Available online at www.easd.procedia.org
EASD Procedia EURODYN (2026) 2081-2104

EASD Procedia

EURODYN 2026
XIII International Conference on Structural Dynamics
M. Beer, C. Papadimitriou, H. Jensen, J. Chen, M. Broggi, J. Behrensdorf (eds.)
Hannover, Germany, 27.09.–01.10 2026

UNCERTAINTY-AWARE MULTI-OBJECTIVE RENDEZVOUS RE-PLANNING FOR MARINE VEHICLES USING A PHYSICS-BASED DIGITAL MODEL

M. Laurenza[1], G. Pepe[2], D. Tonazzi[2], H. M. Gomes[3], and W. J. P. Casas[3]

[1]Department of Engineering, Niccolò Cusano University, Rome, Italy
e-mail: maicol.laurenza@unicusano.it

[2]Department of Mechanical and Aerospace Engineering, Sapienza University of Rome, Rome, Italy
e-mail: gianluca.pepe@uniroma1.it; davide.tonazzi@uniroma1.it

[3]Department of Mechanical Engineering, Universidade Federal do Rio Grande do Sul, Porto Alegre, Brazil
e-mail: herbert@mecanica.ufrgs.br; walter.casas@ufrgs.br

Keywords: Multi-objective optimization, uncertainty quantification, digital twin, autonomous marine vehicles, rendezvous planning, re-optimization.

Abstract. *This paper presents a digital-twin-driven methodology for uncertainty-aware rendezvous planning between a mothership and a marine drone operating in time-varying metocean conditions and in the presence of fixed obstacles. The digital twin combines planar maneuvering dynamics, propulsion-efficiency models, and evolving current, wind, and wave fields. Environmental uncertainty is introduced through stochastic perturbations of intensity and direction with respect to nominal forecast values. A multi-objective optimization based on Particle Swarm Optimization (PSO) constructs the Pareto front between mothership energy, drone energy, and rendezvous time, while a visibility-graph-based layer guarantees obstacle avoidance. The work focuses on the two-vehicle problem, regarded as an elementary building block for broader multi-agent recovery missions, so as to analyze in a physically interpretable way the stability of the optimal geometry under uncertainty. On this basis, a periodic re-optimization procedure is introduced, through which open-loop execution is compared with digital-twin-informed replanning under perturbed environmental realizations. The numerical campaign assesses the deterministic advantage with respect to geometric baselines based on centroid and center of mass, as well as the benefit of periodic replanning in terms of energy cost, rendezvous time, and geometric route dispersion, using indicators of homotopy consistency, path dispersion, and rendezvous-point dispersion. The results show that periodic forecast updates preserve the topological congruence of the solutions and, with respect to frozen-plan execution, provide an average reduction of* 34% *in total energy and* 14% *in rendezvous time, confirming the potential of the digital twin to support robust and computationally sustainable rendezvous decisions in uncertain marine operations.*

ISSN:2311-9020 © 2026 The Authors. Published by EASD Procedia.
Peer-review under responsibility of the Organizing Committee of EURODYN 2026.
doi: 10.47964/1326.10810.204

1 INTRODUCTION

Autonomous Surface Vehicles (ASVs) and Unmanned Surface Vehicles (USVs) are increasingly adopted in maritime operations requiring long endurance, reduced human exposure, and adaptive decision-making capabilities. Typical applications include environmental monitoring, inspection, surveillance, and support for coastal and offshore operations [1, 2]. In many of these missions, small marine drones offer high maneuverability and reduced operating costs, but their limited onboard energy makes periodic interaction with a larger support platform necessary. Mothership–drone architectures equipped with launch and recovery systems (LARS) provide a practical solution, enabling deployment, recovery, recharging, maintenance, and data transfer during the mission [3, 4].

In this context, rendezvous planning becomes a central operational problem. The objective is not only to select a meeting point, but also to determine when and how the mothership and the drone should reach it while accounting for available energy, propulsion efficiency, obstacles, and time-varying metocean conditions. This makes the problem different from classical geometric formulations of rendezvous or aggregation, in which the vehicle model and the environment are often simplified. In marine operations, the same ground-relative trajectory may lead to very different energy costs depending on the relative-to-water velocity, the current direction, the apparent wind, and the efficiency region of the propulsion system.

Several research directions are relevant to this problem. General path planning methods for ASVs include graph search, sampling-based planning, potential-field approaches, and metaheuristic optimization, and have recently been discussed in the context of autonomous marine surface vehicles [1]. Specific rendezvous formulations have been proposed for multiple autonomous marine vehicles operating in current fields [5], for persistent tasks requiring periodic recharging [6], and for AUVs interacting with mobile charging stations in dynamic currents [7]. These works demonstrate the importance of synchronization, endurance, and environmental awareness, but they often rely on simplified kinematic or energetic descriptions. As a result, energy is frequently treated as proportional to distance, travel time, or prescribed speed profiles, rather than being computed through a propulsion-aware model that includes hydrodynamic loads, relative-to-water velocity, actuator limits, and efficiency maps.

A second group of works addresses cooperative rendezvous and formation coordination among multiple vehicles. Recent methods combine optimal assignment and trajectory generation for USV formation rendezvous [8, 9], while learning-based approaches, such as graph-attention-enhanced multi-agent reinforcement learning, have shown promising results in cooperative obstacle avoidance and formation stabilization [10]. These methods are relevant to coordination and real-time decision-making, but their focus is usually formation assembly or cooperative motion control rather than energy-constrained mothership–drone recovery. Moreover, vehicle dynamics are often represented at a kinematic level, and the influence of time-varying metocean forecasts on the stability of the final rendezvous geometry is not their primary concern.

A third relevant direction concerns uncertainty-aware planning and digital twins for maritime autonomy. Uncertainty-aware guidance and planning methods explicitly account for estimation errors, unknown obstacle motion, sensor noise, wind, and model uncertainty [11, 12]. In parallel, digital-twin-based approaches for autonomous maritime systems have been proposed to support model updating, simulation-based decision making, and adaptive control under uncertain operating conditions [13–15]. These contributions show the growing importance of model updating and uncertainty management in marine autonomy. However, compact case studies are still needed in which a physics-based digital twin is used not only to compute a nominal

route, but also to quantify how repeated forecast updates affect the topology, dispersion, and repeatability of the optimized rendezvous path.

This paper isolates the elementary mothership–drone rendezvous problem and turns it into a physically interpretable test case for replanning under uncertainty. Reducing the problem to a single mothership–drone pair removes the combinatorial variability associated with multi-agent recovery order, so that the observed variability can be attributed directly to environmental perturbations and replanning policy. On this elementary block, the paper develops the two elements that are specific to the present study: a periodic, digital-twin-informed re-optimization cycle that compares open-loop execution with replanning on the same perturbed environmental realizations, and a set of homotopy-consistency and path-dispersion metrics that quantify the geometric and topological stability of the rendezvous solution under forecast mismatch.

The underlying propulsion-aware planner, i.e. the vehicle dynamic–propulsive model, the visibility-graph obstacle avoidance, and the PSO-based multi-objective search, is inherited from the authors' previous framework for scalable, energy-efficient rendezvous optimization in multi-agent marine applications [16, 17], reorganized here around the two-vehicle case. That framework targeted fleet-level scalability and recovery sequencing; the present paper instead uses the same physical engine to ask a different question, namely how the geometry and performance of a single rendezvous degrade and recover under forecast uncertainty, while preserving the connection with broader cooperative recovery architectures [9, 18].

The contribution of the paper is therefore threefold. First, a periodic re-optimization strategy is introduced, making it possible to compare open-loop execution with digital-twin-informed replanning on the same perturbed environmental realizations. Second, a targeted uncertainty-quantification campaign is developed, based on repeated solver launches and Monte Carlo perturbations of current and wind fields, and completed by a final multi-obstacle extension designed to highlight the possible emergence of homotopy-class transitions; besides the classical indicators of energy and time, the analysis uses homotopy-consistency, path-dispersion, and rendezvous-point dispersion metrics to evaluate the geometric and operational stability of the solutions under uncertainty. Third, this uncertainty-aware layer is built on a propulsion-aware multi-objective rendezvous planner, driven by a digital twin, for a mothership and a marine drone operating in a time-varying environment with fixed obstacles.

The remainder of the paper is organized as follows. Section 2 introduces the physics-based digital model and the formulation of the rendezvous problem. Section 3 describes the re-optimization cycle under uncertainty and defines the repeatability and robustness indicators. Section 4 presents the numerical campaign and discusses the results, including the deterministic, repeatability, frozen-plan, periodic-replanning, and final multi-obstacle cases. Conclusions and future developments are finally reported in Section 5.

2 PHYSICS-BASED DIGITAL MODEL AND RENDEZVOUS FORMULATION

2.1 Vehicle and environment model

Section 2 summarizes the propulsion-aware planner inherited from [16, 17], reorganized around the two-vehicle case, to the extent needed to make the present paper self-contained; the methodological contribution specific to this paper — periodic re-optimization under uncertainty and the associated robustness metrics — is introduced in Section 3. The physical model used in the planner is recalled here in compact form to make the rendezvous problem formulation self-contained. The description follows the standard Fossen framework for planar maneuvering of marine vehicles [19], already adopted in the reference work [16], but is reorganized around the

elementary case considered in this paper: one mothership and one drone. The goal is to recall, in a compact yet self-contained way, the elements of the digital model required to understand the present study: the generation of rendezvous trajectories, the interaction with current, wind, and waves, the propulsive energy evaluation, and the verification of kinematic, energetic, and geometric constraints.

The calculation chain used at each optimizer evaluation is as follows. The solver proposes a macroscopic choice, namely the mothership direction and speed and the meeting time. From this choice, the rendezvous point and the nominal drone kinematics are derived. If the direct route is not compatible with the obstacles, it is replaced by an admissible polyline.

Each route is then divided into segments, each traversed with local direction, speed, and attitude. On each segment, the twin samples current, wind, and waves, computes the generalized force required to maintain steady motion, and solves a propulsion allocation problem to obtain the minimum power compatible with actuator limits. The total energy is finally obtained by summing the segment powers over the corresponding travel times.

