doi: 10.82124/CEAS-GNC-2026-003 CEAS-GNC-2026-003

Enabling reusability in space transportation using actively controlled parafoils for high-precision landings

Leyre Hernández Palacios Navigation Engineer, Sener Aeroespacial S.A., Tres Cantos, Madrid, Spain, leyre.hernandez@aeroespacial.sener

Joost Veenman Senior Control Engineer and PM, Sener Aeroespacial S.A., Tres Cantos, Madrid, Spain, joost.veenman@aeroespacial.sener

Jesús Ramírez GNC Engineer, Sener Aeroespacial S.A., Tres Cantos, Madrid, Spain, jesus.ramirez@aeroespacial.sener

Ángela Serrano GNC Engineer, Sener Aeroespacial S.A., Tres Cantos, Madrid, Spain, angela.serrano@aeroespacial.sener

Pablo González GNC Engineer, Sener Aeroespacial S.A., Tres Cantos, Madrid, Spain, pablo.gonzalez@aeroespacial.sener

Lourdes Lobera GNC Engineer, Sener Aeroespacial S.A., Tres Cantos, Madrid, Spain, lourdes.lobera@aeroespacial.sener

Antonio Hernández Rodicio-Mateos GNC Engineer, Sener Aeroespacial S.A., Tres Cantos, Madrid, Spain, antonio.hernandez@aeroespacial.sener

Guillermo Alonso GNC Engineer, Sener Aeroespacial S.A., Tres Cantos, Madrid, Spain, guillermo.alonso@aeroespacial.sener

ABSTRACT

The growing demand for reusable space transportation systems underscores the importance of developing reliable recovery technologies for space vehicles, particularly launcher stages. This paper presents an actively controlled parafoil system designed to enhance landing precision, building on the legacy of the Space Rider (SR) mission. The system integrates modern navigation methods, optimization-based guidance, and robust control strategies. It consists of a parafoil attached to a recoverable payload, operating in cooperation with a sea-drone platform. This cooperative scheme reduces the required landing area and mitigates impact energy constraints by using a net-based absorption mechanism. To assess feasibility, a high-fidelity simulation environment was developed, enabling the evaluation of landing accuracy and touchdown velocity, among other requirements. The results confirm the viability of the proposed approach and highlight key challenges for future development and maturation.

Keywords: Space transportation, reusability, guided parafoils, optimization based guidance, robust control, hybrid navigation

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

Acronyms

CoG	Center of Gravity
ConOps	Concept of Operations
DKE	Dynamics, Kinematics and Environment
DoF	Degrees of Freedom
ECMWF	European Centre for Medium-Range Weather Forecasts
EKF	Extended Kalman Filter
ERA5	Fifth generation ECMWF Re-Analysis
ESA	European Space Agency
FDIR	Failure Detection, Isolation and Recovery
FES	Functional Engineering Simulator
FPA	Flight Path Angle
GNC	Guidance, Navigation and Control
GNSS	Global Navigation Satellite System
HALER	High Accurate Landing Equipment for Recovery
HDG	Heading
HDR	Heading Rate
IMU	Inertial Measurement Unit
IXV	Intermediate eXperimental Vehicle
LiDAR	Light Detection And Ranging
LP	Landing Point
LPF	Landing Point Frame
LUT	Lookup Table
LV	Local Vertical
MC	Monte-Carlo
MCI	Mass, Centering and Inertia
MGM	Magnetometer
MPC	Model Predictive Control
OCP	Optimal Control Problem
PYL	Payload
PGNC	Parafoil GNC
PD	Proportional-Derivative
PI	Proportional-Integral
PPP	Precise Point Positioning
RTK	Real-Time Kinematics
SCP	Sequential Convex Programming
SODAR	Sonic Detection And Ranging
SPR	Sea-platform
SR	Space Rider
SSV	Structured Singular Values
SPP	Single Point Positioning
WF	Waypoint Frame
WP	Waypoint

1 Introduction

Two growing needs have been identified by the European Space sector: developing sustainable space technologies, i.e., reusable vehicles; and increasing its autonomy and competitiveness with respect to the global market. One such technology that is low cost and has a high load capacity and manoeuvreability is the parafoil. It is already used in space payload recovery, spacecraft landings, and supply delivery; and has some heritage within Europe, such as the development of the Parafoil GNC (PGNC) for the X-38 by Astrium Aerospace (now Airbus) [1], the Intermediate eXperimental Vehicle (IXV), and the Space Rider (SR) reentry vehicle developed by the European Space Agency (ESA) [2, 3]. Although these systems have demonstrated landing accuracies below a radius of 150 m, they still require a relatively large landing area and controlled environmental conditions to ensure success.

Integrating a mobile vehicle into parafoil-based recovery offers a promising alternative to fixed-point landing by compensating for environmental uncertainties, easing parafoil control, and enabling lighter designs. SpaceX and Rocket Lab have already demonstrated this concept with payload fairing and booster recovery, although efforts remain experimental, as shown by SpaceX's suspension of capture attempts. For Europe, the approach could extend beyond past demonstrations of recovering payloads up to 6 t

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

towards heavier targets such as 10 t upper stages, for which the required larger parafoils would generate higher velocities and impact energies. By allowing the mobile platform to rendezvous with the payload and reduce the relative velocity, and using simple absorption devices such as nets to dissipate residual energy, the risks of structural damage can be mitigated. However, key challenges remain: the complex, wind-sensitive dynamics of parafoils [4] and the stringent timing and coordination demands placed on the mobile vehicle [5].

The High Accurate Landing Equipment for Recovery (HALER) project —which is part of a European Space Agency (ESA) program investigating the feasibility of various reusable space transportation concepts— addresses these challenges by designing, developing, and simulating the following scenario: a 6 t payload from a launcher's first stage re-enters and lands using a parafoil system guided onto a cooperative drone Sea-platform (SPR) equipped with a 60 m × 60 m net. This scenario builds on SR heritage, leveraging existing aerodynamic data for a 6 t parafoil, while introducing autonomous cooperation with the SPR and tightening the landing accuracy requirement from a radius of 150 m to less than 30 m. The HALER system consists of a parafoil attached to the recoverable Payload (PYL), operating in cooperation with the sea-drone platform. It interfaces with the payload's navigation unit (when available, as assumed in the present study case) and integrates the remaining subsystems, distributed between the payload and the SPR, necessary to control the parafoil and accomplish the mission, as illustrated in Fig. 1. Although the system is designed to be compatible with multiple recovery scenarios (e.g., ground landings), this work focuses on the cooperative-based approach.

This paper presents an overview of the HALER project results, beginning with the Concept of Operations (ConOps) in Section 2 outlining the recovery process from deployment to refurbishment. Section 3 then summarizes the Guidance, Navigation and Control (GNC) design, including operational modes, algorithms, and supporting analyses, followed by the results of the Monte-Carlo (MC) campaign in Section 4. Finally, Section 5 provides a feasibility assessment, highlights key challenges, and offers concluding remarks.

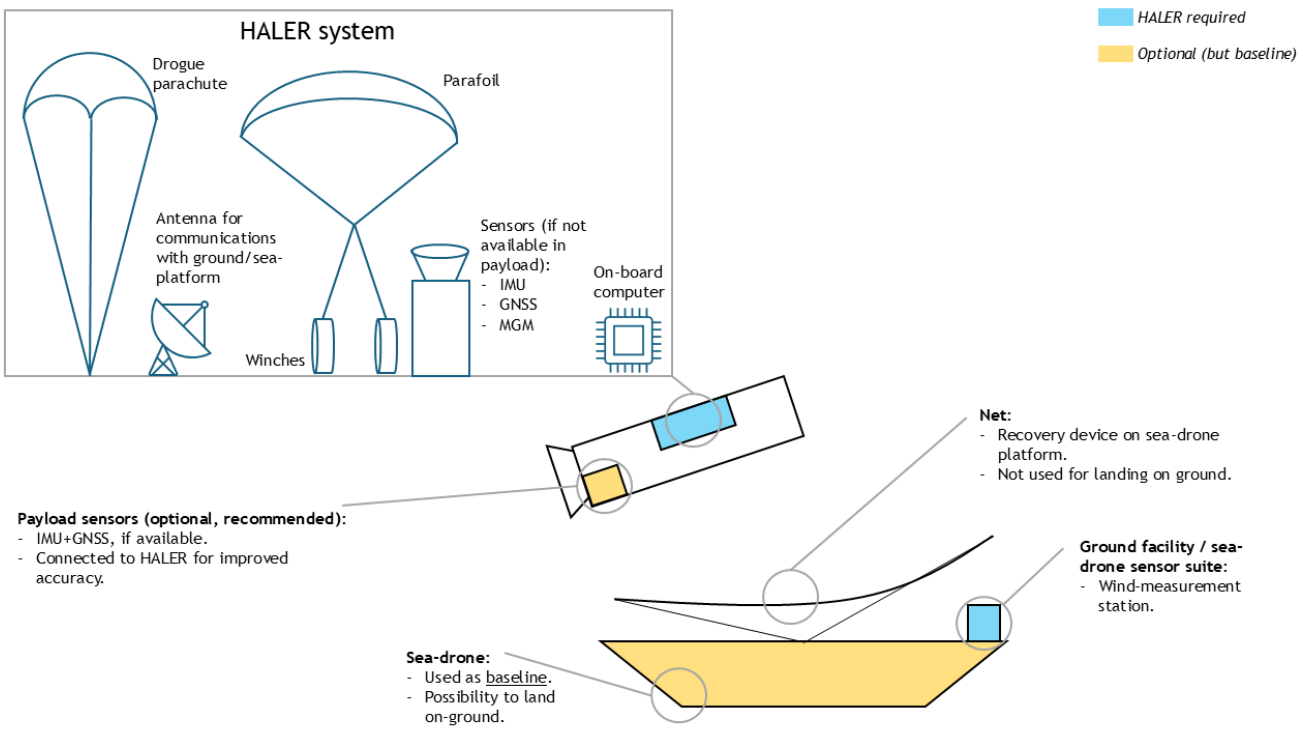

Fig. 1 HALER concept sketch

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

2 Concept of Operations

Coordinating two moving vehicles for a synchronized landing is inherently complex, making a well-defined Concept of Operations (ConOps) essential for mission feasibility. This section outlines the main operational phases, describing the expected behaviour of both vehicles throughout the descent, together with the wind-related operations that are critical for parafoil accuracy. A graphical overview of the ConOps is shown in Fig. 2.

