},e}$ across experiments yields the reported time-to-detect and time-to-confirm distributions, together with their associated range-at-time values under each scenario trajectory.

Following detection confirmation, classification is treated as a second-stage decision. At each dwell, a classification event can occur only if the target has already been confirmed as detected, the classification gate is enabled, and the MD score exceeds the classification threshold. The resulting binary classification event stream is then aggregated across Monte Carlo trials to estimate $P_c(R)$, time-to-classify, and time-to-confirm-classify.

This produces a classification event stream analogous to the detection event stream, but classification can occur only when the MD signature is sufficiently visible. In the present model, visibility is assumed to decrease as the effective classification SNR decreases. A bounded SNR-dependent penalty is therefore applied so that weak or clutter-obscured returns are less likely to exceed the classification threshold, with the maximum reduction limited by `mdPenaltyMax`. The purpose of this term is simply to model degraded MD extraction under poor signal conditions; it is not intended to represent a parameter taken directly from the literature [1,21,30].

Overall, the processor-validation results and the statistical models defined in this chapter establish the simulation-based comparison framework. The single-dwell layer produces

range, Doppler, and CFAR decision observables for each waveform, while the Monte Carlo layer converts those dwell-by-dwell outcomes into the detection and classification metrics used for evaluation. Under this framework, the effects of RCS fluctuation, sea clutter, rain attenuation, and multipath are carried consistently into the reported $\hat{P}_d(R)$, time-to-detect, time-to-confirm, and classification results. Chapter 4 then applies this methodology to the representative UAV scenarios and environmental conditions and presents the resulting comparison between the pulsed LFM, FMCW, and Hybrid operating modes.

Chapter 4

Results

Chapter 4 presents and analyzes the simulation results which are evaluated at starting ranges of 40 km and 10 km which shows the difference between a blue ocean situation and an entering or exiting harbour situation. However, the main analytical focus is restricted to the Attack scenario with a 40 km start range because it is the most operationally demanding and decision-relevant case. This scenario combines low-altitude inbound motion, short warning time, near-horizon operation, and the strongest exposure to sea clutter, rain attenuation, and multipath, and therefore provides the clearest basis for comparing the pulsed LFM, FMCW, and Hybrid operating modes under the common ship-installable constraint set. Rather than reporting every scenario and start range in the main body, this chapter uses the Attack 40 km case as the principal deep-dive and emphasizes interpretation alongside the key tables and figures. The remaining Attack 10 km, Surveillance, and Intelligence results are provided in Appendix A, where they can be used to confirm that the broader trends described here remain consistent across the full scenario set. For scenarios containing more than one UAV, the reported detection and classification probabilities are computed separately for each UAV track and then averaged across the UAVs in that scenario at each range step. The milestone ranges P_{d50} , P_{c50} , and P_{c80} are then extracted from these scenario-averaged probability curves.

The chapter is therefore organized around the results that drive the recommendation. The Attack 40 km case is presented first under Best, Medium, and Worst maritime conditions, followed by interpretation of the operational meaning of the resulting detection and classification milestones. The intent is not to restate every numerical value in paragraph form, but to show which trends matter, why they occur, and how they support the final recommendation.

4.1 Attack Scenarios (40 km Start Range)

The Attack scenario with a 40 km start range is the principal decision case. It represents a low-altitude, multi-UAV inbound geometry that begins close to the radar horizon and is therefore jointly stressed by limited warning time and the most severe maritime interference mechanisms. For that reason, it is the best single case for revealing whether the common-constraint pulsed LFM, FMCW, or Hybrid operating mode is operationally preferable. Table 4-1 summarizes the key confirmed-detection and confirmed-classification milestones for the three waveform modes under Best, Medium, and Worst maritime conditions, while the associated statistics figures show how the cumulative detection and classification behaviour evolves with time and range. The

corresponding results for the remaining scenarios and for the 10 km start cases are provided in Appendix A.

4.1.1 Best-Case Weather (Sea State 1, 0 mm/h rain)

As shown in Table 4-1, under Best conditions, all three operating modes retain detection very close to the 40 km opening range. The pulsed LFM, FMCW, and Hybrid cases all reach $P_{d,50}$ at approximately 38.97 km, which indicates that in the absence of severe environmental penalties the shared aperture, bandwidth, and coherent processing assumptions dominate over waveform-specific differences. The main separation under Best conditions appears not in initial detection, but in the time required to achieve confirmed classification. LFM and FMCW remain closely aligned, reaching $P_{c,50}$ near 33.8 km and $P_{c,80}$ near 32.0 km, whereas Hybrid is substantially delayed, with $P_{c,50}$ and $P_{c,80}$ contracting inward to approximately 28.1 km and 26.7 km, respectively. Thus, in favourable conditions the standalone waveforms are effectively equivalent in the Attack 40 km case, while Hybrid preserves early detection but delays classification because its classification path is only invoked after the LFM-based detection leg has already succeeded.

TABLE 4-1 ATTACK BEST-CASE WEATHER (SEA STATE 1, 0 MM/H RAIN)

Starting Range	Waveform	$P_{d,50}$ time (MM:SS:ss)	$P_{d,50}$ range (km)	$P_{c,50}$ time (MM:SS:ss)	$P_{c,50}$ range (km)	$P_{c,80}$ time (MM:SS:ss)	$P_{c,80}$ range (km)
40 km	LFM	01:16.61	38.97	07:37.52	33.82	09:53.02	31.99
40 km	FMCW	01:16.53	38.97	07:39.62	33.80	09:54.04	31.98
40 km	Hybrid	01:16.46	38.97	14:44.08	28.06	16:23.50	26.72

4.1.2 Medium-Case Weather (Sea State 3, 10 mm/h rain)

As shown in Table 4-2, under Medium conditions, the Attack 40 km case becomes clearly range-limited, but the three operating modes still remain similar at the point of first confirmed detection. LFM, FMCW, and Hybrid all reach $P_{d,50}$ near 20.7 km, showing that by this stage the dominant limitation is the shared rain and clutter. The main distinction again appears in the classification milestones. LFM and FMCW remain nearly coincident, with $P_{c,50}$ occurring near 19.4 km and $P_{c,80}$ near 18.4 km, while Hybrid is shifted farther inward, reaching $P_{c,50}$ near 17.8 km and $P_{c,80}$ near 16.9 km. The interpretation is therefore

that Medium weather compresses the usable warning envelope for all three architectures, but does not yet create a decisive separation between LFM and FMCW. At this level of environmental severity, the shared maritime penalty dominates, while Hybrid remains slower because the LFM detection leg still sets the point at which the FMCW classification leg can begin to contribute.