Two vehicles are considered, denoted by $k \in \mathcal{K} = \{M, D\}$, where M denotes the mothership and D the marine drone. The motion is described in the horizontal plane through a local inertial NED reference frame. The planar position of vehicle k is $\boldsymbol{p}_k = [x_k, y_k]^\top$, while $\boldsymbol{\eta}_k = [x_k, y_k, \psi_k]^\top$ collects position and yaw angle. The velocity vector expressed in the body-fixed frame is $\boldsymbol{\nu}_k = [u_k, v_k, r_k]^\top$, with u_k the surge velocity, v_k the sway velocity, and r_k the yaw rate. The rotation from the body-fixed frame to the NED frame is described by the planar matrix

$$\mathbf{R}(\psi_k) = \begin{bmatrix} \cos\psi_k & -\sin\psi_k \\ \sin\psi_k & \cos\psi_k \end{bmatrix}, \qquad \mathbf{J}(\psi_k) = \begin{bmatrix} \mathbf{R}(\psi_k) & \mathbf{0} \\ \mathbf{0}^\top & 1 \end{bmatrix}. \tag{1}$$

In the digital twin, the planar maneuvering dynamics are described by

$$\begin{aligned} \dot{\boldsymbol{\eta}}_k &= \mathbf{J}(\psi_k)\boldsymbol{\nu}_k, \\ \mathbf{M}_k\dot{\boldsymbol{\nu}}_k + \mathbf{C}_k(\boldsymbol{\nu}_{r,k})\boldsymbol{\nu}_{r,k} + \mathbf{D}_k(\boldsymbol{\nu}_{r,k})\boldsymbol{\nu}_{r,k} &= \boldsymbol{\tau}_k + \boldsymbol{\tau}_k^{\text{wind}} + \boldsymbol{\tau}_k^{\text{wave}}, \end{aligned} \tag{2}$$

where $\mathbf{M}_k$ includes rigid-body mass and added mass, $\mathbf{C}_k$ collects Coriolis, centripetal, and added-damping terms, $\mathbf{D}_k$ represents linear and nonlinear hydrodynamic damping, $\boldsymbol{\tau}_k$ is the vector of forces and moments generated by the propulsion system, $\boldsymbol{\tau}_k^{\text{wind}}$ is the aerodynamic load due to apparent wind, and $\boldsymbol{\tau}_k^{\text{wave}}$ is the wave contribution. The current does not appear as an additive external force, but enters the model through the relative-to-water velocity. Denoting by $\boldsymbol{V}_k = \dot{\boldsymbol{p}}_k$ the ground-relative velocity and by $\boldsymbol{V}_c(\boldsymbol{p}, t)$ the current field, the relative velocity in the inertial plane is

$$\boldsymbol{V}_{r,k}(\boldsymbol{p}_k, t) = \boldsymbol{V}_k - \boldsymbol{V}_c(\boldsymbol{p}_k, t), \tag{3}$$

while the corresponding relative velocity in the body-fixed frame is

$$\boldsymbol{\nu}_{r,k} = \begin{bmatrix} \mathbf{R}^\top(\psi_k)\boldsymbol{V}_{r,k} \\ r_k \end{bmatrix}. \tag{4}$$

This choice is essential for energy evaluation: two trajectories with the same ground-relative speed may require very different power if the alignment with the current changes.

The hydrodynamic part is therefore evaluated on the relative velocity rather than on the absolute velocity. In particular, the mass matrix can be written as $\mathbf{M}_k = \mathbf{M}_{\text{RB},k} + \mathbf{M}_{\text{A},k}$, the same holding for $\mathbf{C}_k$, while the hydrodynamic damping is decomposed into a linear part and a nonlinear part:

$$\mathbf{D}_k(\boldsymbol{\nu}_{r,k})\boldsymbol{\nu}_{r,k} = \mathbf{D}_{\text{L},k}\boldsymbol{\nu}_{r,k} + \mathbf{D}_{\text{NL},k}(\boldsymbol{\nu}_{r,k})\boldsymbol{\nu}_{r,k}. \tag{5}$$

The environmental twin provides, for each position and instant, the fields required to evaluate hydrodynamic resistance, apparent wind, and wave-induced loads. The current and wind fields are described through intensities and directions perturbed with respect to nominal large-scale forecast values:

$$\begin{aligned} U_c(\boldsymbol{p},t;\varepsilon_{U_c}) &= \bar{U}_c(\boldsymbol{p},t)\big(1+\varepsilon_{U_c}\big), & \chi_c(\boldsymbol{p},t;\varepsilon_{\chi_c}) &= \bar{\chi}_c(\boldsymbol{p},t)+\varepsilon_{\chi_c}, \\ U_w(\boldsymbol{p},t;\varepsilon_{U_w}) &= \bar{U}_w(\boldsymbol{p},t)\big(1+\varepsilon_{U_w}\big), & \chi_w(\boldsymbol{p},t;\varepsilon_{\chi_w}) &= \bar{\chi}_w(\boldsymbol{p},t)+\varepsilon_{\chi_w}, \end{aligned} \tag{6}$$

where $\bar{U}_c$, $\bar{U}_w$, $\bar{\chi}_c$, and $\bar{\chi}_w$ are the nominal intensity and direction values, functions of position and time, while ε_{U_c} and ε_{U_w} are relative intensity perturbations, and ε_{χ_c} and ε_{χ_w} are angular direction offsets. These perturbations are interpreted as random parameters associated with the single Monte Carlo realization; the perturbed field therefore preserves the spatio-temporal structure of the nominal forecast while modifying its intensity and orientation in a stochastic way. The corresponding vector fields are reconstructed as

$$\begin{aligned} \boldsymbol{V}_c(\boldsymbol{p},t;\varepsilon_{U_c},\varepsilon_{\chi_c}) &= U_c(\boldsymbol{p},t;\varepsilon_{U_c})\,\boldsymbol{d}_c(\boldsymbol{p},t;\varepsilon_{\chi_c}), \\ \boldsymbol{V}_w(\boldsymbol{p},t;\varepsilon_{U_w},\varepsilon_{\chi_w}) &= U_w(\boldsymbol{p},t;\varepsilon_{U_w})\,\boldsymbol{d}_w(\boldsymbol{p},t;\varepsilon_{\chi_w}), \end{aligned} \tag{7}$$

where

$$\boldsymbol{d}_c = [\cos\chi_c, \sin\chi_c]^\top, \qquad \boldsymbol{d}_w = [\cos\chi_w, \sin\chi_w]^\top, \tag{8}$$

are respectively the propagation directions of current and wind.

The apparent wind seen by the vehicle is

$$\boldsymbol{V}_{a,k}(\boldsymbol{p}_k,t) = \boldsymbol{V}_w(\boldsymbol{p}_k,t) - \boldsymbol{V}_k, \tag{9}$$

and is transformed into the body-fixed frame before calculating $\boldsymbol{\tau}_k^{\mathrm{wind}}$. Denoting by $V_{a,k}$ the apparent wind magnitude, by $\gamma_{a,k}$ its incidence angle with respect to the bow, by ρ_a the air density, by $A_{F,k}$ and $A_{L,k}$ the frontal and lateral exposed areas, and by L_k a characteristic vehicle length, the aerodynamic contribution is written in compact form as

$$\boldsymbol{\tau}_k^{\mathrm{wind}} = \frac{1}{2}\rho_a V_{a,k}^2 \begin{bmatrix} C_X(\gamma_{a,k})A_{F,k} \\ C_Y(\gamma_{a,k})A_{L,k} \\ C_N(\gamma_{a,k})A_{L,k}L_k \end{bmatrix}, \tag{10}$$

where C_X, C_Y, and C_N are angle-dependent aerodynamic coefficients evaluated according to the Blendermann parametrization adopted in [16].

Wave motion is represented through an irregular spectrum with random phases, consistent with the Pierson–Moskowitz formulation. The wave elevation can be expressed as a sum of harmonics

$$\zeta(\boldsymbol{p},t) = \sum_{n=1}^{N_\zeta} A_n \cos\big(\omega_{e,n}t - k_{e,n}\boldsymbol{d}_s^\top \boldsymbol{p} + \varphi_n\big), \tag{11}$$

where N_ζ is the number of harmonic components, $\boldsymbol{d}_s = [\cos\chi_s, \sin\chi_s]^\top$ is the wave propagation direction, $\omega_{e,n}$ and $k_{e,n}$ are respectively the encounter frequency and encounter wave number of the n-th component, φ_n are random phases, and the amplitudes A_n are derived from the assigned spectrum. In the present study, wave motion is not obtained from direct measurements, but is numerically generated from a Pierson–Moskowitz spectrum consistent with the sea state assumed in the numerical campaign. The amplitudes and frequencies of the harmonics define the mean sea condition, while the random phases make it possible to obtain different realizations

of the wave surface under the same nominal sea state. The quantity entering the disturbance model is not the elevation itself, but the local slope in the propagation direction,

$$s_\zeta(\boldsymbol{p}, t) = \boldsymbol{d}_s^\top \nabla\zeta(\boldsymbol{p}, t), \tag{12}$$

which represents the instantaneous inclination of the wave profile perceived by the vehicle. Denoting by $\gamma_{s,k} = \chi_s - \psi_k$ the wave incidence angle with respect to the bow, and by L_k, B_k, and T_k the characteristic length, beam, and draft, respectively, the wave-induced load is approximated as

$$\boldsymbol{\tau}_k^{\text{wave}} = \rho g B_k L_k T_k \begin{bmatrix} \cos\gamma_{s,k}\, s_\zeta \\ -\sin\gamma_{s,k}\, s_\zeta \\ \dfrac{L_k^2 - B_k^2}{24} \sin(2\gamma_{s,k}) \displaystyle\sum_{n=1}^{N_\zeta} k_{e,n}^2 \zeta_n \end{bmatrix}, \tag{13}$$

where ρ is the water density, g the gravitational acceleration, and ζ_n the corresponding contribution in the sum (11). The first and second components describe the equivalent force along and across the hull, while the third one approximates the yaw moment induced by oblique wave incidence. In the planning stage, this contribution does not modify the kinematic structure of the problem, but enters the power required to maintain speed and heading along each segment.

The route Γ_k of vehicle k is not evaluated as a single continuous path, but is discretized into segments $\Gamma_{k,q}$, with $q = 1, \ldots, N_k$. This discretization has two purposes. On the one hand, it naturally represents the polylines generated by obstacle avoidance through the visibility graph; on the other hand, it makes it possible to divide the route into segments short enough that the direction of motion, ground-relative speed, vehicle attitude, and local environmental quantities may be assumed nearly constant. In this way, a quasi-steady condition can be evaluated on each segment and the required propulsion power can be computed consistently.

