

The image shows a fleet of drones flying over a dark, choppy sea at dusk or dawn. In the foreground, a person in an orange life vest is floating in the water. Several drones are positioned around the person, with some casting beams of light downwards. The sky is a deep, dark blue, and the overall scene conveys a sense of a coordinated search and rescue operation.

Energy-Aware Multi-UAV Coordination using Informative Path Planning for Maritime Search & Rescue

Exploiting Flight-Regime Energy Dynamics

Shivesh Damian Bhawan

Energy-Aware Multi-UAV Coordination using Informative Path Planning for Maritime Search & Rescue

Exploiting Flight-Regime Energy Dynamics

by

Shivesh Damian Bhawan

Thesis committee:

Chair:	Dr. Ir. Ewoud J.J. Smeur
Supervisors:	Dr. eng. Marija Popovic Ir. Bart D.W. Remes
External examiner:	Dr. Marilena D. Pavel
Place:	Faculty of Aerospace Engineering, Delft
Project Duration:	June, 2025 - April, 2026
Student number:	5301602

Cover: Generated using ChatGPT, personal prompt

Preface

This thesis represents the final chapter of my studies at Delft University of Technology. Beyond the academic challenges, my time here has been defined by the friendships I've made—bonds that I know will last a lifetime. If my time at Delft has taught me anything, it's that no challenge has a single solution. I've learned that when things inevitably 'go south,' the most powerful tool at your disposal isn't just technical knowledge, but the ability to maintain a level head and dismantle a complex failure into a series of systematic solutions.

Technically, this paper explores how a swarm of autonomous drones can coordinate to explore vast areas. The method I present allows these agents to work together seamlessly; if one drone identifies a potential target—in this context, persons lost in the ocean—it will autonomously lower its altitude to track and confirm the detection with higher certainty. This behavior of precision descent not only improves search reliability but also creates a framework for future practical applications, such as dropping life-saving payloads or medical supplies directly to those in need.

I would like to express my deepest gratitude to my supervisors, Marija Popovic and Bart Remes, for providing invaluable insights week in and week out. Your guidance was instrumental in establishing the theoretical baselines of this work.

A special thank you goes to Erik van der Horst, who served not only as the pilot in command during my field tests but also as the bridge that turned my software concepts into physical reality. I also want to thank Noach Wechtler and Taran Jaswant Singh for their constant support, whether it was helping me squash small bugs or assisting with the logistics of my experiments.

Finally, I want to thank my parents. Your unwavering support has been my foundation throughout this entire journey.

*Shivesh Damian Bhawan
Delft, April 2026*

Contents

Preface	i
Nomenclature	iv
I Preliminary	v
II Scientific Article	viii
1 Introduction	1
2 Related Work	2
2.1 Hybrid UAV Platforms	2
2.2 Multi-UAV Coordination	3
2.3 Energy-Aware Informative Path Planning	3
2.4 Probabilistic Mapping in Dynamic Environments	3
2.5 Summary	4
3 Preliminaries	4
3.1 System Assumptions	4
3.2 Environment and Belief Representation	5
3.3 Synthetic Sensing Model	5
3.4 Platform Models: VSQP and Quadrotor	6
3.4.1 The Variable Skew Quad Plane (VSQP)	6
3.4.2 The Parrot Bebop 2	7
3.4.3 Kinematic Constraints and Wind Compensation	8
3.5 Problem Formulation	8
4 Methodology	9
4.1 System Architecture	9
4.2 Dynamic Environment Modeling	10
4.2.1 Drift Dynamics	10
4.2.2 Conical Diffusion Model	10
4.3 Single Agent Planning: Energy-Aware IPP	11
4.3.1 Motion Primitives	11
4.3.2 Cost-Benefit Evaluation	11
4.3.3 Trajectory Selection	12
4.4 Swarm Coordination: Task Allocation	12
4.4.1 Exploration Mode	13
4.4.2 Task Assignment & Approach	13
4.4.3 Tracking Mode	13
4.4.4 Confirmation	13
4.5 Mission Overview	14
4.6 Summary	14
5 Experiments & Results	15
5.1 Phase I: Algorithmic Validation and Scalability Analysis	16
5.1.1 Simulation Setup	16
5.1.2 Swarm Scalability	17
5.1.3 Environmental Robustness	19
5.1.4 Swarm Sizing and Strategy Comparison	21
5.2 Phase II: Baseline Comparison and Sensitivity Analysis	21
5.3 Phase III: System Integration	24
5.3.1 Network Topology and Middleware Architecture	24
5.3.2 Closed-Loop Control and Kinematic Management	25
5.4 Phase IV: In-field Testing	25

5.4.1	Indoor Verification (TU Delft CyberZoo)	25
5.4.2	Field Validation (Valkenburg Naval Air Base)	27
6	Conclusion	28
7	Recommendations for Future Work	29
7.1	Integrating Real Sensors and Edge Computing	29
7.2	Platform Capabilities and Low-Level Control	29
7.3	Transitioning Towards a Decentralized Swarm Architecture	29
7.4	Hybrid Coarse-to-Fine Search Architecture	30
	References	30
III	Closure	33
A	Declaration of Generative AI Usage	36
B	Test Set-Up and Execution Protocol	36
B.1	Preliminary Hardware and Software Inventory	36
B.2	Network and Infrastructure Configuration	36
B.2.1	Indoor Configuration (TU Delft CyberZoo)	36
B.2.2	Outdoor Configuration (EHVB Field Tests)	38
B.3	Flight Execution Sequence	39
C	Open-Source Contributions to Paparazzi UAV	39
C.1	Repository Information	40
C.2	Directory Structure and Custom Modules	40
C.2.1	Paparazzi Configurations (The SDB Directory)	40
C.2.2	The Centralized Planner (sw/ground_segment/python/Occupancy_Map/)	40
C.3	Execution Guide	40
C.3.1	Build the Paparazzi Environment	41
C.3.2	Launch Paparazzi Center and Select Session	41
C.3.3	Initialize the Swarm	41
C.3.4	Execute the Python Planner	41

Nomenclature

Abbreviation	Definition
CNN	Convolutional Neural Network
ENU	East-North-Up (Coordinate Frame)
EO/IR	Electro-Optical/Infrared
FOV	Field Of View
FSM	Finite State Machine
GCS	Ground Control Station
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
HIL	Hardware-in-the-Loop
IMU	Inertial Measurement Unit
INDI	Incremental Nonlinear Dynamic Inversion
INS	Inertial Navigation System
IP	Internet Protocol
IPP	Informative Path Planning
MCTS	Monte Carlo Tree Search
NED	North-East-Down (Coordinate Frame)
PID	Proportional-Integral-Derivative
PIW	Person in Water
RF	Radio Frequency
RHP	Receding Horizon Planner
RRT	Rapidly-exploring Random Tree
RTK	Real-Time Kinematic
SAR	Search and Rescue
SSID	Service Set Identifier
UAV	Unmanned Aerial Vehicle
UDP	User Datagram Protocol
VSQP	Variable Skew Quad Plane
VTOL	Vertical Take-Off and Landing

Part I

Preliminary

Introduction

In recent years, autonomous aerial systems have advanced rapidly, enabling complex missions in challenging environments. One particularly demanding domain is Search and Rescue (SAR) in maritime environments, where timely victim detection and wide-area coverage are critical. The ocean's vastness and unpredictability make such operations extremely difficult for human crews to perform effectively.

To improve SAR performance, the deployment of autonomous UAV swarms capable of collaboratively detecting drifting victims is highly promising. However, a fundamental bottleneck in multi-UAV maritime operations is the strict endurance limit of battery-powered aircraft. Traditional multirotors excel at low-speed, precise inspection but lack the efficiency for wide-area coverage. Conversely, fixed-wing platforms offer long-range endurance but cannot hover to confirm targets. Hybrid UAVs, such as the Variable Skew Quad Plane (VSQP), bridge this gap by combining Vertical Take-Off and Landing (VTOL) capabilities with the energy-efficient forward flight of a fixed-wing aircraft.

Effectively exploiting these hybrid platforms requires advanced decision-making frameworks. Central to modern autonomous search is Informative Path Planning (IPP), which enables drones to select trajectories that maximize expected information gain rather than relying on rigid, geometric sweep patterns. While IPP has been widely studied, existing approaches largely neglect the non-linear energy dynamics of flight. They rarely penalize energy-intensive hover states during exploration, nor do they explicitly leverage the aerodynamic efficiency of fixed-wing cruise. Furthermore, current IPP frameworks are rarely integrated with swarm coordination and dynamic drift modeling, limiting their applicability to the stochastic realities of ocean search.

Problem Statement

The core problem addressed in this thesis is the efficient coordination of a multi-UAV swarm for maritime SAR under strict energy and kinematic constraints. To maximize the probability of finding a drifting victim before battery depletion, the swarm must continuously balance the conflicting requirements of rapid spatial coverage (exploration) and precise target confirmation (tracking).

Currently, there is a lack of high-level coordination architectures that can dynamically transition a swarm between these phases while explicitly minimizing the aerodynamic penalties of flight. The system must be capable of processing asynchronous telemetry, propagating stochastic ocean drift, and autonomously allocating tasks to prevent redundant searching and inter-agent collisions, all while mathematically biasing trajectories toward the most efficient flight regimes of the hardware.

Research Objective

To develop and deploy an energy-aware swarming algorithm that integrates informative path planning for UAV-based maritime search and rescue missions, with the goal of maximizing victim detection and area coverage efficiency while minimizing total energy consumption in large, dynamic, and partially observable environments.

Research Questions

To achieve the research objective, this thesis addresses the following primary research question:

How can informative path planning be unified with multi-UAV coordination to enable energy-efficient maritime SAR by exploiting distinct flight-regime dynamics?

To systematically answer the primary research question, the following sub-questions are investigated:

1. *What is the minimum swarm size required to explore a bounded maritime environment and achieve a high victim discovery rate within strict endurance and temporal limits?*
This question evaluates the scalability of the proposed framework. It is addressed through extensive Monte Carlo simulations to establish the relationship between operational search area, fleet density, and physical battery limits, ultimately defining the operational constraints of the swarm.
2. *How does embedding an aerodynamic power model into the objective function balance the trade-off between global entropy reduction and endurance?*
This question investigates the core algorithmic contribution of the energy-aware planner. It is answered by benchmarking the proposed Informative Path Planning (IPP) algorithm against traditional Boustrophedon coverage under varying asymmetric wind disturbances.

3. *How can the distinct flight-regime energy dynamics of a hybrid UAV (such as the VSQP) be explicitly exploited during different mission phases to maximize search endurance?*

This question bridges theoretical path planning with physical platform behavior. It is addressed by developing a Finite State Machine (FSM) that dynamically transitions the swarm between high-altitude global exploration and targeted, low-altitude confirmation, and is subsequently validated through Hardware-in-the-Loop (HIL) and real-world multi-UAV flight tests.

Part II

Scientific Article

Energy-Aware Multi-UAV Coordination using Informative Path Planning for Maritime Search & Rescue

Exploiting Flight-Regime Energy Dynamics

Shivesh Damian Bhawan

Autonomous maritime search and rescue faces critical challenges due to limited endurance and dynamic victim drift. This thesis presents a centralized multi-UAV coordination framework based on Energy-Aware Informative Path Planning. The aim is to maximize mission endurance by exploiting the distinct energy dynamics of different flight regimes, such as the efficiency of fixed-wing cruise versus the high cost of hovering. Locally, a Receding Horizon Planner optimizes the trade-off between entropy reduction and energy consumption. The framework provides global coordination by dynamically switching agents between high-altitude exploration and low-altitude target tracking while accounting for the effects of ocean drift. Monte Carlo simulations using the empirical power model of the Variable Skew Quad Plane evaluate the swarm size required to achieve 80% global entropy reduction and successfully confirm multiple victims within a strict 10-minute timespan. The evaluation reveals a fundamental trade-off between spatial clearing speed and environmental robustness—defined as the system’s ability to maintain consistent search performance despite adverse wind disturbances. Although boustrophedon coverage is efficient in optimal conditions, adverse crosswinds cause accelerated energy drain, potentially leading to early mission failure. By dynamically surfing wind gradients, the proposed method extends mission endurance while maintaining sufficient coverage capabilities in these severe crosswind scenarios. Finally, the communication and coordination logic of the framework was validated through real-world flight tests with a swarm of Parrot Bebop 2s.

Keywords: Energy-Aware, Informative Path Planning, Swarming, Dynamic Mapping, Centralized Coordination, Search & Rescue, UAV swarm, Autonomous

1 Introduction

Autonomous multi-UAV systems for search and mapping in challenging environments are gaining increasing attention due to their potential to support applications such as search and rescue (SAR) [1, 2], maritime

monitoring [3, 4], disaster response [5], and large-scale exploration [6–8]. Swarm-based aerial systems offer improved scalability, robustness, and temporal efficiency compared to single-vehicle approaches, making them promising for time-critical SAR missions.

This work focuses on maritime SAR, where vast open-water areas and uncertain victim locations make efficient search a significant challenge. It is a fundamental aerodynamic property that standard rotorcraft consume less energy in moderate forward flight than in stationary hover, as translational lift reduces induced power requirements [9–11]. Hybrid UAV platforms—such as tilt-rotors, tail-sitters—extrapolate this efficiency advantage further. By transitioning from rotary-wing flight to wing-borne cruise, they achieve endurance capabilities far beyond the "sweet spot" of standard multirotor forward flight [12, 13].

Effectively exploiting these distinct regimes requires planning methods that account for their differing energy profiles and transition costs. Our motivation is to develop a framework that leverages the theoretical efficiency of advanced hybrid platforms while remaining adaptable to hardware constraints. To achieve this, this work utilizes the VSQP’s distinct energy dynamics as a representative baseline for large-scale simulations. For physical field trials, standard quadrotors (Parrot Bebop 2) serve as kinematic surrogates to execute regime-optimized trajectories in real-world environments. While prior VSQP studies focus primarily on low-level flight control and transition stability [14, 15], this research addresses the open challenge of high-level planning and cooperative decision-making to fully exploit efficient forward-flight and hover regimes in information-driven search tasks.

Informative Path Planning (IPP) provides a principled approach for guiding UAVs toward areas of high expected information gain [16]. In this context, "information" represents a reduction in positional uncertainty regarding the location of drifting victims. The primary focus of this work is validating the coordination and energy-management algorithms. To enable a systematic evaluation, we employ a modular occupancy-based mapping approach [17, 18].

Centralized coordination strategies have demonstrated strong performance for multi-UAV coverage

and task allocation [19–21]. However, an open challenge remains in combining these strategies with energy-aware IPP that explicitly maximizes efficient fixed-wing cruise for wide-area exploration while strategically reserving hover for precise victim confirmation. Prior work on energy-conscious UAV planning typically focuses on single vehicles [22, 23] or omits hybrid-mode considerations altogether [24].

We contribute to these fields by answering the following research question:

How can informative path planning be unified with multi-UAV coordination to enable energy-efficient maritime SAR by exploiting distinct flight-regime dynamics?

This paper focuses on exploiting flight-regime energy dynamics to optimize the trade-off between rapid information gain and energy consumption during time-critical, swarm-based maritime search. Unlike methods designed for single-mode platforms, our approach relies on mode-dependent behavior—penalizing energy-intensive hover during exploration to prioritize rapid, efficient forward flight, while utilizing holonomic capabilities solely for precise victim inspection. We propose a centralized framework that combines: (i) IPP, (ii) occupancy-based belief mapping, and (iii) coordinated multi-UAV planning that accounts for these energetic trade-offs. To assess this framework, we utilize the hybrid VSQP in large-scale simulations and deploy standard quadrotors in field trials to validate the underlying communication architecture and kinematic execution.

Contribution Statement— This work makes the following contributions:

- **Energy-regime-biased IPP:** We develop a novel path-selection method that mathematically biases UAV behavior toward aerodynamically efficient flight regimes, such as fixed-wing cruise, while penalizing energy-intensive hover states during exploration. We demonstrate the endurance benefits of this approach using the hybrid VSQP as a representative case study in simulation.
- **Integrated multi-UAV coordination:** We introduce a centralized swarm-planning framework that simultaneously reasons about energy usage, belief uncertainty, and inter-UAV coordination.
- **Search efficiency and energy evaluation:** Through large-scale simulation and in-field testing, we demonstrate how flight-regime utilization and IPP jointly optimize the critical tradeoff between rapid information gain and energy consumption.

This report is structured as follows: Chapter 2 reviews relevant literature on UAV SAR, hybrid platforms, dynamic mapping, swarm coordination, and IPP. Chapter 3 defines the sensing assumptions, energy models, and problem formulation. Chapter 4 presents the proposed planning and coordination framework. Chapter 5 describes simulation and in-field experiments, results and discussions. Chapter 6 concludes and outlines future work.

2 Related Work

Research on UAV-based maritime SAR intersects four main domains: hybrid aerial platforms, multi-UAV coordination, informative path planning, and mapping under dynamic and uncertain environments. This section summarizes the most relevant developments in these areas and highlights the gaps addressed in this work.

2.1 Hybrid UAV Platforms

Hybrid UAVs combine the vertical take-off, landing, and hovering capabilities of multirotors with the efficient forward flight of fixed-wing aircraft. Surveys of hybrid and convertible VTOL systems [12, 13] identify common configurations such as tilt-rotor, tilt-wing, tail-sitter, and quadplane architectures. Despite their mechanical differences, these platforms share the same operational principle: a hover mode for precise inspection and low-speed maneuvering, and a fixed-wing mode for long-distance, energy-efficient transit.

The Variable Skew Quad Plane (VSQP) is one such hybrid platform, featuring an oblique rotating wing inspired by the oblique flying wing concept [25]. Prior VSQP work has focused primarily on incremental non-linear dynamic inversion (INDI) based flight control and transition stability [14, 15, 26]. Similar trends hold across the broader hybrid UAV literature, where research is dominated by airframe design and low-level control rather than high-level navigation or mission-planning frameworks for hybrid-mode platforms.

Across hybrid UAV platforms in general, there is a relative lack of high-level navigation frameworks, energy-aware informative planning, and multi-UAV coordination methods that explicitly leverage hybrid-mode capabilities. Most existing coordination frameworks do not account for the non-linear energy trade-offs between efficient forward flight and hover when optimizing informative search paths. As a result, the potential for hybrid aircraft to dynamically transition between these flight regimes to maximize mission endurance remains a significant open challenge in autonomous search and rescue research.

2.2. Multi-UAV Coordination

Swarm coordination strategies generally fall into centralized, decentralized, and behavior-based categories. Centralized architectures assign tasks via a global controller, achieving high coverage uniformity and coordinated exploration [19–21]. These approaches are well suited for maritime SAR, where environments are largely obstacle-free and sector-based allocation is effective, though communication intermittency remains a practical challenge. Master–follower strategies [2] provide a middle ground, reducing communication load while maintaining coherent formation control.

Decentralized approaches emphasize scalability and robustness under intermittent connectivity. Role-based systems [1], Lévy-walk-inspired coordination [7], and distributed consensus behaviors support adaptive exploration, though coverage regularity may suffer. Behavior-based flocking [3, 6] yields cohesive swarm motion but lacks guarantees for systematic SAR coverage.

These decentralized methods are less aligned with this study’s primary objective of maximizing search efficiency under strict battery limits. Decentralized logic often necessitates energy-intensive onboard processing for local map maintenance and frequent inter-agent communication, which directly competes with the aircraft’s propulsion for limited electrical power. Consequently, this work adopts a centralized coordination framework to deliberately offload computationally intensive tasks—such as dynamic drift propagation and global Bayesian sensor fusion—from the energy-constrained UAVs. By prioritizing a link to a Ground Control Station (GCS), the system can leverage centralized task allocation to ensure rigorous spatial deconfliction and optimal coverage without the energetic overhead required for complex inter-agent consensus.

2.3. Energy-Aware Informative Path Planning

Informative Path Planning (IPP) selects trajectories that maximize expected information gain, a task that is inherently complex in maritime SAR due to the vast, featureless search areas and time-sensitive nature of the mission. Foundational approaches, such as submodular optimization [27] and sampling-based planners including RIG-tree [28], RRBT [29], Informed RRT* [30], and BIT* [31], are particularly relevant to this work. These methods provide the algorithmic basis for efficiently searching high-dimensional state spaces where information rewards exhibit diminishing returns—a critical property when determining how much time a swarm should spend in a specific sector before moving on.