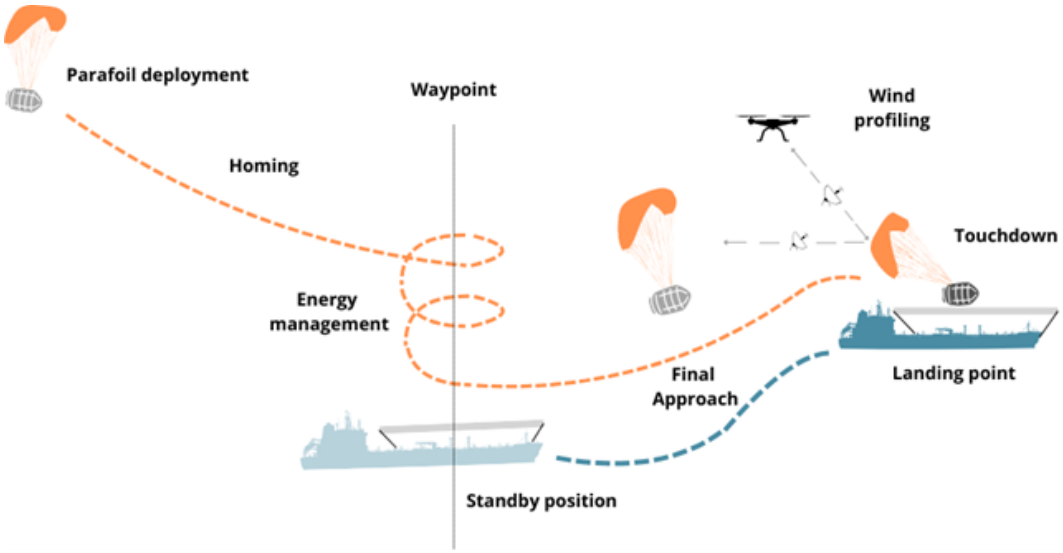

Fig. 2 Concept of Operations sketch for sea-platform recovery

As illustrated, the ConOps covers the entire descent and landing sequence, from parafoil deployment to touchdown on the sea-drone platform. The process consists of five sequential phases:

1) **Parafoil deployment:** Recovery begins with parafoil deployment, typically following drogue chute release under suitable speed, altitude, and distance to Landing Point (LP) conditions. Specific deployment requirements are derived to ensure a safe descent.
2) **Homing:** After deployment, the parafoil flies a straight trajectory toward a designated Waypoint (WP) near the LP, correcting deviations accumulated during re-entry and deployment. The WP is defined geographically without altitude, so the parafoil always points toward it regardless of wind. Placement of the WP must ensure that no trajectory from the deployment area intersects restricted zones. When possible, collocating the WP at the landing point is preferred, as it minimizes sensitivity to wind estimation errors (see Fig. 3)[1].
3) **Energy management:** In this loitering phase, the system dissipates excess altitude before initiating the final approach. This reduces the influence of winds by shortening the final trajectory and provides time for synchronization between the parafoil and the SPR. The maneuver is executed as a spiral with a ground-fixed circular projection, where winds are compensated to maintain a constant horizontal distance from the waypoint.
4) **Final approach:** The parafoil then proceeds toward the landing or interception point within the available time. Traditionally, landings are performed upwind to minimize ground-relative velocity and avoid lateral drift. For SPR recovery, alternative directions may be advantageous, for instance depending on sea conditions. The objective for the SPR is to minimize relative velocity with the parafoil system at the interception point.

[1]The wind creates drifts on ground trajectories depending on the predominant wind direction, requiring from larger or shorter paths depending on the exit point from energy management (different colors in the image show different example of trajectories depending on the triggering point). Without a dynamic WP collocation (preferred to avoid from mission design level) and wind knowledge uncertainty, it is better to situate WP over the LP, such that on average against variety of wind profiles suitable trajectories exists and worst case scenarios are avoided.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

5) **Touchdown:** The final maneuver aims to minimize impact forces at landing, achieved using the flare Model Predictive Control (MPC) [6].

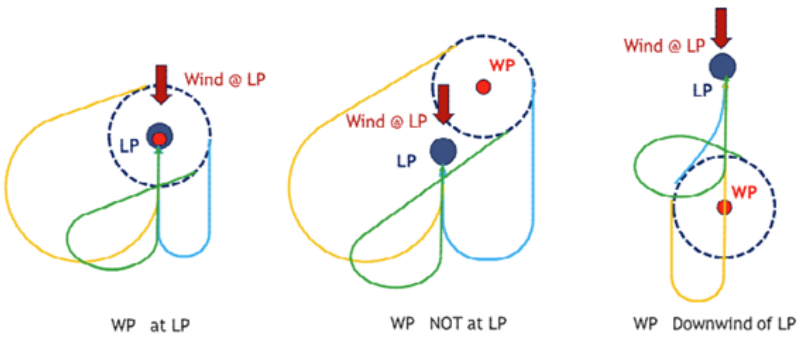

Fig. 3 Waypoint positioning effect on typical trajectories

In addition to the descent and landing phases, accurate wind characterization is essential for high-accuracy parafoil landings. With relatively low airspeeds, parafoils are highly sensitive to wind conditions, which can vary significantly and lead to large trajectory errors if not properly measured. Long-term atmospheric measurements should be collected at the landing site months in advance to build reliable forecasting models, using data sources such as European Centre for Medium-Range Weather Forecasts (ECMWF). Closer to the landing event, updated measurements —taken hours or minutes beforehand— are required to refine these forecasts. Moreover, during recovery, the SPR will be equipped with a weather station and wind-profiling sensors (e.g., Light Detection And Ranging (LiDAR) or Sonic Detection And Ranging (SODAR)) to provide real-time measurements. Alternatively, airborne assets such as UAVs or balloons can be deployed to characterize atmospheric conditions.

Finally, for reuse, the parafoil must undergo inspection and refurbishment to repair potential damage such as cut or weakened chords. Repacking the parafoil is a time-consuming process that requires specialized facilities. To minimize downtime and maintain launch cadence, the HALER system is designed to be modular: a used parafoil can be removed for refurbishment while a replacement unit is immediately installed.

3 GNC design

This section provides a high-level overview of the GNC design, covering both the overall architecture and the individual functional components. To guide the parafoil during re-entry, a dedicated GNC has been developed as an evolution of the SR system [3]. The HALER GNC consists of the functional blocks illustrated in Fig. 4. Some of these functions are distributed across the SPR; for example, navigation sensors and wind-related tasks (such as wind profiling and estimation) are hosted on the SPR, yet remain integral to the HALER system.

The interfaces between the HALER system and the SPR are not depicted in the previous figure. The interface consists of a radio-frequency communication link through which both systems exchange information regarding sensor measurements, navigation estimates, and other G&C parameters including the landing point and agreed time and velocity. Most data is exchanged at 1 Hz, with the exception of the uplink wind and Global Navigation Satellite System (GNSS) Precise Point Positioning (PPP) messages, which are sent every 5 min, due to limited measurement frequency and bandwidth.