Considering therefore a generic segment $\Gamma_{k,q}$, the digital twin samples the local current, the apparent wind, and the wave contribution along that segment, and associates with it the ground-relative speed and the attitude $\psi_{k,q}$. Neglecting the short transient phases of acceleration and realignment between consecutive segments, the generalized force that the propulsion system must realize is

$$\boldsymbol{\tau}_{k,q}^{\text{req}}(\psi_{k,q}) = \mathbf{C}_k(\boldsymbol{\nu}_{r,k,q})\boldsymbol{\nu}_{r,k,q} + \mathbf{D}_k(\boldsymbol{\nu}_{r,k,q})\boldsymbol{\nu}_{r,k,q} - \boldsymbol{\tau}_{k,q}^{\text{wind}} - \boldsymbol{\tau}_{k,q}^{\text{wave}}. \tag{14}$$

where $\boldsymbol{\tau}_{k,q}^{\text{req}} = [X_{k,q}^{\text{req}}, Y_{k,q}^{\text{req}}, N_{k,q}^{\text{req}}]^\top$ is the generalized vector of longitudinal force, lateral force, and yaw moment that must be balanced by the propulsion system on segment q.

This vector does not, however, coincide with the elementary actions produced by the actuators. It represents an equivalent demand at the vehicle body level, which must then be realized through a physically admissible combination of thrusts, azimuthal orientations, rotation rates, and local attitude. The role of propulsion allocation is therefore to map the generalized demand $\boldsymbol{\tau}_{k,q}^{\text{req}}$ into the actual commands of the available actuators. This mapping is necessary because the same generalized resultant may be obtained through different combinations of propulsive actions, but these combinations do not have the same absorbed power nor the same efficiency.

The index ℓ distinguishes the individual propulsors of vehicle k. In the case considered here, a conventional marine propulsion system based on propellers is assumed, with azimuth propulsors for the mothership and a simpler underactuated configuration for the drone. This choice is not incidental: propellers exhibit maximum-efficiency regions that depend on the advance ratio and

the rotation rate, and the planner may therefore prefer routes and speeds that keep the propulsion line under more favorable operating conditions.

For propulsor ℓ of vehicle k, the thrust $T_{k,\ell}$ and the shaft torque $Q_{k,\ell}$ are computed by means of Wageningen B characteristic curves:

$$T_{k,\ell} = \rho n_{k,\ell}^2 d_{k,\ell}^4 K_T(J_{k,\ell}), \qquad Q_{k,\ell} = \rho n_{k,\ell}^2 d_{k,\ell}^5 K_Q(J_{k,\ell}), \tag{15}$$

where ρ is the water density, $n_{k,\ell}$ the rotation rate, $d_{k,\ell}$ the propeller diameter, K_T and K_Q the thrust and torque coefficients, and $J_{k,\ell}$ the advance ratio, defined from the local incident flow velocity on the propeller. More details on the curve parametrization, the adopted coefficients, and the actuator geometry are reported in the reference work [16]. At the segment level, the twin solves a thrust-allocation problem to determine intensities, azimuthal orientations, and local heading that satisfy the required resultant with minimum power:

$$\mathcal{P}_{k,q} = \min_{\mathbf{a}_{k,q} \in \mathcal{A}_{k,q}} \sum_{\ell=1}^{N_{T,k}} \frac{2\pi n_{k,\ell} Q_{k,\ell}}{\eta_{m,k,\ell}(n_{k,\ell}, \sigma_{k,\ell})} \qquad \text{s.t.} \qquad \mathbf{B}_k(\boldsymbol{\beta}_{k,q})\mathbf{T}_{k,q} = \boldsymbol{\tau}_{k,q}^{\text{req}}. \tag{16}$$

In Eq. (16), q denotes the route segment, $N_{T,k}$ is the number of propulsors, and $\boldsymbol{a}_{k,q}$ collects the admissible action variables on the segment. In particular, in the present implementation $\boldsymbol{a}_{k,q}$ contains the local heading $\psi_{k,q}$ and, for each propulsor, the rotation rate $n_{k,\ell}$ and the azimuth angle $\beta_{k,\ell}$; the thrusts $T_{k,\ell}$ and torques $Q_{k,\ell}$ then follow from the curves (15). Moreover, $\boldsymbol{\beta}_{k,q}$ is the vector of azimuth angles, $\mathbf{B}_k$ the allocation matrix, and $\boldsymbol{\tau}_{k,q}^{\text{req}}$ the generalized force required to compensate drag, wind, waves, and segment kinematic constraints. The function $\eta_{m,k,\ell}$ represents the motor or drivetrain efficiency at the considered operating point, while $\sigma_{k,\ell}$ indicates the corresponding load level. Limits on rotation rates, azimuth angles, reconstructed thrusts, power, and relative velocities are included in the admissible set $\mathcal{A}_{k,q}$. It follows that the energy cost is not a surrogate for the traveled distance, but results from the compatibility between relative-to-water motion, local environment, and propulsion efficiency.

2.2 Pairwise rendezvous formulation and decision-variable parametrization

The operational problem consists in determining a meeting point $\boldsymbol{p}_R$ and a meeting time t_R such that the mothership and the drone simultaneously reach the same position starting from the initial conditions $\boldsymbol{p}_{M,0}$ and $\boldsymbol{p}_{D,0}$. Consistently with the two-vehicle formulation recalled in the reference work [16, 17], the objective vector simultaneously minimizes mothership energy, drone energy, and rendezvous time:

$$\min \ \boldsymbol{J}(\boldsymbol{x}) = \min \left[E_M(\boldsymbol{x}),\ E_D(\boldsymbol{x}),\ t_R(\boldsymbol{x})\right]^\top. \tag{17}$$

The energies E_M and E_D are evaluated along the actual routes traversed, after possible reconstruction through the visibility graph and after the environmental discretization of the segments. The rendezvous condition imposes instead

$$\boldsymbol{p}_M(t_R) = \boldsymbol{p}_D(t_R) = \boldsymbol{p}_R. \tag{18}$$

If planar motion is described through the five quantities $(V_M, \chi_M, V_D, \chi_D, t_R)$, where V_k is the ground-relative speed and χ_k the direction of motion of vehicle k, the vector condition (18) introduces two scalar constraints. Three degrees of freedom therefore remain. As in the journal paper, the independent choice concerns the mothership motion and the meeting time; the drone speed and direction are then recovered analytically from synchronization.

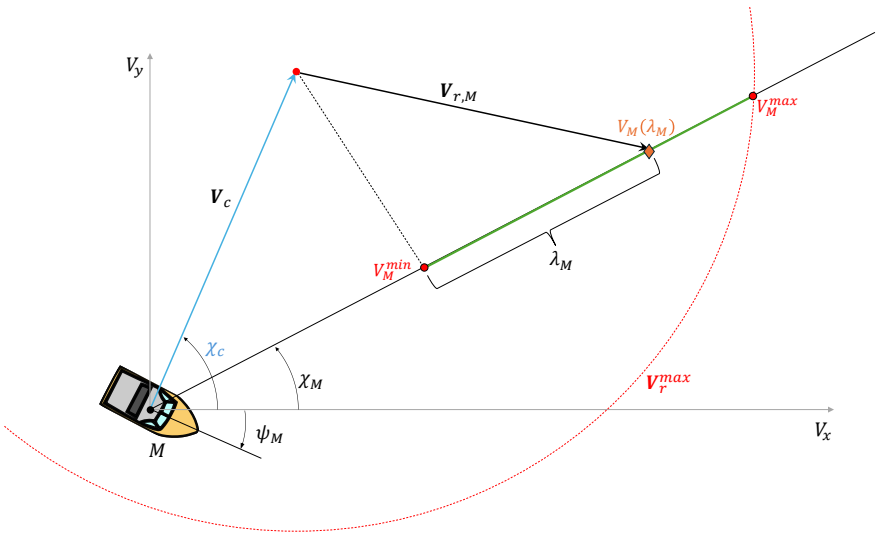


Figure 1: Velocity-space interpretation of the parametrization used for the mothership. The blue vector $\boldsymbol{V}_c$ represents the local current, oriented according to χ_c. The red dashed circle, centered at the tip of $\boldsymbol{V}_c$, denotes the locus of ground-relative velocities compatible with the maximum relative-to-water speed limit $V_{r,M}^{\max}$. For a fixed direction χ_M, the admissible ground-relative speed is restricted to the segment between $V_M^{\min}$ and $V_M^{\max}$, and the normalized parameter λ_M selects the operating point $V_M(\lambda_M)$ on that segment. The figure also reports the vehicle yaw angle ψ_M, to distinguish the body heading from the direction of motion adopted in the rendezvous parametrization.

The main difference with respect to a purely kinematic parametrization is that the mothership speed is not searched directly as an absolute magnitude. The current field shifts the set of admissible ground-relative velocities: for a fixed direction of motion, the speed interval compatible with the relative-to-water speed limit changes with the intensity and direction of the current. To avoid poorly conditioned search regions that depend on direction, the decision vector

$$\boldsymbol{x} = [\chi_M,\ \lambda_M,\ t_R]^\top, \tag{19}$$

is introduced, where χ_M is the mothership direction of motion and $\lambda_M \in [0, 1]$ is a normalized parameter interpolating within the admissible speed interval. The geometric meaning of this parametrization is shown in Fig. 1: the current shifts the center of the admissible ground-relative velocity domain, while λ_M selects an operating point along the half-line of direction χ_M.

More precisely, denoting by V_c and χ_c the magnitude and direction of the current sampled at the mothership position, the projection of the current along the direction of motion is $V_c \cos(\chi_M - \chi_c)$. The relative-to-water speed limit therefore defines the maximum ground-relative speed magnitude reachable along the same direction:

$$V_M^{\max} = V_c \cos(\chi_M - \chi_c) + \sqrt{\max\left[\left(V_{r,M}^{\max}\right)^2 - V_c^2 \sin^2(\chi_M - \chi_c), 0\right]}. \tag{20}$$

The minimum admissible value along the motion half-line is taken as the non-negative part of the current projection,

$$V_M^{\min} = \max[V_c \cos(\chi_M - \chi_c), 0]\,. \tag{21}$$

The mothership ground-relative speed is finally reconstructed as

$$V_M = V_M^{\min} + \lambda_M\left(V_M^{\max} - V_M^{\min}\right), \qquad \lambda_M \in [0, 1]. \tag{22}$$

This choice keeps λ_M within a fixed dimensionless interval while preserving the physical constraint on the relative-to-water speed. The corresponding longitudinal component of the relative speed can be read as $V_{r,M} = V_M - V_c \cos(\chi_M - \chi_c)$; possible numerical round-off errors at the edge of the admissible interval are treated as constraint violations and aggregated in the indicator ϕ described in Eq. (33).