A notable single-UAV IPP framework is designed by Popovic et al. [16], which couples sampling-based planning with mutual information metrics and Gaus-

sian process mapping. While effective for informative coverage, it relies on a multirotor platform, static GP fields, and vision-based sensing, and thus does not address multi-UAV coordination, hybrid flight modes, or energy-aware behavior required in maritime SAR.

To handle the larger scales and evolving uncertainties typical of maritime drift, advanced IPP methods have moved toward non-myopic planning. TIGRIS, for example, optimizes long-range fixed-wing paths in large search spaces [32], while POMDP-based approaches using Monte Carlo Tree Search [33] provide a framework for decision-making under high uncertainty. These planners are specifically relevant to SAR because they address the "curse of dimensionality" inherent in sprawling ocean environments and provide the non-myopic foresight required to balance immediate information gain with the long-term goal of returning to base before battery depletion.

2.4. Probabilistic Mapping in Dynamic Environments

Effective SAR operations rely on probabilistic representations of survivor locations to guide search efforts. Occupancy grid mapping [17, 18] provides a computationally efficient foundation for this task, with multi-resolution extensions such as probabilistic quadrees [34] addressing scalability in large areas. In the maritime domain, however, static representations are insufficient due to wind and current forces. Consequently, standard occupancy grids must be augmented with drift dynamics to maintain accurate belief states over time.

This temporal evolution is critical for planning; drift causes the probability distribution of a victim to shift and expand, leading to "information decay" in previously searched regions. If the planner treats the environment as static, it may fail to account for victims drifting from unobserved sectors into "cleared" areas, resulting in sub-optimal search allocation. By integrating drift dynamics, the planning framework can strategically schedule revisits to high-probability zones and adjust sweep patterns to intercept the evolving probability density.

Belief propagation models, such as the Leeway Field Method [35], are the standard for simulating these dynamics, using stochastic ensembles to model survivor drift. While distributed mapping frameworks exist to support cooperative exploration [36, 37], these approaches rarely integrate dynamic ocean evolution with the specific flight-regime energy constraints of UAVs. This work addresses that gap by incorporating dynamic belief updates into a unified, energy-aware planning framework.

2.5. Summary

The literature provides robust foundations for hybrid platforms, swarm coordination, and informative path planning; however, these elements are typically treated in isolation. As illustrated in Table 2.1, current frameworks lack an integrated approach that simultaneously reasons about multi-UAV coordination, mode-dependent energy dynamics, and the temporal uncertainty of maritime drift. This research addresses this fragmented landscape by developing an IPP framework that moves beyond standard spatial coverage to explicitly prioritize aerodynamically efficient flight regimes. This positioning directly motivates the design choices presented in Chapter 4, specifically the selection of a centralized architecture. This choice is necessitated by the need to offload the computational burden of high-fidelity drift-aware mapping while enforcing energy-optimal behavior across the swarm, ensuring that mission endurance is maximized without compromising individual agent processing limits.

Table 2.1: Comparison of SAR and non-SAR (multi-)UAV path planning methods relevant to this study.

Method	Multi-UAV	Centralised	Scalable	Energy-aware	Hybrid platform	Dynamic occup. map	Real-world tested
This work	✓	✓	✓	✓	✓	✓*	×
Popovic et al. [16]	×	/	×	×	×	✓	×
Cho et al. [21]	✓	✓	✓	×	×	×	×
Yan et al. [20]	✓	✓	✓	×	×	×	×
Horyna, Baca, Walter et al. [3]	✓	×	✓	×	×	×	✓
Sojo et al. [7]	✓	×	✓	×	×	×	✓
Datsko et al. [24]	✓	✓	×	✓	×	×	✓

* Synthetic dynamic occupancy field with time-varying victim probabilities.

3 Preliminaries

This chapter establishes the foundational models and mathematical formulation for the energy-aware multi-UAV search problem. It defines the operational environment, the probabilistic sensing mechanism, and the specific kinematic and energetic constraints of the hybrid VSQP platform and the Parrot Bebop 2. These elements lead to the formulation of the formal problem statement, which defines the optimization objective and constraints addressed by the methodology in Chapter 4.

3.1. System Assumptions

The following assumptions define the operational scope and limitations of this research. They establish the technical boundary between the proposed

high-level coordination framework and the underlying hardware, perception, and control architectures.

1. **Environment:** The mission is executed in a bounded two-dimensional search region with no static obstacles. Although the physical environment is static, survivor locations evolve over time according to a drift process that emulates wind- and current-driven motion.
2. **State Estimation:** Each UAV has access to reliable position and attitude estimates obtained from GPS and onboard IMU sensors. State information is assumed to be sufficiently accurate for planning at the spatial scales considered.
3. **Communication Model:** All UAVs maintain communication with a centralized Ground Control Station (GCS), which aggregates their state information and maintains a global probabilistic belief map. Inter-UAV communication occurs implicitly through the GCS.
4. **Mapping Representation:** The search region is discretized into a 2D occupancy-based belief grid. The belief map evolves due to drift dynamics, sensor updates, and Bayesian filtering. While UAVs dynamically change altitude during victim confirmation phases, the spatial discretization of the search environment remains strictly two-dimensional.
5. **Sensing Model:** To isolate the planning and coordination algorithms from the complexities of computer vision, victim detection is modeled using a synthetic proximity-based sensor. A detection is triggered probabilistically when a UAV enters a defined sensing radius around a victim location. This abstraction ensures that the framework remains modular and compatible with future visual or multimodal detectors.
6. **Energy-Aware Optimization:** The planner evaluates motion primitives by optimizing an objective function that incorporates a continuous, velocity-dependent energy model. By seeking to minimize total energy consumption, the mathematical optimization naturally biases the generated velocity commands toward the hardware’s most efficient speeds without the planner requiring explicit awareness of discrete flight modes.
7. **Low-Level Flight Control:** For both the VSQP and the Bebop 2, low-level stabilization and attitude regulation are handled by onboard Incremental Non-linear Dynamic Inversion (INDI) flight controllers. This allows the high-level planner to treat the aircraft as a kinematic executor that can reliably track commanded velocity vectors.

3.2. Environment and Belief Representation

The maritime search region is modeled as a bounded two-dimensional domain $\mathcal{H} \subset \mathbb{R}^2$. To enable computational tractability, we discretize \mathcal{H} into a regular Cartesian grid $\mathcal{C} = \{c_1, \dots, c_{|\mathcal{C}|}\}$ with a uniform resolution parameter Δs , where each cell $c \in \mathcal{C}$ represents a square area of Δs^2 . Each cell stores an occupancy probability

$$p_t(c) \in [0, 1], \quad (3.1)$$

representing the likelihood that a victim occupies the corresponding location at time t . The complete belief map is denoted as the set of all cell probabilities:

$$M_t = \{p_t(c) \mid c \in \mathcal{C}\}. \quad (3.2)$$

To bridge this theoretical formulation with our experimental implementation, we define a specific operational domain for the search task. The mission area corresponds to the EHVb test site, whose polygonal boundary is integrated as a Paparazzi flight plan. Figure 3.1 shows the satellite imagery with the boundary waypoints used to define the search domain.



Figure 3.1: Real-world test area at Valkenburg (EHVB), showing the soft geofence boundary (S1–S9). This polygon defines the searchable maritime region \mathcal{A} .

With no prior information about the victim location, the belief is initialized uniformly [18],

$$p_0(c) = 0.5, \quad \forall c \in \mathcal{C}, \quad (3.3)$$

reflecting maximum initial uncertainty. The initial grid structure is shown in Figure 3.2.



Figure 3.2: Initial occupancy-based belief map M_0 . All grid cells begin with uniform prior probability $p_0(c) = 0.5$.

As the mission progresses, the belief evolves through two mechanisms: (i) drift propagation caused by ocean-like motion of victims, and (ii) Bayesian updates from UAV observations. Formally, the recursive update is:

$$M_{t+\Delta t}^{\text{pred}} = \mathcal{D}(M_t), \quad M_{t+\Delta t} = \mathcal{U}(M_{t+\Delta t}^{\text{pred}}, z_t). \quad (3.4)$$

Here, \mathcal{D} represents the prediction step, which transports probability mass according to the environmental drift model. The specific formulation of this advection-diffusion process is detailed in the Methodology (Chapter 4). \mathcal{U} represents the update step that incorporates new sensor measurements z_t to reduce uncertainty.

To quantify the uncertainty of the belief map, we employ Shannon Entropy. The entropy of a single cell c is given by [16]:

$$H(c) = -p_t(c) \log_2(p_t(c)) - (1 - p_t(c)) \log_2(1 - p_t(c)). \quad (3.5)$$

The total map uncertainty is the sum over all cells: $H(M_t) = \sum_{c \in \mathcal{C}} H(c)$. The Information Gain $I(\pi)$ used in the problem formulation is defined as the expected reduction in this total entropy resulting from the observations collected along a trajectory π .

3.3. Synthetic Sensing Model

The UAVs are equipped with a downward-facing sensor modeled as a conical Field of View (FOV) with half-angle θ_{FOV} . Unlike simple proximity models, this formulation captures the realistic trade-off between coverage area and sensing fidelity. For a UAV at altitude h_t , the sensing footprint $\Omega_i(t)$ is a disk of radius $R_{\text{FOV}}(h_t) = h_t \tan(\theta_{\text{FOV}})$. We define a binary detection variable $z_t \in \{0, 1\}$, where $z_t = 1$ indicates a positive detection. The true positive probability

$P_D(h_t) = P(z_t = 1 | x_c = 1)$ decays with altitude:

$$P_D(h_t) = P_0 \cdot \min\left(1, \frac{h_{\text{ref}}}{h_t}\right), \quad (3.6)$$

where P_0 is the nominal detection rate and h_{ref} is the reference altitude for maximum resolution.

Given a sensor measurement z_t obtained at altitude h_t , the belief map is updated via the Bayesian filter \mathcal{U} (Eq. 3.4). For any cell $c \in \Omega_i(t)$, the occupancy probability is updated as follows:

$$p_{t+\Delta t}(c) = \frac{P(z_t | x_c = 1) p_t^{\text{pred}}(c)}{P(z_t | x_c = 1) p_t^{\text{pred}}(c) + P(z_t | x_c = 0) [1 - p_t^{\text{pred}}(c)]}, \quad (3.7)$$

where $P(z_t = 1 | x_c = 0) = P_F$ represents the constant false alarm rate. For cells outside the sensor footprint ($c \notin \Omega_i(t)$), the occupancy probability simply retains the predicted value $p_t^{\text{pred}}(c)$ generated by the drift propagation model. Although this setup establishes a consistent baseline for evaluating coordination, it represents an abstraction that removes real-world perception failure modes, such as environmental clutter, vision-specific false positives, and viewpoint-dependent detection probabilities.

3.4. Platform Models: VSQP and Quadrotor

This research utilizes two distinct UAV platforms to bridge the gap between the theoretical performance of hybrid aircraft and the practical requirements of real-world swarm coordination. The Variable Skew Quad Plane (VSQP) [14, 15, 26] serves as the primary baseline for large-scale simulations to evaluate the endurance benefits of exploiting hybrid flight regimes. In parallel, standard quadrotors (Parrot Bebop 2) are employed in physical field trials as kinematic surrogates to validate the framework's communication, coordination, and trajectory execution architecture.

While these platforms possess vastly different physical airframes, the high-level planner interfaces with them identically. Both utilize onboard INDI flight controllers for low-level stabilization, and both receive continuous velocity commands from the centralized Ground Control Station (GCS) via the Ivy software bus. The primary distinction between the platforms—and the driving factor behind the planner's behavior—lies in their respective flight-regime energy dynamics.

3.4.1. The Variable Skew Quad Plane (VSQP)

The Variable Skew Quad Plane (VSQP) is a hybrid VTOL platform that combines the maneuverability of a multicopter with the aerodynamic efficiency of a

fixed-wing aircraft. As depicted in Figure 3.3, it relies on four rotors for vertical lift and a pusher-propeller for forward flight. Because of its extreme energetic asymmetry, it serves as the primary theoretical model to demonstrate the endurance benefits of regime-aware informative path planning.




Figure 3.3: The VSQP in take-off, landing, and hover mode [15]

Empirical Energy Model Derivation

To optimize mission endurance, the planner utilizes a continuous, mode-dependent power model derived from high-frequency telemetry logs collected during Navy field trials of the VSQP platform¹.

To isolate the steady-state aerodynamic power consumption, the raw data required synchronization and filtering. Flight logs capturing instantaneous electrical power ($P_{elec} = V \cdot I$), True Airspeed (v_{tas}), and vertical velocity (v_z) were synchronized using nearest-neighbor interpolation with a 0.2 s tolerance. To ensure the model reflects purely equilibrium aerodynamic costs (where Lift \approx Weight), data points representing ground idling ($P_{elec} < 100$ W) and transient altitude corrections ($|v_z| > 0.25$ m/s) were filtered out.

The resulting filtered data was segmented into three kinematic regimes based on airspeed, yielding the continuous piecewise function $P(v)$ for forward airspeed $v(t)$:

$$P(v) = \begin{cases} 1751 \text{ W} & \text{if } 0 \leq v < 12 \text{ m/s} \\ 154.5v - 104.6 \text{ W} & \text{if } 12 \leq v < 17 \text{ m/s} \\ 323 \text{ W} & \text{if } v \geq 17 \text{ m/s} \end{cases} \quad (3.8)$$

¹The specific telemetry logs utilized in this study are confidential property of the Royal Netherlands Navy.




Figure 3.4: Derivation of the Empirical Energy Model from field telemetry. The scatter plot represents filtered power consumption data collected during multiple flight trials. The continuous piecewise model (thick lines) is fitted to these clusters: constant power for Hover (Red), linear interpolation for the Transition phase (Orange Dashed), and constant efficient power for Fixed-Wing Cruise (Green).

As visualized in Figure 3.4 and quantified by Equation 3.8, the VSQP operates across three distinct regimes:

- **Hover Mode** ($0 \leq v < 12 \text{ m/s}$): Lift is generated entirely by the four vertically oriented rotors. This is the most energy-intensive state, modeled as a constant 1751 W draw derived by averaging steady-state hover data from low-speed flight tests.
- **Transition Mode** ($12 \leq v < 17 \text{ m/s}$): The mechanical reconfiguration phase where the wing rotates. Modeled via linear regression, this phase exhibits a steep positive power penalty (+154.5 W per m/s). To ensure planning robustness, this dataset incorporates high-wind trials where the vehicle operated at high skew angles, resulting in a conservatively high power estimate.
- **Fixed-Wing Cruise Mode** ($v \geq 17 \text{ m/s}$): The wing is fully rotated ($\Lambda = 90^\circ$), and the UAV flies as a conventional aircraft. Data extracted from fixed-wing cruise logs confirms a massive 5.4 \times efficiency gain over hover, requiring a constant baseline of only 323 W. This highly efficient aerodynamic state is the primary regime for wide-area coverage.

The total energy cost $E(\pi_i)$ for a trajectory π_i of duration T is computed as the time integral of this instantaneous power profile:

$$E(\pi_i) = \int_0^T P(v(t)) dt. \quad (3.9)$$

By utilizing this continuous formulation rather than discrete penalties, the path planner can mathematically exploit specific cost gradients—for example,

minimizing time spent in hover during inspection tasks, or accelerating quickly through the inefficient transition band to reach the cruise regime. This formulation will also be used for the derived energy model of the Parrot Bebop 2, described in the next section.

3.4.2. The Parrot Bebop 2

The Parrot Bebop 2 is a standard commercial quadrotor utilized in this work to physically validate the planner’s algorithms in real-world conditions. While lacking a fixed-wing cruise mode, standard rotorcraft still exhibit distinct energy dynamics, consuming less power in moderate forward flight than in a stationary hover due to translational lift [9]. By supplying the planner with the Bebop 2’s specific empirical energy model, the platform acts as a kinematic surrogate, proving that the centralized coordination logic can successfully optimize trajectories for any given hardware profile.

Empirical Energy Model Derivation

To adapt the planner for the physical field trials, a continuous power curve was established for the Bebop 2 using onboard flight telemetry. For this flight test², we first hovered, and we flew at speeds of 8 and 13 m/s, as can be seen in figure 3.5. Similar to the VSQP methodology, the raw data was synchronized and rigorously filtered to isolate steady-state equilibrium flight. Data points exhibiting transient altitude corrections ($|v_z| > 0.25 \text{ m/s}$) or significant horizontal acceleration ($|a| > 0.2 \text{ m/s}^2$) were discarded to ensure the model reflects steady cruise.

As visualized in Figure 3.5, the filtered data exhibits the characteristic parabolic power curve of a multirotor [9, 11]. A second-degree polynomial regression was applied to formulate the continuous instantaneous power function $P_{bebop}(v)$ for forward airspeed v :

$$P_{bebop}(v) = 0.219v^2 - 1.626v + 81.300 \quad (3.10)$$

²The video of part of this test can be seen here: <https://youtube.com/shorts/nYUfg9NWGbU?is=aS6RKNLwSdgwzyCs>




Figure 3.5: Empirical Aerodynamic Power Curve for the Parrot Bebop 2. The scatter points represent filtered, steady-state telemetry data. The solid red line represents the quadratic fit, clearly illustrating the power reduction from hover to the optimal forward speed due to translational lift.

This continuous quadratic model allows the objective function to mathematically exploit aerodynamic "sweet spots" without requiring pre-defined flight modes. It successfully captures the three key phases of multirotor energy dynamics [9, 11]:

- **Hover Phase ($v \approx 0$ m/s):** Consumes approximately 81.3 W to maintain stationary lift. By embedding this in the cost-benefit evaluation, the planner naturally penalizes static hovering during exploration.
- **Translational Lift Phase (Optimal Cruise):** As forward speed increases, the rotor disk processes undisturbed air, increasing efficiency. The theoretical minimum power draw occurs at the vertex of the parabola ($v_{opt} \approx 3.71$ m/s), dropping the required power to approximately 78.3 W. The IPP framework uses this cost gradient to bias the swarm's trajectories toward this energy-optimal speed.
- **Parasitic Drag Phase:** Beyond the optimal speed, the quadratic increase in aerodynamic profile drag overtakes the benefits of translational lift, causing power demands to rise steeply at higher velocities. This penalty ensures that the hardware's physical limits are respected through mathematical optimization rather than rigid behavioral rules.

3.4.3. Kinematic Constraints and Wind Compensation

Let the state of each UAV i be defined as $\mathbf{x}_i = [x_i, y_i, z_i, \psi_i]^T$, where $\mathbf{p}_i = [x_i, y_i, z_i]^T$ represents the inertial position and ψ_i is the heading angle. The planner generates a desired ground-relative velocity vector $\mathbf{v}_{plan} \in \mathbb{R}^3$, which drives the positional dynamics:

$$\dot{\mathbf{p}}_i(t) = \mathbf{v}_{plan}(t). \quad (3.11)$$

Kinematic Limits

To ensure feasible trajectories across different hardware, the motion primitives must respect state-dependent constraints:

- **Horizontal Speed:** Bounded by $v_{min} \leq \|\mathbf{v}_{plan}\| \leq v_{max}$, with a nominal energy-optimal speed dictated by the platform's specific aerodynamic profile (e.g., $v_{opt} \approx 19$ m/s for the VSQP, and $v_{opt} \approx 3.71$ m/s for the Bebop 2).
- **Turning Rate:** Kinematic constraints dynamically adapt to the flight regime. During high-speed forward flight (e.g., VSQP wing-borne cruise), the heading evolution is non-holonomic and constrained by a maximum turn rate $\Delta\psi_{max}$ (set to $10^\circ/s$). Conversely, at low speeds or during Tracking phases, while multirotor configurations theoretically allow for holonomic flight, the planner deliberately avoids commanding instantaneous directional changes. Sudden velocity vector reversals require aggressive pitch and roll maneuvers that induce severe controller oscillation. To maintain aerodynamic stability and a consistent sensor footprint, the change in velocity is strictly bounded by a maximum acceleration limit $a_{max} = 2.0$ m/s².
- **Vertical Rates:** Climb and descent rates are restricted to physically safe bounds, typically $v_z \in [-1, 2]$ m/s in hover and $v_z \in [-2, 2]$ m/s in forward flight.

Wind Compensation and Airspeed

A critical aspect of maritime operations is the strong wind field at flight altitude. While the planner dictates the ground-relative velocity \mathbf{v}_{plan} required to track a path, the UAV must fly relative to the moving air mass. The required air-relative command velocity \mathbf{v}_{cmd} is given by the vector difference:

$$\mathbf{v}_{cmd} = \mathbf{v}_{plan} - \mathbf{v}_{wind}, \quad (3.12)$$

where \mathbf{v}_{wind} is the local wind vector.