The GNC modes are shown in Fig. 5. With the exception of the NAV_ONLY mode, each corresponds directly to a phase of the ConOps bearing the same name. It should be noted that these represent the nominal GNC modes; additional modes related to Failure Detection, Isolation and Recovery (FDIR), standby, or passivation are not included here.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

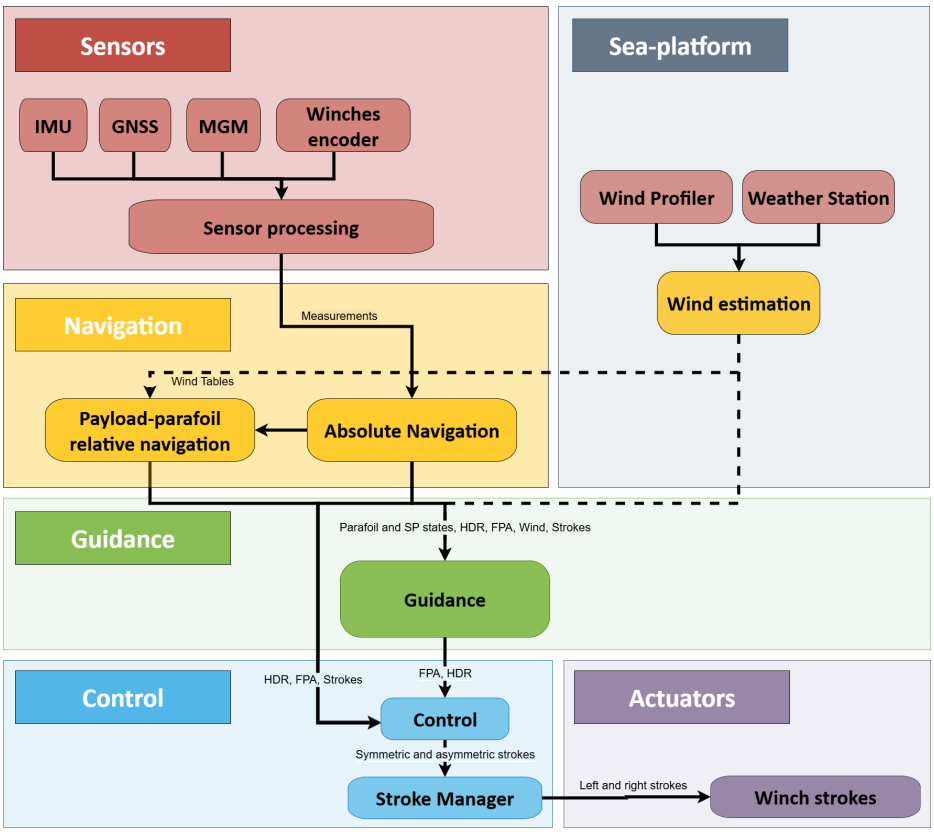

Fig. 4 HALER GNC subsystem's functional architecture

3.1 Guidance

Given the weak (dynamic) coupling between the longitudinal and lateral dynamics of the parafoil system, the GNC is structured to handle these motions independently (with the exception of the final approach, where some coupling may be beneficial). This modular approach simplifies algorithm design, as each algorithm focuses on a single state.

- **Longitudinal algorithms:** These generate a reference Flight Path Angle (FPA). Depending on the mission phase, the commanded FPA may be constant or set to achieve the glide slope required to reach a target. During touchdown, commands are issued directly to the winch speeds via MPC.
- **Lateral algorithms:** These generate a reference Heading Rate (HDR). The commanded rate can be derived either from a heading-tracking logic or from an optimal trajectory generator.

To avoid issuing unfeasible command pairs, i.e., those leading to unattainable responses or exceeding actuator limits, protection functions are applied. Requested FPA and heading rates are saturated, along with their rate of change, according to the design parameters tuned for each GNC mode. This ensures safe operation while prioritizing different behaviors depending on the system's capability and the chosen design philosophy. Importantly, these constraints are explicitly embedded in the optimization schemes for the final approach and touchdown.

For each phase defined in the ConOps, the guidance system employs a tailored strategy. The overall objective is to generate feasible trajectories in terms of FPA, Heading (HDG), and HDR, updating them as a function of the current position and velocity. The only exception is the touchdown phase, where guidance and control are merged into an MPC framework [6]. The strategies for each mode are as follows:

1) **Parafoil deployment:** Guidance remains passive, issuing no commands.
2) **Homing:** The vehicle follows a straight path from the deployment point to the waypoint, maintaining constant ground speed and correcting heading deviations due to wind. The FPA is held constant, selected to maximize range while retaining sufficient heading-rate authority.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

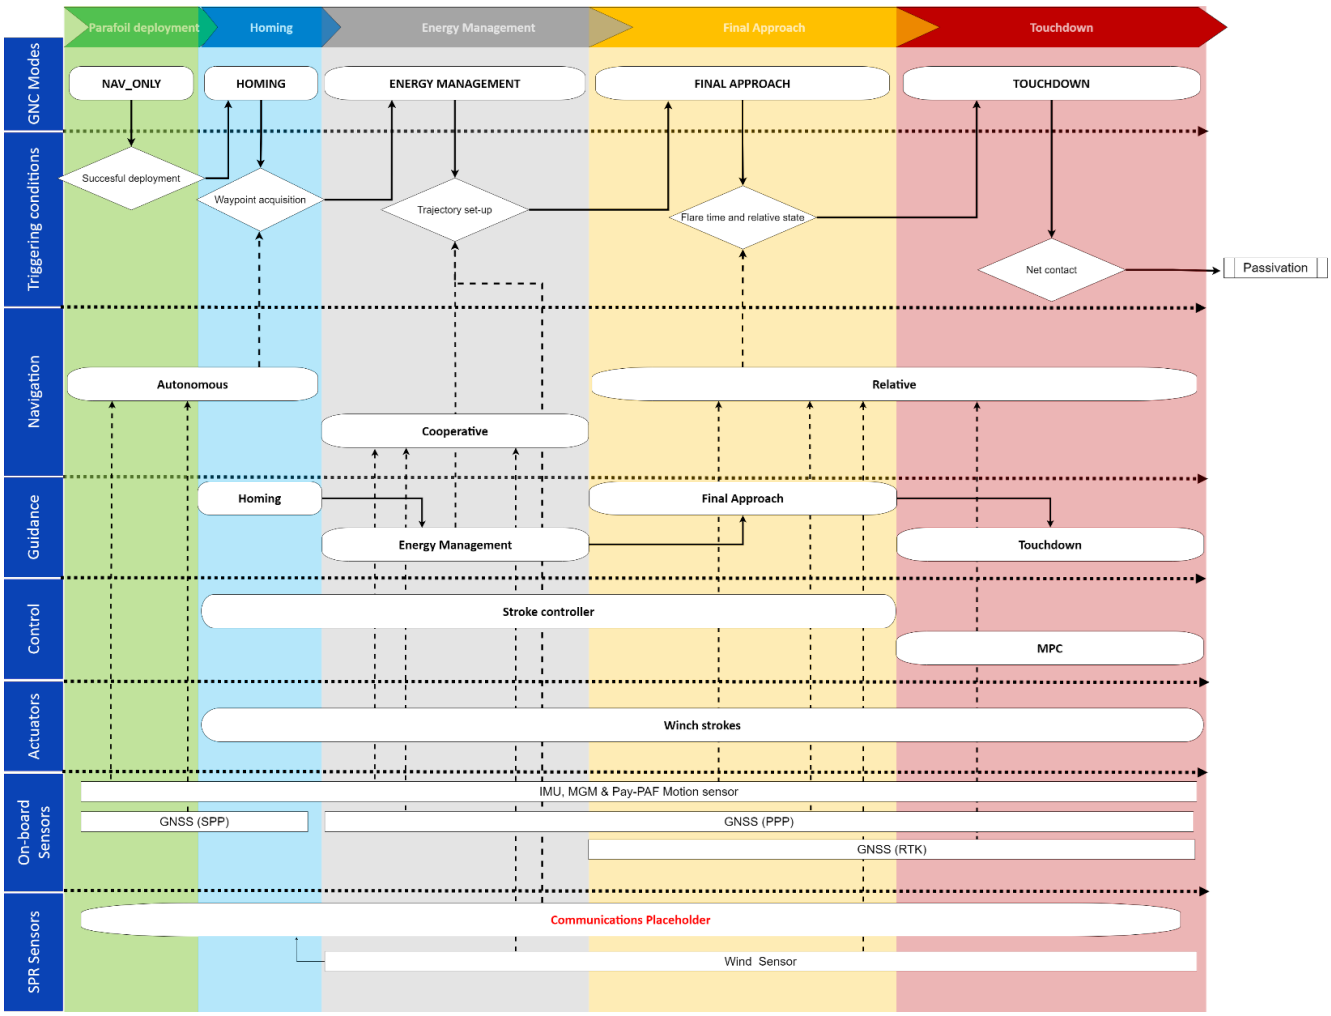

Fig. 5 GNC modes architecture and relation with ConOps

3) **Energy management:** The trajectory tracks a spiral around the waypoint, dissipating altitude and enabling synchronization with the SPR.
4) **Final approach:** An optimization-based trajectory generator [7] computes synchronized landing paths against the prevailing wind direction.
5) **Touchdown:** A closed-loop flare manoeuvre [6] minimizes touchdown velocity (vertical and horizontal trade-off). Heading and heading rate are set to zero, while the FPA is commanded through the MPC.

The final approach strategy relies on on-board convex optimization, executed in real time to merge multiple objectives and adapt the trajectory to wind estimates and sea-platform motion updates. Two cooperative schemes, with option (b) being the one that was finally selected, were explored during the design process:

(a) **Cooperative capture:** The parafoil trajectory is predicted onboard, and the sea platform adapts its own trajectory accordingly. The parafoil does not directly control the SPR, but provides updated landing objectives (point and velocity) to minimize errors and relative velocities.
(b) **Chaser capture:** The parafoil aligns its trajectory with the sea platform, nominally within a corridor aligned with the wind direction. The SPR adjusts its velocity and small lateral displacements to match the predicted landing point, feeding these corrections back to the parafoil system for the next guidance update.