Once χ_M, λ_M, and t_R have been assigned, the rendezvous point is determined by the mothership kinematics:

$$\boldsymbol{V}_M = V_M \begin{bmatrix} \cos\chi_M \\ \sin\chi_M \end{bmatrix}, \qquad \boldsymbol{p}_R = \boldsymbol{p}_{M,0} + t_R \boldsymbol{V}_M. \tag{23}$$

The ground-relative velocity required for the drone to reach the same point at the same time is therefore

$$\boldsymbol{V}_D = \frac{\boldsymbol{p}_R - \boldsymbol{p}_{D,0}}{t_R}, \qquad V_D = \|\boldsymbol{V}_D\|, \qquad \chi_D = \operatorname{atan2}(V_{D,y}, V_{D,x}). \tag{24}$$

The drone kinematic feasibility is verified by comparing the relative-to-water speed $\|\boldsymbol{V}_D - \boldsymbol{V}_{c,D}\|$ with $V_{r,D}^{\max}$, where $\boldsymbol{V}_{c,D}$ is the current sampled along the drone route or, in a first evaluation, at its initial position. If synchronization requires a relative speed above the limit, the particle is not discarded arbitrarily: it is retained in the PSO population but marked as infeasible through the violation indicator. In this way, geometric baselines and uncertainty-perturbed solutions can also be compared under the same feasibility logic.

Expressions (20)–(24) define the nominal kinematics before the possible deviation induced by obstacles. If the direct path is not admissible, the rendezvous point remains the same, but the route geometry is reconstructed through the obstacle-avoidance layer described later and the segment velocities are updated to preserve the arrival time t_R.

2.3 Energy evaluation and operational constraints

For each candidate solution, the route of vehicle k is denoted by Γ_k and subdivided into N_k segments $\Gamma_{k,q}$, $q = 1, \ldots, N_k$. On each segment, the direction of motion and the reference speed are assumed constant. The local environmental quantities, namely current, wind, and wave motion, are instead evaluated by the digital model at representative points and times of the segment and then treated as mean values over that portion. The required energy is computed as the sum of the minimum segment powers:

$$E_k = \sum_{q=1}^{N_k} \mathcal{P}_{k,q}\, \Delta t_{k,q}, \qquad k \in \{M, D\}, \tag{25}$$

where $\mathcal{P}_{k,q}$ is obtained from the solution of the allocation problem (16) and $\Delta t_{k,q}$ is the travel time of the segment. If $\ell_{k,q}$ is the segment length and $V_{k,q}$ the ground-relative speed assigned to the segment, then $\Delta t_{k,q} = \ell_{k,q}/V_{k,q}$ and $\sum_q \Delta t_{k,q} = t_R$ for both vehicles. In this way, synchronization is preserved even when the direct route is replaced by an obstacle-avoidance polyline.

The admissibility of a solution is verified with respect to kinematic, energetic, temporal, and geometric constraints. In compact form:

$$0 \le \|\boldsymbol{V}_{r,k,q}\| \le V_{r,k}^{\max}, \qquad E_k \le E_k^{\text{avail}}, \qquad 0 < t_R \le t^{\max}, \qquad \Gamma_k \subset \Omega_{\text{free}}, \qquad k \in \{M, D\}, \tag{26}$$

where Ω_{free} is the navigable domain, obtained by subtracting the set of fixed obstacles from the operational domain. The constraints therefore do not represent obstacle avoidance alone, but also include kinematic and propulsive limits: maximum relative-to-water speed, available

energy, maximum time, and thrust allocation compatibility. For the solutions generated by the optimizer, the geometric constraint is largely enforced constructively by the visibility graph. For the geometric baselines and for the re-evaluations under uncertainty, however, energetic or kinematic violations may arise even when the route does not intersect the obstacle. Any such violations are aggregated in the indicator ϕ, defined in Section 3. This indicator distinguishes feasible from infeasible solutions and quantifies the normalized magnitude of the constraint exceedance.

2.4 Path construction with obstacle avoidance

Fixed obstacles are modeled as closed polygons O_m, $m = 1, \dots, N_O$, and the free domain is $\Omega_{\text{free}} = \Omega \setminus \cup_m O_m$. For each vehicle, it is first checked whether the direct segment between the initial position and the rendezvous point belongs to Ω_{free}. If no intersections are present, the direct segment is retained. Otherwise, a visibility graph is constructed, whose nodes are the initial vehicle position, the point $\boldsymbol{p}_R$, and the obstacle vertices. An edge is admitted only if the segment connecting two nodes does not intersect any obstacle. The admissible route is then obtained as the shortest path on the graph.

The resulting polyline is subsequently refined from the environmental point of view. Given two consecutive points of the route, additional subdivision points are inserted whenever the variation of an environmental indicator, built from current and wind intensity and direction, exceeds a prescribed threshold. This operation produces segments on which local conditions are sufficiently uniform to apply the power calculation (16) consistently. If obstacle avoidance increases the path length, the segment speeds are recomputed so that the total travel time remains equal to t_R; if this requires a relative-to-water speed above the admissible limit, the solution is marked as constrained or infeasible through ϕ.

The obstacle-handling and segmentation steps described above complete the calculation chain outlined at the beginning of this section. The resulting Pareto front therefore represents a trade-off between mission time and the energy consumption of the two vehicles, already filtered with respect to geometric and operational feasibility. Within this structure, the PSO does not directly optimize the low-level actuator commands, but selects the geometry and synchronization of the rendezvous; the physical–propulsive model then translates each candidate choice into required power, energy cost, and constraint margins.

3 RE-OPTIMIZATION UNDER UNCERTAINTY AND ROBUSTNESS METRICS

3.1 Periodic re-optimization cycle

The planner described in Section 2 is used as a deterministic module within a digital-twin-driven re-optimization policy. At the initial time t_0, the twin receives the positions of the two vehicles, constructs the environmental forecast $\widehat{\mathcal{E}}_0$, and computes a set of Pareto-optimal solutions. An operational selection rule, denoted by $\mathcal{S}(\cdot)$, extracts from the front the solution to be executed.

During the mission, the actual vehicle state and the environmental forecast are updated at times

$$t_k = t_0 + k\Delta t_r, \qquad k = 1, 2, \dots, \tag{27}$$

unless the rendezvous has already been completed. The interval Δt_r is a parameter of the numerical campaign and controls the trade-off between adaptation promptness and computational cost. At each update time t_k, the rendezvous problem is reconstructed from the current mothership and drone positions, $\boldsymbol{p}_M(t_k)$ and $\boldsymbol{p}_D(t_k)$, taken as new initial conditions, and from the updated

environmental forecast $\widehat{\mathcal{E}}_k$ over the remaining horizon. The updated problem is then solved again through MOPSO, a multi-objective formulation of Particle Swarm Optimization:

$$\mathcal{P}_k = \mathrm{MOPSO}\left(\mathcal{X}(t_k),\, \boldsymbol{J}(\boldsymbol{x}; \widehat{\mathcal{E}}_k)\right), \qquad \boldsymbol{x}_k^\star = \mathcal{S}(\mathcal{P}_k), \tag{28}$$

where $\mathcal{X}(t_k)$ denotes the set of rendezvous decisions still admissible at update time t_k. In particular, it contains the decision vectors $\boldsymbol{x} = [\chi_M, \lambda_M, t_R]^\top$ compatible with the current positions of the two vehicles, the remaining mission time, the relative-to-water speed limits, the remaining available energy, and the geometric constraints imposed by the obstacles. The symbol $\mathcal{P}_k$ denotes the Pareto front approximated by the multi-objective solver, while $\boldsymbol{x}_k^\star$ is the solution selected and followed until the next update.

This setup clearly separates uncertainty modeling from the optimizer. Each individual call to the planner remains deterministic, since it uses a fixed environmental forecast. The effect of uncertainty is instead evaluated by repeating the mission over perturbed realizations of the metocean fields and comparing two policies: open-loop execution, in which the nominal plan is maintained until rendezvous, and periodic replanning, in which the digital twin updates the forecast and progressively corrects the rendezvous geometry. In this way, a full stochastic reformulation of the problem is avoided, while still preserving a direct measure of the operational value of forecast updates.

3.2 Repeatability and robustness indicators

The metrics used in the numerical campaign have two purposes. The first is to verify the numerical repeatability of the solver, that is, to establish whether different random seeds of the PSO produce consistent solutions on the same deterministic scenario. The second is to measure operational robustness, that is, to quantify how energy, time, route geometry, and rendezvous point change when the environmental fields are perturbed or updated.

To compare multiple routes geometrically, each path $\Gamma_k^{(j)}$ of vehicle $k \in \{M, D\}$, obtained in execution or realization j, is resampled on a common grid of normalized curvilinear abscissa $s_m \in [0, 1]$, $m = 1, \ldots, N_s$. The vector $\boldsymbol{r}_k^{(j)}(s_m) = [x_k^{(j)}(s_m), y_k^{(j)}(s_m)]^\top$ indicates the planar position of the vehicle along the resampled path, at normalized curvilinear distance s_m from the initial position toward the rendezvous point. The ensemble mean path is

$$\bar{\boldsymbol{r}}_k(s_m) = \frac{1}{N_r} \sum_{j=1}^{N_r} \boldsymbol{r}_k^{(j)}(s_m), \tag{29}$$

where N_r is the number of considered runs. The normalized path dispersion is then

$$D_{\Gamma,k} = \frac{1}{N_r N_s L_{\mathrm{ref},k}} \sum_{j=1}^{N_r} \sum_{m=1}^{N_s} \left\| \boldsymbol{r}_k^{(j)}(s_m) - \bar{\boldsymbol{r}}_k(s_m) \right\|, \tag{30}$$

where $L_{\mathrm{ref},k}$ is a reference length, chosen as the straight-line distance between the initial vehicle position and the nominal rendezvous point. Small values of $D_{\Gamma,k}$ indicate that the routes remain concentrated within a narrow geometric tube, whereas large values indicate higher spatial sensitivity.

Euclidean distance alone does not necessarily distinguish routes belonging to different topological families. For this reason, each path is also assigned to a homotopy class h_j with respect to the set of obstacles. In the case of multiple islands, these classes identify route families that

pass through the same free corridors or bypass the obstacles according to the same topological sequence. Operationally, in the numerical campaign the class h_j is identified by the free corridor crossed by the route with respect to the island configuration.

The homotopy-consistency rate is defined as

$$\eta_h = \max_{h \in \mathcal{H}} \frac{1}{N_r} \sum_{j=1}^{N_r} I(h_j = h), \tag{31}$$

where h_j is the homotopy class of the j-th solution, $I(\cdot)$ is the indicator function, and $\mathcal{H}$ is the set of admissible classes in the free domain. The value $\eta_h = 1$ indicates that all runs belong to the same topological class; lower values instead signal the presence of solutions distributed among different corridors or bypass families.