TABLE 4-2 ATTACK MEDIUM-CASE WEATHER (SEA STATE 3, 10 MM/H RAIN)

Starting Range	Waveform	$P_{d,50}$ time (MM:SS:ss)	$P_{d,50}$ range (km)	$P_{c,50}$ time (MM:SS:ss)	$P_{c,50}$ range (km)	$P_{c,80}$ time (MM:SS:ss)	$P_{c,80}$ range (km)
40 km	LFM	23:46.02	20.75	25:29.07	19.36	26:41.64	18.38
40 km	FMCW	23:44.30	20.77	25:26.35	19.39	26:42.36	18.37
40 km	Hybrid	23:45.80	20.75	27:24.25	17.80	28:29.13	16.93

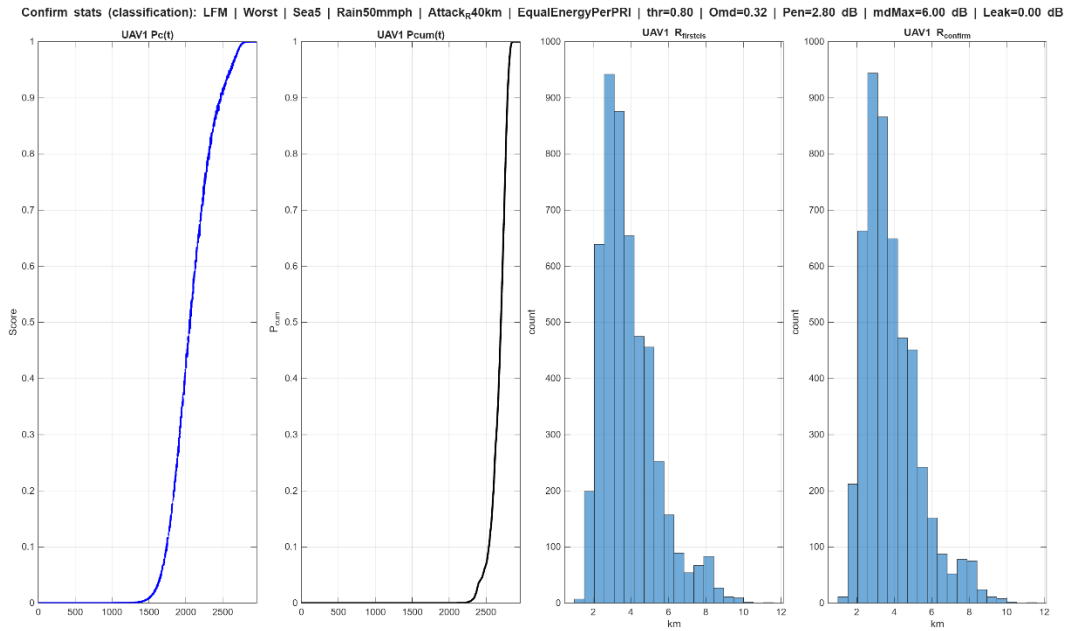
4.1.3 Worst-Case Weather (Sea State 5, 50 mm/h rain)

As shown in Table 4-3, the Attack 40 km case under Worst conditions is the decisive result. Here, the target remains effectively undetectable over most of the inbound track, and the radar does not reach $P_{d,50}$ until the target has closed to very short range. In the LFM case, $P_{d,50}$ is reached only at 3.56 km, with $P_{c,50}$ at 3.54 km and $P_{c,80}$ at 2.66 km. Hybrid behaves almost identically, reaching $P_{d,50}$ at 3.52 km, $P_{c,50}$ at 3.46 km, and $P_{c,80}$ at 2.59 km, which confirms that the limiting mechanism is still the LFM-based detection path. FMCW shows the same qualitative trend, but the key difference is that all three milestones occur farther out: $P_{d,50}$ at 4.71 km, $P_{c,50}$ at 4.68 km, and $P_{c,80}$ at 3.94 km. Relative to LFM, this corresponds to gains of approximately 1.15 km at $P_{d,50}$, 1.14 km at $P_{c,50}$, and 1.28 km at $P_{c,80}$. These gains are not large fractions of the full 40 km design window, but in the hardest operational case they are decisive because they directly convert into additional reaction time in the most clutter-limited portion of the engagement. This is the result that drives the recommendation.

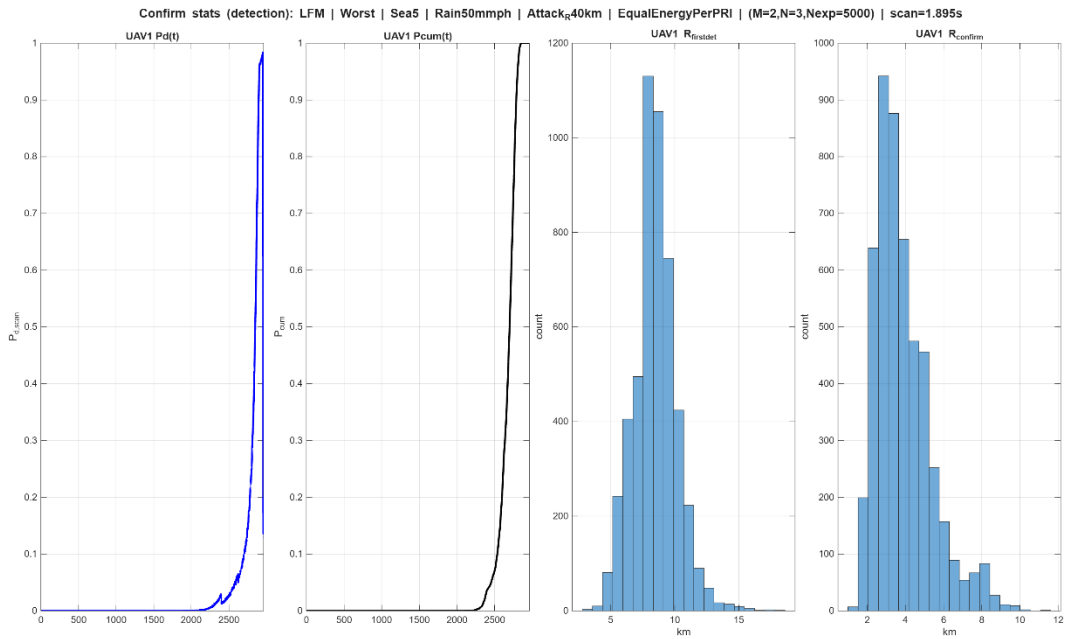
TABLE 4-3 ATTACK WORST-CASE WEATHER (SEA STATE 5, 50 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
40 km	LFM	44:59.00	3.56	45:00.92	3.54	46:05.82	2.66
40 km	FMCW	43:33.99	4.71	43:36.01	4.68	44:31.20	3.94
40 km	Hybrid	45:02.53	3.52	45:06.32	3.46	46:11.01	2.59

Fig.4-1 demonstrates the worst-case Attack 40 km behaviour of the LFM mode by showing both the detection and classification statistics in time and in range. In panel (a), the first plot shows the probability of classification as a function of scan time, while the second plot shows the cumulative confirmed classification probability after the M -of- N rule is applied. The third histogram gives the distribution of the first range at which classification is achieved across the Monte Carlo trials, and the fourth gives the distribution of the corresponding confirmed classification range. Panel (b) is interpreted in the same way for detection: the first plot shows probability of detection versus time, the second shows cumulative confirmed detection probability, the third shows the distribution of first detection range, and the fourth shows the distribution of confirmed detection range. Taken together, the figure shows that under the worst maritime condition the LFM mode remains below useful performance over most of the inbound track and only rises sharply after substantial target closure, so occasional early events are not enough for reliable declaration; instead, both confirmed detection and confirmed classification occur only at short range.



(a)



(b)

Figure 4-1 (a) Classification Stats and (b) Detection Statistics for LFM radar during Attack Scenario under Worst Case Scenario

The worst-case Attack 40 km results show that severe sea clutter, rain, and low-grazing geometry compress the usable warning envelope much more strongly for the LFM-based modes than for FMCW. Figure 4-1 shows that both detection and classification remain weak over most of the approach and only become reliably confirmable after substantial target closure, while Table 4-1 shows that FMCW preserves a consistent outward margin over LFM and Hybrid at the main confirmed milestones. This makes the worst-case Attack 40 km geometry the most informative case and provides the basis for the broader cross-condition interpretation presented in the next subsection.