The underlying Optimal Control Problem (OCP) is convex except for the parafoil dynamics, whose non-convexities can be efficiently addressed through Sequential Convex Programming (SCP) [8]. At

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

(a) Cooperative capture **(b) Chaser capture**

Fig. 6 Parafoil final approach guidance schemes

Sener, dedicated tools for auto-coding the OCP formulation and SCP implementation have already been developed and validated, focusing on feasibility within the computational constraints of on-board systems. The SCP-based guidance is explicitly tailored for computational efficiency, as detailed in [7]. Specifically, the optimization problem is solved at 2.5 Hz, with computational times in user computer (i7-12800H (2.40 GHz) processor) in the order of 0.5 seconds. The scope of the work for the SCP guidance was functional verification, and further work is expected towards on-board implementation and real-time execution times by code tailoring process. A similar process was followed for the MPC used in the flare [6], heritage from Space Rider, for which maximum computational time of 3.5 ms was demonstrated in a GR740 LEON 4 space graded processor. Fig. 7 depicts generated guidance profiles for random initial conditions all successfully reaching the SPR.

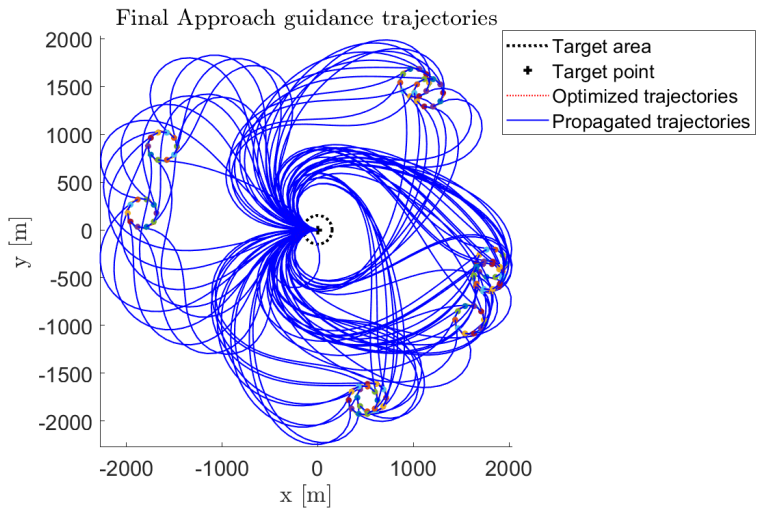

Fig. 7 Parafoil final approach guidance initial conditions robustness

3.2 Navigation

To achieve the required performance, the navigation solution adopts a collaborative architecture, combining sensors located both onboard the payload and on the sea platform. Sensor measurements are exchanged via the data link between PYL and SPR, allowing each system to remain informed of the other's state while also sharing additional information about the environment. The overall system architecture is shown in Fig. 4.

The navigation modes are determined by the availability of sensors and communication links, which in turn define the algorithms employed. Four modes are defined:

1) **Re-entry (REE):** Active before parafoil deployment, mainly relies on inertial navigation.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

2) **Autonomous (AUT):** Initiated after parafoil deployment, it operates without direct communication with the SPR. Some information (e.g., wind table initialization) must be preloaded or transmitted via the ground station. Consequently, wind estimation is limited, resulting in reduced state-estimation performance.
3) **Cooperative (COP):** Once the system is within communication range of the SPR, data exchange via the link becomes active. This enables reception of wind tables from the wind profiler and PPP-enhanced GNSS data, significantly improving navigation performance.
4) **Relative (REL):** In the final mission phase, when the target is the recovery net on the SPR, the Real-Time Kinematics (RTK) augmentation technique is employed to maximize navigation accuracy and minimize landing error.

The corresponding operational phases and sensors used in each mode are shown in Table 1.

Table 1 Navigation modes with corresponding operational phases and active sensors

Nav. mode	Operational phases	On Payload					On Sea-Platform	
		IMU	GNSS (SPP)	GNSS (PPP)	GNSS (RTK)	MGM	Wind profiler	Weather station
REE	Pre-parafoil	x	x*			x		
AUT	Parafoil Deployment and Homing	x	x			x		
COP	Loitering	x		x		x	x	x
REL	Final Approach, Flare and Touchdown	x		x	x	x	x	x

*Depending on availability (GNSS blackout during re-entry).

Absolute navigation is achieved through the hybridization of the Inertial Measurement Unit (IMU), GNSS, and Magnetometer (MGM) sensors. The IMU provides high-frequency but noisy measurements, which are corrected using the lower-rate yet absolute measurements from the GNSS (for linear states) and the MGM (for angular states). The hybridization algorithm implementing this process is illustrated in Fig. 8.

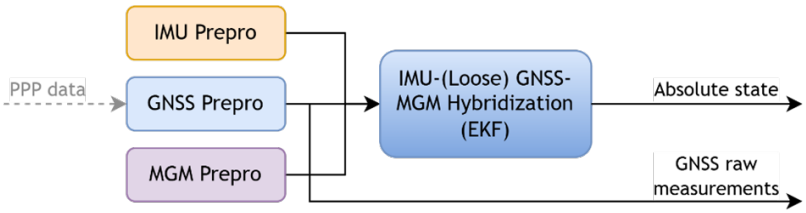

Fig. 8 Absolute navigation algorithm diagram

First, sensor measurements are preprocessed to ensure proper time-tagging, scaling, and rotation alignment. In cooperative or relative modes, the GNSS operates with PPP corrections received via the data link; otherwise, it defaults to Single Point Positioning (SPP) mode. Certain raw measurements, such as those from the GNSS, are also output for transmission through the data link. The core hybridization algorithm then fuses the IMU, GNSS (in loose coupling), and MGM measurements using an Extended Kalman Filter (EKF) [9]. The navigation mode determines the effective GNSS precision (SPP, PPP, or RTK), which is incorporated into the filter and directly affects the accuracy of the absolute state estimates as shown in Figs. 9 and 10.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

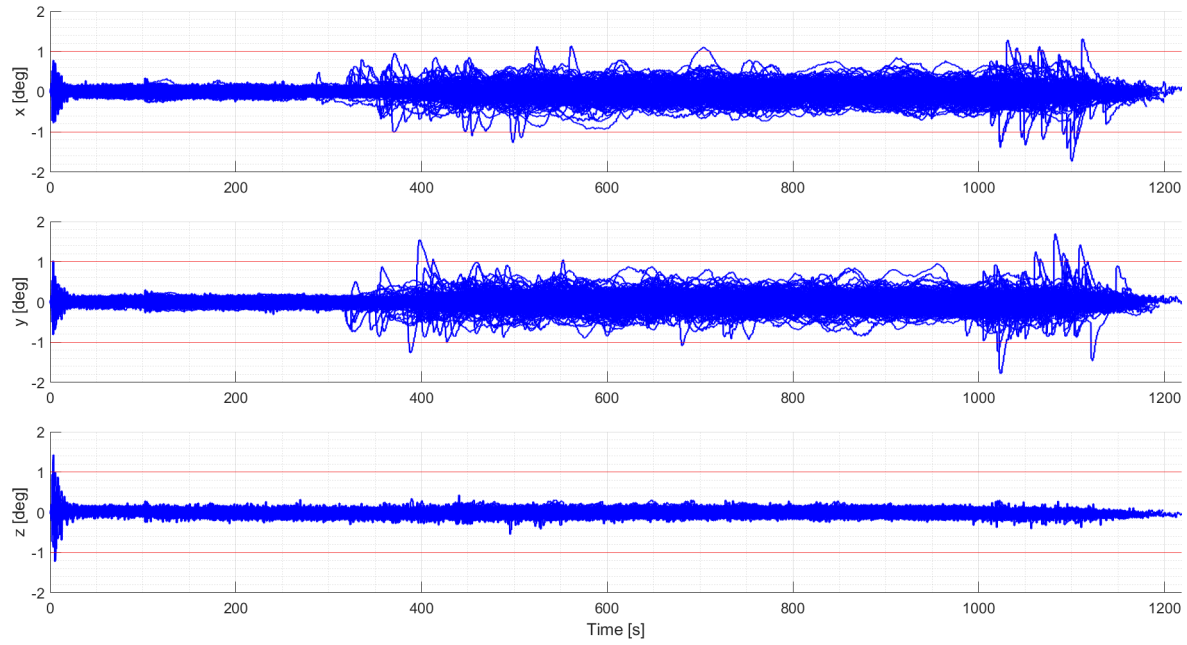

Fig. 9 Attitude knowledge error for 250 MC cases (blue) with requirement lines (red)