Flying against the wind requires a higher command magnitude $\|\mathbf{v}_{cmd}\|$ to maintain ground progress, significantly increasing power consumption. Conversely, flying downwind allows for energy savings. The continuous energy models defined previously explicitly account for this physical reality by computing power as a function of the air-relative speed $\|\mathbf{v}_{cmd}\|$.

3.5. Problem Formulation

We frame the multi-UAV maritime SAR mission as a constrained trajectory optimization problem. Let $i \in \{1, \dots, N\}$ denote an individual agent within a swarm of N UAVs operating within the bounded search region $\mathcal{H} \subset \mathbb{R}^2$. The objective is to compute a joint

search policy $\pi = \{\pi_1, \dots, \pi_N\}$, where each π_i represents a time-indexed trajectory (a sequence of velocity vectors) for UAV i over a planning horizon.

The global objective function $J(\pi)$ is defined as the sum of individual utilities:

$$\max_{\pi} J(\pi) = \sum_{i=1}^N (I(\pi_i) - \lambda E(\pi_i)), \quad (3.13)$$

where $I(\pi_i)$ is the expected information gain—defined as the reduction in Shannon entropy $H(M_t)$ as formulated in Section 3.2—resulting from observations along trajectory π_i . The term $E(\pi_i)$ represents the platform-specific energy cost derived via the time integral of the power curve (Eq. 3.9), and λ is a weighting parameter that balances the trade-off between information gain and energy efficiency.

This maximization is subject to the following hard constraints for every UAV i :

$$E_{\text{used},i}(t) + E(\pi_i) \leq E_{\text{budget}}^{(i)} \quad (3.14)$$

$$\forall t : \pi_i(t) \in \mathcal{H}_{\text{safe}} \quad (3.15)$$

$$\pi_i \in \Pi_{\text{feasible}} \quad (3.16)$$

Here, Eq. (3.14) ensures that the energy required for the planned trajectory does not exceed the agent’s remaining battery budget $E_{\text{budget}}^{(i)}$, accounting for energy already expended, $E_{\text{used},i}(t)$. Eq. (3.15) enforces geofencing by requiring the UAV’s position $\pi_i(t)$ to remain within the safe region $\mathcal{H}_{\text{safe}}$, while Eq. (3.16) ensures the trajectory remains within the set of kinematically feasible motions Π_{feasible} defined by the flight envelope. Note that inter-agent collision avoidance is not explicitly modeled as a hard constraint in this local optimization; instead, it is handled at the coordination level through the centralized task allocation and the repulsive nature of the entropy map.

4 Methodology

This chapter presents the proposed Energy-Aware Informative Path Planning (IPP) framework, designed to unify multi-UAV coordination with the strategic exploitation of flight-regime energy dynamics. While the system utilizes a centralized architecture to manage the global belief map and offload heavy computation, the core research contribution is the development of a planning logic that mathematically couples path selection with non-linear power profiles to actively prioritize energy-efficient trajectories and maximize mission endurance.

The key aspect of this approach is the explicit integration of energy-intensive flight states—such as hovering during wide-area exploration—into the IPP

objective function to naturally bias the swarm toward aerodynamically efficient regimes. To achieve this, a Ground Control Station (GCS) maintains the global occupancy probabilities and assigns tasks, while executing per-agent receding-horizon optimization that accounts for both information gain and predicted power draw. To minimize unnecessary onboard energy expenditure, the physical UAVs operate as kinematic executors, receiving continuous velocity commands over the Ivy software bus. By centralizing the optimization process, the framework enables the swarm to navigate the complex trade-off between rapid victim detection and strict battery constraints, providing a systematic solution for large-scale maritime search.

4.1. System Architecture

The system architecture abstracts computations from the energy-constrained real-time execution hardware. As illustrated in Figure 4.1, the centralized server includes three primary logical components, while the physical UAVs act as a separate kinematic execution layer:

1. **Dynamic Mapping & Sensor Model (Server):** The central server acts as the global knowledge hub. It executes the computationally intensive drift propagation and global Bayesian sensor fusion. By centralizing these tasks, we completely offload heavy processing from the physical UAVs.
2. **Energy-Aware IPP (Server-Side Node):** Rather than calculating trajectories onboard, a dedicated instance of the Energy-Aware IPP algorithm runs on the ground server for each individual UAV. Using the continuously updated global belief map M_t , this layer generates and scores candidate trajectories to maximize the objective function subject to the kinematic constraints of the UAVs, outputting a continuous velocity command $\mathbf{v}_{cmd}(t)$.
3. **Swarming & Task Allocation (Server):** This supervisory layer monitors the belief map for high-probability detection events. It dynamically reassigns agents from *Exploration* to *Tracking* mode based on wind-adjusted energetic suitability, seamlessly passing these task assignments down to the respective per-agent IPP nodes.

Ultimately, the physical UAVs operate strictly as kinematic executors within a continuous real-time control loop. They stream their state telemetry—specifically position and heading $(\mathbf{p}_i(t), \psi_i(t))$ —to the server over the Ivy software bus. In return, the server computes and transmits back the optimized $\mathbf{v}_{cmd}(t)$ vectors. The onboard flight controllers then execute these commands by relying on their internal state $\mathbf{x}_i(t)$ to close the physical control loop.

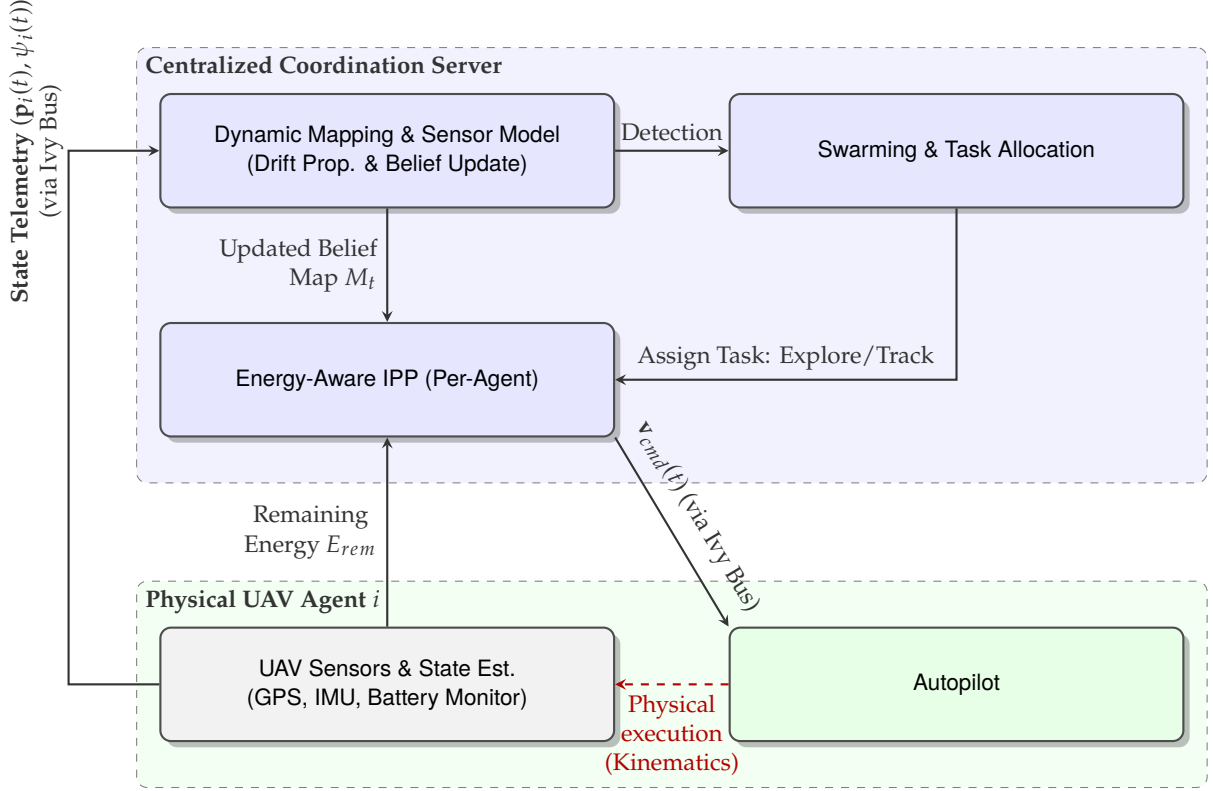


Figure 4.1: The closed-loop system architecture illustrating the continuous real-time control cycle. The physical UAV acts as a kinematic executor, streaming state telemetry via the Ivy bus to the server. The server updates the dynamic map, triggers tracking modes upon detection, and computes the continuous energy-aware velocity command $\mathbf{v}_{cmd}(t)$ to send back to the autopilot, which executes the command using internal state feedback.

4.2. Dynamic Environment Modeling

Victims drifting at sea are transported by environmental forces such as the ocean drift and the wind. To enable predictive path planning, we model this motion as a superposition of deterministic advection and stochastic diffusion.

4.2.1. Drift Dynamics

The mean drift velocity of the victim, \mathbf{v}_{drift} , is modeled as the vector sum of the local ocean current and the wind-induced leeway:

$$\mathbf{v}_{drift} = \mathbf{v}_{current}^{mean} + \mathbf{v}_{leeway}(\mathbf{W}_{10}), \quad (4.1)$$

where $\mathbf{v}_{current}^{mean}$ represents the mean surface current and \mathbf{v}_{leeway} accounts for the wind's drag on the floating object, driven by the 10 m wind vector \mathbf{W}_{10} [35].

This velocity drives the advection of the belief map over time. On the discrete grid \mathcal{C} , the probability mass is transported using a stochastic transition kernel $P(c | c')$. The predicted belief at the next time step is given by:

$$p_{t+\Delta t}(c) = \sum_{c' \in \mathcal{C}} P(c | c') p_t(c'). \quad (4.2)$$

In the system implementation, this transition kernel

$P(c | c')$ is realized in two stages to capture both physical phenomena: first, a deterministic spatial shift of the grid matrix proportional to $\mathbf{v}_{drift}(\Delta t)$ models the advection; second, a Gaussian convolution filter is applied to the shifted map to model the stochastic diffusion. This ensures that the probability mass is conserved while correctly capturing the growing uncertainty of the target's exact location as time progresses.

4.2.2. Conical Diffusion Model

While Equation 4.1 describes the mean trajectory, stochastic environmental turbulence introduces positional uncertainty that grows over time. Following a detection event at location \mathbf{p}_d and time t_d , we model the region of possible victim locations using a *Conical Diffusion Model*. Intuitively, this conical geometry reflects the linear expansion of positional uncertainty over time; the "fan-out" represents how a single point of certainty at the moment of detection evolves into a widening area of occupancy probability as the cumulative effects of variable wind and current forces increase.

The predicted search region $\mathcal{R}(t)$ expands anisotropically along the mean drift vector. Geometrically, this is defined as a cone with its origin anchored at the

initial detection point \mathbf{p}_d :

$$\mathcal{R}(t) = \left\{ \mathbf{x} \in \mathbb{R}^2 \mid \|\mathbf{x} - \mathbf{p}_d\| \leq R(t), \right. \\ \left. |\angle(\mathbf{x} - \mathbf{p}_d) - \psi_{\text{drift}}| \leq \frac{\theta_{\text{cone}}}{2} \right\} \quad (4.3)$$

where ψ_{drift} is the heading of the drift vector and θ_{cone} represents the angular uncertainty due to current variability. The search radius $R(t)$ bounds the stochastic uncertainty and expands linearly over time:

$$R(t) = R_0 + \gamma \|\mathbf{v}_{\text{drift}}\| (t - t_d), \quad (4.4)$$

where R_0 is the base uncertainty radius at detection and γ is an expansion coefficient derived from the estimated variance of the flow field.

Within this bounded region $\mathcal{R}(t)$, the victim occupancy probability $p_t(\mathbf{x})$ evolves via a discrete advection-diffusion process. First, the probability mass advects deterministically along the drift vector. Subsequently, a Gaussian smoothing kernel $G(\mathbf{x}; \sigma_{\text{diff}})$ is convoluted with the grid to simulate environmental stochasticity:

$$p_{t+1}(\mathbf{x}) = p_t^{\text{adv}}(\mathbf{x}) * G(\mathbf{x}; \sigma_{\text{diff}}) \quad (4.5)$$

where $*$ denotes the discrete convolution operator and p_t^{adv} represents the advected belief state. This diffusive spreading naturally causes the peak probability intensity to decay over time, as the total probability mass is conserved across the increasingly dispersed belief cloud. This results in a dynamic, high-entropy region that the IPP framework must prioritize for tracking, as visualized in Figure 4.2.

4.3. Single Agent Planning: Energy-Aware IPP

The primary goal of the single-agent planner is to identify a flight path that maximizes the discovery of new information about victim locations while keeping the aircraft in its most energy-efficient flight states. To solve the global maximization problem defined in Eq. (3.13), a Receding Horizon Planner (RHP) is implemented on the centralized coordination server. At each control epoch t , the planner projects the individual UAV's state forward over a fixed prediction horizon of N_{steps} . It generates a set of candidate trajectories Π_{cand} by sampling sequential spatial states at a fixed step distance d_{step} , evaluates the cumulative utility of these paths against the current occupancy belief map, and computes the optimal continuous velocity command $\mathbf{v}_{\text{cmd}}(t)$ to be transmitted to the UAV for execution.

4.3.1. Motion Primitives

To efficiently sample the search space while respecting platform-specific limits, the planner generates candidate trajectories using state-dependent spatial primitives. The geometry of the primitive set is governed by

a platform-dependent transition velocity $v_{\text{trans}}^{(i)}$, which differentiates between holonomic and non-holonomic flight regimes for platform $i \in \{\text{VSQP}, \text{Bebop}\}$.

- **Adaptive Heading Primitives:** The planner samples N_{dir} candidate endpoints on a circle of radius d_{step} . The angular sampling aperture Φ is a function of the current airspeed v_a :
 - *Holonomic Regime* ($v_a \leq v_{\text{trans}}^{(i)}$): The sampling is omnidirectional ($\Phi = 2\pi$). For the Bebop 2, this transition speed is set high to exploit quadrotor agility. For the VSQP, this corresponds to the hover-heavy regime where the tilted rotors provide full XY translation.
 - *Non-Holonomic Regime* ($v_a > v_{\text{trans}}^{(i)}$): The sampling is restricted to a forward-facing cone ($\Phi < \pi$) centered on the current heading. For the VSQP, this enforces the aerodynamic constraints of wing-borne flight, preventing stalls or high-drag sideslips. For the Bebop 2, this constraint ensures the aircraft remains within the small-angle assumptions of the steady-state energy model by preventing aggressive, high-acceleration reversals.
- **Altitude Primitives:** For every radial heading, the planner evaluates N_h discrete altitude levels. This allows the system to optimize the sensing trade-off: lower altitudes provide higher detection certainty (P_D) but smaller coverage area (Ω), while higher altitudes provide a wider footprint at the cost of resolution.
- **Environmental Feasibility:** Every primitive is subjected to a wind-resolution check. The planner calculates the ground velocity \mathbf{v}_g required to reach a sample, then resolves the corresponding airspeed $\mathbf{v}_a = \mathbf{v}_g - \mathbf{v}_{\text{wind}}$. A primitive is discarded if the required \mathbf{v}_a exceeds the platform's maximum airspeed $v_{\text{max}}^{(i)}$ or requires an acceleration $\dot{\mathbf{v}}_g$ beyond the hardware's thrust-to-weight limits.

4.3.2. Cost-Benefit Evaluation

For each candidate path $\pi \in \Pi_{\text{cand}}$, the planner computes a scalar utility score that couples expected information gain with the geometrically resolved energy costs. Because information gain (measured in bits of entropy reduction) and energy consumption (measured in Joules) exist on different numerical scales, the framework utilizes normalization constants to balance their influence within the objective function.

1. **Information Gain (I):** The planner simulates the sensor footprint along π and computes the expected reduction in global map entropy $H(M_t)$.




Figure 4.2: The Victim Drift and Uncertainty Model parameters. **(a)** Initial detection at \mathbf{p}_d , where the concentrated red region represents the high initial **occupancy probability**. **(b)** Advection-Diffusion process: The progression from red to blue illustrates the temporal evolution of the victim's location. As the belief moves along the mean drift vector $\mathbf{v}_{\text{drift}}$, the map is smoothed using a Gaussian kernel (σ_{diff}); the reader can observe the "peak" probability fading and spreading out, representing increasing positional uncertainty over time. **(c)** Tracking Model: The search cone $\mathcal{R}(t)$ provides a geometric bound for this expansion, fanning out linearly with a coefficient γ to encompass the growing area where the victim is likely to be located.

This value is scaled by a factor representing the maximum theoretically achievable information gain over the planning horizon, ensuring the term remains bounded and comparable to the energy cost.

- Energy Cost (E):** The energy consumption is estimated by evaluating the continuous power model $P(\|\mathbf{v}_a\|)$ over the trajectory segment. This cost is normalized by the platform's maximum power draw to map the energy expenditure into a dimensionless range.

The weighting parameter λ then acts as a dimensionless "priority factor." This normalization prevents the larger raw magnitudes of the energy model from overwhelming the information gain, allowing the optimizer to find a stable equilibrium between the two. By adjusting λ , the mission can be tuned to prioritize rapid occupancy probability reduction in high-stakes scenarios or to favor energy-efficient aerodynamic regimes when endurance is the primary constraint. This approach ensures that the planner's "bias" toward specific speeds or altitudes is a result of calculated mission trade-offs rather than an artifact of numerical scaling.

4.3.3. Trajectory Selection

The planner selects the locally optimal trajectory π^* by solving a finite-horizon instance of the global objective function. This selection represents a locally optimal decision over the defined prediction window N_{steps} , rather than a global mission optimum. Finding a globally optimal path for the entire mission duration is computationally intractable in a dynamic maritime environment where the underlying occupancy probabilities are constantly shifting due to environmental drift.

To determine π^* , the planner evaluates all candidate paths using the weighted utility:

$$\pi^* = \arg \max_{\pi \in \Pi_{\text{cand}}} J(\pi). \quad (4.6)$$

The process of generating and scoring these can-

didates is visualized in Figure 4.3. The planner simulates the spatial sensor footprint along each path, accumulating the expected information gain from the underlying grid cells while simultaneously integrating the predicted power draw.

We adopt a **Receding Horizon** approach for three primary reasons: (i) *Reactivity*: It allows the system to respond immediately to stochastic detections or unexpected changes in environmental drift; (ii) *Computational Efficiency*: It keeps the search space bounded, enabling the centralized server to manage high-frequency updates for multiple agents; (iii) *Feedback Loop Stability*: By discarding all but the first segment and re-planning, the framework closes the control loop with the UAV's current telemetry, ensuring the issued $\mathbf{v}_{\text{cmd}}(t)$ remains grounded in the aircraft's actual physical state. Consequently, the server derives and transmits only the first continuous velocity command of π^* before restarting the planning cycle.




Figure 4.3: Visualization of the Trajectory Selection process. The planner generates multiple candidate paths (π_1, π_2, \dots) using motion primitives. For each candidate, it simulates the expected sensor footprints (dotted circles). The path π^* (solid blue) is selected because it maximizes the objective function $J(\pi)$, offering the best balance of information gain (entropy reduction) and energy efficiency compared to the rejected alternatives (gray dashed).

4.4. Swarm Coordination: Task Allocation

While the single-agent IPP framework (Section 4.3) ensures that each UAV operates at its local aerodynamic and informational optimum, the swarm's collective efficiency depends on how these individual

trajectories are distributed across the mission area. The coordination layer serves as a high-level supervisor that prevents redundant coverage and manages mission priorities. By utilizing the centralized belief map as a shared reference, the system can partition the search space, ensuring that the "repulsive" nature of reduced entropy in one agent's path naturally pushes other agents toward unobserved regions.

The high-level behavior of the swarm is governed by the Finite State Machine (FSM) illustrated in Figure 4.4. This FSM translates global mission requirements into specific local objectives for each agent, distinguishing between two primary roles:

- **Exploration Role (Default):** The primary state where UAVs utilize the full energy-aware IPP logic to reduce overall map uncertainty. In this mode, the weight λ is tuned to favor long-endurance, wide-area coverage at optimal cruise speeds.
- **Tracking Role (Active):** Triggered upon a victim detection event. To ensure the most efficient response, the coordination server reassigns the agent with the smallest effective energy-distance (D_{eff}) to the detection site. This reassignment modifies the agent's local objective function to prioritize loitering over the detected occupancy probability peak, even if this requires higher-power hover states.