4.1.4 Interpretation of the Attack 40 km Results

Across Best, Medium, and Worst conditions, the Attack 40 km case shows a consistent pattern. Under Best and Medium conditions, LFM and FMCW are nearly indistinguishable in the main detection milestones, while Hybrid preserves similar detection behaviour but reaches classification later. More specifically, the Hybrid mode uses the same FMCW micro-Doppler classifier as the standalone FMCW case, but only after the LFM detection leg has first declared a confirmed track and handed that track to the classifier. Hybrid therefore inherits the shorter-range LFM detection gate before classification can begin, which is why its P_{c50} and P_{c80} milestones occur later than in standalone FMCW. Under Worst conditions, however, FMCW retains a meaningful detection and classification advantage over the two LFM-based modes. This is consistent with the broader scoring results which include the scenarios in Appendix A, which show that although FMCW is not universally dominant in every case, it is the most consistent waveform and avoids the most severe degradation in the hardest cases. As such, the FMCW mode performs the best.

The physical explanation follows directly from the common-constraint design basis established in Chapter 3. Both candidates are normalized to the same energy per measurement and per dwell, so the comparison does not favour one waveform by assigning it more energy on target. Even so, FMCW retains the larger waveform-processing term because each chirp occupies the full 0.5 ms repetition interval, whereas the pulsed candidate is limited to a total transmit on-time of 50 μ s within that same interval. Moreover, the pulsed LFM design must divide its transmit energy across short staggered segments in order to preserve close-range access, longer-range energy on target, and ambiguity control within the same waveform structure, whereas FMCW maintains continuous illumination throughout the dwell. This is why the LFM and FMCW results are similar in favourable conditions but separate in the hardest low-altitude, clutter-limited Attack case. Hybrid does not recover the FMCW margin because its detection leg remains tied to the LFM front end, and the FMCW classification leg cannot help until detection has already been achieved.

4.2 Remaining Engagement Time Results

To express the results in operational terms, the ranges reported in the preceding sections and appendix are converted to remaining engagement time before ship intercept. For a detection reached at range R_m in a scenario with constant inbound radial speed v_r , the available reaction time is

$$T_{\text{react}} = \frac{R_m}{v_r} \quad (4.1)$$

In the present scenarios, $v_r = 13.5$ m/s for Attack, $v_r = 10$ m/s for Surveillance, and $v_r = 3$ m/s for Intelligence. Because the 10 km start cases already lie inside the short-range envelope and most of their milestones remain close to 10 km, Table 4-4 focuses on the 40 km start cases, which are the more relevant warning-range conditions for shipboard response.

TABLE 4-4 REMAINING ENGAGEMENT TIME - SUMMARY

Scenario	Weather	T_{react} at $P_{d,50}$, LFM (MM:SS)	T_{react} at $P_{d,50}$, FMCW (MM:SS)	T_{react} at $P_{d,50}$, Hybrid (MM:SS)	T_{react} at $P_{c,80}$, LFM (MM:SS)	T_{react} at $P_{c,80}$, FMCW (MM:SS)	T_{react} at $P_{c,80}$, Hybrid (MM:SS)
Attack	Best	49:18	49:18	49:18	39:33	39:31	33:02
Attack	Medium	25:37	25:39	25:37	22:41	22:41	20:54
Attack	Worst	04:24	05:49	04:21	03:17	04:52	03:12
Surveillance	Best	66:31	66:31	66:31	65:24	65:24	65:00
Surveillance	Medium	65:43	65:42	65:43	58:23	58:23	57:07
Surveillance	Worst	59:27	59:28	59:25	58:04	58:03	56:20
Intelligence	Best	222:10	222:10	222:10	210:03	210:00	207:17
Intelligence	Medium	152:23	152:17	152:23	65:27	65:40	43:20
Intelligence	Worst	71:30	71:33	71:30	68:37	68:40	68:23

Table 4-4 shows that the operational effect of waveform choice depends strongly on both scenario and environmental severity. The Attack case is the most demanding. Under Best weather, all three operating modes provide substantial remaining engagement time at both the $P_{d,50}$ and $P_{c,80}$ milestones, with little practical difference between LFM and FMCW at

first detection, while Hybrid is slower at high-confidence classification. Under Medium weather, the ship has less time to respond, but all three modes still provide more than 20 minutes of remaining engagement time in the Attack case. The clearest separation occurs under Worst weather. At the $P_{d,50}$ milestone, the remaining engagement time is (04:24) for LFM, (05:49) for FMCW, and (04:21) for Hybrid. All three modes remain above the nominal 4 min Action Stations requirement at first detection, but LFM and Hybrid exceed the requirement by only a small margin. If action is delayed until a higher-confidence classification is reached, represented here by the $P_{c,80}$ milestone, LFM and Hybrid no longer meet the 4 min requirement, with remaining times of (03:17) and (03:12), respectively, whereas FMCW remains above the requirement at (04:52). This is the clearest operational conclusion from the worst-case Attack 40 km result: FMCW does not simply detect the UAV at longer range, but provides the ship with more usable time to react in the most demanding low-altitude maritime engagement.

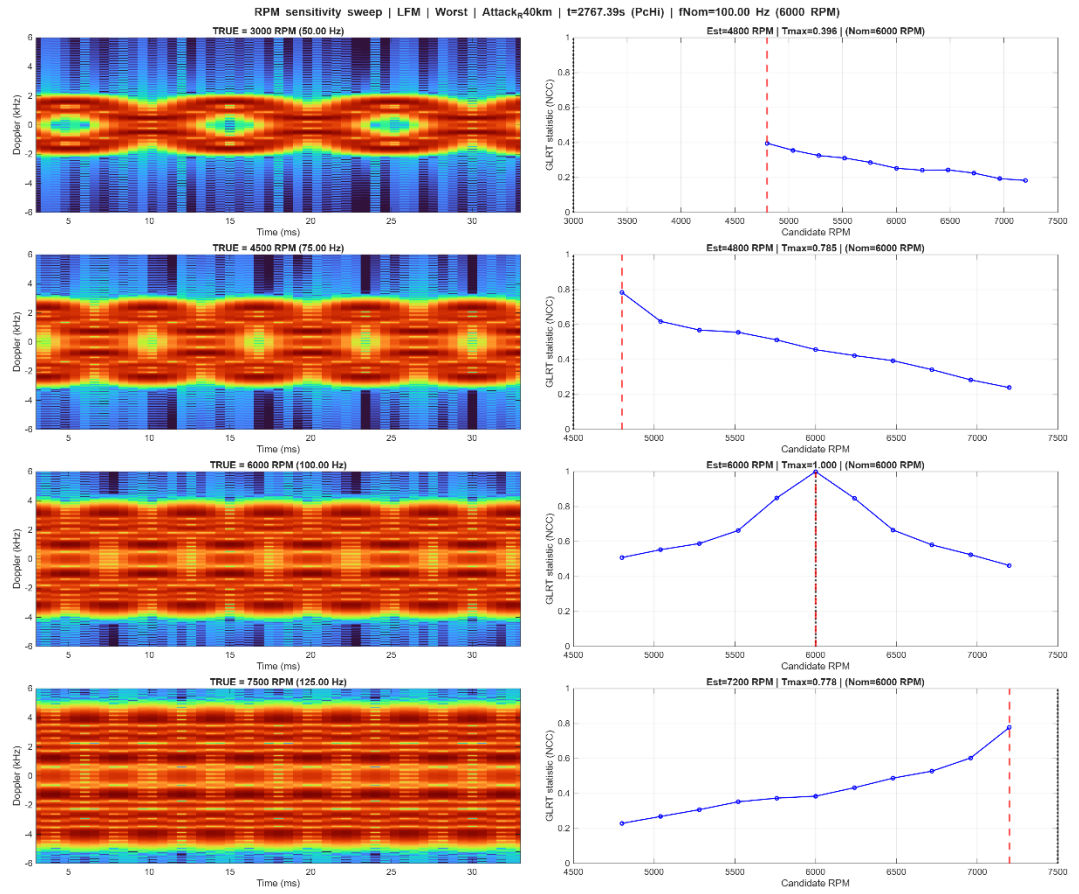
The Surveillance and Intelligence results found in Appendix A provide important context for that conclusion. In the Surveillance case, all three modes retain substantially larger engagement margins, even in Worst weather. At $P_{d,50}$, the remaining time stays close to one hour, ranging from (59:25) to (59:28), while at $P_{c,80}$ it remains between (56:20) and (58:04). These values show that for the medium-altitude Surveillance geometry, the environmental penalties are not severe enough to reduce the engagement timeline to the point where waveform choice becomes operationally decisive, even though Hybrid remains consistently slower in classification than the two standalone modes. In the Intelligence case, the margins are larger still because of the slower target speed. Under Best weather the remaining engagement times extend to several hours, with 222 minutes at $P_{d,50}$ and approximately 210 minutes at $P_{c,80}$. Even in Worst weather, the Intelligence case still retains more than an hour of remaining engagement time at both milestones for all three modes. The one notable exception is the Medium-weather Hybrid classification result, which decreases to (43:20) at $P_{c,80}$, compared with (65:27) for LFM and (65:40) for FMCW, again reflecting the poorer performance of the Hybrid mode.