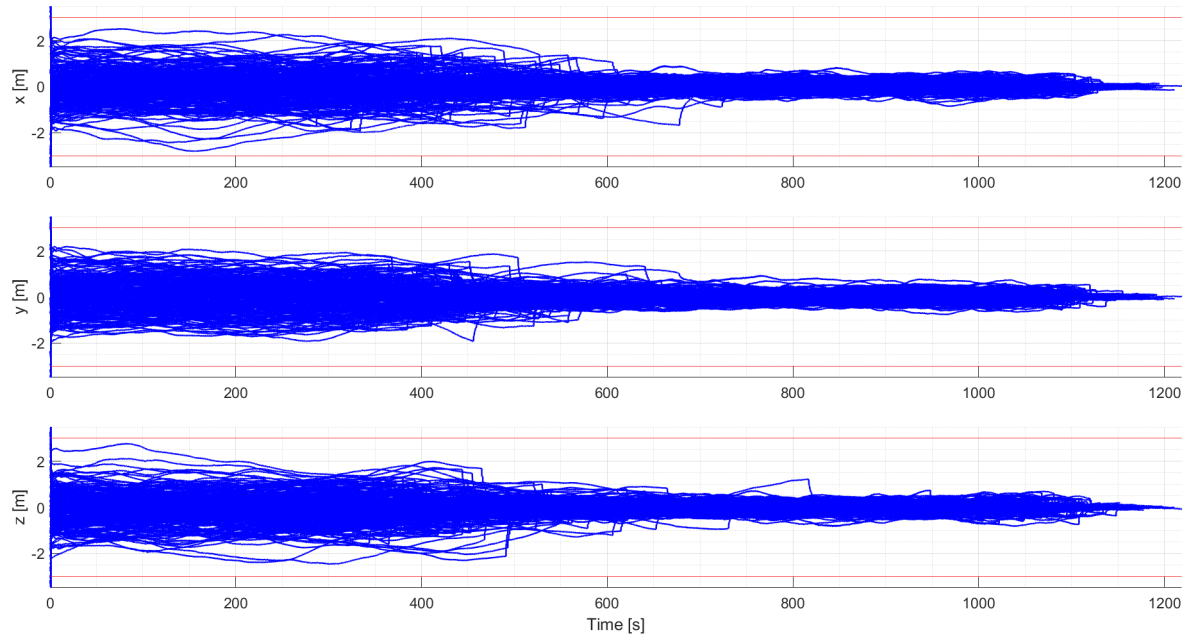

Fig. 10 Position knowledge error for 250 MC cases (blue) with requirement lines (red)

Since parafoil flight takes place in a dynamic environment dominated by wind and controlled through aerodynamic actions, it is essential to estimate the system states relative to the air, i.e., the air states. Two methods are considered for their estimation:

- **Look-up tables:** The Flight Path Angle (FPA) and airspeed V_{AIR} are estimated using precomputed Lookup Tables (LUTs) derived from open-loop simulations.
- **Wind measurements:** The velocity vector of the parafoil with respect to the air in the local frame is obtained by subtracting wind speed from the ground velocity measured by GNSS. From this vector, FPA and Heading (HDG) are derived as the angles with respect to the horizontal plane and the north direction, respectively.

With appropriate wind sensors and expected aerodynamic uncertainties, the second method is expected to be more accurate than the LUTs, since having measurements is a more adaptable and robust

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

system than relying on precomputed or simulated data. Both factors, accuracy and robustness, are critical to meeting the landing requirements. And therefore, the second method is the one used in the simulations presented in Section 4.

The Heading Rate (HDR) is updated under the assumption of negligible sideslip, i.e., the airspeed lies in the system's longitudinal symmetry plane. This allows the rate of change of the airspeed vector to be interpreted as an attitude change, measurable by onboard sensors. The heading rate is then approximated by projecting the yaw rate onto the Local Vertical (LV) reference frame. To ensure robustness, a windowed Hampel filter [10] is applied to this signal for outlier detection and replacement.

3.3 Control

The guided parafoil control problem consists of actuating two winches that pull or release the parafoil lines, thereby modifying the dynamic properties of the system. These inputs affect both the flight path and the heading rate. Asymmetric deflections are used to control the HDR, and thus the heading, while symmetric deflections control the FPA and airspeed, resulting in changes to the vertical and horizontal velocity components.

The dynamics model previously developed for SR can also be applied in the present activity [3]. This validated model describes the full 6-Degrees of Freedom (DoF) motion of the parafoil assembly, augmented by the relative yaw angle between the parafoil and the payload. The result is a 7-DoF model of the parafoil–payload system, also employed in Functional Engineering Simulator (FES) simulations. Based on this model, three controllers are designed, namely: 1) the HDG feedback controller; 2) the HDR feedback + feedforward controller; and 3) the FPA feedback + feedforward controller. These controllers are synthesized using an H_∞ design approach and subsequently assessed for robustness, supported by a sensitivity analysis based on Morris' method [11, 12].

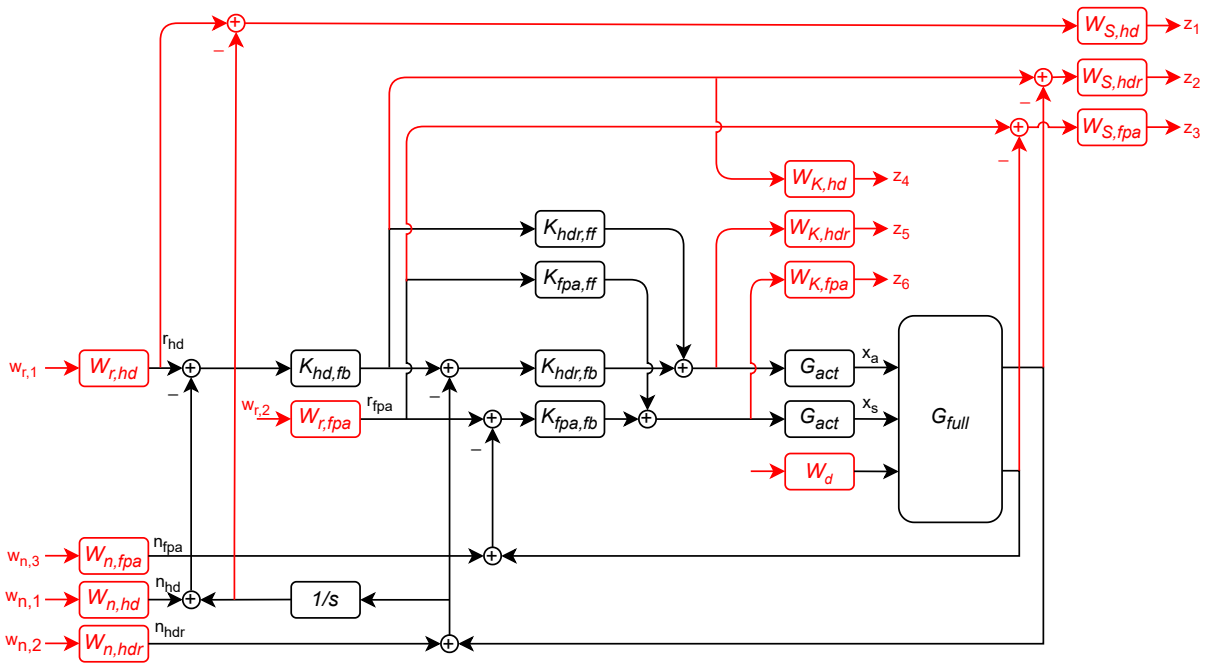

Fig. 11 Overall combined (heading rate + flightpath angle) control interconnection

An actuator manager is then included to allocate compatible stroke commands, ensuring command feasibility and prioritizing specific channels according to the flight phase. Finally, during the touchdown and landing phase, the control functions are replaced by a dedicated MPC designed for parafoil flare execution [6].

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

Given the synthesized controllers, the next objective is to demonstrate that the overall control system remains robustly stable and preserves performance in the presence of uncertainties and disturbances. The corresponding uncertainty model, associated with the system interconnection in Fig. 11, encompasses 48 uncertain parameters including aerodynamic coefficients, Mass, Centering and Inertia (MCI) parameters, parafoil/payload geometric properties, and air density. Lower- and upper-bound computations confirm that the overall control interconnection is robustly stable and performs satisfactorily under the modeled uncertainties. Due to the large number of uncertainties and their repetitions in the model, a full-scale robustness analysis is computationally intractable [13]. Instead, a sensitivity analysis using Morris' method [12] is first performed. This sampling-based method quantifies the effect of uncertainty variations on performance outputs. The Results are shown in Fig. 12.