These two indicators capture different, operationally distinct failure modes. A drop in η_h means that a small forecast update can switch the selected corridor entirely: the mothership and drone would be committed to a materially different route, with implications for traffic separation near the obstacle, for the reliability of the communication link with the drone, and for the safety margins assumed in mission planning. A large $D_{\Gamma,k}$, by contrast, does not change the route family but widens it: it quantifies the lateral buffer and, indirectly, the energy reserve that must be budgeted around the nominal corridor for the plan to remain feasible under forecast mismatch. In this sense, η_h close to one together with small $D_{\Gamma,k}$ is the operationally desirable regime: the mission planner can commit to a route family in advance and treat replanning as a fine correction rather than a source of route uncertainty.

The terminal cloud of rendezvous points is summarized by the standard deviations σ_x and σ_y of the Cartesian coordinates of the points $\boldsymbol{p}_R^{(j)}$. These quantities measure the dispersion of the final meeting point along the two directions of the operational domain and are reported separately to highlight whether the terminal variability is distributed in a nearly isotropic way or prevails along one specific direction.

In addition to the geometric quantities, the variability of the main scalar mission outputs is evaluated. Denoting generically by z one of the quantities

$$z \in \{E_M, E_D, E_{\text{tot}}, t_R\},$$

the coefficient of variation is defined as

$$\mathrm{CV}_z = \frac{\sigma_z}{\mu_z}, \tag{32}$$

where μ_z and σ_z are respectively the mean and standard deviation of z computed over the ensemble of N_r runs. The same definition therefore yields CV_{E_M}, CV_{E_D}, $\mathrm{CV}_{E_{\text{tot}}}$, and CV_{t_R}. The coefficient of variation is useful because it separates the mean value of the performance from its relative dispersion: a policy may reduce the mean energy while increasing the variability across realizations, or may produce a slightly worse mean performance but one that is more repeatable.

Finally, feasibility is quantified through a scalar constraint-violation indicator. Writing the normalized inequality constraints in the form $g_\ell(\boldsymbol{x}) \le 0$, with $\ell = 1, \dots, n_g$, one defines

$$\phi^{(j)} = \sum_{\ell=1}^{n_g} \max\left(g_\ell(\boldsymbol{x}^{(j)}), 0\right), \tag{33}$$

where $\boldsymbol{x}^{(j)}$ is the solution obtained in the j-th run. Therefore, $\phi^{(j)} = 0$ indicates that all constraints are satisfied, whereas $\phi^{(j)} > 0$ measures the aggregated and normalized amount of constraint exceedance.

In the calculation of ϕ, the constraints g_ℓ include the normalized margins associated with maximum relative-to-water speed, available energy, maximum time, propulsion-allocation limits, and route membership in the free domain Ω_{free}. Consequently, a positive value of ϕ does not necessarily indicate a collision with an obstacle, but more generally the exceedance of at least one of the operational limits imposed on the rendezvous problem. The feasibility rate of the ensemble is therefore

$$\rho_f = \frac{1}{N_r} \sum_{j=1}^{N_r} I\left(\phi^{(j)} = 0\right). \tag{34}$$

The metrics $\phi^{(j)}$ and ρ_f therefore make it possible to distinguish a simple cost variation from a true loss of operational feasibility under environmental perturbations.

4 NUMERICAL CAMPAIGN AND RESULTS

4.1 Reference scenario, environmental uncertainty, and numerical setup

The reference scenario includes a single island represented by a polygonal obstacle, selected so that the direct line-of-sight route between the two vehicles is blocked and the visibility-graph-based layer becomes active. The initial positions of the mothership and the drone are kept fixed throughout the campaign, so that all comparisons are performed on the same mission geometry and can therefore be attributed to solver variability, environmental uncertainty, or replanning policy, rather than to changes in the underlying scenario. The choice of a single island is deliberate: it makes the homotopy class non-trivial while preserving an interpretable connection between environmental perturbations and route dispersion. The resulting reference scenario, together with the representative rendezvous trajectories used in the following, is illustrated in Fig. 2. The nominal metocean configuration includes time-varying current and wind fields and a synthetic irregular sea consistent with the wave model adopted in the digital twin. In particular, the nominal field is defined by a mean current of 1.10 m/s oriented at $-25°$, a mean wind of 6.5 m/s oriented at 155°, and a synthetic irregular sea associated with a sea state corresponding to Beaufort 4. These values do not represent extreme conditions, but a sufficiently energetic environment to make both propulsion cost and the effect of forecast perturbations on route feasibility and geometry relevant.

In the uncertainty-aware cases, perturbed environmental realizations are generated according to the parametrization introduced in Eq. (6), without modifying the structure of the nominal model. In the results reported here, perturbed environmental realizations are generated by assuming

$$\varepsilon_{U_c}, \varepsilon_{U_w} \sim \mathcal{N}(0, 0.10^2), \qquad \varepsilon_{\chi_c}, \varepsilon_{\chi_w} \sim \mathcal{N}(0, (10°)^2). \tag{35}$$

This corresponds to a 10% variability in current and wind intensity and a standard deviation of 10° in the corresponding directions. These values are chosen to represent a plausible forecast mismatch: sufficiently limited to preserve a clear link with the nominal forecast, yet large enough to generate appreciable effects on energy, feasibility, and route structure.

In the replanning case, the same environmental model is updated consistently with the global mission time, so that each replanning instant uses a forecast recalculated on an evolving metocean field. In this sense, the comparison between the frozen-plan and replanning cases

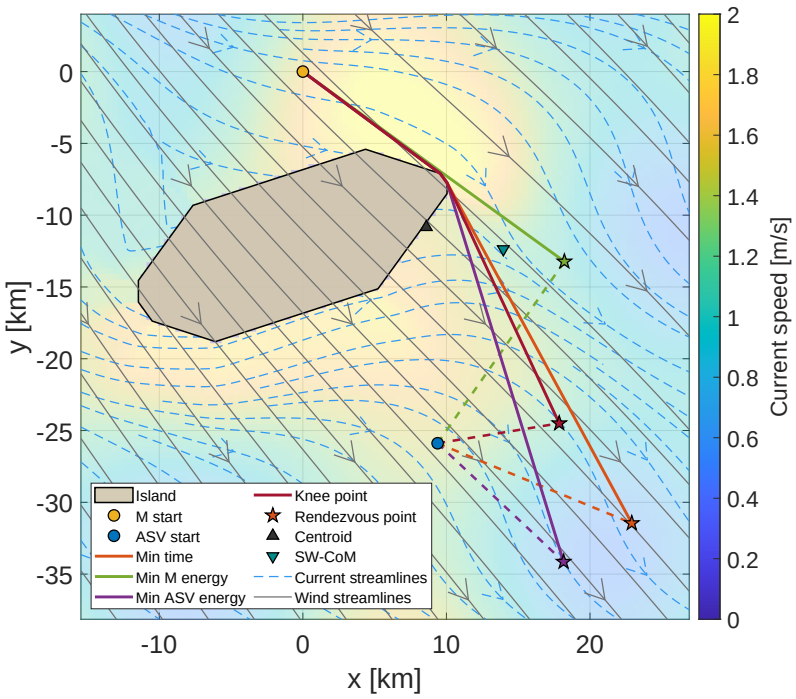


Figure 2: Reference mission scenario used in the numerical campaign. The figure shows the island obstacle, the initial positions of the mothership and the drone, the nominal current field, and the representative rendezvous trajectories associated with the selected Pareto-optimal solutions and geometric baselines.

does not measure only the effect of a static perturbation, but also the ability of the strategy to adapt to an environment that continues to evolve while the mission is underway.

The campaign is organized into five cases. Case A is the deterministic multi-objective reference, in which the rendezvous problem is solved in Pareto form to expose the trade-off among mothership energy, drone energy, and rendezvous time; selected representative Pareto-optimal solutions are then compared with two geometric baselines, namely a centroid-based construction and a speed-weighted center-of-mass construction. Case B addresses deterministic repeatability by repeatedly solving the same optimization problem on the same fixed scenario while varying only the PSO seed. This case is introduced to verify that the optimizer returns geometrically and energetically consistent solutions even under stochastic solver initialization.

Case C evaluates frozen-plan execution under environmental uncertainty. A nominal deterministic solution selected from Case B is kept fixed and its performance is re-evaluated over multiple perturbed environmental realizations. In this way, the campaign quantifies the energy penalty, feasibility degradation, and route dispersion induced by forecast mismatch when no replanning is allowed. Case D considers the same class of environmental perturbations as Case C, but activates the digital-twin-informed replanning policy of Eq. (28) at regular time intervals. At each update instant, the current mission state and the updated forecast are used to solve the rendezvous problem again over the remaining horizon. The comparison between Cases C and D therefore isolates the operational value of periodic forecast updating and re-optimization.

Finally, Case E is a multi-obstacle extension of Case D. In this final test, the domain contains multiple islands and is designed so that several topologically distinct corridors toward the rendezvous region are simultaneously admissible. The purpose is not to introduce a new full energetic benchmark, but to verify whether periodic replanning may induce not only local corrections of route geometry, but also changes in homotopy class under the effect of environmental perturbations and forecast updates.

Uncertainty and campaign		Nominal environment	
Parameter	Value	Parameter	Value
Monte Carlo realizations $N_{\mathrm{MC}}^{C,D}$	100	Mean current speed $\bar{U}_c$	1.10 m/s
Monte Carlo realizations N_{MC}^{E}	10	Mean current direction $\bar{\chi}_c$	−25°
Replanning interval Δt_r (Cases D–E)	0.5 h	Mean wind speed $\bar{U}_w$	6.5 m/s
Std. dev. of relative perturbation $\varepsilon_{U_c}, \varepsilon_{U_w}$	0.10	Mean wind direction $\bar{\chi}_w$	155°
Std. dev. of angular perturbation $\varepsilon_{\chi_c}, \varepsilon_{\chi_w}$	10°	Sea state	Beaufort 5

Table 1: Essential parameters of the uncertainty campaign and nominal metocean conditions, consistent with the parametrization introduced in Eq. (6).

Table 1 summarizes the essential parameters of the numerical campaign. In particular, the left column reports the number of Monte Carlo realizations, the replanning interval, and the dispersion levels assumed for the stochastic variables ε_{U_c}, ε_{U_w}, ε_{χ_c}, and ε_{χ_w} introduced in Eq. (6), whereas the right column collects the corresponding nominal values of the current and wind fields, together with the sea state adopted in the simulations.