By switching between these roles based on D_{eff} , the coordination layer ensures that the swarm maintains a balance between discovering new victims and maintaining a high-fidelity track on known ones, while always selecting the most energy-capable agent for time-sensitive tasks.

4.4.1. Exploration Mode

In the absence of detections, the swarm operates in Exploration Mode with the objective of maximizing global map entropy reduction. Candidate trajectory segments are sampled from the reachable space within the geofence \mathcal{A} .

To prevent "flocking behavior"—where multiple agents converge on the same high-utility region—the framework employs a Sequential Planning Strategy with Anticipated Belief Updates. The coordination server triggers the planning cycle sequentially for agents $i = 1 \dots N$. When Agent i selects its optimal trajectory π_i^* , the expected information gain is immediately applied to a shared "planning belief map" as a *virtual* update. Consequently, when Agent $i + 1$ begins its optimization step, it perceives the area covered by Agent i as already explored. This mechanism naturally incentivizes orthogonal search trajectories and ensures spatial deconfliction without the need for complex inter-agent communication.

4.4.2. Task Assignment & Approach

The transition to the Tracking state is triggered by a discrete sensor detection event $z_{t_d} = 1$. Upon detection, the system records the initial target coordinates \mathbf{p}_d corresponding to the local belief peak:

$$\mathbf{p}_d = \arg \max_{c \in \text{FOV}_i} p_{t_d}(c) \quad (4.7)$$

The Central Server immediately evaluates the cost metric D_{eff} to assign the most suitable agent to intercept this location.

Wind-Adjusted Effective Distance

To minimize the energetic cost of interception, the server assigns the tracker based on a wind-adjusted effective distance metric. For a UAV i at position $\mathbf{p}_i(t)$, the cost to reach \mathbf{p}_d is:

$$D_{\text{eff}}(i) = \|\mathbf{r}_i\| + \gamma_{\text{wind}} \left| \mathbf{v}_{\text{wind}} \cdot \frac{\mathbf{r}_i}{\|\mathbf{r}_i\|} \right|, \quad (4.8)$$

where $\mathbf{r}_i = \mathbf{p}_d - \mathbf{p}_i(t)$ is the displacement vector. This ensures that agents positioned upwind are prioritized for downwind interceptions, leveraging the environment to conserve battery.

Approach and Stabilization

The selected agent enters the Approach Phase, flying towards the detection coordinate \mathbf{p}_d . Upon arrival, the agent enters an Initial Lock period (T_{lock}), holding its position relative to the drifting centroid $\mathbf{c}(t)$ as defined in Equation 4.3. This stabilization step ensures the search region $\mathcal{R}(t)$ is centered before engaging the constrained IPP.

4.4.3. Tracking Mode

Following stabilization, the UAV switches to Cone Tracking. The IPP planner is now constrained to generate candidate trajectories strictly within the predicted conical diffusion region $\mathcal{R}(t)$. To prevent "boundary lock-up" at the edge of the cone, a 10 m soft-constraint buffer is applied, allowing the RHP to generate smooth re-entry maneuvers.

To mitigate "ghost chasing" from residual belief peaks, a Gating Filter is implemented. The tracker rejects visual belief centroids that diverge by more than 20 m from the theoretical drift model $\mathbf{c}(t)$, ensuring the UAV remains anchored to the physics-based prediction.

4.4.4. Confirmation

The transition to the Confirmation phase (initiating descent) serves as a logical gate that prevents the UAV from abandoning its efficient search altitude until there is sufficient evidence to justify a high-energy commitment. Intuitively, this trigger fires when the swarm

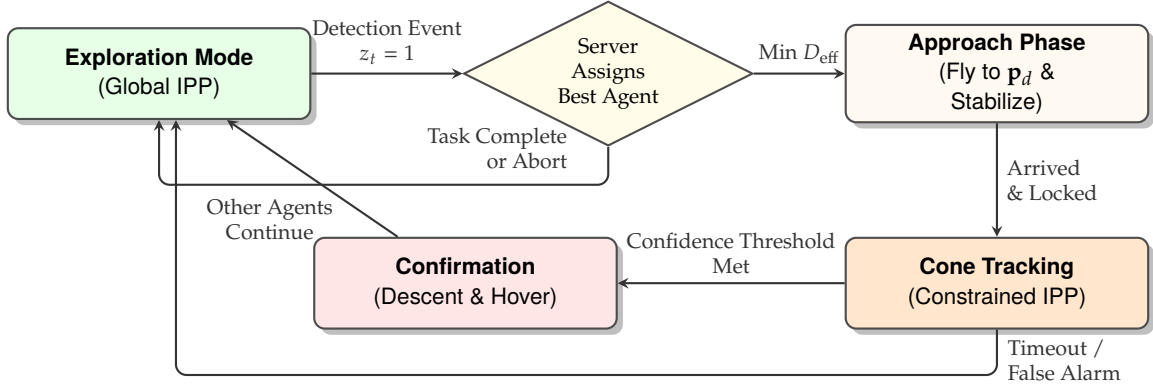


Figure 4.4: Finite State Machine (FSM) governing the swarm coordination. When a detection occurs, the Central Server executes a **Task Allocation** routine (Section 4.4) to assign the optimal agent based on the minimum effective energy-distance (D_{eff}). This metric ensures the responding UAV is selected based on its actual energy-to-arrival cost under current wind conditions rather than Euclidean distance. The selected agent then enters an **Approach Phase** to stabilize the drift estimate before initiating **Cone Tracking** once the confidence threshold defined in Eq. 4.9 is met.

has "seen enough" to transition from a wide-area scan to a localized dive. This is quantified by evaluating the search region $\mathcal{R}(t)$ through three perspectives: the "sharpness" of a specific sighting (p_{max}), the "consensus" across the entire area (\bar{p}), or the "density" of high-likelihood cells (ρ_{conf}). Formally, the transition is triggered when any of the following conditions are met:

$$(\rho_{conf} \geq \rho_{th}) \vee (\bar{p} \geq p_{mean_th}) \vee (p_{max} \geq p_{max_th}) \quad (4.9)$$

where ρ_{conf} is the fraction of cells exceeding the confidence threshold p_{conf} , \bar{p} is the mean occupancy probability, and p_{max} is the peak probability detected within the cone.

Once triggered, the tracker executes a descent to visually verify the target. To maximize mission efficiency, a dynamic confirmation ceiling allows the system to confirm a target at higher altitudes (e.g., 35 m) if the mean probability \bar{p} is exceptionally high (≥ 0.90), bypassing the need for a full descent. During the dive, the system utilizes a "Stop-and-Center" safety logic: if the horizontal alignment error exceeds a graduated threshold (tightening from 10 m to 2 m as altitude decreases), the vertical velocity v_z is inhibited. This ensures the UAV re-aligns with the occupancy probability peak before continuing its descent, preventing "fly-over" errors during the high-resolution confirmation.

4.5. Mission Overview

The interaction between the Energy-Aware IPP planner, the Dynamic Mapping layer, and the Centralized Swarm Coordinator is summarized in the operational sequence illustrated in Figure 4.5.

The mission begins in **Exploration Mode** (a), where the swarm operates at a high altitude to maximize the

sensor footprint and global map entropy reduction using the wind-compensated IPP strategy. Upon a discrete sensor detection event $z_t = 1$ (b), the system registers the initial target coordinate \mathbf{p}_d . The Central Server executes the task allocation logic, identifying the agent to break formation based on the D_{eff} metric. The selected agent transitions to the **Approach Phase** (c), descending toward \mathbf{p}_d and entering the initial lock to center the theoretical drift estimate $\mathbf{c}(t)$, while the unassigned agents maintain their high-altitude exploration.

Once stabilized, the assigned agent initiates **Cone Tracking** (d). It maintains a lower altitude for higher-resolution sensing, constraining its IPP primitives strictly to the predicted conical region $\mathcal{R}(t)$. Finally, if the probabilistic confidence criteria defined in Eq. 4.9 are met, the agent enters **Confirmation** (e), executing a final descent to h_{min} to conclusively validate the target before the server clears the local belief map and the agent climbs to resume global exploration.

4.6. Summary

This chapter detailed the implementation of a unified framework for energy-aware multi-UAV coordination in dynamic maritime environments. By formalizing the mathematical bridge between non-linear power dynamics and probabilistic search, the methodology provides the operational backbone for the comparative evaluations presented in Chapter 5. Specifically, the framework components enable the following experimental capabilities:

- **Quantifiable Energy-Search Trade-offs:** The integration of the continuous power model $P(\|\mathbf{v}_a\|)$ into the RHP allows Chapter 5 to quantitatively demonstrate how the planner sacrifices immediate information gain to stay within efficient aerodynamic envelopes. This enables the simulation-



Figure 4.5: Visualization of cooperative swarming behavior. (a) Distributed high-altitude exploration. (b) Detection phase: High-flying UAV₁ identifies p_d and triggers task allocation. (c) Approach phase: The server assigns UAV₂ to intercept, as its wind-adjusted effective distance (D_{eff}) is lower than that of UAV₁, making it the energetically optimal choice. (d) Tracking phase: UAV₂ maintains a low-altitude "close-up" on the advecting target within the cone $\mathcal{R}(t)$. (e) Confirmation: Target validation via descent to h_{min} . Note that icon scale (and sensor radius) visualizes relative flight altitude.

based comparison of the VSQP's hybrid endurance against traditional multirotor benchmarks.

- **Drift-Adaptive Search Resilience:** The *Conical Diffusion Model* and the drift-aware mapping layer allow the experiments to test the system's resilience against stochastic environmental forces. Operationally, this enables the swarm to maintain a high-probability track on a target even when the initial detection coordinates become obsolete due to current and wind, a key metric for the success rates reported in the following chapter.
- **Autonomous Swarm Management:** The Finite State Machine (FSM) and the *Wind-Adjusted Effective Distance* (D_{eff}) metric provide the logic required for the autonomous multi-UAV field trials. By automating the transition between *Exploration* and *Tracking*, the coordination layer removes the need for manual piloting or hard-coded waypoints, enabling the physical Bebop 2 trials to validate the framework's real-time

communication and task-switching stability.

By coupling flight-regime constraints with probabilistic environment modeling, this methodology establishes a systematic approach for long-duration maritime search. The following chapter evaluates this framework through a series of large-scale Monte Carlo simulations and physical multi-UAV field deployments to validate its effectiveness in real-world SAR scenarios.

5 Experiments & Results

To rigorously evaluate the proposed Energy-Aware Coordination framework, the experimental campaign is structured into four progressive phases. **Phase I** validates the framework's fundamental logic and scalability in a simulation environment. After the validation is completed, the method is tested against baselines in **Phase II**. In **Phase III** it will be explained how a bridge is formed between the high-level planner, the

Paparazzi simulation, and the autopilot. Finally, **Phase IV** involves real-world flight tests.

5.1. Phase I: Algorithmic Validation and Scalability Analysis

The initial validation phase focuses on the logical correctness of the Swarm Coordination FSM, the drift-tracking mechanism, and the energy-aware IPP. This phase evaluates the framework’s fundamental logic and operational limits within a controlled simulation environment, systematically evaluating performance across three distinct campaigns:

1. **Scalability Analysis:** Evaluating coordination efficiency, computational overhead, and map entropy reduction as the swarm size and search area dimensions are scaled.
2. **Environmental Robustness:** Assessing search resilience and tracking stability under varying drift velocities and wind noise to simulate unpredictable maritime conditions.
3. **Swarm Sizing and Strategy Comparison:** Determining the optimal swarm configurations required to meet high-recall mission objectives across diverse search areas and altitude strategies.

This abstraction allows for rapid iteration and scalability testing without the real-time constraints and hardware risks associated with physical field trials. In this phase, the simulation uses the VSQP model exclusively to isolate the coordination logic from the variables of heterogeneous fleet management.

5.1.1. Simulation Setup

To ensure consistency with the theoretical models established in previous chapters, the simulation imports the kinematic and energetic constraints directly from Chapter 3 and Chapter 4. The maritime search region \mathcal{H} is discretized into a 2D occupancy-based belief grid with a uniform resolution of $\Delta s = 10$ m, a value selected to balance sensing granularity with the computational requirements of the real-time planning server. While the physics of flight and energy consumption are model-derived, the operational behavior of the swarm is governed by a set of specific algorithmic coefficients and environmental parameters.

Table 5.1 details the nominal parameters used to initialize the experiments; these values are held constant across all trials unless explicitly stated otherwise.

The reference altitude $h_{ref} = 20$ m is calibrated to the “Recognition” standard defined in aerial SAR guidelines [38, 39], which is operationally sufficient for maritime rescue initiation. While forensic “Identification” requires a hazardous descent to sub-10 m, such proximity introduces unacceptable latency and risk in

time-critical maritime scenarios. At 20 m, the optical resolution allows the distinction of human survivors from the ocean.

The Field of View parameter $\theta_{FOV} = 45^\circ$ (90° full angle) is selected to approximate the optical characteristics of primary wide-angle visual sensors used in commercial search operations. Industry specifications for standard enterprise platforms confirm that visual search cameras typically range between 84° (e.g., DJI Mavic 3 Enterprise [40]) and 93° (e.g., Skydio X10 [41]).

The confirmation logic parameters are calibrated to enforce a high-sensitivity search phase followed by a high-precision validation phase. To prioritize Recall during the initial search, the confidence threshold $p_{conf} = 0.60$ is set marginally above the uniform prior (0.5) to filter stochastic noise while maintaining a low barrier for candidacy. Similarly, the spatial fraction threshold $\rho_{th} = 0.25$ is intentionally relaxed. This ensures that fragmented belief clusters—characteristic of sensor noise at 50 m—trigger an inspection descent rather than being discarded as false negatives.

Robustness and precision are enforced through temporal and probabilistic constraints. Task abortion is governed by a temporal window ($T_{timeout} = 30$ s) rather than transient fluctuations in occupancy probability, preventing early abandonment of valid targets due to sensor noise. Finally, during the confirmation phase, the spatial fraction trigger is disengaged in favor of a peak probability threshold ($p_{max} \geq 0.75$). This shift ensures that final confirmation is grounded in the identification of a specific, high-certainty detection rather than the diffuse probability mass sufficient for initial detection.

To ensure results are statistically significant and not artifacts of specific random seeds, every experimental configuration is evaluated across at least 20 independent Monte Carlo trials. A trial is classified as a “Success” if the swarm achieves a global entropy reduction of $\geq 80\%$ within the mission constraints. This threshold is chosen based on classic search theory [42], which identifies an asymptotic relationship between search effort and area coverage; the 80% mark represents the completion of the “Coarse Search” phase, beyond which the marginal information gain per unit of energy diminishes exponentially. This prevents the evaluation from being skewed by the “asymptotic tail” of the search process.

Table 5.1: Detailed simulation parameters for Phase I. The *Grid Bounding Box* defines the extent of the meshgrid before it is masked by the EHVB Polygon. Values are categorized as model-derived (*mod.*) or experimental/tunable (*exp.*).

Parameter	Value
<i>Grid & Environment (Sec. 3.2)</i>	
Test Area Boundary (\mathcal{H})	EHVB Polygon (mod.)
Grid Bounding Box	$\approx 500 \times 500$ m (exp.)
Grid Resolution (Δs)	10 m (exp.)
<i>Drift & Diffusion Dynamics (Sec. 4.2)</i>	
Cone Expansion Factor (γ)	1.0 (exp.)
Belief Smoothing Factor (σ_{diff})	1.0 grid/step (exp.)
<i>Sensing (Sec. 3.3)</i>	
Nominal Detection Rate (P_0)	0.9 (exp.)
FOV Half-Angle (θ_{FOV})	45 deg (exp.)
Reference Altitude (h_{ref})	20 m (exp.)
<i>IPP Planner Settings (Sec. 4.3)</i>	
Prediction Horizon (N_{steps})	3 (Explore) / 2 (Track) (exp.)
Step Length (d_{step})	40 m (exp.) / 20 m (Trk.) (exp.)
Heading Samples (N_{dir})	16 (exp.)
Energy Norm. (λ) - Explore	0.01 (exp.)
Energy Norm. (λ) - Track	0.05 (exp.)
<i>FSM Constraints (Sec. 4.4)</i>	
Initial Lock Duration	8.0 s (exp.)
Cone Search Timeout	30.0 s (exp.)
Re-tasking Cooldown	30.0 s (exp.)
Constraint Buffer	10.0 m (exp.)
Max Acceleration (a_{max})	6.0 m/s ² (mod.)
Max Speed (v_{max})	20.0 m/s (mod.)
Battery Capacity VSQP	300 Wh (mod.)
Hover Confirmation Time	10.0 s (exp.)
Task Allocation Wind Factor (γ_{wind})	5.0 (exp.)
<i>Confirmation Logic (Eq. 4.9)</i>	
Cell Confidence Level (p_{conf})	0.60 (exp.)
Cell Fraction Threshold (ρ_{th})	0.25 (exp.)
Mean Prob. Threshold ($p_{\text{mean_th}}$)	0.40 (exp.)
Peak Prob. Threshold ($p_{\text{max_th}}$)	0.75 (exp.)

5.1.2. Swarm Scalability

Experimental Setup - This campaign evaluates the coordination framework’s performance as both fleet size N and geographical scale L increase, testing the limits of centralized task allocation and the efficiency of the energy-aware planner. The swarm sizes range from $N \in \{1, 3, 5, 7, 10, 15, 25, 35, 50, 75, 100, 125, 150, 175\}$ agents across four maritime domains with side lengths $L \in \{500, 2000, 5000, 10000\}$ m. To maintain computational tractability for large-scale swarms ($N > 50$) while preserving search fidelity, the grid resolution is dynamically scaled according to:

$$\Delta s = \max(10.0, L/250) \text{ [m]}. \quad (5.1)$$

All agents utilize the VSQP platform model and the empirical power curves defined in Section 3.4. Missions are subject to two primary constraints: a hard 20% physical battery reserve and a dynamic time limit $T_{\text{max}} = \min(3600, (L/20) \cdot 10)$ seconds. The diverse range of N and L configurations is designed to

identify parameter sensitivity—specifically the saturation point where adding more agents to a fixed search area no longer yields proportional gains in occupancy probability reduction.

Results- The scalability analysis characterizes the physical and algorithmic boundaries of information-driven coordination. The primary findings are summarized in Figure 5.1.

The results illustrate a distinct transition in swarm dynamics as the operational scale increases. In the 500 m and 2000 m maps, the system is primarily search-limited. The steep slopes across all swarm sizes indicate that the bottleneck is the physical area available to be covered; we observe a rapid saturation point at $N \approx 15$ for the 2000 m map and at $N \approx 5$ for the 500 m map, where adding more agents yields negligible gains in search speed. At this density, the repulsive nature of the IPP objective function cannot overcome the physical crowding of sensor footprints, leading to redundant coverage.

In contrast, the 10,000 m domain reveals an energy-limited regime. Here, the entropy reduction curves exhibit near-linear slopes, indicating that the search area is vast enough to prevent significant agent overlap, allowing for efficient, independent exploration. However, a clear failure threshold emerges: for $N < 100$, the swarm is unable to reach the 80% occupancy probability reduction target before exhausting the T_{max} or battery constraints. This highlights that in large-scale maritime areas, the aerodynamic efficiency of the individual VSQP must be supported by a minimum "critical density" of agents to ensure mission success.

Figure 5.1 illustrates the "acceleration effect" of the swarm as the mission area expands. In the 500 m map (top-left), all swarms (even a single agent, $N = 1$) successfully achieve the 80% Coarse Search threshold within the mission time, with the curves becoming nearly vertical as N increases. In the 2000 m and 5000 m maps (top-right and bottom-left), a clear linear exploration phase emerges; larger swarms maintain a steeper, constant slope for longer, while smaller swarms exhibit a shallow gradient that eventually terminates prematurely due to battery exhaustion. The 10,000 m map (bottom-right) provides a visual proof of the *Battery Wall*: swarms with $N \leq 100$ fail to reach the threshold, as the energy cost of the initial transit consumes the usable battery capacity before sufficient coverage can occur. However, because the Coarse Search metric stops at 80%, this barrier is more flexible than strict total coverage: at $N = 125$, the swarm finally achieves the search capacity required to hit the 80% threshold just before energy depletion.