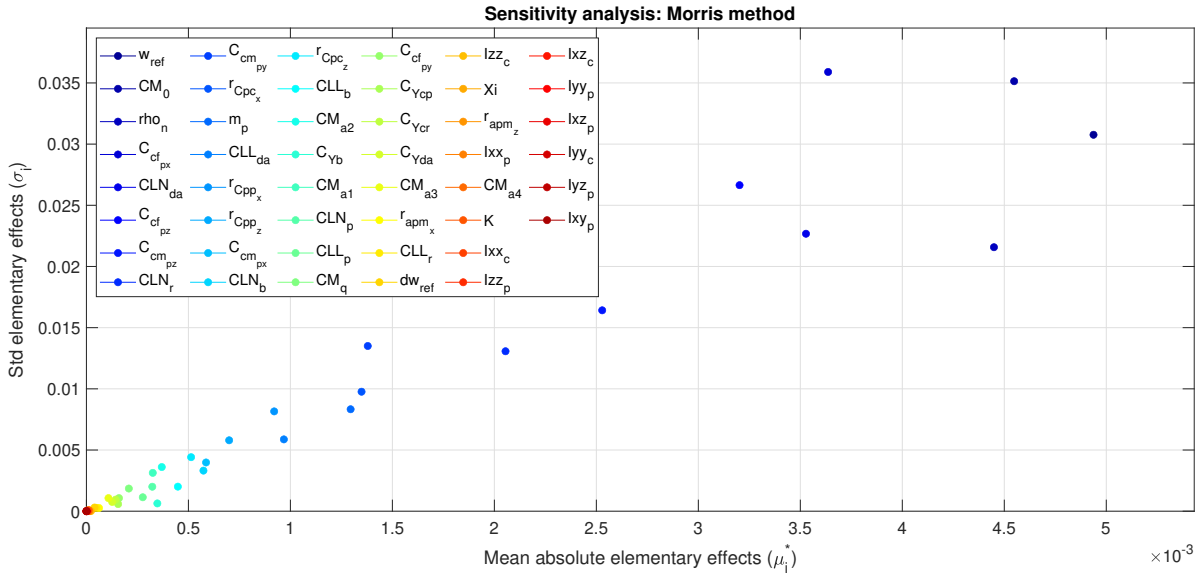

Fig. 12 Sensitivity analysis using Morris' method

As can be seen, the mean absolute elementary effects (μ_i^*) are very small, indicating that the system is largely insensitive to parametric uncertainties. Similarly, the standard deviation of the elementary effects (σ_i), which measures correlation with other parameters, is also small. For the robustness analysis, the first eight most influential parameters are selected for a μ-robustness study. Since the normalized air density (ρ_n) is repeated 52 times in the closed-loop weighted uncertain plant, the computational burden for a μ-analysis remains high. To address this, robustness analyses are performed over a grid of ρ_n values. Figure 13 shows the Structured Singular Values (SSV) computed for different air densities, allowing determination of the robust stability margin.

The results indicate that the system is robustly stable, tolerating up to 180% of the modeled uncertainties. Robust performance is then verified by computing upper bounds of the worst-case gains for different ρ_n values, assuming 100% of modeled uncertainty in the performance channels. The worst-case gains from reference inputs to tracking error outputs are guaranteed to remain below 1.05, 0.8, and 1.05 for the heading, heading rate, and flightpath angle channels, respectively, for all selected uncertainties.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

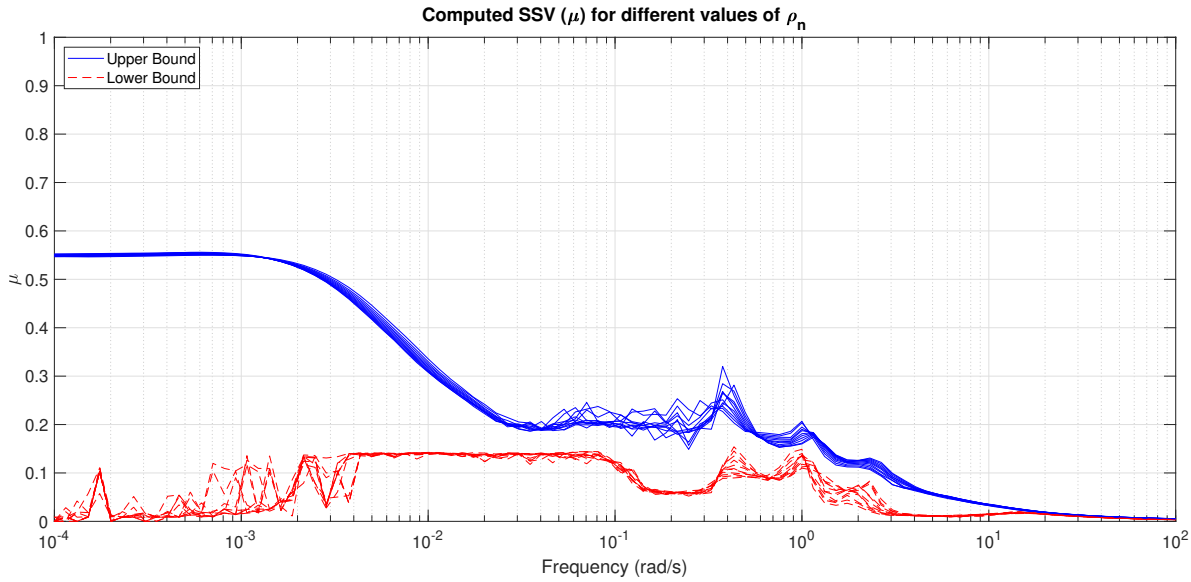

Fig. 13 Robust stability: Computed SSV (μ) upper (blue) and lower (red) bounds

4 Test results

This section presents the results obtained from different test scenarios, consisting of MC campaigns with 250 shots each. In every simulation run, the DKE conditions were varied, including initial states, payload MCI, atmospheric and wind conditions, as well as sensor and actuator noise, among others. Across all tests, the objective was to land within the 60 m by 60 m net on the SPR, while ensuring maximum relative velocities at touchdown do not exceed 15 m/s horizontally, 2.5 m/s vertically, and 1 m/s laterally.

4.1 Simulation environment

The test results were obtained using a high-fidelity simulation environment comprising the Dynamics, Kinematics and Environment (DKE), sensor and actuator models, the parafoil GNC, and the SPR with a dedicated autopilot. The parafoil dynamics were characterized by a 7-DoF model, previously validated for the SR mission, including "apparent mass" induced by the trapped air. The environmental model features gravity (up to J2), wind and atmosphere based on ECMWF Fifth generation ECMWF Re-Analysis (ERA5) data (wind, density, pressure, temperature), and an air data module deriving aerodynamic angles, Mach number, and dynamic pressure. Sensors and actuators are represented by behavioural models: on-board units include an IMU (STIM380H), GNSS receiver (NavRIX), MGM (PEGASUS-1), and winches (RBE(H)-O6212-A). The SPR integrates the navigation unit, plus two wind-sensing LiDARs (ZX300M and Windcube Explore).

On the other hand, the SPR dynamics are captured with a simplified 3-DoF model (surge, sway, yaw), combining hydrostatics, hydrodynamics (added mass, drag), and environmental effects from wind, currents, and waves. A kinematics module maps body-fixed to earth-fixed states. For GNC, the SPR uses a behavioural navigation model based on the Seapath 385-R2, providing motion and position data. Guidance commands from HALER specify the desired landing position, velocity direction, and time to land. A geometric strategy switches automatically between four modes: standby, straight, turn, and catch-up. The control is achieved through thrust commands to the main propellers, modelled symmetrically around the ship's axis. The SPR guidance provides the desired speed magnitude and yaw angle, which are tracked by a 2-DoF controller: a Proportional-Integral (PI) loop generates the longitudinal force from the velocity error, and a Proportional-Derivative (PD) loop computes the yaw torque from the heading

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

error. These commands form the generalized force vector $\boldsymbol{\tau} = (F_x, T_z)^T$, which is then mapped into individual thruster forces through a static allocation matrix based on each propeller's lateral distance to the ship's Center of Gravity (CoG).

4.2 Test 1: Stationary SPR

The first test was conducted with a stationary SPR, serving as a baseline for subsequent experiments to evaluate the benefits of a cooperative approach compared to a case like SR. The results show that, although the landing accuracy improves compared to SR, the required precision cannot be achieved with a static SPR. This is illustrated in Fig. 14, which highlights the dispersion of the landing error.

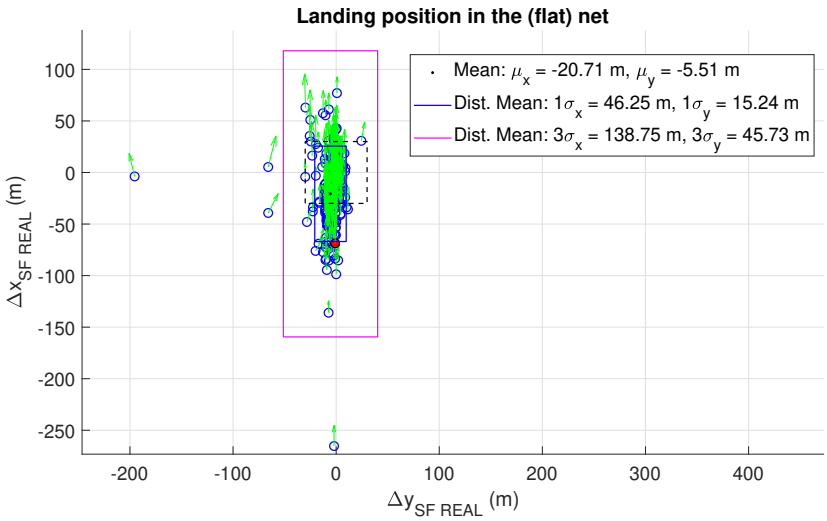

Fig. 14 Landing position (blue) in the flat net (dashed black line) with relative velocity arrows (green) and 3σ error bounds (magenta) of test with fixed landing location for wind speeds up to $10\,\mathrm{m/s}$

The dispersion is most pronounced in the X-direction (along-wind) due to two main factors: (a) the guidance system is tuned to minimize Y-direction error, leaving X-errors to be corrected by the SPR through acceleration or deceleration; and (b) certain aerodynamic–wind interactions cause transient variations in the FPA, which accumulate into errors in both time-to-land and trajectory length. While factor (a) could be mitigated, doing so would increase dispersion along Y. Factor (b) is more sensitive to wind intensity, since reduced ground speed leaves less margin for error correction.