4.2 Case A: deterministic Pareto structure

The deterministic benchmark reveals a clear trade-off among total rendezvous time t_R, mothership energy E_M, and drone energy E_D. Even in the simplified pairwise setting, the problem of selecting the rendezvous point cannot be reduced to a purely geometric construction. Fast interception requires the mothership, and to a lesser extent the drone, to operate at high effective propulsive effort, whereas energy-saving solutions exploit a more favorable current alignment and accept longer synchronization times. The pairwise setting makes this interpretation particularly transparent, since the trade-off is not masked by sequencing effects.

The representative Pareto-optimal solutions reported in Table 2 make this structure explicit. The minimum-time solution achieves $t_R = 151.69$ min, but at the price of a very high mothership expenditure, $E_M = 484.40$ MJ, while the drone contribution remains limited, $E_D = 0.138$ MJ. The minimum-drone-energy solution lies in the same fast-rendezvous region, preserving the same rendezvous time while reducing the drone effort to only 0.022 MJ at the cost of an even larger mothership energy, $E_M = 652.17$ MJ. At the opposite end, the minimum-mothership-energy solution reduces the mothership expenditure to only 5.72 MJ, but requires a significantly longer rendezvous time, $t_R = 189.39$ min, together with a larger drone effort, $E_D = 1.369$ MJ. The knee point provides the most balanced compromise: it preserves the same rendezvous time as the fast solutions while reducing the mothership expenditure to 133.40 MJ, with a moderate drone cost of 0.210 MJ. In practical terms, the main design freedom therefore concerns not whether rendezvous can be achieved, but how the energetic burden is redistributed between mothership and drone under synchronization constraints. The pairwise projections of Fig. 3 make these trade-offs particularly easy to interpret. In the (t_R, E_M) plane, the Pareto set is widely spread, indicating that small gains in rendezvous time may require very large increases in mothership energy. In the (t_R, E_D) plane, the spread is narrower, suggesting that the dominant energetic burden of fast interception is carried primarily by the mothership. Finally, the (E_M, E_D)

Point	t_R [min]	E_M [MJ]	E_D [MJ]	E_M [kWh]	E_D [kWh]	ϕ [-]
Min time	151.69	484.40	0.138	134.55	0.038	0
Min E_M	189.39	5.72	1.369	1.59	0.380	0
Min E_D	151.69	652.17	0.022	181.16	0.006	0
Knee	151.69	133.40	0.210	37.06	0.058	0
Centroid baseline	175.96	34.16	2.971	9.49	0.825	1.062
Weighted CoM baseline	240.16	120.46	1.695	33.46	0.471	0.176

Table 2: Representative deterministic solutions for Case A in absolute quantities. The Pareto-representative points are feasible ($\phi = 0$), whereas both geometric baselines violate the energy constraints ($\phi > 0$); the violation indicator ϕ is defined in Eq. (33).

projection highlights a clear inverse trend: strong mothership-energy savings are generally obtained by increasing drone expenditure and accepting a delayed rendezvous. The knee point lies in the central portion of this trade-off region and emerges as the most balanced operating condition among the representative solutions. Figure 2 explains the physical origin of these trends

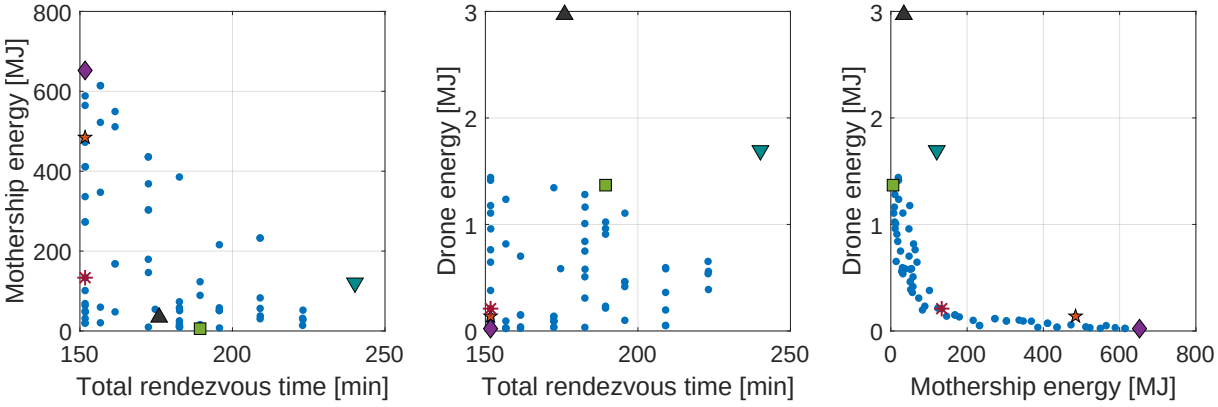


Figure 3: Pairwise projections of the deterministic Pareto front in absolute quantities.

by showing the corresponding trajectories. In the minimum-time and minimum-drone-energy solutions, the rendezvous point is located far downstream and close to the lower-right portion of the domain, so that the drone can follow a relatively favorable path while the mothership performs the energetically dominant maneuver around the island. By contrast, the minimum-mothership-energy solution shifts the rendezvous point upward and closer to the drone, thereby reducing the mothership propulsive effort at the price of a longer powered transfer for the drone. The knee solution remains geometrically close to the fast-rendezvous family, but moderates the aggressiveness of the mothership maneuver and therefore achieves a substantial reduction in E_M without increasing t_R.

A further observation from Fig. 2 is that all representative Pareto solutions belong to the same dominant homotopy class with respect to the island. This indicates that, in the deterministic case, the main degree of freedom is the precise placement of the rendezvous point along a dynamically favorable corridor, rather than the choice of a completely different obstacle-bypass topology. Hence, deterministic optimization does not generate radically different topological routes, but rather redistributes the energetic effort within the same physically meaningful route family.

Table 2 also allows a direct comparison with the two geometric baselines. For both con-

Indicator	Value
Feasibility rate ρ_f [-]	1
Time coefficient of variation CV_{t_R} [-]	0.0020
Mothership-energy coefficient of variation CV_{E_M} [-]	0.0055
Drone-energy coefficient of variation CV_{E_D} [-]	0.0213
Mothership path dispersion $D_{\Gamma,M}$ [km]	0.0378
Drone path dispersion $D_{\Gamma,D}$ [km]	0.0503
Dominant homotopy rate η_h	1.0000

Table 3: Repeatability indicators for the deterministic PSO campaign of Case B.

structions, the geometric rendezvous point is first computed from the initial positions of the vehicles. If this point falls inside the obstacle polygon, it is relocated to an admissible position outside the island using the same visibility-graph-based obstacle-handling logic adopted by the planning methodology; the corrected point is selected among candidates generated just outside the obstacle boundary so as to minimize the combined length of the obstacle-avoiding paths from the mothership and the drone. The corresponding baseline rendezvous configuration is then constructed directly from this corrected target, without any further optimization step.

Both the centroid-based and the speed-weighted center-of-mass constructions remain infeasible, since their aggregate constraint-violation indicator is positive. The centroid baseline yields a relatively low mothership energy, $E_M = 34.16$ MJ, but requires a very large drone expenditure, $E_D = 2.971$ MJ, and therefore significantly violates the energy constraints, with $\phi = 1.062$. The speed-weighted center-of-mass baseline performs better than the centroid in terms of overall balance, but still remains inferior to the optimized solutions: it produces the longest rendezvous time, $t_R = 240.16$ min, together with $E_M = 120.46$ MJ and $E_D = 1.695$ MJ, and remains infeasible with $\phi = 0.176$. These results confirm that simple geometric constructions, even when made consistent with the obstacles, do not properly account for obstacle-induced path elongation, current alignment, synchronization requirements, and nonlinear propulsive demand, and therefore cannot be regarded as reliable surrogates of the physically optimal rendezvous point.

4.3 Case B: deterministic repeatability

The deterministic repeatability campaign was carried out by solving the same fixed rendezvous problem 20 times while varying only the PSO seed. The results, summarized in Table 3, confirm that the optimizer is highly stable with respect to stochastic initialization.

All runs are feasible and remain confined to the same obstacle-bypassing homotopy class, with dominant homotopy rate $\eta_h = 1.00$. The coefficients of variation are very small for all scalar outputs: $\mathrm{CV}_{t_R} = 0.0020$, $\mathrm{CV}_{E_M} = 0.0055$, and $\mathrm{CV}_{E_D} = 0.0213$. The path-dispersion indicators also confirm geometric stability, with $D_{\Gamma,M} = 0.0378$ km and $D_{\Gamma,D} = 0.0503$ km. Overall, these values show that stochastic solver initialization induces only minor local variations around the same deterministic optimum, without triggering topological route changes.

Since seed sensitivity is negligible, the variability observed in Cases C and D can be attributed to environmental perturbations and forecast updates rather than to solver artifacts. This level of repeatability is also what makes the solver usable inside an onboard replanning loop: a policy that returned a materially different route at every re-optimization, even under an unchanged forecast, would undermine the very continuity of the plan that replanning is meant to preserve.

Metric	Nominal	Mean	Std. dev.	CV_z
Rendezvous time t_R [min]	138.70	138.70	0.00	0.0000
Mothership energy E_M [MJ]	155.72	183.92	52.03	0.2829
Drone energy E_D [MJ]	0.250	1.703	4.426	2.5991
Total energy E_{tot} [MJ]	155.97	185.62	53.46	0.2880
Constraint violation ϕ [-]	0	0.855	2.862	–
Robustness and geometric indicators				
Feasibility rate ρ_f [-]		0.86		
Mean total-energy increase $\overline{\Delta E}_{\text{tot}}$ [MJ]		29.66		
Mean total-energy increase $\overline{\Delta E}_{\text{tot}}$ [%]		19.01		
Mothership path dispersion $D_{\Gamma,M}$ [km]		1.0433		
Drone path dispersion $D_{\Gamma,D}$ [km]		1.0579		
Dominant homotopy rate η_h [-]		1.00		

Table 4: Summary statistics for Case C, i.e., execution of the nominal frozen rendezvous plan under perturbed environments.

4.4 Case C: frozen-plan execution under environmental uncertainty

Case C evaluates the sensitivity of a fixed nominal rendezvous plan to environmental perturbations. A deterministic solution selected from Case B is kept unchanged, and the same spatio-temporal rendezvous target is re-evaluated over 100 perturbed environmental realizations, thereby isolating the effect of forecast mismatch when replanning is not allowed.