The mean time to reach the 80% coverage threshold demonstrates a power-law decay relative to swarm size across all successful map scales, as summarized in Figure 5.2. For small-scale operations (500 m), the search

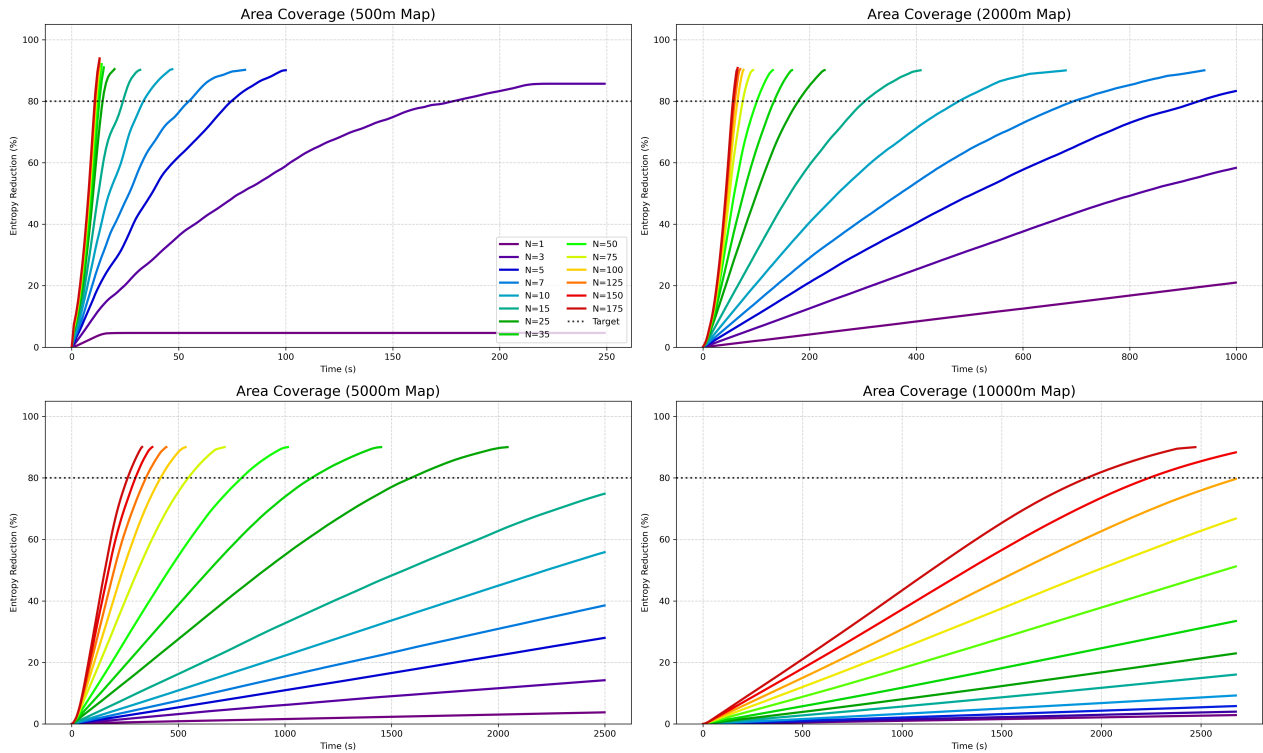


Figure 5.1: Temporal evolution of global entropy reduction across four map scales (500 m to 10,000 m). As search area increases, the swarm transitions from a saturation-prone regime where additional agents provide diminishing returns, to an endurance-constrained regime where mission completion is strictly dependent on achieving a minimum fleet density.

efficiency plateaus rapidly beyond $N = 15$, suggesting that the map becomes "spatially saturated" where additional agents provide negligible marginal utility. Conversely, larger scales (5000 m) continue to show significant performance gains up to $N = 100$, highlighting the framework's ability to effectively distribute high-density swarms.

required to clear an area before individual batteries are depleted.

- **Success Jump:** For 2000 m and 5000 m maps, a sharp transition from 0% to 100% success occurs as N increases, representing the point where the collective search rate outpaces the individual energy drain rate.
- **Physical Limit at 10 km:** At the 10,000 m scale, the success rate remains at 0% for $N \leq 100$. A sharp transition to success is now observed at $N = 125$, identifying a formidable "Battery Wall" where transit energy heavily taxes the usable battery capacity, requiring massive parallelization to clear the Coarse Search phase in time.

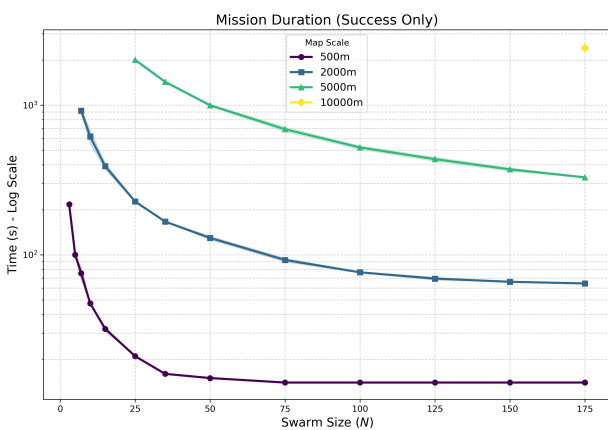


Figure 5.2: Search Efficiency: Mean Time to 80% Coverage across various swarm and map sizes.

The success probability curves in Figure 5.3 reveal a critical Minimum Viable Swarm (MVS) threshold

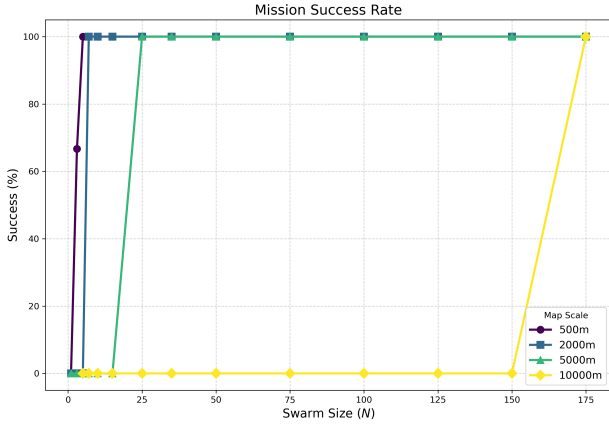


Figure 5.3: Mission success rate as a function of swarm size, highlighting the physical endurance limits at all environment scales.

The temporal evolution of global entropy reduction (Figure 5.1) illustrates the efficiency of the IPP allocator.

- **Linearity vs. Asymptote:** High-density swarms achieve nearly linear entropy reduction during early mission phases, indicating optimal spatial deconflation with minimal redundant searching. By concluding the mission at the 80% threshold, the swarm avoids the highly inefficient asymptotic tail of the search curve.
- **Speed Multiplication:** On a 5000 m map, a 100-UAV swarm completes the 80% coverage task approximately four times faster than a 15-UAV swarm, reducing the critical "window of uncertainty" from roughly 2000 seconds to under 500 seconds.

Based on the transition points observed in Figure 5.3, a generalized sizing rule can be established for the VSQP platform. By terminating at the 80% Coarse Search threshold and avoiding the asymptotic tail, the data suggests that for large-scale operations, the swarm achieves a highly efficient effective search rate of approximately 1.1 km^2 per UAV per hour. To achieve 80% Coarse Search coverage of an area A in a target time T , the required swarm size N can be estimated as:

$$N \approx \frac{A [\text{km}^2]}{0.00031 \cdot T [\text{s}]} \quad (5.2)$$

This heuristic is strictly bounded by the platform's physical endurance limits:

- **Individual Search Capacity:** Accounting for transit overhead, each agent can efficiently clear a maximum of $\approx 0.82 \text{ km}^2$ per charge.
- **Endurance Bound:** The maximum continuous mission time is capped at ≈ 44 minutes (2666 s)

due to the 323 W cruise power drain and the 300 Wh battery capacity.

- **Operational Ceiling:** Swarms must satisfy the density $N \geq A/0.82$ for mission success. At the 10 km scale (100 km^2), this implies a theoretical minimum of $N \approx 122$ agents required to clear the Coarse Search phase. This mathematical limit aligns perfectly with the experimental results, where the $N = 100$ configuration failed, but the $N = 125$ configuration successfully reached the 80% threshold just before battery depletion.

5.1.3. Environmental Robustness

Experiment - The primary aim of this experiment is to evaluate the framework's operational boundaries when subjected to high wind and wind-induced victim drift, and high-task-density scenarios, meaning scenarios where a lot of victims need to be detected in a bounded area within a constrained time frame. We subject a fixed swarm of $N = 3$ UAVs (at 50 m altitude) to three environmental severity levels: *Low*, *Medium*, and *High*. To ensure the results are not biased by specific flight headings, the wind direction is randomized across $[0, 2\pi)$ for every trial. The victim drift follows the standard Person-in-Water (PIW) leeway model, where velocity is $\approx 3\%$ of the 10 m wind speed W_{10} [35] (Table 5.2).

To stress-test the coordination logic and identify the system's "saturation point" regarding task management, we vary the number of victims $M \in \{1, \dots, 10\}$. This configuration tests the interplay between the UAV's aerodynamic efficiency against headwinds and the planner's ability to maintain the **occupancy probability** of multiple rapidly moving targets within the $500 \times 500 \text{ m}$ search domain.

Table 5.2: Environmental severity settings for Campaign B. The 3% leeway factor represents the standard drift rate for a Person in Water (PIW) under influence of wind and waves [35].

Severity	Wind Speed (W_{10})	Victim Drift (v_{drift})
Low	2.0 m/s	0.06 m/s
Medium	6.0 m/s	0.18 m/s
High	10.0 m/s	0.30 m/s

Results - The experimental results (Figure 5.4) demonstrate that the proposed framework achieves high environmental Invariance, maintaining consistent performance metrics across varying wind severities. For each severity level, 25 trials were performed. The system effectively manages aerodynamic disturbances by dynamically adjusting its flight profile and relying on robust algorithmic re-tasking logic.

As illustrated in Figure 5.4A, the mission success rate—defined as the successful confirmation of 100% of victims—is near-perfect when the victim count is within the fleet's capacity ($M \leq 3$). In these scenarios,




Figure 5.4: Environmental Robustness Results ($N = 3$). A) Success rates remain consistent across wind severities, declining primarily due to swarm capacity and time limits. B) Tracking instability (Aborts) is minimal (Median < 10), and increases with high severity level C) Mission duration for successful runs is consistent for all severity levels

the swarm processes targets in parallel. However, excess victims force the system into a serialized mission state progression (*Exploration* \rightarrow *Tracking* \rightarrow *Confirmation*). Because this cycle requires a finite time to satisfy the confidence requirements, the 600 s mission endurance permits only a limited number of sequential acquisitions, causing success rates to decline beyond $M = 5$. While the framework is resilient, high-severity wind conditions introduce an additional temporal penalty; the reduced ground speed during upwind transits causes the success rate to degrade slightly faster than in calmer conditions.

The tracking stability analysis (Figure 5.4B) highlights the efficacy of the framework's probabilistic validation gates (defined in Eq. 4.9). Across all severities, the median number of "aborts" stabilizes at approximately 9 to 10. Rather than indicating a failure of the control system, these aborts represent a conservative and intentional re-evaluation: if the local occupancy probability drops below strict thresholds due to sensor noise, the FSM triggers a reset to avoid wasting energy on "ghost" targets. Crucially, this behavior is environmentally invariant; the near-identical abort rates across all wind levels prove that tracking resets are driven by the internal logic of the framework rather than physical instability caused by the wind.

Finally, the Mission Duration metric (Figure 5.4C) demonstrates predictable and resilient scaling for successful runs. Mission times scale smoothly upward from roughly 150 s for a single victim to approximately 500 s for ten victims. The overlapping trajectories and confidence intervals across all severity levels confirm the aerodynamic resilience of the planner. By enforcing kinematic constraints—specifically a minimum guaranteed upwind ground speed—the UAVs successfully penetrate high-severity wind fields to capture upwind targets. As a result, the swarm reliably completes all required search legs without stalling, ensuring that overall mission durations remain robust and consistent

regardless of the ambient wind conditions.

Validation of Kinematic Asymmetry

To empirically validate that the observed "survivorship bias" is driven by ground speed limits, a limit-case simulation was conducted with a wind severity of 20 m/s. Given the UAVs' maximum cruise airspeed of 20 m/s, this scenario represents the theoretical limit where upwind progression becomes impossible ($v_{gs} \rightarrow 0$).

As visualized in Figure 5.5, the swarm exhibits Kinematic Saturation, effectively hitting a "Reachability Wall" at the upwind boundary. The UAVs are pinned to the downwind sector, capable of flying North-South but physically unable to move Westwards. This confirms that the statistical skew in mission duration (Figure 5.4C) is indeed an artifact of kinematic asymmetry: upwind legs are not inefficient, but significantly slower, causing upwind missions to systematically time out while downwind missions succeed rapidly.

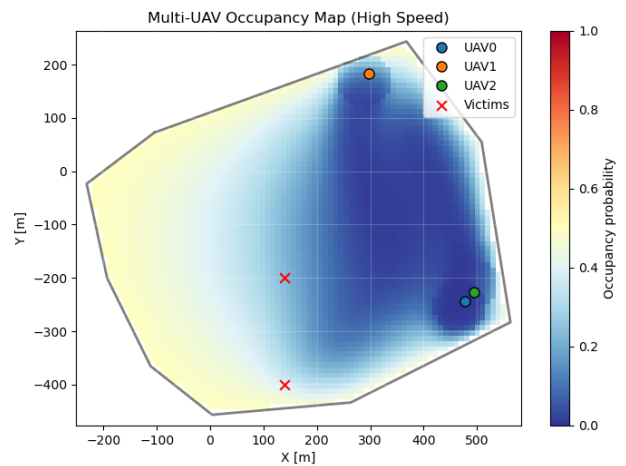


Figure 5.5: Limit Case Analysis (20 m/s Wind). The swarm faces a "hard wall" where $v_{wind} > v_{airspeed}$, rendering the upwind sector (West) kinematically unreachable. This proves that high-wind operations are constrained by ground speed asymmetry.

5.1.4. Swarm Sizing and Strategy Comparison

Experiment - Having validated the environmental robustness of the framework, we now investigate the scalability of the total system. The objective of this experiment is to determine the minimum swarm size required to achieve a high Victim Discovery Rate—defined as a recall $\geq 90\%$ —across varying operational areas (from 500×500 m up to 5000×5000 m) and victim densities $\{3, 5, 10, 15, 25\}$. To ensure the evaluation focuses on search efficiency rather than boundary effects, victims are consistently initialized at least 150 m within the search area perimeters.

Furthermore, we evaluate the impact of sensor diversity by comparing two distinct altitude strategies:

- **Uniform Strategy:** A homogeneous swarm where all UAVs maintain a constant search altitude of 50 m. This serves as a baseline for maximum coverage rate.
- **Mixed Strategy:** A heterogeneous altitude allocation where half of the fleet operates at 35 m and the other half at 50 m.

Results - The experimental results reveal a critical inversion in swarm strategy effectiveness based on the scale of the operational area. While the *Mixed* altitude strategy is highly efficient for smaller search areas, it suffers a severe performance degradation on massive maps, where the *Uniform* strategy proves to be much more robust and scalable.

As illustrated in Figure 5.6A, for smaller maps (e.g., 500 m to 2000 m), the Mixed swarm (dashed lines) and the Uniform swarm (solid lines) perform equally well, meaning that the Mixed swarm is preferred since it uses less energy. The swarm dynamically evaluates travel cost versus information gain (4.8); if a 50 m UAV detects a weak signal, the IPP tasks the UAV that is available and has the smallest effective distance D_{eff} to the target area to investigate. On a small map, a 35 m UAV is frequently nearby, allowing the swarm to operate more efficiently, since confirmation height is only 5 m down.

The lookup tables (Figures 5.6B and 5.6C) demonstrate that this low-high UAV advantage collapses on the massive 5000×5000 m map. In these scenarios, the mixed swarm requires significantly more UAVs to reach the 90% target recall compared to the Uniform swarm. This degradation is not a failure of coordination, but rather a mathematical consequence of area coverage deficit.

In a massive 25 km^2 operational area with a strict 600 s endurance limit, the primary mission bottleneck shifts from target confirmation to pure spatial exploration. The Uniform swarm (all UAVs at 50 m) maximizes the collective Field of View (FOV) and total sweep rate of the fleet. Conversely, the Mixed swarm sacrifices the FOV of half its fleet by restricting them to 35 m. While

these 35 m UAVs generate high-quality local data, their small FOV make them highly inefficient at exploring vast stretches of empty ocean. Consequently, the 50 m UAVs in the Mixed swarm are forced to shoulder the majority of the exploration burden, effectively halving the swarm’s spatial clearing rate.

The rare instances where a small Mixed swarm achieved high recall on the 5000 m map were statistical anomalies driven by “lucky spawns,” where the targets randomly spawned directly in the path of the UAVs, bypassing the need for map-wide exploration.

While heterogeneous altitude distribution accelerates search times in constrained environments by optimizing confirmation confidence, large-scale Maritime Search and Rescue operations are strictly bottlenecked by volumetric sweep limits, mandating homogeneous, high-altitude swarms to maximize area coverage.

5.2. Phase II: Baseline Comparison and Sensitivity Analysis

Having validated the core mechanics in Phase I, this evaluation focuses on benchmarking the framework’s performance against industry-standard SAR protocols. The primary objective is to identify the specific environmental crossovers where the proposed energy-aware adaptability provides a measurable advantage over rigid geometric patterns. Furthermore, we conduct a sensitivity analysis on the energy-weighting parameter λ to evaluate its effectiveness in balancing mission endurance against spatial coverage.

Experiment - To validate the efficiency of the energy-aware IPP, a swarm of $N = 5$ VSQP agents was deployed. The proposed method is benchmarked against two baselines: (i) a non-energy-aware version of the IPP ($\lambda = 0$), and (ii) a *Boustrophedon Coverage* strategy [43]. The latter represents traditional SAR protocols where the swarm partitions the search area into parallel sweep corridors, disregarding the underlying **belief map** (Figure 5.7).

The comparison was conducted under two distinct 6.0 m/s wind vectors, selected from a preliminary sweep of all wind angles (Figure 5.8). We specifically compare the aerodynamically optimal wind angle (290°) against the most unfavorable regime (310°), where crosswinds drive the UAVs toward the search area boundaries.

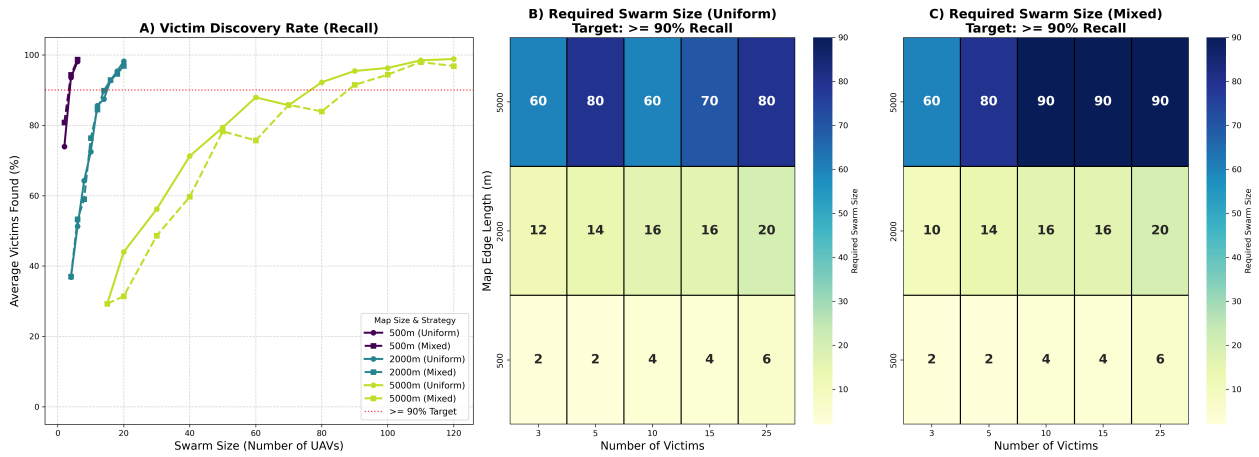


Figure 5.6: Scalability and Strategy Comparison: Uniform vs. Mixed Swarms. (A) Average Victim Discovery Rate as a function of swarm size. The Mixed altitude strategy (dashed lines) consistently reaches the 90% target threshold (red dotted line) with similar UAVs compared to the Uniform strategy (solid lines), except for the 5000 by 5000 m area. (B) Required Swarm Size lookup table for the Uniform strategy (all UAVs at 50 m), showing that (C) Required Swarm Size lookup table for the Mixed strategy (UAVs half at 35, and 50 m).