4.3 Micro-Doppler Results

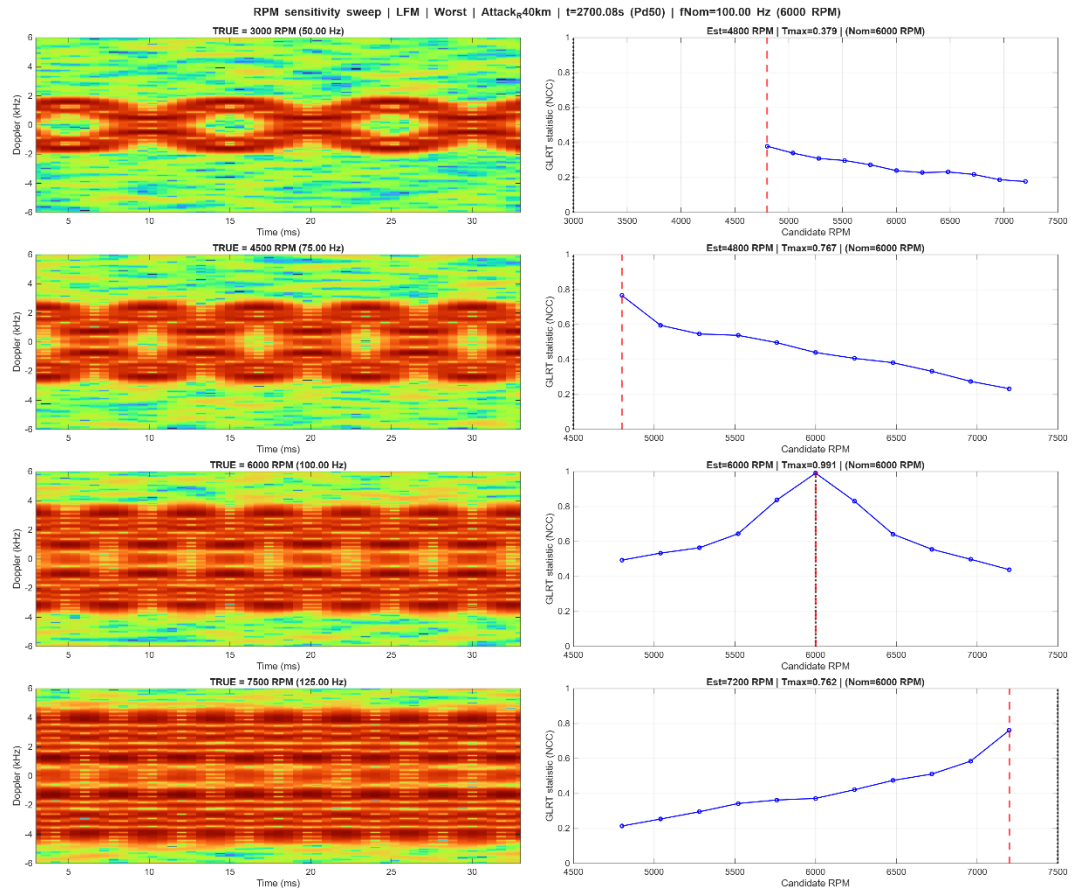
To quantify sensitivity of the MD classification stage to rotor-speed mismatch, a Revolution per Minute (RPM) sweep was performed in which the true rotor speed was set to 3000, 4500, 6000, and 7500 RPM, while the template bank remained centred at the nominal 6000 RPM with the configured $\pm 20\%$ uncertainty band, that is, 4800 to 7200 RPM. Fig. 4-2 should be read row by row. Each row corresponds to one true rotor speed. In both panels, the left-hand plot in each row is the observed MD spectrogram over the selected time window, whereas the right-hand plot shows the corresponding Generalized Likelihood Ratio Test/Normalized Cross-Correlation (GLRT/NCC) score evaluated across

the candidate RPM values in the template bank. Here, the GLRT provides a statistical measure of how strongly the observed micro-Doppler data support a given RPM-template hypothesis relative to the others, while the NCC quantifies the normalized similarity between the observed spectrogram and each candidate template. The spectrograms show the expected physical trend: increasing RPM increases the micro-Doppler extent and changes the repetition structure over the dwell. The score plots then show how that changing rotor signature maps onto the template bank used for classification. In this analysis, T_{max} denotes the maximum GLRT statistic across all RPM templates, equivalently the peak NCC between the observed micro-Doppler spectrogram, within the selected region of interest, and the template set. As such, T_{max} provides a unitless best-match measure of how well the observation aligns with any candidate RPM. The GLRT is therefore used here as a best-match selector across the RPM bank, while NCC provides the associated normalized similarity measure of template match quality.

The four rows show the classifier behaviour under increasing mismatch. For the 3000 RPM case, the true rotor speed lies well below the lower bank boundary of 4800 RPM. The best match therefore saturates at 4800 RPM and the peak statistic is weak, with $T_{max} \approx 0.38$ to 0.40. For the 4500 RPM case, the true rotor speed is still outside the bank but closer to its lower edge, so the estimate again clips at 4800 RPM, but the peak remains moderately high, at approximately $T_{max} \approx 0.76$ to 0.79. For the nominal 6000 RPM case, which lies at the centre of the bank, the estimator locks correctly at 6000 RPM and the match becomes near-ideal, with $T_{max} \approx 0.991$ at $P_{d,50}$ and $T_{max} \approx 1.000$ at the later high-confidence classification snapshot. For the 7500 RPM case, the true rotor speed lies above the upper bank boundary, so the estimate saturates at 7200 RPM and the score remains moderately high because the true RPM is close to the upper edge of the assumed template interval rather than far outside it. Thus, the figure shows both effects of mismatch clearly: first, clipping of the estimated RPM to the nearest allowable template, and second, reduction in peak match quality as the true rotor speed moves farther outside the assumed uncertainty band.



(a)



(b)

Figure 4-2 RPM sensitivity sweep for the LFM micro-Doppler classifier in the worst-case Attack 40 km condition: (a) later high-confidence classification snapshot and (b) earlier $P_{d,50}$ snapshot. Each row corresponds to a different true rotor speed, with the left-hand panel showing the observed spectrogram and the right-hand panel showing the GLRT/NCC score across the candidate RPM templates.