4.3 Test 2: Cooperative SPR with perfect auto-pilot

For this test, the landing location is intentionally varied, assuming that the SPR can perfectly reach the HALER-commanded LP. The guidance system continuously adapts the commanded LP based on environmental conditions and the manoeuvrability limits of the vehicle, updating the target landing coordinates and timing during approach as prediction accuracy improves. The test was performed under different maximum wind conditions (i.e., $4\,\mathrm{m/s}$, $6\,\mathrm{m/s}$, $8\,\mathrm{m/s}$ and $10\,\mathrm{m/s}$) and with random winds including turbulence, as shown in Fig. 15.

Comparing the different wind intensities, we find that landing accuracy is not significantly degraded by stronger winds. However, higher wind speeds induce greater variability in the LP commands. Notably, for wind intensities above $8\,\mathrm{m/s}$, the agility demanded of the SPR exceeds the real system's capabilities, a limitation further analyzed in Section 4.4. Despite these effects, LP command variations are generally contained within $50\,\mathrm{m}$ in most simulations, and below $100\,\mathrm{m}$ in all cases except for the $10\,\mathrm{m/s}$ scenario. These variations are primarily aligned with the wind direction, deliberately tuned to favor SPR acquisition of the LP through acceleration and deceleration rather than large trajectory adjustments.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

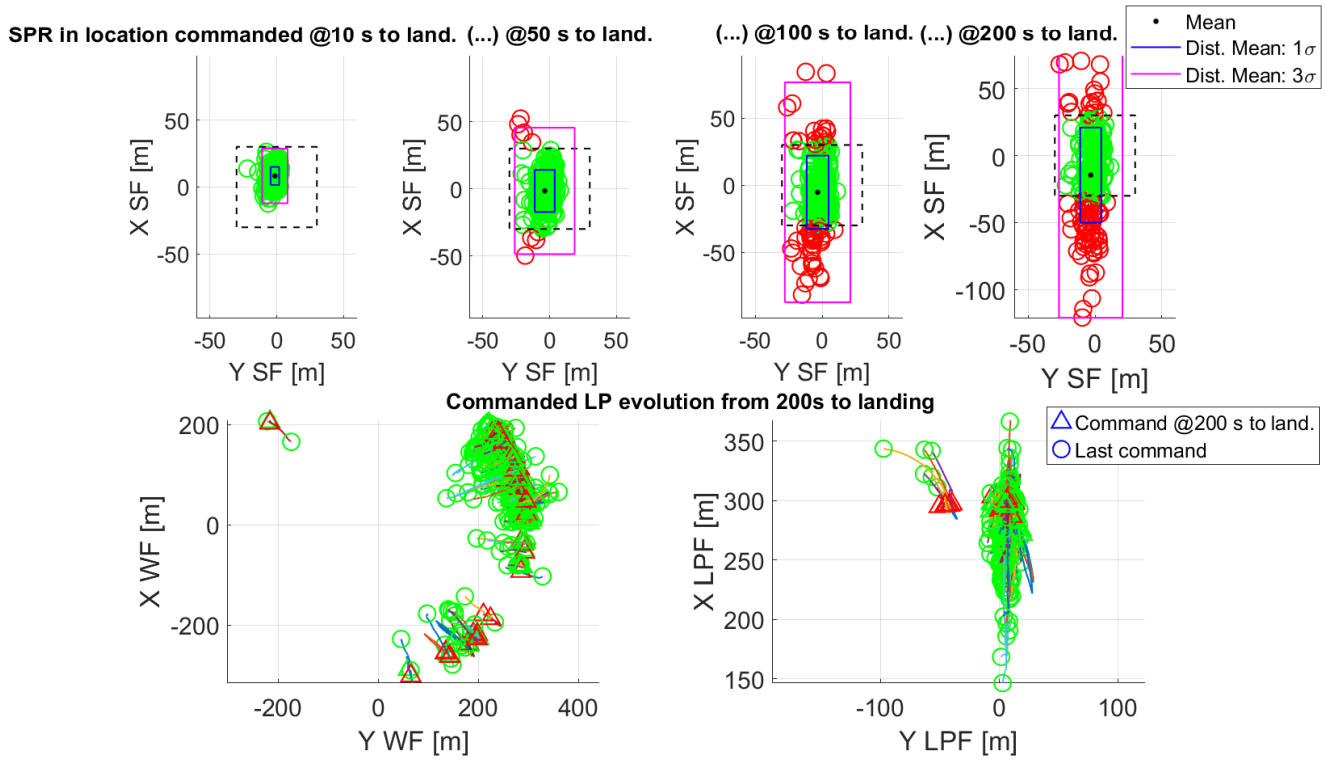

Fig. 15 Landing accuracy prediction (top) —green circles denote success cases (inside the net, i.e., within the dashed black bounds) while red cases land outside the net, magenta bounds represent the 3σ error in landing— and landing point evolution (bottom) —colored lines connect the first (triangle) and last (cicle) commanded LPs in Waypoint Frame (WF) (left) and Landing Point Frame (LPF) (right)— for winds up to $10\,\mathrm{m/s}$

4.4 Test 3: Cooperative SPR with representative autopilot

For this test, the simplified drone ship autopilot was included in the loop, thereby accounting for the effect of SPR agility on landing precision. The results are shown in Fig. 16.

As can be seen, the required landing performance is statistically achieved, as the flare-induced bias in the wind direction is effectively compensated [6]. Moreover, the outcomes are consistent with those obtained under the assumption of an ideally varying landing location, indicating that the HALER GNC properly accounts for SPR limitations when commanding LP adjustments during the Final Approach. Outlier cases were not analyzed within the scope of this project. However, it is expected that a more detailed SPR GNC design and additional tuning could eliminate these occurrences and further improve overall performance. Conversely, higher-fidelity simulations may reveal some degradation once simplifying assumptions, such as flat sea conditions, are removed.

Regarding touchdown conditions, both vertical and horizontal velocities remain within the specified requirements, as shown in Fig. 17, where the limits are indicated by the red line. Nevertheless, some outliers appear in the lateral velocity, suggesting imperfect alignment between the parafoil and SPR trajectories. This is a consequence of prioritizing landing accuracy over directional alignment. Further development of the SPR GNC is expected to mitigate this effect. Finally, with respect to the touchdown attitude, the roll angle satisfies the requirements, with a 3σ value below $0.6°$. In contrast, the touchdown pitch shows a bias towards negative values, ranging from $-5°$ to $1.5°$. This bias is independent of wind intensity and cannot be mitigated through GNC adjustments, since the flare MPC already operates at its maximum capacity. A straightforward solution would be to correct it via the parafoil rigging angle. Alternatively, the recovery net may be able to absorb the impact sufficiently, making a positive landing pitch unnecessary.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

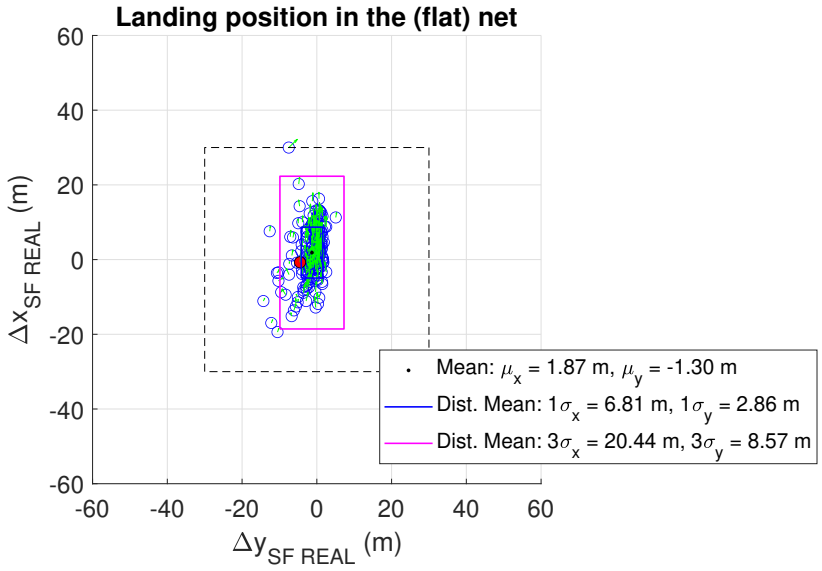

Fig. 16 Landing position (blue) in the flat net (dashed black line) with relative velocity arrows (green) and 3σ error bounds (magenta) of test with SPR autopilot for winds up to $10\,\mathrm{m/s}$