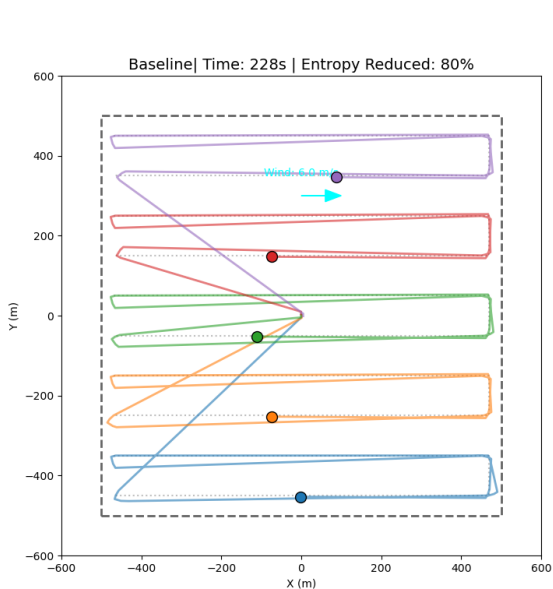


Figure 5.7: Flight paths of the Baseline (Boustrophedon) geometric search strategy with $N = 5$ agents. Unlike the proposed probabilistic IPP, this method rigidly partitions the operational area into parallel sweep corridors, disregarding underlying belief maps. In this specific trial, the swarm successfully achieved the 80% entropy reduction target in 228 seconds under a 6.0 m/s wind vector.




Figure 5.8: Evaluation of the baseline efficiency across all wind angles, identifying 290° as the optimal aerodynamic condition and 310° as the worst condition.

Results - Under the optimal environmental conditions (290° wind), the traditional Lawnmower baseline outperformed the Proposed Method in both mission duration and energy conservation. Because the baseline executes deterministic, straight-line trajectories, it maximizes time spent in the highly efficient fixed-wing cruise regime. Conversely, the Energy-Aware IPP incurred a temporal and energetic "cost of adaptability" due to the kinematic overhead of continuously computing curved, collision-avoiding trajectories. As seen in Figure 5.9, under favorable and structured wind conditions, deterministic straight lines are highly efficient, and the swarm easily met the 80% coverage threshold across all tested λ values.

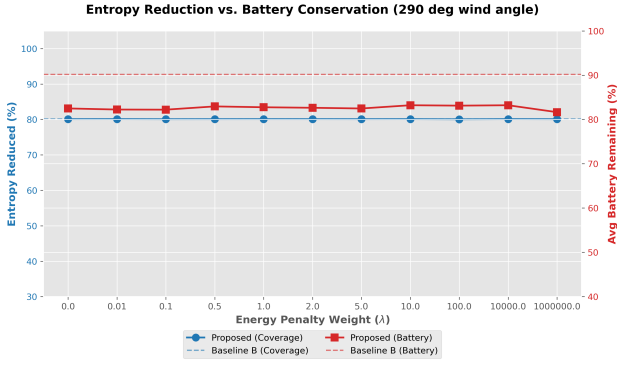


Figure 5.9: Sensitivity analysis under the optimal 290° wind condition. The baseline (dashed lines) sets a highly efficient benchmark that the proposed method approaches but does not strictly surpass in perfect weather.

However, maritime SAR environments are highly stochastic, and this optimal efficiency is fragile. When exposed to adverse aerodynamic conditions—specifically the aerodynamic cliff at 310° where crosswinds drive the UAVs into map corners—the rigid geometry of the baseline forced the swarm to continuously fight severe aerodynamic drag, resulting in significant battery depletion. In these realistic and unpredictable scenarios, the Energy-Aware IPP framework proves its value of adaptability. As shown in Figure 5.10 (filtering for successful missions), the proposed planner autonomously adapts its flight paths to utilize wind gradients, effectively outperforming the geometric baseline in battery conservation.

Analysis of Figure 5.10 demonstrates that λ acts as a critical pacing factor for the swarm. A low λ (prioritizing information only) leads to high-power, erratic maneuvers, while $\lambda \approx 0.5$ identifies the operational sweet spot for the VSQP platform by penalizing high-draw hover states (1751 W) in favor of cruise states. While the specific value of $\lambda = 0.5$ is tuned to the power curves of the VSQP, the underlying principle—using a normalized weight to balance the information-to-energy ratio—is robust across different hardware profiles, as it allows the planner to respect the physical discontinuity between hover and wing-borne flight.

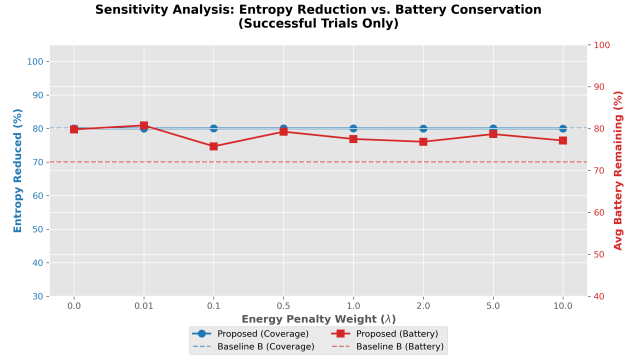


Figure 5.10: Performance frontier under the adverse 310° wind condition (Successful runs only). The proposed method ($\lambda \approx 0.5$) successfully dynamically avoids the aerodynamic cliff, conserving more battery than the rigid geometric baseline.

Despite this aerodynamic superiority, the unfiltered dataset for the 310° wind condition (Figure 5.11) exposes a key limitation of the framework. Under severe environmental stress, the IPP swarm occasionally failed to reach the 80% coverage threshold before the 600-second deadline. This is a direct consequence of spatial myopia inherent to Receding Horizon planners. As the swarm approached the final stages of the coverage search, the remaining unknown areas shrank into isolated pockets. When forced into upwind map corners by the 310° wind, local candidate trajectories returned zero Information Gain. Without a global geometric map to pull them out of these local minima, the UAVs occasionally became trapped, loitering inefficiently in terms of information gain, until the mission timed out.

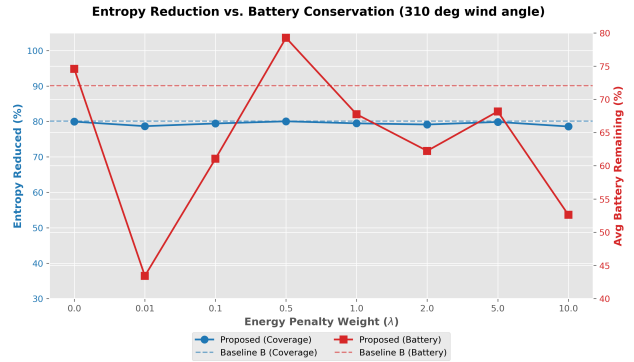


Figure 5.11: Sensitivity analysis under the adverse 310° wind condition (All runs, including timeouts). The erratic coverage variance at extreme λ values highlights mission failures caused by the swarm trapping in local minima during upwind maneuvers.

Ultimately, the evaluation demonstrates a fundamental operational trade-off: the proposed framework trades absolute spatial clearing speed for improved aerodynamic efficiency. While the probabilistic IPP was never faster than the geometric baseline, it successfully mitigated wind-induced drag in adverse con-

ditions. By dynamically adapting its flight paths, the energy-aware approach resulted in lower average battery consumption to reach the 80% threshold in the majority of sub-optimal trials compared to the rigid baseline. However, despite this energetic advantage, its reliance on a receding horizon leads to diminishing returns and occasional local minima as global entropy depletes. This indicates that to achieve full map clearance, the IPP should be utilized for the Coarse Search phase before handing off to a deterministic Fine Search algorithm.

5.3. Phase III: System Integration

While Phases I and II established the theoretical efficacy and algorithmic limits of the framework, Phase III transitions the focus toward system integration readiness. The objective of this final validation stage is not to replicate the large-scale statistical analysis of previous sections, but to demonstrate the technical feasibility of the real-world deployment. Specifically, this phase serves as a proof-of-concept for the hardware-in-the-loop (HIL) architecture, proving that the centralized coordination server, the communication middleware, and the autopilot stack can synchronize reliably under real-time constraints. By deploying the framework at the CyberZoo and Valkenburg Naval Air Base (EHVB), we validate the system’s resilience to asynchronous data streams, network latencies, and hardware stochasticity—confirming its readiness for operational field missions.

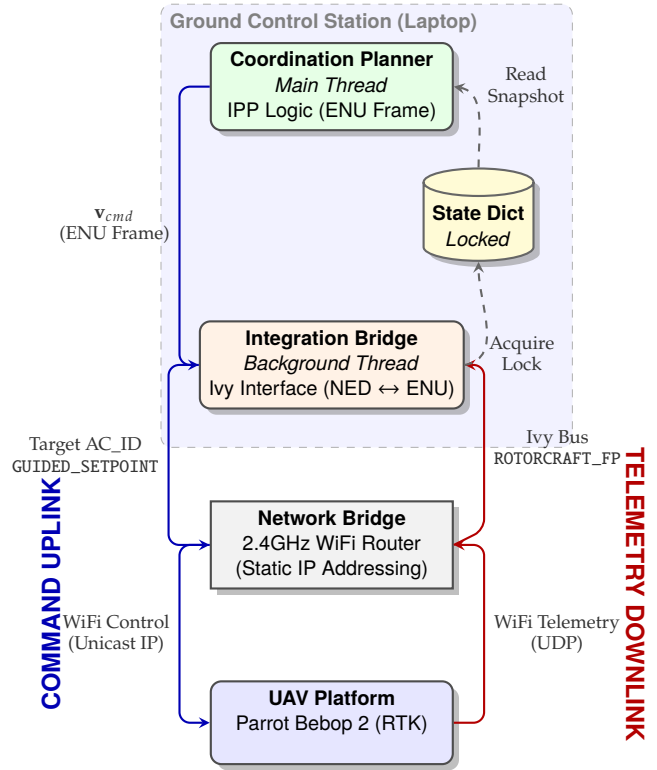


Figure 5.12: Hardware-in-the-Loop (HIL) Communication Architecture. The data flow is partitioned into a **Command Uplink** channel (left, blue) and a **Telemetry Downlink** channel (right, red). Node styling denotes functional domains: Green rounded rectangles represent custom application logic, Orange denotes translation middleware, the Yellow cylinder highlights shared thread-safe memory, the Gray sharp rectangle denotes passive network infrastructure, and Blue represents the embedded flight hardware.

5.3.1. Network Topology and Middleware Architecture

To facilitate real-time swarm coordination, we developed an integration bridge (depicted within the GCS block in Figure 5.12) using the Ivy software bus. This architecture enables a single centralized server to manage a distributed swarm of heterogeneous UAVs by abstracting the communication and coordinate transformation layers—a capability that was previously impossible within the standard standalone autopilot framework. The physical communication layer utilizes a centralized star topology, where a ground-based router acts as the bridge between the GCS and the distributed fleet.

Asynchronous Threading Logic

In contrast to the synchronized simulation environment, real-world telemetry arrives with varying latencies and frequencies. The bridge architecture utilizes a multi-threaded approach to decouple the high-frequency state updates of the aircraft from the lower-frequency planning cycle (≈ 5 Hz). By implementing a thread-safe "State Dictionary," the system ensures

that the planner always operates on a consistent spatial snapshot of the entire swarm, preventing the race conditions that typically emerge during multi-agent data ingestion.

Addressing and Frame Transformation

The bridge dynamically manages agent identification through Aircraft IDs (AC_ID), routing command packets to specific agents while simultaneously handling frame transformations. Because the autopilot operates in a North-East-Down (NED) frame and the IPP planner utilizes an East-North-Up (ENU) grid, the bridge automatically scales and transforms incoming telemetry into metric ENU coordinates. This abstraction allows the high-level planner to remain platform-agnostic, focusing on search optimization rather than low-level coordinate arithmetic.

5.3.2. Closed-Loop Control and Kinematic Management

The integration phase validates the functional handshake between the IPP algorithm and the physical flight controllers. As illustrated by the Command Uplink data flow in Figure 5.12, the bridge consumes optimal velocity vectors \mathbf{v}_{cmd} and transmits them as hardware-compatible setpoints.

To ensure aerodynamic stability during real-world maneuvers, the bridge enforces a kinematic yaw rate limiter. While the desired heading is derived from the velocity vector, angular updates are clamped to a maximum rotation rate ($20^\circ/\text{s}$). This prevents the aggressive roll oscillations that can occur when a high-speed aircraft attempts to track the instantaneous heading changes of a stochastic path planner. By maintaining independent, un-buffered control channels for each aircraft, this architecture demonstrates that a centralized server can achieve stable, closed-loop kinematic control over a physical swarm in real-time.

By successfully closing the loop between the IPP server and the physical hardware, Phase III confirms the framework’s integration readiness. This stage proves that the bottleneck of communication and coordinate transformation has been resolved. With the hardware-software handshake validated, the framework is prepared for Phase IV, which will evaluate the autonomous search performance and victim discovery rates in unconstrained maritime field environments.

5.4. Phase IV: In-field Testing

This final phase presents the physical validation of the coordination framework through two distinct experimental tiers: a full indoor proof-of-concept and a partial outdoor system validation. First, we utilize the TU Delft CyberZoo facility to evaluate the end-to-end autonomy of the swarm, including target detection, tracking, and inter-agent coordination within a con-

trolled environment. Subsequently, we present results from the Valkenburg Naval Air Base (EHVB), which serves to validate the system’s operational robustness in an unconstrained outdoor setting. Due to the unavailability of the custom VSQP prototypes at the time of testing, the Parrot Bebop 2 quadrotor serves as the common physical flight platform. These trials demonstrate that the energy-aware coordination logic is platform-agnostic and capable of maintaining stable search patterns despite the real-world constraints of sensor noise and communication latencies.

5.4.1. Indoor Verification (TU Delft CyberZoo)

Experiment- The initial system integration tests are performed in the TU Delft Cyberzoo, a controlled $10 \times 10 \times 7$ m flight arena. Since the environment is GNSS-denied, the system utilizes an OptiTrack Motion Capture system. The Integration Bridge connects physically to the Cyberzoo’s local router to ingest real-time rigid-body data, injecting it into the autopilot’s INS to emulate high-precision GPS coordinates. To ensure safety within the confined volume, the planner’s kinematic constraints are strictly clamped to a maximum velocity of $v_{max} = 0.5$ m/s, and a maximum acceleration of $a_{max} = 0.3$ m/s². Furthermore, the energy-aware components of the cost function are disabled ($w_{energy} = 0$), forcing the planner to optimize purely for information gain.

To accommodate the scaled-down indoor geometry, several high-level coordination parameters were aggressively tightened compared to the open-water methodology. The dynamic confirmation ceiling (detailed in Section 4.4.4) was disabled, enforcing a hardcoded confirmation descent to 1.0 m. The stop-and-center descent safety logic (Section 4.4.4) was reduced to a strict 0.8 m horizontal tolerance, and the visual gating filter (Section 4.4.3) was tightened from 20 m to 2.0 m to prevent the UAV from locking onto spurious probability peaks near the geofence boundaries.

In terms of sensing, the reference altitude h_{ref} is set to 1 m, and the FOV Half-Angle θ_{FOV} to 25° . The optimization step length is set to 2.0 m for global exploration and reduced to 1.5 m during localized cone-tracking.

In the first experiment, single-UAV IPP was tested. The setup consisted of one UAV acting as the “victim” and one acting as the “explorer”. The victim UAV drifted with a constant velocity of 0.1 m/s across the CyberZoo while maintaining a 0.5 m altitude, while the explorer conducted its search from a 2.0 m altitude.

In the second experiment, multi-UAV IPP was tested using two active “explorer” UAVs and a passive physical proxy serving as the “victim”. To safely simulate a surface-level target without introducing severe rotor-downwash interference or mid-air OptiTrack marker occlusion during the confirmation descent,




Figure 5.13: Single-Agent Validation Results in the CyberZoo. Yellow boxes indicate the UAV Agent, while the red box indicates the UAV victim proxy. Arrows originating from the boxes indicate the current motion trajectory of the respective UAV. From left to right: (1) Agent registers an initial detection at \mathbf{p}_d via Bayesian sensor fusion. (2) Agent transitions to \mathbf{p}_d and initiates cone tracking. (3) Upon satisfying the confidence criteria (Eq. 4.9), the agent descends to the 1 m h_{ref} to finalize confirmation. (4) Agent begins ascent to initial scanning altitude. (5) UAV reaches scanning altitude and resumes Exploration IPP.

the victim was represented by a tracked piece of cardboard. This proxy was pulled via a cable to simulate ocean drift. The two explorers maintained respective search altitudes of 2.0 m and 3.0 m to guarantee spatial deconfliction. To manage this heterogeneous mix of physical UAVs and passive markers, the centralized Integration Bridge utilized strict ID-based message routing.

Results of Single Agent Test - The flight data confirms the logical correctness of IPP and the FSM on physical hardware. As illustrated in Figure 5.13, the system successfully executed the full search-and-confirm lifecycle with a single agent.

First, regarding IPP and trajectory tracking, the UAV tracked the high-level velocity setpoints generated by the planner in the ENU frame. The Integration Bridge correctly transformed these into valid NED commands for the Paparazzi autopilot, resulting in smooth coverage of the search area (Figure 5.13, top).

Second, regarding detection and transition, the agent correctly triggered the *Detection Event* upon entering the proximity radius of the victim proxy, transitioning the FSM from *Exploration* to *Tracking* mode without oscillation.

Finally, the most critical validation is visible in the altitude profile of Figure 5.13 (bottom). The agent stabilized over the target coordinates and executed a controlled descent from $h = 2$ m to the confirmation altitude ($h_{ref} = 1$ m). Once the “Confirmed” state was logged, the UAV maintained a stable hover over the victim, validating the vertical control logic.¹

Discussion- While the flight logic proved robust, the experimental campaign revealed significant challenges regarding the physical communication layer. The CyberZoo facility presents a spectrally congested environment, with multiple overlapping 2.4 GHz and 5 GHz networks used for building infrastructure and motion capture systems.

Consequently, the commercial-grade router used in the Integration Bridge struggled to broadcast its SSID on a stable channel. Network association times were highly stochastic, often requiring up to 30 minutes to establish connection with the UAVs. This “connectivity lottery” suggests that while the Star Topology (Section 5.3.1) is sufficient for open-field operations (Phase IV), it is vulnerable in spectrally dense indoor environments. However, it is important to note that once the connection was established, the Integration Bridge maintained a stable connection, and the asynchronous threading architecture successfully handled the telemetry stream without packet loss during the flight.

Results of Multi Agent Test - The flight data confirmed the successful execution of distributed Multi-UAV IPP, dynamic task allocation, and safe spatial deconfliction². Initially, the two search agents successfully partitioned the search space to perform exploration without explicit inter-agent communication, driven entirely by the shared global entropy map (as illustrated in Figure 5.14, top-left).

Upon the cardboard victim entering the simulated Field of View of one of the agents, the Central Server correctly triggered the Detection Event. The server evaluated the Wind-Adjusted Effective Distance (D_{eff}) and autonomously assigned the energetically optimal UAV to break off from exploration and intercept the target (Figure 5.14, top-right).

The assigned tracker transitioned seamlessly into Cone Tracking mode, maintaining a lower altitude over the moving proxy. Once the probabilistic confidence criteria were met, the agent executed a stable confirmation descent over the cardboard target (Figure 5.14, bottom-left). Unlike previous configurations utilizing an active UAV as a victim—which suffered from severe OptiTrack marker occlusion and identity merging—the low-profile physical proxy allowed the tracking cameras to maintain distinct rigid-body iden-

¹The video of this test can be seen here: https://youtu.be/VugWa_E22gU?si=xdQb0mNW-4E4sW_g

²The video of this test can be seen here: https://youtu.be/ZNlnYty_aHc?si=QtyGBKjyiu3BPiPW




Figure 5.14: Sequential evolution of the multi-UAV occupancy map demonstrating successful task allocation and target confirmation. **Top-Left:** Initial detection of the physical victim proxy (red ‘x’) by the swarm triggers a localized rise in belief probability. **Top-Right:** The Central Server evaluates the effective distance and autonomously assigns UAV2 (green) to intercept the target. UAV2 transitions into Cone Tracking mode, restricting its path planning to the predicted diffusion region (green dashed boundary) where confidence rapidly peaks (dark red). **Bottom-Left:** Upon satisfying the probabilistic confidence criteria, UAV2 executes a stable confirmation descent directly over the target. **Bottom-Right:** Following successful validation, the server clears the local belief map (probability and entropy reduced to near-zero), and UAV2 safely ascends to resume distributed global exploration.

tities throughout the vertical maneuver. Following confirmation, the agent safely ascended and resumed global exploration, while the server cleared the local belief map of the confirmed target (Figure 5.14, bottom-right).