The geometric behavior of the ensemble is shown in the left panel of Fig. 4. Despite the environmental perturbations, all trajectories preserve the same dominant obstacle-bypassing topology, as confirmed by a dominant homotopy rate $\eta_h = 1.00$. However, the route tube becomes significantly wider than in Case B, especially in the terminal portion of the mission, where synchronization constraints amplify environmental variability. The path-dispersion indicators reach $D_{\Gamma,M} = 1.0433$ km and $D_{\Gamma,D} = 1.0579$ km, about twenty times larger than the corresponding values in Case B, and the rendezvous-point cloud exhibits standard deviations of 1.447 km and 2.032 km along the x and y directions, respectively. The frozen plan therefore remains topologically meaningful, but loses geometric precision as environmental mismatch accumulates during execution. The energetic consequences are substantial, as reported in Table 4. The rendezvous time remains fixed at $t_R = 138.70$ min by construction, but the energetic outputs become highly variable. The mean mothership energy increases from the nominal value of 155.72 MJ to 183.92 MJ, with $CV_{E_M} = 0.2829$. The drone energy is even more sensitive: its nominal value of 0.250 MJ rises to a mean of 1.703 MJ, with $CV_{E_D} = 2.5991$ and extreme cases close to 30 MJ. This asymmetry is physically expected, since the nominal drone expenditure is small and therefore highly sensitive to adverse current and wind perturbations. As a consequence, the mean total energy increases by 29.66 MJ, i.e., about 19% with respect to the nominal plan.

The operational impact is equally significant. The feasibility rate drops to 0.86, meaning that 14 realizations out of 100 violate the energy constraints. The mean value of the constraint-violation indicator ϕ (Eq. (33)) rises to 0.855, with worst-case values up to 19.7. A deterministic plan that is fully feasible under the nominal forecast may therefore become infeasible when executed open loop under perturbed metocean conditions. These degradations, namely route-tube

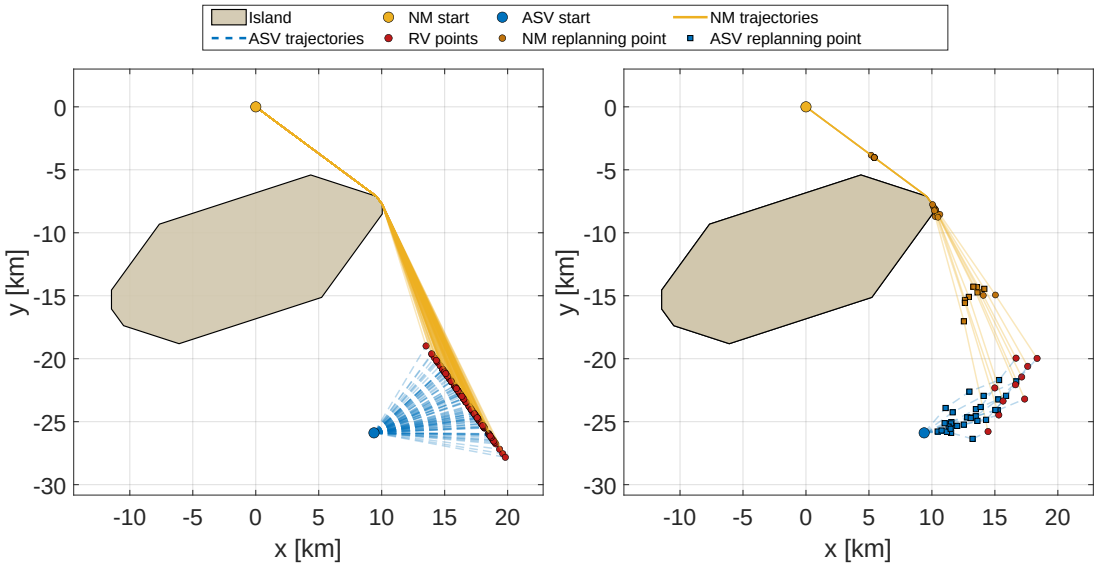


Figure 4: Trajectory overlays under environmental uncertainty for the two execution policies. Left: Case C, in which the nominal rendezvous plan is frozen and executed in open loop. Right: Case D, in which the rendezvous plan is periodically recomputed using updated vehicle states and updated environmental information. In both cases the dominant obstacle-bypassing topology is preserved, but periodic replanning produces a visibly narrower route tube and a more compact terminal approach.

widening, energy increase, and feasibility loss, all stem from the same cause: the synchronization constraint forces both vehicles to absorb forecast mismatch through additional propulsive effort rather than through geometric adaptation.

4.5 Cases C and D: frozen-plan execution versus periodic replanning

Figure 4 and Table 5 directly compare the frozen-plan policy of Case C with the digital-twin-informed periodic replanning policy of Case D under the same class of perturbed environmental realizations. The two cases differ only in whether the rendezvous plan is kept fixed or is periodically recomputed from updated vehicle states and updated environmental information; the comparison therefore isolates the operational contribution of replanning. From a geometric viewpoint, both policies preserve the same dominant obstacle-bypassing topology across the whole Monte Carlo ensemble, with dominant homotopy rate equal to one in both cases. The difference lies in the compactness of the resulting trajectory family. In the frozen-plan case, the route tube progressively widens toward the terminal portion of the mission, where synchronization constraints amplify the effect of environmental mismatch. With periodic replanning, the ensemble remains more compact and the final approach appears visibly more organized, as shown in Fig. 4. Replanning therefore does not induce topological route changes, but corrects the rendezvous geometry within the same dominant route family. Table 5 quantifies these gains: a 13.9% reduction in mean rendezvous time, a 34.2% reduction in mean total energy, and comparable reductions in the coefficients of variation of E_M and E_{tot}. The time gain arises because the digital twin lets the vehicles adapt the meeting geometry dynamically, rather than converging on a fixed spatio-temporal target agreed upon before the mismatch was known. The energy gain is carried almost entirely by the mothership, whose path dispersion $D_{\Gamma,M}$ also drops substantially, whereas the drone path dispersion $D_{\Gamma,D}$ remains nearly unchanged: since the mothership performs the longer and energetically dominant maneuver around the island, it is also the vehicle whose trajectory has the most room to be corrected by an updated forecast, while

Metric	Case C	Case D	Reduction [%]
Rendezvous time t_R [min]	138.70	119.43	13.9
Mothership energy E_M [MJ]	183.92	120.76	34.3
Drone energy E_D [MJ]	1.703	1.394	18.1
Total energy E_{tot} [MJ]	185.62	122.15	34.2
Mothership-energy coefficient of variation CV_{E_M} [-]	0.2829	0.1772	37.4
Total-energy coefficient of variation $\text{CV}_{E_{\text{tot}}}$ [-]	0.2880	0.1924	33.2
Standard deviation of the RV cloud along x, σ_x [km]	1.447	1.264	12.7
Standard deviation of the RV cloud along y, σ_y [km]	2.032	1.925	5.3
Mothership path dispersion $D_{\Gamma,M}$ [km]	1.0433	0.6715	35.6
Drone path dispersion $D_{\Gamma,D}$ [km]	1.0579	1.0334	2.3
Dominant homotopy rate η_h [-]	1.00	1.00	–
Mean number of replanning events $\overline{N}_{\text{rep}}$ [-]	–	4.00	–

Table 5: Comparison between frozen-plan execution (Case C) and periodic replanning (Case D) under environmental uncertainty. Positive reductions indicate better performance of Case D with respect to Case C.

the drone's short, near-direct approach offers little geometric margin for replanning to act on.

In operational terms, these gains are not interchangeable with a simple safety margin added at the planning stage. A 34% reduction in total energy at fixed mission geometry represents propulsive budget that can instead be spent on station-keeping, sensor operation, or an additional recovery cycle before the drone must return to the mothership for recharging; conversely, absorbing the same forecast mismatch open loop would require that energy to be reserved up front, on every mission, whether or not the adverse realization actually occurs. The 14% reduction in rendezvous time has a comparable operational reading: it shortens the interval during which the drone operates near its energy limit and the two vehicles are exposed to evolving sea and wind conditions before meeting, which is precisely the phase in which the frozen plan (Case C) shows the largest feasibility loss. Both gains are obtained without a change of route topology, which is itself operationally relevant: a mission planner can commit to a corridor and a rough schedule in advance, and treat replanning as a correction of that plan rather than as a source of last-minute route changes.

Taken together, these results show that periodic digital-twin-informed replanning is not a cosmetic correction of the nominal plan, but a substantial improvement acting simultaneously on mission time, energy expenditure, and output robustness. The digital twin therefore acts not as a route-redesign tool, but as a forecast-assimilation layer that progressively corrects the rendezvous geometry within the same physically meaningful corridor.

4.6 Case E: multi-obstacle extension

To highlight a behavior that does not emerge in the single-island scenario, a variant of Case D is considered in a multi-obstacle environment, designed so that several topologically distinct corridors toward the rendezvous region are simultaneously feasible. The purpose of this test is not to introduce a new complete energy benchmark, but rather to assess whether digital-twin-driven periodic replanning can induce not only local corrections of the route geometry, but also actual transitions between alternative topological families of solutions.

Fig. 5 shows the ensemble of trajectories obtained in the multi-island case. Unlike the previous cases, the trajectories do not remain confined within a single, merely deformed route

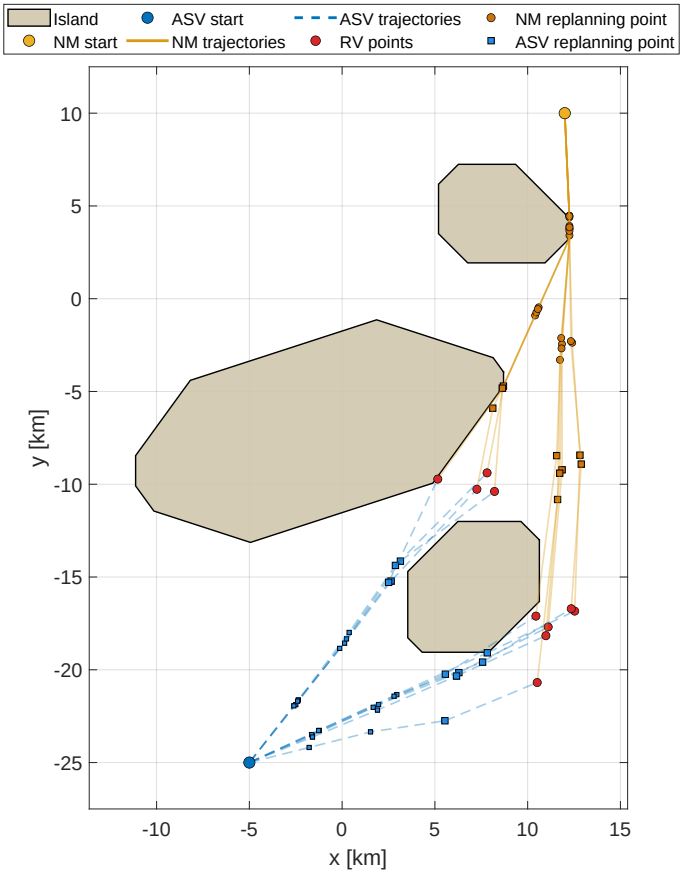


Figure 5: Multi-obstacle case with periodic replanning. The ensemble of mothership and drone trajectories, together with the terminal rendezvous-point cloud and the replanning instants, shows the presence of alternative corridors and multiple topological families of solutions.