The difference between panel (a) and panel (b) is not the qualitative trend, but the strength of the evidence available to the classifier at the time the snapshot is taken. Panel (a), taken at the later high-confidence classification point, produces slightly stronger peak statistics than panel (b), taken at the earlier $P_{d,50}$ point. This is consistent with classification occurring later in the run when the micro-Doppler evidence is cleaner and the classifier has greater margin. Even so, both panels support the same conclusion: when the true RPM lies inside the template bank, the classifier remains highly accurate, whereas when the true RPM lies outside the bank, the estimate clips to the nearest bank boundary and the match statistic falls. The $\pm 20\%$ RPM uncertainty therefore directly determines the rotor-speed

interval over which template-based micro-Doppler classification remains reliable before becoming limited by template mismatch.

4.4 Overall Performance Comparisons

Across all scenario-environment scenarios, each waveform is evaluated using the same operational figures of merit and then reduced to a single weighted score. The detection figures of merit are the 80% confirmed-detection range, $R_{d,80}$, the 80% confirmed-detection time, $T_{d,80}$, and the median confirmed-detection range, $\tilde{R}_{\text{conf,det}}$. The classification figures of merit are defined in the same manner: the 80% confirmed-classification range, $R_{c,80}$, the 80% confirmed-classification time, $T_{c,80}$, and the median confirmed-classification range, $\tilde{R}_{\text{conf,cls}}$. Let $\hat{R}_{d,80}$, $\hat{T}_{d,80}$, and $\hat{\tilde{R}}_{\text{conf,det}}$ denote the case-normalized versions of the three detection metrics, and let $\hat{R}_{c,80}$, $\hat{T}_{c,80}$, and $\hat{\tilde{R}}_{\text{conf,cls}}$ denote the corresponding normalized classification metrics. The detection and classification subscores are then defined as

$$S_{\text{Det}} = \frac{\hat{R}_{d,80} + \hat{T}_{d,80} + \hat{\tilde{R}}_{\text{conf,det}}}{3} \quad (4.2)$$

$$S_{\text{Cls}} = \frac{\hat{R}_{c,80} + \hat{T}_{c,80} + \hat{\tilde{R}}_{\text{conf,cls}}}{3} \quad (4.3)$$

The total waveform score is formed as

$$S_{\text{Total}} = w_{\text{Det}}S_{\text{Det}} + w_{\text{Cls}}S_{\text{Cls}} \quad (4.4)$$

where w_{Det} is the detection-weight coefficient and w_{Cls} is the classification-weight coefficient. In the present study,

$$w_{\text{Det}} = 0.65, w_{\text{Cls}} = 0.35, w_{\text{Det}} + w_{\text{Cls}} = 1 \quad (4.5)$$

This weighting reflects the design assumption that reliable detection is operationally more important than early classification.

Fig. 4-3(a) summarizes all of the results. The rows correspond to the three weather cases, namely Worst (Sea 5, Rain 50 mm/h), Medium (Sea 3, Rain 10 mm/h), and Best (Sea 1, Rain 0). The columns correspond to the six scenario/start-range combinations: Attack R40 km, Attack R10 km, Intel R40 km, Intel R10 km, Surveillance R40 km, and Surveillance R10 km. Each cell is colour-coded by the winning radar candidate, with green for FMCW, blue for LFM, and orange for Hybrid. Within each cell, the first line identifies the winner, the second gives the total weighted score S together with the score difference compared to the runner-up, and the next lines report the winning values of $R_{d,80}$ and $R_{c,80}$. The Δ values show the winner's range margin over the runner-up. The purpose of the figure is therefore not only to show which waveform wins each case, but also how strong that win is and whether it is driven mainly by detection range, classification range, or both.

Read in that way, the figure shows that FMCW is the dominant winner across the matrix, especially in the most important Attack 40 km column, where it wins under Best, Medium, and Worst weather. That column is the main operational decision case, because it represents the hardest low-altitude warning geometry. The figure also shows that LFM still wins several cells, particularly in some Intelligence and Surveillance cases, and that Hybrid wins only one cell. This is important because the recommendation is not based on FMCW winning every case, but on its ability to remain competitive across almost all cases while also winning the most demanding one. In other words, the winner map shows both case-by-case variation and the fact that the broadest robustness belongs to FMCW.

Scenario	Attack R40 km	Attack R10 km	Intel R40 km	Intel R10 km	Surveil R40 km	Surveil R10 km
Worst Sea 5 Rain 50	FMCW S = 1.00 Δ = +0.99 R.d80 = 3.994 km R.c80 = 3.968 km Δ = +1.328 km	FMCW S = 1.00 Δ = +0.99 R.d80 = 3.970 km R.c80 = 3.943 km Δ = +1.303 km	LFM S = 0.99 Δ = +0.51 R.d80 = 12.405 km R.c80 = 12.376 km Δ = -0.003 km	FMCW S = 0.80 Δ = +0.10 R.d80 = 9.960 km R.c80 = 9.948 km Δ = +0.000 km	LFM S = 0.74 Δ = +0.26 R.d80 = 16.797 km R.c80 = 14.133 km Δ = -0.004 km	FMCW S = 0.74 Δ = +0.29 R.d80 = 9.965 km R.c80 = 9.908 km Δ = +0.068 km
Medium Sea 3 Rain 10	FMCW S = 0.82 Δ = +0.20 R.d80 = 19.005 km R.c80 = 18.358 km Δ = +1.417 km	FMCW S = 0.80 Δ = +0.30 R.d80 = 9.828 km R.c80 = 9.794 km Δ = +0.003 km	LFM S = 1.00 Δ = +0.32 R.d80 = 26.244 km R.c80 = 20.720 km Δ = +0.000 km	LFM S = 0.76 Δ = +0.30 R.d80 = 9.989 km R.c80 = 9.979 km Δ = +0.012 km	HYBRID S = 0.65 Δ = +0.11 R.d80 = 27.959 km R.c80 = 17.449 km Δ = -2.113 km	LFM S = 0.80 Δ = +0.46 R.d80 = 9.966 km R.c80 = 9.932 km Δ = +0.001 km
Best Sea 1 Rain 0	FMCW S = 0.80 Δ = +0.45 R.d80 = 39.846 km R.c80 = 31.999 km Δ = -0.016 km	LFM S = 0.80 Δ = +0.30 R.d80 = 9.953 km R.c80 = 9.918 km Δ = +0.001 km	FMCW S = 0.47 Δ = +0.02 R.d80 = 39.965 km R.c80 = 35.056 km Δ = +6.277 km	FMCW S = 0.80 Δ = +0.35 R.d80 = 9.990 km R.c80 = 9.982 km Δ = +0.009 km	FMCW S = 0.80 Δ = +0.11 R.d80 = 39.890 km R.c80 = 32.668 km Δ = +0.021 km	FMCW S = 0.80 Δ = +0.35 R.d80 = 9.966 km R.c80 = 9.940 km Δ = +0.030 km

Legend

- FMCW
- LFM
- HYBRID

S = Detection score (0 to 1)
 R.d80 = 80% detection range (km)
 R.c80 = 80% confirmed-classification range (km)
 Δ = Margin over runner-up (km)

Figure 4-3 Cross-case synthesis results of all cases

Fig. 4-4 then collapses the 18 individual results into an overall summary across all cases. Here, the mean score is the average weighted score across the full 18-case matrix, so it indicates the waveform’s typical overall performance. The minimum score is the lowest weighted score achieved in any single case, so it indicates the waveform’s worst-case performance or robustness floor. On that basis, FMCW remains the strongest overall option. It achieves 11 case wins, compared with 6 for LFM and 1 for Hybrid. It also has the highest mean score, 0.68, and the highest minimum score, 0.35. By comparison, LFM has a mean score of 0.53 and a minimum score of 0.01, while Hybrid has a mean score of 0.25 and a minimum score of 0.00. This means that FMCW is not only the most frequent winner, but also the most stable overall performer, because it is strongest on both the average-case measure and the worst-case measure.