Discussion— The multi-agent flight test successfully validated the complete Energy-Aware IPP and Swarm Coordination framework in a controlled indoor setting. The physical hardware proved capable of dynamically transitioning between distributed, high-altitude exploration and localized, low-altitude target confirmation without inter-agent collisions or data-link degradation. By utilizing a passive physical proxy alongside active searchers, the system demonstrated robust integration of asynchronous telemetry streams and localized sensor fusion.

Overall, the validation within the CyberZoo served

as a successful controlled proof-of-concept for the centralized architecture. Nevertheless, an indoor environment cannot fully validate the aerodynamic robustness of the swarm. Indoor testing inherently lacks unpredictable real-world variables, such as wind gusts, full-scale RTK inaccuracies, and long-range communication constraints. Therefore, transitioning to the outdoor testing phase at Valkenburg Naval Air Base remains essential to evaluate the framework’s energy-awareness against physical wind-induced drag.

5.4.2. Field Validation (Valkenburg Naval Air Base)

Following successful indoor verification, the system was deployed to the Valkenburg Naval Air Base (EHVB) for outdoor testing. The primary focus of this campaign was the structural validation of the framework’s

deployment stack—specifically the robustness of the asynchronous network architecture, the precision of the Real-Time Kinematic (RTK) localization, and the reliability of the cross-frame kinematic command mapping under real-world conditions.

- **Network and Communication Architecture:** We tested the ability of the integration bridge to maintain stable, low-latency links with a fleet of three Parrot Bebop 2 UAVs. The 2.4 GHz star topology was evaluated for its ability to sustain telemetry downlink and velocity command uplink at operational ranges of approximately 500 m.
- **RTK Localization:** Outdoor state estimation transitioned from motion capture to a RTK GPS architecture. Each UAV was equipped with an onboard RTK rover antenna receiving differential corrections from a central base RTK station. This was tested to ensure centimeter-level positional accuracy, which is a prerequisite for the high-resolution mapping required by the drift models.

Results- The EHVB campaign served as a rigorous stress-test for the system’s plumbing. While the mission infrastructure was successfully validated, specific airframe instabilities prevented the execution of the full autonomous search. The outcomes are categorized below:

1. **Integration Validation (Success):** The Integration Bridge successfully maintained simultaneous communication with three active UAVs. The fleet achieved consistent RTK "FIX" status, and the ENU-to-NED coordinate transformations functioned correctly under real-world network jitter. This confirms that the communication and localization backbone is ready for large-scale deployment.
2. **Kinematic Execution (Partial Success):** The UAVs successfully executed multi-waypoint geometric patterns (e.g., synchronized squares) commanded via the flight plan from the GCS. However, the agents were following pre-defined trajectories to test command uptake rather than actually performing the IPP methodology.
3. **Operational Bottleneck:** Under outdoor conditions, the default Paparazzi airframe configuration for the Bebop 2 lacked the tuning required for stable hovering. Since the *Confirmation* phase (Section 4.4.4) requires a precise, low-altitude hover to verify a target, it was deemed unsafe to hand over full closed-loop control to the autonomous planner. Consequently, the real-time entropy-driven search remains validated in simulation and indoor environments only.

The field trials successfully bridged the gap between simulation and the physical world by validating the asynchronous network scalability and RTK precision. While the physical tuning of the Bebop 2 limited the execution of the final FSM search states, the campaign proved that the underlying architecture—from the coordination server to the kinematic bridge—is robust and capable of managing a physical swarm in a maritime-representative environment.

6 Conclusion

This thesis addressed the critical challenge of autonomous maritime Search and Rescue (SAR) by investigating how to effectively balance the rapid spatial coverage required to find drifting victims against the strict energy limitations of UAVs. The central research question of this work was: *How can informative path planning be unified with multi-UAV coordination to enable energy-efficient maritime SAR by exploiting distinct flight-regime dynamics?*

To answer this, a centralized, energy-aware coordination framework was developed. By integrating IPP with a continuous, flight-regime-specific power model, the system mathematically optimized UAV flight behavior based on operational necessity. The planner prioritized aerodynamic efficiency during the broad-area search—naturally biasing the swarm toward energy-efficient, fixed-wing cruise speeds—while selectively utilizing the platform’s high-maneuverability capabilities only when required for localized, high-precision tasks.

The framework’s effectiveness was validated across three core contributions:

1. **Flight-Regime-Aware IPP:** Through the implementation of a Receding Horizon Planner utilizing the empirical power curves of the VSQP, the system demonstrated a distinct energetic advantage in adverse and stochastic environments. By integrating continuous aerodynamic costs into the objective function, the planner successfully mitigated the energy penalties associated with high-power maneuvers. While deterministic patterns remain competitive in steady, optimal winds, the energy-aware IPP proved essential for sustaining mission endurance in unpredictable maritime settings by dynamically exploiting favorable wind gradients.
2. **Integrated Multi-UAV Coordination:** The coordination framework enabled the swarm to autonomously balance broad-area exploration with localized target custody through an energetically-informed task allocation logic. The primary result was the successful demonstration of the wind-adjusted effective distance (D_{eff})

metric, which ensured that agent assignments were based on the actual energy-to-arrival cost rather than simple Euclidean proximity. This strategy maintained swarm-level efficiency and prevented task-stranding even under severe environmental disturbances.

- 3. Search Efficiency and Comparative Evaluation:** Large-scale simulations established a highly efficient Coarse Search capacity of approximately 1.1 km² per UAV per hour, while identifying a transit-induced "Battery Wall" at the 10 km scale. Crucially, benchmarking against a static Boustrophedon baseline revealed a fundamental operational trade-off: the framework trades absolute spatial clearing speed for environmental robustness. While geometric paths were faster in optimal weather, the proposed IPP dynamically surfed wind gradients to prevent the severe battery depletion suffered by rigid baselines in adverse crosswinds. However, its reliance on a receding horizon exposed a susceptibility to local minima as map entropy depleted, confirming that the framework is optimally deployed as a wide-area coarse search engine prior to a deterministic fine search hand-off.

Bridging the gap between simulation and the real world, the integrated coordination logic and asynchronous middleware were first successfully verified through indoor flight campaigns at the TU Delft CyberZoo. Subsequent hardware-in-the-loop (HIL) field testing at the Valkenburg Naval Air Base provided a partial validation of the system's outdoor architecture. The centralized server successfully established asynchronous, multi-threaded communication with a physical swarm over a 2.4 GHz network, verifying the RTK localization. However, because the Parrot Bebop 2 drones were unstable during hover, it was deemed unsafe to execute the autonomous tracking and confirmation phases in real-world wind conditions. Thus, while the communication and localization infrastructure was successfully validated outdoors, the full autonomous path planner requires further low-level control refinement before field deployment.

7 Recommendations for Future Work

While the proposed framework successfully unifies energy-aware IPP with swarm coordination, the simulated constraints and physical field tests revealed several avenues for future improvement. Future research should prioritize upgrading the sensing modalities and transitioning the network topology toward a distributed system to move beyond the experimental

star configuration. Furthermore, the final stages of map entropy reduction must be robustified against the spatial myopia and local minima observed during the high-wind simulation campaigns in Phase II.

7.1. Integrating Real Sensors and Edge Computing

To isolate the coordination logic and establish a mathematical baseline, this research utilized a synthetic, proximity-based probabilistic sensor model (Section 3.3). The logical next step is to integrate actual multimodal sensor payloads, such as Electro-Optical/Infrared (EO/IR) cameras. However, the EHVB field trials (Phase IV) confirmed that the current architecture is highly scalable precisely because it relies on the transmission of lightweight telemetry.

Transmitting raw, high-resolution video streams from multiple UAVs would quickly saturate the 2.4 GHz network bandwidth observed at EHVB. To resolve this, future implementations must leverage edge computing. This transition would allow the UAV to perform onboard object detection (e.g., via YOLO [44]), project those detections into global ENU coordinates (X, Y), and transmit only a "micro-packet" of metadata back to the server. This transition would validate the assumption made in Chapter 4—that the centralized planner can remain vision-agnostic by treating sensor outputs as pure probabilistic updates.

7.2. Platform Capabilities and Low-Level Control

The ultimate bottleneck during the EHVB field campaign (Section 5.4) was the physical platform's inability to safely execute commanded trajectories in real-world wind conditions. Because the *Confirmation* phase of the FSM requires a precise, low-altitude hover, the default airframe configurations proved insufficient for full autonomous deployment in unconstrained environments.

Future work must bridge this gap between high-level planning and low-level execution. Furthermore, testing the IPP algorithm on a physical hybrid platform (such as the actual VSQP) would validate the planner's ability to smoothly manage the power-discontinuity between fixed-wing cruise and holonomic hover—a transition currently validated only in the Phase II simulations.

7.3. Transitioning Towards a Decentralized Swarm Architecture

The field validation campaign highlighted the range limitations of the centralized Star Topology assumed in Section 5.3.1. While the integration bridge successfully managed the swarm up to 500 m from the GCS, true maritime SAR operations require scales of several kilometers. Relying on a single centralized

server creates both a communication bottleneck and a single point of failure.

Transitioning toward a decentralized architecture is a critical next step for operational deployment rather than a long-term extension. By moving toward distributed belief mapping—where agents share local map updates via long-range, low-bandwidth RF datalinks (e.g., LoRa)—and auction-based task allocation, the swarm would become resilient to the intermittent connectivity typical of maritime environments. This would relax the perfect communication assumption established in Chapter 3 and allow the D_{eff} metric to be calculated locally by each agent.

7.4. Hybrid Coarse-to-Fine Search Architecture

Phase II results exposed a vulnerability in the "asymptotic tail" of the search process. As global entropy dropped below the 80% threshold, the receding horizon planner frequently suffered from spatial myopia, becoming trapped in local minima when the remaining high-probability cells were isolated in up-wind corners.

To ensure 100% map clearance, future iterations should implement a hybrid coarse-to-fine architecture. The Energy-Aware IPP should be deployed specifically as the coarse search engine to rapidly clear the initial 80% of the environment while maximizing battery life. Once the entropy reduction rate plateaus, the central server should autonomously trigger a deterministic frontier-based exploration algorithm [45]. This hybrid approach would perfectly leverage the aerodynamic efficiency of the proposed IPP for the majority of the flight, while utilizing the mathematical guarantees of geometric exploration to efficiently clear the final, isolated pockets of uncertainty.

References

- [1] Ulrich Dah-Achinanon et al. "Search and rescue with sparsely connected swarms". In: *Autonomous Robots* 47.6 (2023), pp. 849–863. doi: 10.1007/s10514-022-10080-7. url: <https://doi.org/10.1007/s10514-022-10080-7>.
- [2] Soon-Jo Chung et al. "A survey on aerial swarm robotics". In: *IEEE Transactions on Robotics* 34.4 (2018), pp. 837–855.
- [3] J. Horyna et al. "Decentralized swarms of unmanned aerial vehicles for search and rescue operations without explicit communication". In: *Autonomous Robots* 47 (2023), pp. 77–93. doi: 10.1007/s10514-022-10066-5.
- [4] Ignacio Martinez-Alpiste et al. "Search and rescue operation using UAVs: A case study". In: *Expert Systems with Applications* 178 (2021), p. 114937. issn: 0957-4174. doi: <https://doi.org/10.1016/j.eswa.2021.114937>. url: <https://www.sciencedirect.com/science/article/pii/S095741742100378X>.
- [5] Martin Messmer et al. "UAV-Assisted Maritime Search and Rescue: A Holistic Approach". In: *arXiv preprint* 2403.14281 (2024). arXiv:2403.14281.
- [6] Faqihza Mukhlish and John Page. "Alignment Method for Flocking Swarm on Search and Rescue Mission". In: *2023 8th International Conference on Instrumentation, Control, and Automation (ICA)*. 2023, pp. 241–246. doi: 10.1109/ICA58538.2023.10273116.
- [7] Antonio Sojo et al. "Multi-UAV Planning in Search and Rescue Missions using Optimal Search Effort Allocation". In: *2025 International Conference on Unmanned Aircraft Systems (ICUAS)*. North Carolina, USA, May 2025.
- [8] Ming Yan et al. "Task allocation and route planning of multiple UAVs in a marine environment based on an improved particle swarm optimization algorithm". In: *EURASIP Journal on Advances in Signal Processing* 2021.1 (2021), pp. 1–12. doi: 10.1186/s13634-021-00804-9.
- [9] J. Gordon Leishman. *Principles of Helicopter Aerodynamics*. 2nd. Cambridge University Press, 2006.
- [10] Federal Aviation Administration. *Helicopter Flying Handbook (FAA-H-8083-21B)*. U.S. Department of Transportation. Washington, D.C., 2019.
- [11] Mariusz Jacewicz et al. "Quadrotor Model for Energy Consumption Analysis". In: *Energies* 15.19 (2022). issn: 1996-1073. doi: 10.3390/en15197136. url: <https://www.mdpi.com/1996-1073/15/19/7136>.
- [12] Adnan S. Saeed et al. "A survey of hybrid Unmanned Aerial Vehicles". In: *Progress in Aerospace Sciences* 98 (2018), pp. 91–105. issn: 0376-0421. doi: <https://doi.org/10.1016/j.paerosci.2018.03.007>. url: <https://www.sciencedirect.com/science/article/pii/S0376042117302233>.
- [13] Guillaume J.J. Ducard and Mike Allenspach. "Review of designs and flight control techniques of hybrid and convertible VTOL UAVs". In: *Aerospace Science and Technology* 118 (2021), p. 107035. issn: 1270-9638. doi: <https://doi.org/10.1016/j.ast.2021.107035>. url: <https://www.sciencedirect.com/science/article/pii/S1270963821005459>.

- [14] T.M.L. De Ponti et al. "Incremental Nonlinear Dynamic Inversion controller for a Variable Skew Quad Plane". In: *2023 International Conference on Unmanned Aircraft Systems (ICUAS)*. 2023, pp. 241–248. doi: 10.1109/ICUAS57906.2023.10156289.
- [15] Tomaso Maria Luigi De Ponti et al. "Unified-Actuator Nonlinear Dynamic Inversion Controller for the Variable Skew Quad Plane". In: *Journal of Guidance, Control, and Dynamics* 48.5 (May 2025). Technical Notes. doi: 10.2514/1.G008659. URL: <https://doi.org/10.2514/1.G008659>.
- [16] Marija Popovic et al. *An informative path planning framework for UAV-based terrain monitoring*. 2020. arXiv: 1809.03870 [cs.RO]. URL: <https://arxiv.org/abs/1809.03870>.
- [17] Sonia Waharte and Niki Trigoni. "Supporting Search and Rescue Operations with UAVs". In: *Proceedings of the 2010 International Conference on Emerging Security Technologies*. EST '10. USA: IEEE Computer Society, 2010, pp. 142–147. ISBN: 9780769541754. doi: 10.1109/EST.2010.31. URL: <https://doi.org/10.1109/EST.2010.31>.
- [18] Alberto Elfes. "Occupancy Grids: A Probabilistic Framework for Robot Perception and Navigation". PhD thesis. Carnegie Mellon University, 1989.
- [19] Aryo Jamshidpey et al. *Centralization vs. decentralization in multi-robot coverage: Ground robots under UAV supervision*. Aug. 2024. doi: 10.48550/arXiv.2408.06553.
- [20] Ming Yan et al. "Task allocation and route planning of multiple UAVs in a marine environment based on an improved particle swarm optimization algorithm". In: *EURASIP Journal on Advances in Signal Processing* 2021.1 (2021), p. 94. doi: 10.1186/s13634-021-00804-9. URL: <https://doi.org/10.1186/s13634-021-00804-9>.
- [21] Sung-Won Cho et al. "Multi-UAV Coverage Path Planning Based on Hexagonal Grid Decomposition in Maritime Search and Rescue". In: *Mathematics* 10.1 (2022). ISSN: 2227-7390. doi: 10.3390/math10010083. URL: <https://www.mdpi.com/2227-7390/10/1/83>.
- [22] E. Balasubramanian et al. "Optimal energy efficient path planning of UAV using hybrid MACO-MEA* algorithm: theoretical and experimental approach". In: *Journal of Ambient Intelligence and Humanized Computing* 14 (2023), pp. 13847–13867. doi: 10.1007/s12652-022-04098-z. URL: <https://doi.org/10.1007/s12652-022-04098-z>.
- [23] Wenlong Meng et al. "Advances in UAV Path Planning: A Comprehensive Review of Methods, Challenges, and Future Directions". In: *Drones* 9.5 (2025). ISSN: 2504-446X. doi: 10.3390/drones9050376. URL: <https://www.mdpi.com/2504-446X/9/5/376>.
- [24] Denys Datsko et al. "Energy-Aware Multi-UAV Coverage Mission Planning With Optimal Speed of Flight". In: *IEEE Robotics and Automation Letters* 9.3 (Mar. 2024), pp. 2893–2900. ISSN: 2377-3774. doi: 10.1109/lra.2024.3358581. URL: <http://dx.doi.org/10.1109/LRA.2024.3358581>.
- [25] W. H. Andrews et al. "AD-1 Oblique Wing Aircraft Program". In: *SAE Transactions* 89 (1980), pp. 3571–3577. doi: 10.4271/801180. URL: <https://doi.org/10.4271/801180>.
- [26] D. C. van Wijngaarden and B. D. W. Remes. "INDI Control for the ObliqueWing-Quad Plane Drone". In: *13th International Micro Air Vehicle Conference*. Final published version. 2022, pp. 119–126. URL: <https://www.imavs.org/tag/imav2022/>.
- [27] Chandra Chekuri and Martin Pal. "A Recursive Greedy Algorithm for Walks in Directed Graphs". In: *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*. FOCS '05. USA: IEEE Computer Society, 2005, pp. 245–253. ISBN: 0769524680. doi: 10.1109/SFCS.2005.9. URL: <https://doi-org.tudelft.idm.oclc.org/10.1109/SFCS.2005.9>.
- [28] Geoffrey A. Hollinger and Gaurav S. Sukhatme. "Sampling-based robotic information gathering algorithms". In: *The International Journal of Robotics Research* 33.9 (2014), pp. 1271–1287. doi: 10.1177/0278364914533443. eprint: <https://doi.org/10.1177/0278364914533443>. URL: <https://doi.org/10.1177/0278364914533443>.
- [29] Jur Van Den Berg et al. "LQG-MP: Optimized path planning for robots with motion uncertainty and imperfect state information". In: *The International Journal of Robotics Research* 30.7 (2011), pp. 895–913.
- [30] Jonathan D Gammell et al. "Informed RRT*: Optimal sampling-based path planning focused via direct sampling of an admissible ellipsoidal heuristic". In: *2014 IEEE/RSJ International Conference on Intelligent Robots and Systems*. IEEE, 2014, pp. 2997–3004.
- [31] Jonathan D Gammell et al. "Batch informed trees (BIT*): Sampling-based optimal planning via the heuristically guided search of implicit random geometric graphs". In: *2015 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2015, pp. 3067–3074.