tube. Instead, at least two distinct families of paths clearly emerge, associated with different ways of traversing the obstacle field. The rendezvous-point cloud is also more fragmented, with terminal clusters consistent with the presence of alternative corridors. This behavior is consistent with the role of replanning in a scenario with multiple competing passages: small variations in the updated forecast and in the vehicle states may alter the relative convenience of the different corridors and trigger a switch in homotopy class. Table 6 quantitatively confirms this interpretation. The dominant homotopy rate drops to $\eta_h = 0.60$, and two distinct homotopy classes are observed in the Monte Carlo ensemble. The dominant signature is $[+1, +1, +1]$, where each element indicates the bypass side (left $+1$ or right -1) with respect to each obstacle polygon. However, this signature accounts for only 60% of the realizations, while the remaining 40% follow an alternative topological configuration. Since comparing the energetic performance of different classes is not the objective of this test, scalar outputs such as CV_{E_M} and path-dispersion indicators such as $D_{\Gamma,k}$ are not reported here. Indeed, the dominant source of variability in this case is topological rather than geometric, and it is adequately captured by η_h and by the observed class distribution. From a methodological standpoint, this test completes the picture provided by the previous cases. In Cases B and C, and also in Case D for the single-island scenario, homotopy coherence remains essentially equal to one, and replanning mainly acts as a mechanism for geometric and energetic correction. In the multi-obstacle case, by contrast, the richer domain makes multiple comparable bypass strategies available, and the sensitivity of the digital twin to updated environmental conditions translates into genuine competition between

Metric	Value
Dominant homotopy rate η_h [-]	0.60
Observed homotopy classes [-]	2
Dominant homotopy signature	$[+1, +1, +1]$
Frequency of dominant signature [%]	60
Average number of replanning events $\overline{N}_{\text{rep}}$ [-]	4.00

Table 6: Topological indicators for the multi-obstacle test with periodic replanning.

homotopy classes. In this sense, the multi-island test reveals a behavior that is closer to what may be expected in realistic operational environments, where robustness concerns not only the cost of the selected trajectory, but also the stability of its topological structure.

5 CONCLUSIONS

This paper has examined the pairwise rendezvous problem between a mothership and a marine drone within a digital-twin-based framework, with the explicit goal of isolating this elementary block from the broader multi-agent recovery sequence. This choice makes the uncertainty-quantification narrative physically interpretable, since the observed variability can be attributed directly to environmental perturbations, solver initialization, or replanning policy, rather than to combinatorial changes in a higher-level recovery order.

The deterministic benchmark has shown that rendezvous-point selection cannot be reduced to a purely geometric construction. The multi-objective optimization reveals a structured trade-off among mothership energy, drone energy, and rendezvous time under synchronization, propulsion, and environmental constraints, while the geometric baselines based on centroid and speed-weighted center of mass are infeasible in the tested scenario. Even within the same dominant homotopy class, the location of the rendezvous point along a dynamically favorable corridor has a strong effect on how the energetic burden is distributed between the two vehicles.

The repeatability campaign also confirms that the PSO-based solver is highly stable with respect to stochastic initialization. All runs in the deterministic case are feasible, the coefficients of variation of energy and rendezvous time remain very small, and the dominant homotopy rate is equal to one. This result is important because it shows that the variability observed in the uncertainty-aware cases is primarily due to environmental perturbations rather than to numerical instability of the optimizer.

The open-loop uncertainty analysis shows that a nominal deterministic plan may degrade substantially under perturbed metocean conditions. Although the dominant route topology is preserved, the route tube widens significantly, the total energy increases on average by about 19%, and the feasibility rate drops to 0.86. This confirms that a plan that is fully feasible under the nominal forecast may lose operational robustness when executed without updates under forecast mismatch.

Periodic digital-twin-informed replanning recovers a large fraction of this lost performance. With an average of four replanning events, the mean total energy is reduced by 34% and the mean rendezvous time by 14% with respect to the frozen-plan policy, together with a reduction in output dispersion and a more compact trajectory family. Importantly, in the single-island scenario this improvement is not obtained through a change of route topology, but through a progressive correction of the rendezvous geometry within the same physically meaningful corridor. In this sense, the digital twin acts as a forecast-assimilation and adaptation layer, rather

than as a simple replacement of the nominal route.

The final multi-obstacle extension completes this picture by showing that, in richer environments, replanning may also induce transitions between distinct homotopy classes. In that case, the dominant homotopy rate decreases to 0.60, revealing genuine competition between alternative corridors. This result indicates that, in more complex operational domains, robustness concerns not only the cost of the selected trajectory but also the stability of its topological structure.

These findings suggest that the pairwise rendezvous block is a promising foundation for uncertainty-aware recovery strategies in more complex cooperative marine operations.

REFERENCES

[1] Yubing Wu, Tao Wang, and Shuo Liu. "A Review of Path Planning Methods for Marine Autonomous Surface Vehicles". In: *Journal of Marine Science and Engineering* 12.5 (2024). ISSN: 2077-1312. DOI: 10.3390/jmse12050833. URL: https://www.mdpi.com/2077-1312/12/5/833.

[2] Jiri de Vos, Robert G. Hekkenberg, and Osiris A. Valdez Banda. "The Impact of Autonomous Ships on Safety at Sea – A Statistical Analysis". In: *Reliability Engineering & System Safety* 210 (2021), p. 107558. ISSN: 0951-8320. DOI: https://doi.org/10.1016/j.ress.2021.107558.

[3] Maicol Laurenza, Gianluca Pepe, Davide Tonazzi, et al. "Surface Unmanned Multipurpose Research Marine Vehicle: SUNMARE Project". In: *Progress in Marine Science and Technology* 6 (2022), pp. 522–529. DOI: 10.3233/PMST220062.

[4] Edoardo I. Sarda and Manhar R. Dhanak. "Launch and Recovery of an Autonomous Underwater Vehicle From a Station-Keeping Unmanned Surface Vehicle". In: *IEEE Journal of Oceanic Engineering* 44.2 (2019), pp. 290–299. DOI: 10.1109/JOE.2018.2867988.

[5] Z. Zeng et al. "Rendezvous Path Planning for Multiple Autonomous Marine Vehicles". In: *IEEE Journal of Oceanic Engineering* 43.4 (2017), pp. 907–925. DOI: 10.1109/JOE.2017.2723058.

[6] N. Mathew, S. L. Smith, and S. L. Waslander. "Multirobot Rendezvous Planning for Recharging in Persistent Tasks". In: *IEEE Transactions on Robotics* 31.1 (2015), pp. 128–142. DOI: 10.1109/TRO.2014.2380593.

[7] B. Li et al. "Rendezvous Planning for Multiple AUVs With Mobile Charging Stations in Dynamic Currents". In: *IEEE Robotics and Automation Letters* 4.2 (2019), pp. 1653–1660. DOI: 10.1109/LRA.2019.2896899.

[8] Y. Zhang, P. Peng, and J. Wei. "Research on Formation Rendezvous for Unmanned Surface Vehicles Considering Arbitrary Initial Positions and Headings". In: *Electronics* 13.14 (2024), p. 2867. DOI: 10.3390/electronics13142867.

[9] L. Nesi et al. "A new tow maneuver of a damaged boat through a swarm of autonomous sea drones". In: *IFAC-PapersOnLine* 52.21 (2019), pp. 360–366. DOI: 10.1016/j.ifacol.2019.12.333.

[10] C. Zhu et al. "A formation rendezvous method for multiple unmanned surface vehicles based on GAT-modMADDPG". In: *Ocean Engineering* 341 (2025), p. 122483. DOI: 10.1016/j.oceaneng.2025.122483.

[11] Xuliang Yao et al. “Path planning for autonomous underwater vehicles based on interval optimization in uncertain flow fields”. In: *Ocean Engineering* 234 (2021), p. 108675. DOI: 10.1016/j.oceaneng.2021.108675.

[12] Karim Ahmadi Dastgerdi et al. “Uncertainty aware path planning and collision avoidance for marine vehicles”. In: *IFAC-PapersOnLine*. Vol. 58. 20. 2024, pp. 235–240. DOI: 10.1016/j.ifacol.2024.10.060.

[13] Minahil Raza et al. “Towards Integrated Digital-Twins: An Application Framework for Autonomous Maritime Surface Vessel Development”. In: *Journal of Marine Science and Engineering* 10.10 (2022), p. 1469. DOI: 10.3390/jmse10101469.

[14] Nuwan Sri Madusanka et al. “Digital Twin in the Maritime Domain: A Review and Emerging Trends”. In: *Journal of Marine Science and Engineering* 11.5 (2023), p. 1021. DOI: 10.3390/jmse11051021.

[15] Henrik Stokland Berg et al. “Digital twin syncing for autonomous surface vessels using reinforcement learning and nonlinear model predictive control”. In: *Scientific Reports* 15 (2025), p. 9344. DOI: 10.1038/s41598-025-93635-9.

[16] Maicol Laurenza et al. “Toward Scalable Energy-Efficient Rendezvous Optimization for Multiagent Marine Applications”. In: *IEEE Journal of Oceanic Engineering* 51.2 (2026), pp. 1106–1128. DOI: 10.1109/JOE.2026.3655631.

[17] Maicol Laurenza et al. “Adaptive Rendezvous Planning Algorithm for Energy-Efficient Marine Operations”. In: *Progress in Marine Science and Technology* 10 (2025), pp. 775–782. DOI: 10.3233/PMST250092. URL: https://www.scopus.com/inward/record.uri?eid=2-s2.0-105016120178&doi=10.3233%2fPMST250092&partnerID=40&md5=50c56a74c6f506b35c0c12a2182ec73a.

[18] Nicola Roveri et al. “Safe and Secure Control of Swarms of Vehicles by Small-World Theory”. In: *Energies* 13.5 (2020), p. 1043. DOI: 10.3390/en13051043.

[19] Thor I. Fossen. *Guidance and Control of Ocean Vehicles*. Wiley, 1994.