This distinction between mean score and minimum score is important for interpretation. A high mean score shows that a waveform performs well on average across the full set of scenarios and environments. A high minimum score shows that even in its weakest case, the waveform still retains useful performance. In other words, the mean score measures overall consistency across the case set, while the minimum score measures resilience in the hardest individual case. Fig. 4-4 therefore shows that FMCW is preferred not simply because it wins the largest number of cases, but because it also avoids the severe collapse seen in the weakest LFM and Hybrid outcomes.




Figure 4-4 Overall summary across all cases

Taken together, the results of Chapter 4 show that the recommendation is driven by consistency and worst-case resilience rather than by universal dominance in every individual case. The detailed Attack 40 km analysis identified the low-altitude, clutter-limited attack geometry as the decisive operational case, while the cross-case synthesis showed that FMCW provides the strongest aggregate balance of detection and classification performance across the 18 evaluated scenario–environment combinations, with the highest mean score and the strongest worst-case floor. The remaining Surveillance and Intelligence results support that interpretation by showing that waveform differences are less operationally decisive when the engagement timeline is longer. Chapter 5 therefore draws these results together into a final design-oriented recommendation, then defines the modelling limits within which that recommendation should be interpreted, and finally outlines the most important directions for future work, including richer UAV behaviour, implementation-level radar effects, and more advanced classification methods. Ultimately, the detailed Attack 40 km results and the cross-case synthesis support the final recommendation that FMCW is the preferred architecture within the tested constraint set. This recommendation is driven not by universal dominance in every case, but by FMCW’s stronger aggregate balance of detection and classification performance and, most importantly, by its superior worst-case robustness in the most operationally relevant maritime conditions.

Chapter 5

Conclusion

5.1 Summary

In Chapter 2, the key radar concepts needed for the comparison were reviewed. This included the basic operation of coherent radar systems, AESA beamsteering, pulsed LFM and FMCW waveform fundamentals, target radar cross section, maritime clutter and rain effects, the detection processing chain, and the representative UAV mission profiles used to define the problem. Chapter 2 also reviewed relevant literature examples to show how practical naval and UAV-detection radars are parameterized and what physical and operational trade-offs must be carried into a fair comparison.

In Chapter 3, the simulation-based comparison framework was developed under a common set of naval design constraints so that the pulsed and FMCW radar candidates could be assessed on an equal basis. This chapter defined the shared radar parameters, the shipboard search strategy, the dwell and CPI budgeting, the Range-Doppler and CFAR detection processing, and the Monte Carlo confirmation logic used to evaluate performance.

In Chapter 4, the results of the comparison were presented with analytical emphasis on the Attack 40 km case as the main operational decision geometry, supported by the broader scenario set summarized across the appendix cases. These results showed how sea clutter, rain, and start range compress detection and classification performance, and they provided the basis for comparing LFM, FMCW, and Hybrid operation using common figures of merit and a cross-case weighted synthesis.

5.2 Conclusions and Discussion

The results show that FMCW is the preferred radar architecture within the adopted ship-installable constraint set, but the reason is not that it dominates every individual case. Rather, FMCW is preferred because it provides the most consistent performance across the mission envelope and retains the strongest margin in the hardest operational conditions, particularly the low-altitude, clutter-limited Attack 40 km geometry. In the favourable cases, the differences between the waveform candidates remain small. The more important result is that as environmental severity increases, FMCW degrades less severely than the LFM-based modes and therefore preserves more useful warning and classification margin under the conditions that matter most for shipboard counter-UAV response.

The main reason for this behaviour lies in the waveform itself. Within the common comparison framework, both candidates were normalized to the same energy per measurement and per dwell, so the comparison did not favour one waveform by assigning it more energy on target. Even so, FMCW retained the larger waveform-processing term because each chirp occupied the full 0.5 ms repetition interval, giving $BT = 2500$, whereas the pulsed LFM candidate was limited to a total transmit on-time of 50 μ s per repetition interval, giving $BT = 250$. This corresponds to a nominal 10 dB advantage in the waveform-processing term for FMCW. In addition, the pulsed design must divide its available energy into short, staggered transmit segments in order to preserve close-range access, long-range energy on target, and ambiguity control within the same repetition interval, whereas FMCW spreads that same interval energy continuously across the full chirp. As a result, the LFM design must balance several competing waveform constraints simultaneously, while FMCW maintains continuous illumination throughout each dwell. This explains why FMCW produced the more even detection and classification margin across the scenarios examined in this study.

The Hybrid operating mode did not recover that advantage. Although it benefited from an FMCW-based micro-Doppler classification path, its primary detection leg remained tied to the LFM front end. As a result, Hybrid preserved little of the FMCW robustness in the hardest cases while also delaying classification relative to the standalone waveforms. The Hybrid concept therefore remained useful as an architectural test case, but it did not provide enough improvement over the standalone candidates to justify the added operating complexity within the present framework.

This performance advantage, however, is not free. In the present comparison, the FMCW solution assumes a two-panel implementation, whereas the LFM solution assumes a single-panel implementation. FMCW is therefore preferred on performance grounds, but that preference must be weighed against the added hardware burden, cost, integration effort, and Size, Weight and Power (SWaP) associated with the second panel. For a shipboard counter-UAV application, this means that the recommendation in favour of FMCW is conditional on the platform being able to absorb that additional implementation burden. If the available installation volume, topside weight margin, cooling capacity, or integration complexity is tightly constrained, then a pulsed LFM solution may still remain attractive as a simpler, lower-burden alternative, even if it is not the strongest performer in the hardest cases. The design question is therefore not only which waveform performs best in principle, but whether the available shipboard SWaP margin is sufficient to realize the preferred architecture in practice.

5.3 Future Work

The experiments in this thesis were conducted under a common set of assumptions for the pulsed and FMCW radar candidates so that the comparison remained focused on waveform-dependent performance rather than unequal baseline design choices. Nevertheless, several modelling limitations remain. The UAV trajectories were simplified to straight-line inbound motion toward the ship, and although statistical RCS fluctuation was included to account for aspect-dependent variation in target reflectivity, this does not fully reproduce the kinematic and aspect complexity associated with realistic UAV manoeuvres. In addition, target classification was based primarily on micro-Doppler extraction and therefore did not exploit other potentially informative behavioural or kinematic signatures. The present framework also did not explicitly include radar implementation errors or RF front-end non-idealities. That omission is particularly relevant for FMCW, where practical receiver sensitivity, transmitter–receiver leakage, finite isolation, phase noise, and limited dynamic range can elevate the effective noise floor and reduce post-processing SNR.

Future work should therefore extend the framework in four main directions. First, the target and scenario models should be made more realistic by including manoeuvring, loitering, crossing, and evasive UAV trajectories, together with richer time-varying target signatures. This would allow the classification stage to incorporate not only micro-Doppler information, but also motion history, trajectory shape, acceleration patterns, and other behavioural features that may help distinguish UAVs from non-UAV targets. Second, the framework should be expanded to evaluate additional radar architectures, including multiple-input multiple-output (MIMO) radar concepts and a broader set of waveform families and hybrid operating modes [32]. Third, future work should move beyond pure simulation toward prototyping and implementation-oriented validation, including hardware-in-the-loop testing or experimental validation with representative RF front ends, so that the comparison can be tested against practical leakage, isolation, sensitivity, and dynamic-range limitations. Fourth, with those richer motion and signal features available, machine-learning-based classification methods could then be introduced to reinforce the radar’s interpretation of target behaviour and potentially extend reliable classification to greater ranges [33]. In that setting, machine learning would not replace the present framework, but would build on it by exploiting a broader feature set than MD alone.

The present work should therefore be interpreted as a controlled baseline comparison of waveform performance, with future extensions focused on more realistic target dynamics, broader radar architectures, prototype-based validation, hardware non-idealities, and more advanced machine-learning-assisted classification.