- [32] Brady Moon et al. "TIGRIS: An Informed Sampling-based Algorithm for Informative Path Planning". In: *2022 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, Oct. 2022, pp. 5760–5766. doi: 10.1109/iros47612.2022.9981992. URL: <http://dx.doi.org/10.1109/IROS47612.2022.9981992>.
- [33] Roman Marchant et al. "Sequential Bayesian optimisation for spatial-temporal monitoring". In: *Proceedings of the Thirtieth Conference on Uncertainty in Artificial Intelligence*. UAI'14. Quebec City, Quebec, Canada: AUAI Press, 2014, pp. 553–562. ISBN: 9780974903910.
- [34] Gerhard K. Kraetzschmar et al. "Probabilistic quadrees for variable-resolution mapping of large environments". In: *IFAC Proceedings Volumes 37.8 (2004)*. IFAC/EURON Symposium on Intelligent Autonomous Vehicles, Lisbon, Portugal, 5-7 July 2004, pp. 675–680. ISSN: 1474-6670. doi: [https://doi.org/10.1016/S1474-6670\(17\)32056-6](https://doi.org/10.1016/S1474-6670(17)32056-6). URL: <https://www.sciencedirect.com/science/article/pii/S1474667017320566>.
- [35] Øyvind Breivik et al. "Wind-induced drift of objects at sea: The leeway field method". In: *Applied Ocean Research* 33.2 (Apr. 2011), pp. 100–109. ISSN: 0141-1187. doi: 10.1016/j.apor.2011.01.005. URL: <http://dx.doi.org/10.1016/j.apor.2011.01.005>.
- [36] Ragesh K. Ramachandran et al. "Information Correlated Lévy Walk Exploration and Distributed Mapping Using a Swarm of Robots". In: *IEEE Transactions on Robotics* 36.5 (Oct. 2020), pp. 1422–1441. ISSN: 1941-0468. doi: 10.1109/tro.2020.2991612. URL: <http://dx.doi.org/10.1109/TRO.2020.2991612>.
- [37] Siyuan Wu et al. *Decentralized Multi-Agent Trajectory Planning in Dynamic Environments with Spatiotemporal Occupancy Grid Maps*. 2024. arXiv: 2404.15602 [cs.RO]. URL: <https://arxiv.org/abs/2404.15602>.
- [38] DJI Enterprise. *SAR Altitude Guide Results: Standing Up, 50m, Slanted*. Accessed: 2026-02-04. DJI Enterprise Insights. 2022. URL: <https://enterprise-insights.dji.com/saraltitudeguideresult10>.
- [39] DJI Enterprise and SkyBound Rescuer. *Drone SAR Altitude Guide*. Accessed: 2026-02-04. DJI Enterprise Insights. 2022. URL: <https://enterprise-insights.dji.com/saraltitudeguide>.
- [40] DJI Enterprise. *Mavic 3 Enterprise Series: Technical Specifications*. Accessed: 2026-02-04. SZ DJI Technology Co., Ltd. 2024. URL: <https://enterprise.dji.com/mavic-3-enterprise/specs>.
- [41] Skydio, Inc. *Skydio X10 Technical Specifications*. Accessed: 2026-02-04. Skydio, Inc. 2024. URL: <https://www.skydio.com/x10/technical-specs>.
- [42] B. O. Koopman. "The Theory of Search. II. Target Detection". In: *Operations Research* 4.5 (1956), pp. 503–531. doi: 10.1287/opre.4.5.503. eprint: <https://doi.org/10.1287/opre.4.5.503>. URL: <https://doi.org/10.1287/opre.4.5.503>.
- [43] Howie Choset. "Coverage of Known Spaces: The Boustrophedon Cellular Decomposition". In: *Autonomous Robots* 9 (2000), pp. 247–253. doi: 10.1023/A:1008958800904.
- [44] Jianhao Xu et al. "YoloOW: A Spatial Scale Adaptive Real-Time Object Detection Neural Network for Open Water Search and Rescue From UAV Aerial Imagery". In: *IEEE Transactions on Geoscience and Remote Sensing* 62 (2024), pp. 1–15. doi: 10.1109/TGRS.2024.3395483.
- [45] Brian Yamauchi. "A frontier-based approach for autonomous exploration". In: *Proceedings 1997 IEEE International Symposium on Computational Intelligence in Robotics and Automation CIRA'97. 'Towards New Computational Principles for Robotics and Automation'*. IEEE. 1997, pp. 146–151.

Part III

Closure

Conclusion

This thesis addressed the critical challenge of autonomous maritime SAR by investigating how to effectively balance the rapid spatial coverage required to find drifting victims against the strict energy limitations of UAVs. To solve this, a centralized, energy-aware coordination framework was developed, integrating IPP with flight-regime-specific energy dynamics.

The core objective of this research was driven by one primary research question, supported by three sub-questions. The findings and corresponding conclusions for each are detailed below.

Sub-Question 1: What is the minimum swarm size required to explore a bounded maritime environment and achieve a high victim discovery rate within strict endurance and temporal limits?

Through extensive Monte Carlo simulations (Phase I), the research established that the minimum swarm size is strictly bounded by a transit-induced "Battery Wall" and volumetric sweep limits. For large-scale operations (e.g., a $10,000 \times 10,000$ m area), individual energy drain outpaces spatial coverage if the swarm is too small; a theoretical minimum of $N \approx 125$ VSQP agents is required to successfully achieve the 80% "Coarse Search" entropy reduction threshold over 100 km^2 before battery depletion. By prioritizing this 80% threshold and avoiding the inefficient asymptotic tail of the search process, the framework achieved a highly efficient effective search rate of approximately 1.1 km^2 per UAV per hour.

Furthermore, when optimizing specifically for a high victim discovery rate ($\geq 90\%$ recall), the evaluation exposed a critical inversion in optimal altitude strategies based on the operational scale. While a heterogeneous "Mixed" altitude swarm (combining 35 m and 50 m search altitudes) proved highly efficient for smaller regions (up to 2000 m) by accelerating localized confirmations, it suffered severe performance degradation in massive domains. For expansive areas (5000 m and beyond), operations are strictly bottlenecked by the collective field of view. In these scenarios, a homogeneous, high-altitude (uniform) swarm is mandated to maximize wide-area coverage and ensure the minimum number of agents can successfully achieve the target recall.

Sub-Question 2: How does embedding an aerodynamic power model into the objective function balance the trade-off between global entropy reduction and endurance?

Benchmarking the Energy-Aware IPP against traditional Boustrophedon coverage (Phase II) revealed a fundamental operational trade-off: the proposed framework trades absolute spatial clearing speed for environmental robustness. By coupling the IPP objective function with a continuous power model, the swarm mathematically penalized energy-intensive maneuvers. While rigid geometric paths proved faster in perfect weather, the energy-aware IPP actively utilized wind gradients, avoiding the severe battery depletion suffered by the baseline in adverse crosswinds. Sensitivity analysis identified that setting the energy penalty weight to $\lambda \approx 0.5$ established the operational "sweet spot," effectively balancing the need for rapid entropy reduction with the necessity of preserving battery life in stochastic environments.

Sub-Question 3: How can the distinct flight-regime energy dynamics of a hybrid UAV (such as the VSQP) be explicitly exploited during different mission phases to maximize search endurance?

The distinct energy dynamics of hybrid platforms were successfully exploited by coupling the local IPP planner with a global Finite State Machine (FSM). The framework naturally biased the swarm toward highly efficient, fixed-wing cruise speeds during the high-altitude *Exploration* phase. Energy-intensive holonomic flight (hovering) was strictly reserved for the localized, low-altitude *Confirmation* phase. Furthermore, the transition between these phases was governed by a newly introduced wind-adjusted effective distance (D_{eff}) metric. This ensured that task allocation—specifically the assignment of a UAV to track a detected target—was based on the actual aerodynamic energy-to-arrival cost rather than simple Euclidean proximity, thereby preserving swarm-level endurance.

Primary Research Question: How can informative path planning be unified with multi-UAV coordination to enable energy-efficient maritime SAR by exploiting distinct flight-regime dynamics?

This thesis demonstrates that IPP and multi-UAV coordination can be successfully unified through a centralized, energy-aware architecture that treats the physical reality of flight as a core constraint of intelligent search. By offloading complex advection-diffusion mapping to a Ground Control Station and utilizing asynchronous middleware, the framework allowed physical UAVs to act as highly efficient kinematic executors.

The physical field trials (Phase III and IV) at the TU Delft CyberZoo and Valkenburg Naval Air Base confirmed the structural readiness of this architecture. While airframe instabilities of the Bebop 2 platform

limited the execution of full autonomous hovering outdoors, the campaigns successfully validated the real-time coordination logic, the asynchronous network scalability, and the precision of the RTK integration. Ultimately, this research proves that shifting path planning from purely spatial algorithms to flight-regime-aware models provides a critical energetic advantage, laying the groundwork for the deployment of fully autonomous, hybrid UAV swarms in unpredictable maritime environments.

A Declaration of Generative AI Usage

During the research, development, and writing of this thesis, Google Gemini was used as a supplementary AI assistant. To ensure full academic transparency in accordance with university guidelines, the specific applications of this tool throughout the project lifecycle are detailed below:

- **Software Implementation and Debugging:** Generative AI was utilized as a pair-programming assistant to accelerate the development of the simulation environment. This included translating my mathematical cost functions into Python, optimizing data structures for Monte Carlo simulations, and troubleshooting runtime errors and integration bugs.
- **Data Analysis and Visualization:** After conducting flight tests and running simulations, the AI assisted in drafting Python scripts (using libraries such as Pandas and Matplotlib) to process large datasets. It also served as a sounding board for structuring and formatting the analytical figures to best communicate my findings.
- **Literature Study:** Gemini was used as an interactive search and summarization tool to assist in navigating academic literature. It helped brainstorm conceptual frameworks which I then independently evaluated and synthesized to structure the methodology presented in this thesis.

B Test Set-Up and Execution Protocol

To ensure the reproducibility of the physical experiments, this appendix details the practical hardware configurations, network routing, and chronological flight execution protocols utilized during both the indoor TU Delft CyberZoo verification and the outdoor Valkenburg Naval Air Base (EHVB) field trials.

B.1. Preliminary Hardware and Software Inventory

Before deploying the swarm, the following core components were established as the baseline Hardware-in-the-Loop (HIL) architecture:

- **Ground Control Station (GCS):** A HP ZBook Studio G5 laptop with an Intel Core i7-9750H CPU @ 2.60GHz \times 12 was running the central Python coordination script (Informative Path Planning, FSM, and advection-diffusion mapping), and the Paparazzi GCS v7.0.
- **Swarm Router:** A Linksys WRT1200AC router. This acted as the central communication hub for the entire fleet, handling all code uploads, telemetry downlinks, and command uplinks.
- **Aerial Platforms:** A homogeneous fleet of Parrot Bebop 2 quadrotors modified with a u-Blox M8P GNSS chip, capable of RTK. Each UAV was flashed with the Paparazzi open-source autopilot.
- **Software:** The Paparazzi UAV software, Python 3.12, and pprzlink for communication between the GCS and the UAVs.
- **Power Supply:** For outdoor testing, an external portable power station is required to supply continuous power to the Swarm Router and the RTK Base Station.

B.2. Network and Infrastructure Configuration

The physical network topology and localization infrastructure required distinct configurations depending on whether the test was conducted indoors or outdoors. In both scenarios, the GCS laptop was connected directly to the Swarm Router via a hardwired Ethernet connection to ensure minimum latency and maximum reliability.

B.2.1. Indoor Configuration (TU Delft CyberZoo)

Because the CyberZoo is a GNSS-denied environment, localization relied entirely on an external OptiTrack Motion Capture system. The setup procedure was as follows:

1. **Network Bridging:** To ensure minimal latency and reliable data ingestion, the Ground Control Station (GCS) laptop was hardwired to the Swarm Router via an Ethernet connection. The Swarm Router was subsequently bridged to the internal CyberZoo router using a secondary Ethernet link (the orange cable). This physical bridge allowed the Swarm Router to receive the high-frequency localization packets broadcast by the MoCap system. This set-up can be seen in Fig. B.1.

2. **OptiTrack Initialization (UAVs):** Reflective markers were physically attached to each Bebop 2. Within the OptiTrack tracking software, a "Rigid Body" was created for each UAV.
3. **Victim Proxy Tracking (Cardboard):** To safely simulate a surface-level target without introducing severe rotor-downwash interference or mid-air marker occlusion, a flat piece of cardboard was utilized as a physical proxy. Reflective OptiTrack markers were attached to the cardboard, and a tow cable was affixed to allow manual pulling, physically simulating ocean drift. A corresponding Rigid Body was created in OptiTrack so the centralized server could track its dynamic motion across the arena. This configuration can be seen in Fig. B.2
4. **NatNet Configuration:** It was critical that the Rigid Body in the OptiTrack software was named exactly after the UAV's Paparazzi AC_ID (e.g., 121, 219, 123). This specific naming convention ensured that the NatNet protocol correctly parsed the coordinates and injected the rigid-body data directly into the corresponding UAV's INS to emulate GPS. More information on how to do this can be found here: <https://github.com/tudelft/UnifiedMocapRouter>.



Figure B.1: Hardware networking setup at the Ground Control Station. The laptop executing the centralized planner is connected via Ethernet to the Swarm Router. The Swarm Router is then bridged to the CyberZoo's internal network (via the orange Ethernet cable) to ingest the high-frequency OptiTrack telemetry.




Figure B.2: Physical flight setup inside the CyberZoo. The two active Bebop 2 exploration agents are visible alongside the passive cardboard victim proxy. Both the UAVs and the cardboard are equipped with reflective OptiTrack markers for precise localization, and the proxy is tethered to a cable to simulate continuous environmental drift during the multi-agent search.

B.2.2. Outdoor Configuration (EHVB Field Tests)

During the outdoor trials at Valkenburg Naval Air Base we relied on Real-Time Kinematic (RTK) GPS. The take-off set-up can be seen in figure B.3, which adds the RTK base station to the set-up compared to the CyberZoo set-up.

1. **Base Station Initialization:** An RTK base station antenna (GGSFTP.50.7.A.08) was connected to the GCS. Before flight operations could commence, the base station had to be initialized and surveyed-in using the u-center GNSS evaluation software to establish a highly accurate absolute "HOME" coordinate. The steps that need to be followed for the set-up can be found here: <https://wiki.paparazziuav.org/wiki/Disco>
2. **Correction Broadcast:** Once initialized, differential corrections are routed through the GCS and broadcast over the Swarm Router to the RTK rover modules mounted on each Bebop 2, allowing the fleet to achieve centimeter-level positional fixes.



Figure B.3: Three UAVs in their take-off position accompanied with the RTK base-antenna on a stable platform.

B.3. Flight Execution Sequence

Regardless of the testing environment, the procedure for launching the autonomous swarm followed a strict chronological sequence to ensure safety and system stability:

1. **Code Upload:** While the UAVs were powered on and connected to the Swarm Router, the compiled Paparazzi firmware (containing the specific AC_ID, flight plan and airframe file) was uploaded to each UAV via the network.
2. **Physical Placement:** The UAVs were carried into the flight arena (the CyberZoo or the EHVB field) and placed on the ground with sufficient spatial separation to prevent aerodynamic interference or collision during takeoff.
3. **Sequential Takeoff:** Using the Paparazzi Center GCS interface, the UAVs were commanded to take off one by one. Upon reaching their default altitude, they automatically entered a "Standby" hovering state.
4. **Guided Mode Activation:** Once the entire fleet was airborne and stable, the flight mode for all UAVs was switched to Guided Mode, which explicitly allows the autopilot to bypass its internal waypoint navigation and accept external velocity commands.
5. **Planner Execution:** Finally, the custom Python coordination script was executed on the GCS laptop. The script immediately established connections over the Ivy bus, read the current telemetry, and began transmitting the optimized continuous velocity commands.

For the EHVB test, we never got past a stable hover, and we could thus not move to step 5. This is because there is a magnetometer heading issue with the Bebop 2 on Paparazzi. These issues need to be fixed, before performing a swarming test outside.

C Open-Source Contributions to Paparazzi UAV

The framework developed in this research heavily leverages and extends the Paparazzi open-source UAV autopilot system. To ensure the reproducibility of both the Monte Carlo simulations and the physical flight

tests, all custom airframes, flight plans, and the Python-based centralized planner have been integrated into a dedicated branch of the TU Delft Paparazzi repository.

This appendix serves as the technical documentation and execution guide for future researchers wishing to utilize or extend the Energy-Aware IPP framework.

C.1. Repository Information

- **Host Organization:** TU Delft Micro Air Vehicle Lab (MAVLab)
- **Repository URL:** <https://github.com/tudelft/paparazzi>
- **Target Branch:** VSQP_NAV_OK

C.2. Directory Structure and Custom Modules

The repository contains two primary areas of contribution:

C.2.1. Paparazzi Configurations (The SDB Directory)

The hardware configurations and waypoint logic required to interface with the Bebop 2 and the simulated VSQP platforms are housed under dedicated SDB (Shivesh Damian Bhawan) directories within the Paparazzi configuration tree:

- **Airframes (conf/airframes/SDB/):** Contains the XML configuration files specifying the kinematic limits, PID tuning, and active telemetry modules for the surrogate platforms. Crucially, these files instantiate the guided module, which allows the autopilot to accept external velocity vectors. Note that the airframe files with animal names are for the Bebop 2s. They are sufficient for flying in the CyberZoo with their given modifications specified in the files, but do not use them outside because of the magnetometer issue! Also when using a rotational wing file, it is advised to ask MavLab staff for the latest version.
- **Flight Plans (conf/flight_plans/SDB/):** Contains the XML state-machine logic for the UAVs. These flight plans strip away complex geometric navigation, reducing the UAV's native behavior to a simple sequence: Takeoff → Standby → Guided Mode (listening to the Python planner). The CyberZoo flight plans are sufficient, but the EHVB ones for the Bebop 2 need additional safety measures; hard kill at the outer boundaries. It is advised to discuss this with MavLab staff.

C.2.2. The Centralized Planner (sw/ground_segment/python/Occupancy_Map/)

The entirety of the high-level coordination framework is contained within the root-level Occupancy_Map folder. This folder acts as the "Ground Control Station" logic and it also contains critical simulations used in the paper:

- **The IPP and Mapping integration:** Python scripts executing the Receding Horizon Planner, calculating the continuous aerodynamic power model ($P(\|v_a\|)$), and selecting optimal spatial trajectories, while also keeping and updating the Dynamic Mapping logic with the Conical Diffusion Model. `multi_uav_cyberzoo_faster` was used for the CyberZoo tests. It runs a simulation if nothing is streamed via the Ivy bus. `muti_uav_sim_faster_EHVB` is the Valkenburg version, which was only used in simulation, and not tested yet outside.
- **Paparazzi Bridge:** One ivy-python middleware script responsible for unpacking ROTORCRAFT_FP telemetry and broadcasting GUIDED_SETPOINT_NED commands, is given as example for communication.
- **Monte Carlo Scripts:** Any script containing "Monte Carlo" in its name is used for the Monte Carlo simulations. The files with the added "visualization" then process the data of the Monte Carlo simulations and make the plots.
- **Baseline Scripts:** Anything containing "Baseline" is for the baseline comparison. The one with "optimal wind direction" is for determining the energy consumption for a configuration in all wind directions.
- **Data Parsing:** `mutli_uav_real_test_plot` is used to parse data from flights and to create and visualize its 3D path. `map_and_sat` is used to create a figure of the Valkenburg test area with an overlapping probability map. `data_parsing_energy_model_emp` is used to create

C.3. Execution Guide

To deploy the Energy-Aware swarm, follow this chronological sequence using the pre-configured Paparazzi sessions:

C.3.1. Build the Paparazzi Environment

Clone the TU Delft repository, checkout the specific branch, and compile the base system:

```
git clone https://github.com/tudelft/paparazzi.git
cd paparazzi
git checkout VSQP_NAV_OK
echo bubba && git submodule sync && git submodule update && make clean && make
```

C.3.2. Launch Paparazzi Center and Select Session

Start the graphical GCS interface:

```
./paparazzi
```

Within the Paparazzi Center GUI, first build the airframe files you want in your swarm, one-by-one. Then go to the "Operation" tab and select the appropriate custom Session from the top drop-down menu depending on the deployment environment (from my custom made control panel `userconf/SDB/sdb_control_panel.xml`):

- **SDB EHVB Simulation:** Used for software-in-the-loop testing.
- **sdb control cyberzoo:** Used for physical indoor flights with the OptiTrack MoCap system.
- **SwarmyRTK:** Used for physical outdoor field tests relying on the RTK GPS architecture.

C.3.3. Initialize the Swarm

Once the *GCS* and *Server* modules are running:

1. Use the GCS interface to power on the aircraft and command a sequential takeoff.
2. Wait for the aircraft to reach the default holding altitude and stabilize.
3. Switch the flight block to **Guided Mode**. The aircraft can now receive external velocity commands over the Ivy bus.

C.3.4. Execute the Python Planner

Open the python IPP script, and execute the main coordination script. The script will immediately bind to the Ivy bus, ingest the current UAV telemetry, and begin transmitting the continuous Energy-Aware IPP velocity commands to the swarm, while also receiving and integrating real-time positional and velocity data.