5.4 Contributions to the field

This thesis makes three principal contributions to the field of shipboard counter-UAV radar evaluation. First, it develops a common-constraint MATLAB-based simulation framework for the fair comparison of pulsed AESA and FMCW AESA shipborne radar candidates under the same naval design assumptions. In this sense, the contribution is not only the final recommendation, but also the construction of an auditable comparison method in which frequency band, aperture, steering sector, range resolution, search burden, and energy allocation are held on a common basis. Second, the thesis integrates sea clutter, rain attenuation, multipath, and fluctuating UAV radar cross section into a single Monte Carlo evaluation chain, thereby linking waveform choice to operationally meaningful detection, confirmation, and classification outcomes in a maritime environment. Third, the thesis produces decision-oriented figures of merit and an evidence-based recommendation for an AOPV-relevant counter-UAV radar architecture, while also establishing a reusable baseline framework that can extend to future studies.

References

- [1] S. Li, Y. Chai, M. Guo, and Y. Liu, "Research on Detection Method of UAV Based on Micro-Doppler Effect," in *2020 39th Chinese Control Conference (CCC)*, Jul. 2020, pp. 3118–3122. doi: 10.23919/CCC50068.2020.9189414
- [2] M. I. Skolnik, Ed., *Radar Handbook*, 3rd ed. New York: McGraw-Hill, 2008.
- [3] D. Oliver, "UKRAINE'S UNMANNED AIR WAR: UAVs are being used at the centre of air operations conducting ISTAR missions on a daily basis.," *Armada Int.*, no. 3, pp. 26–28, Jun. 2022.
- [4] A. Cheranyov and E. Dukhan, "Methods of Detecting Small Unmanned Aerial Vehicles," in *2021 Ural Symposium on Biomedical Engineering, Radioelectronics and Information Technology (USBREIT)*, May 2021, pp. 0218–0221. doi: 10.1109/USBREIT51232.2021.9455043.
- [5] Fisheries and Oceans Canada, Government of Canada, "National Shipbuilding Strategy and the Canadian Coast Guard." Accessed: Mar. 17, 2025. [Online]. Available: <https://www.ccg-gcc.gc.ca/fleet-flotte/nss-sncn-eng.html>
- [6] Treasury Board of Canada Secretariat, "Directive on the Management of Projects and Programmes." Accessed: Mar. 17, 2025. [Online]. Available: <https://www.tbs-sct.canada.ca/pol/doc-eng.aspx?id=32594>.
- [7] M. C. Budge and S. R. German, *Basic Radar Analysis*, Second Edition. Artech House, 2020
- [8] B. R. Mahafza, *Radar Signal Analysis and Processing using MATLAB*. Boca Raton: CRC Press, 2009.
- [9] R. Bil, "Modern phased array radar systems in Germany," in *Proc. IEEE Int. Symp. Phased Array Syst. Technol.*, 2016, doi: 10.1109/ARRAY.2016.7832544.
- [10] P. Eudeline, "Survey of active electronically scanned antenna in Thales radars," in *Proc. IEEE Int. Symp. Phased Array Syst. Technol.*, 2013, doi: 10.1109/ARRAY.2013.6731793.
- [11] S. Harman, "The performance of a novel three-pulse radar waveform for marine radar systems," in *Proc. European Radar Conf.*, Amsterdam, Netherlands, 2008, pp. 160–163
- [12] W. Hao, G. Chen, Y. Li and Y. Hou, "Analysis of RCS Characteristics in Dynamic Flight of Stealth UAV," *2023 IEEE 3rd International Conference on Electronic Technology, Communication and Information (ICETCI)*, Changchun, China, 2023, pp. 1208-1212.
- [13] Thales, "SMART-S Mk2 3D Surveillance Radar," Accessed Jan 2026: <https://www.thalesdsi.com/wp-content/uploads/2018/12/Smart-S-Data-Sheet.pdf>

- [14] T. Moon, "Beyond FFT: Precision Vibration Tracking with FMCW Radar and Kalman Estimators," 2022 IEEE Radar Conference (RadarConf22), pp. 1–6, May 2024, doi: <https://doi.org/10.1109/radarconf2458775.2024.10548102>
- [15] M. Ezuma, E. Ozdemir, C. Anjinappa, and I. Guvenc, "Micro-UAV detection with a low-grazing-angle millimeter wave radar," in Proc. 2019 IEEE Radio and Wireless Symp. (RWS), Orlando, FL, USA, Jan. 2019, pp. 1–4.
- [16] S. Harman, "A comparison of staring radars with scanning radars for UAV detection: Introducing the Alarm™ staring radar," in Proc. 12th European Radar Conf. (EuRAD), Paris, France, 2015, doi: 10.1109/EuRAD.2015.7346268.
- [17] V. Gregers-Hansen and R. Mital, "An empirical sea clutter model for low grazing angles," in Proc. IEEE Radar Conf., Pasadena, CA, USA, May 2009, pp. 1–5.
- [18] MATLAB "Simulating Radar Returns from Moving Sea Surfaces - MATLAB & Simulink," Mathworks.com, 2022. <https://www.mathworks.com/help/radar/ug/simulating-radar-returns-from-moving-sea-surfaces.html> (accessed Jan. 22, 2026).
- [19] S. Mohanty, C. Singh and V. Tiwari, "Estimation of rain attenuation losses in signal link for microwave frequencies using ITU-R model," in Proc. IEEE Int. Geoscience and Remote Sensing Symp. (IGARSS), Beijing, China, Jul. 2016, pp. 532–535.
- [20] B. Widrow and S. D. Stearns, Adaptive Signal Processing. Prentice Hall, 1985.
- [21] S. Li, Y. Chai, M. Guo, and Y. Liu, "Research on Detection Method of UAV Based on Micro-Doppler Effect," in 2020 39th Chinese Control Conference (CCC), Jul. 2020, pp. 3118–3122. doi: 10.23919/CCC50068.2020.9189414.
- [22] D. Hambling, "Moving targets: Implications of the Russo-Ukrainian war for drone terrorism," CTC Sentinel, Jul. 2025, pp. 1–8.
- [23] T. Ryan, "Modern combat: Canadian forces confronts the drone threat both at sea and on land," Canadian Forces (magazine article, PDF provided), vol. 32, no. 9, n.d.
- [24] DrDrone.ca, "DJI Phantom 4 Pro Specification," DrDrone.ca, 2022. <https://drdrone.ca/pages/phantom-4-pro-technical-specifications> (accessed Jan. 22, 2026)
- [25] Y. Ma, J. Dai, H. Jing, and D. Jing, "On the radar detection analysis under the interference environment based on the Monte Carlo method," in Proc. 11th Int. Symp. on Antennas, Propagation and EM Theory (ISAPE), Guilin, China, Oct. 2016, pp. 492–495, doi: 10.1109/ISAPE.2016.7834022.
- [26] International Telecommunication Union, "Attenuation by atmospheric gases," Recommendation ITU-R P.676-11, Sep. 2016
- [27] Government of Canada "Halifax class," www.canada.ca, Apr. 19, 2013. <https://www.canada.ca/en/navy/corporate/fleet-units/surface/halifax-class.html>

- [28] P. M. Djuric and Y. Huang, "Estimation of a Bernoulli parameter p from imperfect trials," *IEEE Signal Processing Letters*, vol. 7, no. 6, pp. 160–163, Jun. 2000, doi: 10.1109/97.844638.
- [29] R. Guay, "The dynamic measurement of unmanned aerial vehicle radar cross section," M.A.Sc. thesis, Royal Military College of Canada, Kingston, ON, Canada, Apr. 2016.
- [30] X. Fang and G. Xiao, "Rotor Blades Micro-Doppler Feature Analysis and Extraction of Small Unmanned Rotorcraft," in *IEEE Sensors Journal*, vol. 21, no. 3, pp. 3592-3601, Feb.1, 2021
- [31] Government of Canada "Halifax class Ship Standing Orders" Apr. 20, 2020. [https://collaboration-navy.forces.mil.ca/sites/SeaTrainingGroup/Shared%20Documents/Published%20References%20\(Public\)/SSOs_Fleet%20SOPs/SSOs%20AL%2011/SSO%20AL11%20%2025%20April%202020.pdf#search=Halifax%20Class%20Ship%20Standing%20orders](https://collaboration-navy.forces.mil.ca/sites/SeaTrainingGroup/Shared%20Documents/Published%20References%20(Public)/SSOs_Fleet%20SOPs/SSOs%20AL%2011/SSO%20AL11%20%2025%20April%202020.pdf#search=Halifax%20Class%20Ship%20Standing%20orders) (Only on DWAN)
- [32] J. Klare, O. Biallawons and D. Cerutti-Maori, "UAV detection with MIMO radar," *2017 18th International Radar Symposium (IRS)*, Prague, Czech Republic, 2017, pp. 1-8
- [33] W. Nie *et al.*, "UAV Detection and Localization Based on Multi-Dimensional Signal Features," in *IEEE Sensors Journal*, vol. 22, no. 6, pp. 5150-5162, March 15, 2022

Appendix A Supplementary Scenario Results

Appendix A provides the remaining scenario results that are not discussed in detail in the main body of Chapter 4, including the Attack 10 km, Surveillance, and Intelligence cases under Best, Medium, and Worst maritime conditions. These results are retained for completeness and to show that the principal interpretation developed from the Attack 40 km case is representative of the broader scenario set. The results are presented in form of tables.

TABLE A-1 ATTACK BEST-CASE WEATHER (SEA STATE 1, 0 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	00:02.87	9.96	00:05.04	9.93	00:06.11	9.92
10 km	FMCW	00:02.87	9.96	00:05.07	9.93	00:06.21	9.92
10 km	Hybrid	00:02.87	9.96	00:07.12	9.90	00:09.12	9.88

TABLE A-2 ATTACK MEDIUM-CASE WEATHER (SEA STATE 3, 10 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	00:05.34	9.93	00:08.02	9.89	00:15.49	9.79
10 km	FMCW	00:05.61	9.92	00:08.33	9.89	00:15.69	9.79
10 km	Hybrid	00:05.34	9.93	00:11.25	9.85	00:18.54	9.75

TABLE A-3 ATTACK WORST-CASE WEATHER (SEA STATE 5, 50 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	07:59.37	3.53	08:01.33	3.50	09:06.32	2.62
10 km	FMCW	06:30.64	4.73	06:32.63	4.70	07:28.66	3.94
10 km	Hybrid	07:58.40	3.54	08:02.19	3.49	09:07.75	2.61

TABLE A-4 SURVEILLANCE BEST-CASE WEATHER (SEA STATE 1, 0 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	00:02.91	9.97	00:05.63	9.94	00:07.51	9.92
10 km	FMCW	00:02.91	9.97	00:05.59	9.94	00:07.52	9.92
10 km	Hybrid	00:02.91	9.97	00:07.86	9.92	00:09.76	9.90
40 km	LFM	00:08.79	39.91	07:04.72	39.29	07:32.95	39.24
40 km	FMCW	00:08.58	39.91	07:05.63	39.29	07:32.74	39.24
40 km	Hybrid	00:08.81	39.91	10:02.07	39.03	10:22.74	39.00

TABLE A-5 SURVEILLANCE MEDIUM-CASE WEATHER (SEA STATE 3, 10 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	00:03.78	9.96	00:06.82	9.93	00:08.92	9.91
10 km	FMCW	00:03.88	9.96	00:06.74	9.93	00:08.84	9.91
10 km	Hybrid	00:03.78	9.96	00:09.39	9.91	00:11.53	9.88
40 km	LFM	00:57.09	39.43	11:59.79	35.68	13:55.77	35.03
40 km	FMCW	00:58.29	39.42	12:01.69	35.67	13:55.74	35.03
40 km	Hybrid	00:56.60	39.43	14:25.38	34.86	16:10.71	34.27

TABLE A-6 SURVEILLANCE WORST-CASE WEATHER (SEA STATE 5, 50 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	00:31.23	9.69	01:19.44	9.21	01:36.46	9.04
10 km	FMCW	00:28.14	9.72	01:06.68	9.33	01:22.30	9.17
10 km	Hybrid	00:31.23	9.69	01:46.51	8.94	02:02.28	8.78
40 km	LFM	12:00.80	35.67	13:21.47	35.22	14:30.10	34.84
40 km	FMCW	11:58.52	35.68	13:21.17	35.23	14:30.58	34.83
40 km	Hybrid	12:02.55	35.65	15:56.17	34.36	17:35.41	33.80

TABLE A-7 INTELLIGENCE BEST-CASE WEATHER (SEA STATE 1, 0 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	00:02.85	9.99	00:05.62	9.98	00:07.74	9.98
10 km	FMCW	00:02.85	9.99	00:05.58	9.98	00:07.73	9.98
10 km	Hybrid	00:02.85	9.99	00:08.24	9.98	00:10.42	9.87
40 km	LFM	00:04.95	39.99	08:39.69	38.44	12:11.03	37.81
40 km	FMCW	00:04.67	39.99	08:43.28	38.43	12:13.43	37.80
40 km	Hybrid	00:04.87	39.99	11:41.25	37.90	14:57.85	37.31

TABLE A-8 INTELLIGENCE MEDIUM-CASE WEATHER (SEA STATE 3, 10 MM/H RAIN)

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	00:02.90	9.99	00:05.67	9.98	00:07.96	9.98
10 km	FMCW	00:02.90	9.99	00:05.59	9.98	00:07.92	9.98
10 km	Hybrid	00:02.90	9.99	00:08.81	9.98	00:11.40	9.97
40 km	LFM	69:52.26	27.43	153:06.31	12.44	156:44.68	11.78
40 km	FMCW	69:56.26	27.41	152:47.37	12.49	156:30.62	11.82
40 km	Hybrid	69:53.50	27.43	175:56.74	8.33	179:13.54	7.80

TABLE A-9 INTELLIGENCE WORST-CASE WEATHER (SEA STATE 5, 50 MM/H RAIN) — SUMMARY

Starting Range	Waveform	P_d 50 time (MM:SS:ss)	P_d 50 range (km)	P_c 50 time (MM:SS:ss)	P_c 50 range (km)	P_c 80 time (MM:SS:ss)	P_c 80 range (km)
10 km	LFM	00:05.61	9.98	00:09.61	9.97	00:17.19	9.95
10 km	FMCW	00:05.61	9.98	00:09.56	9.97	00:17.05	9.95
10 km	Hybrid	00:05.61	9.98	00:14.24	9.96	00:23.90	9.93
40 km	LFM	150:42.94	12.87	150:52.89	12.84	153:35.39	12.35
40 km	FMCW	150:41.59	12.88	150:53.26	12.84	153:31.70	12.36
40 km	Hybrid	150:43.97	12.87	151:28.13	12.74	153:49.55	12.31

Overall, the appendix results support the main conclusion of the thesis. Because the Surveillance and Intelligence scenarios place the UAV at higher altitude, the targets are less strongly affected by low-grazing sea clutter and the resulting environmental penalty on P_d and P_c is reduced. This allows detection and classification to occur earlier and at greater range than in the Attack geometry. The supplementary cases therefore confirm that the principal waveform separation is most important in the hardest low-altitude case, while the broader scenario set remains consistent with the recommendation developed in the main body.