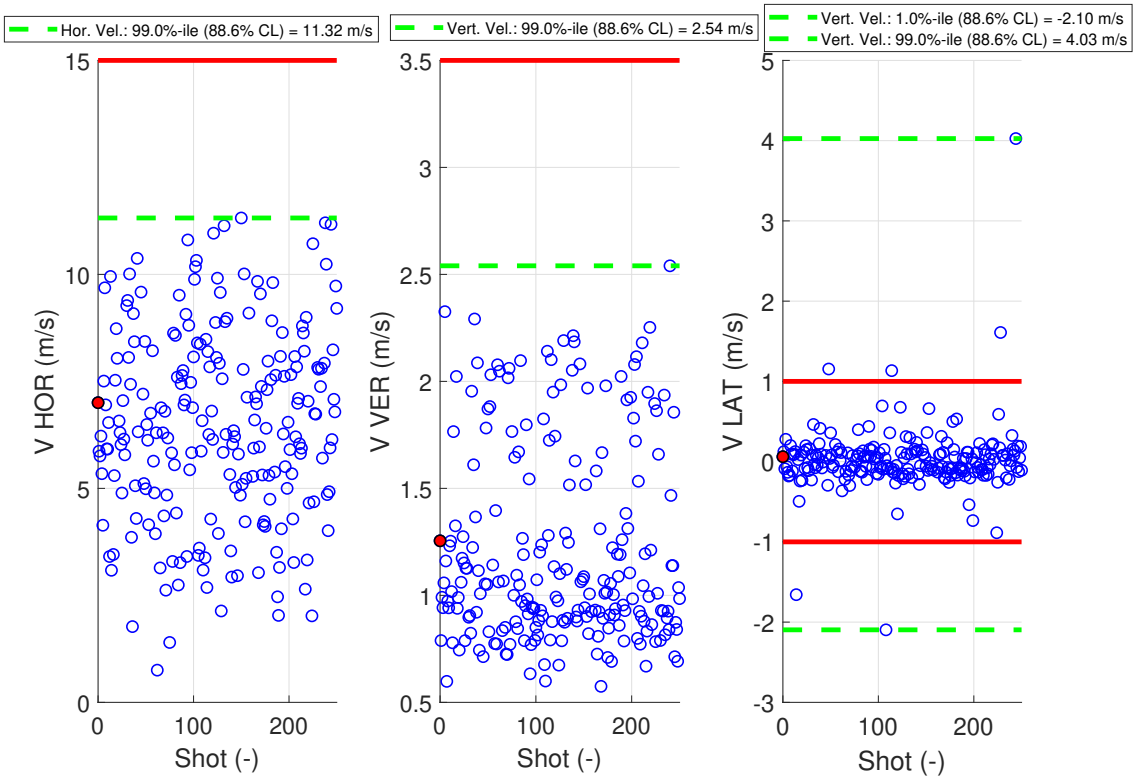

Fig. 17 Touchdown velocity (blue) with 3σ bounds (dashed green) and requirement (red) lines for winds up to 10 m/s

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

5 Feasibility assessment and conclusions

Being a generalization and evolution of the SR design, the HALER building block has been developed with a state-of-the-art parafoil GNC solution integrating modern navigation techniques, an optimization-based guidance, and a robust control scheme. A detailed simulation environment was implemented to evaluate the GNC design for various scenarios (Sections 3 and 4). From a GNC point-of-view, the HALER concept demonstrates feasibility regarding landing accuracy, touchdown velocities, and attitude. Key conclusions from the simulation campaigns are:

- **Landing execution:** The GNC successfully guides the parafoil assembly to a designated landing site, autonomously updating the trajectory under wind variations and other uncertainties. The landing performance is largely invariant to wind intensity, supporting precise landings even at speeds of 10 m/s.
- **Navigation and control:** The navigation system reliably estimates the system-state with minimal error, while wind measurements are sufficiently accurate for optimized guidance. The controller exhibits robust performance and sufficient bandwidth to execute commanded trajectories.
- **Interaction with the sea drone:** In scenarios where the drone ship is stationary, the net is insufficient to guarantee a successful capture, highlighting the need for active drone-parafoil coordination. When the drone follows commanded positions, the guidance anticipates its manoeuvreability limitations, allowing successful capture in most cases. Simplified autopilot implementations sometimes produced unrealistic behaviour, but these issues are expected to be resolved with a more comprehensive SPR GNC.

Several areas for refinement and further study were identified:

- **Requirement definition:** Some legacy requirements (e.g., maximum wind speed and touchdown pitch) could be revised, as performance is invariant to wind intensity and net impact absorption may mitigate pitch biases.
- **Environmental modeling:** Higher-fidelity models, particularly of the ship dynamics and net interaction, are recommended to accurately assess landing performance.
- **Flare-induced wind bias:** This can be compensated by adjusting LP commands and accounting for vertical velocity effects on predicted timing.
- **SPR manoeuvreability and agility:** Limits in agility affect the drone's ability to reach the commanded landing location, impacting the net design and the go/no-go criteria for wind conditions.

A few outliers were observed in lateral velocity and pitch, due to prioritizing landing accuracy over alignment. These could be mitigated through improved GNC tuning or minor mechanical adjustments, such as parafoil rigging. Higher-fidelity simulations may reveal performance degradation if simplifying assumptions (e.g., flat sea) are removed, indicating areas for continued study. In summary, the HALER building block demonstrates robust, precise, and feasible performance under realistic operational conditions, while identified challenges point to clear avenues for future improvement and refinement.

Acknowledgments

The authors are grateful for the funding provided by the European Space Agency (ESA) in the scope of the project "Actively Controlled Parafoil Building Block for Increased Precision Landing", contract no. 4000144651/24/FR/MA, which was carried out by Sener Aeroespacial S.A. (Spain) in collaboration with IDS (Italy), CIRA (Italy) and CIMSA (Spain).

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.

Declaration of Use of Artificial Intelligence

Artificial intelligence was used to prepare the manuscript to improve the readability of the text in terms of language and grammar.

References

[1] Jenny Stein, Chris Madsen, and Alan Strahan. An Overview of the Guided Parafoil System Derived from X-38 Experience. In *18th AIAA Aerodynamic Decelerator Systems Technology Conference and Seminar*, 05 2005. ISBN: 978-1-62410-054-3. doi: 10.2514/6.2005-1652.

[2] Roberto Angelini and Angelo Denaro. IXV re-entry demonstrator: Mission overview, system challenges and flight reward. *Acta Astronautica*, 124:18–30, 2016. ISSN: 0094-5765. Intermediate Experimental Vehicle (IXV) – Special Publications from the 6th EUCASS. doi: 10.1016/j.actaastro.2016.02.015.

[3] F. Cacciatore, R. Haya Ramos, L. Tarabini Castellani, A. Figueroa, J. Veenman, S. Ramírez, C. Recupero, M. Kerr, and J.A. Bejar. The Design of the GNC of the Re-entry Module of Space Rider. In *8th EUropean Conference for Aeronautics and Space Sciences (EUCASS)*, 2019. doi: 10.13009/EUCASS2019-1016.

[4] Antonio Figueroa-González, Francesco Cacciatore, and Rodrigo Haya-Ramos. Landing Guidance Strategy of Space Rider. *Journal of Spacecraft and Rockets*, 58(4):1220–1231, July 2021. ISSN: 1533-6794. doi: 10.2514/1.a34957.

[5] Davide Bonetti, Giovanni Medici, Gonzalo Blanco Arnao, Samuele Salvi, Andrea Fabrizi, and Murray Kerr. Reusable Payload Fairings: Mission Engineering and GNC Challenges. In *8th EUropean Conference for Aeronautics and Space Sciences (EUCASS)*, 2019. doi: 10.13009/EUCASS2019-638.

[6] Jesús Ramírez, Jorge Cardín, H. Gutierrez, F. Cacciatore, and V Preda. Embedded Optimization for Space Rider Reentry Module Parafoil GNC. *International Astronautical Congress (IAC)*, 2023. doi: 10.5270/esa-gnc-icatt-2023-073.

[7] Jesús Ramírez, Antonio Hernández-Rodicio, and Guillermo Alonso. Parafoil Terminal Landing Guidance with On-Board Convex Optimization. In *11th EUropean Conference for Aeronautics and Space Sciences (EUCASS)*, 2025. doi: 10.13009/eucass2025-335.

[8] Jesús Ramírez and Lukas Hewing. Sequential convex programming for optimal line of sight steering in agile missions. In *9th EUropean Conference for Aeronautics and Space Sciences (EUCASS)*, 2022. doi: 10.13009/EUCASS2022-7361.

[9] Paul Groves. *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*. Artech House, Boston, 03 2008. ISBN: 978-1-58053-255-6.

[10] Hancong Liu, Sirish Shah, and Wei Jiang. On-line outlier detection and data cleaning. *Computers & Chemical Engineering*, 28(9):1635–1647, 2004. ISSN: 0098-1354. doi: 10.1016/j.compchemeng.2004.01.009.

[11] Joost Veenman, Carsten W. Scherer, Carlos Ardura, Samir Bennani, Valentin Preda, and Bénédicte Girouart. IQClab: A new IQC based toolbox for robustness analysis and control design. *IFAC-PapersOnLine*, 54(8):69–74, 2021. ISSN: 2405-8963. 4th IFAC Workshop on Linear Parameter Varying Systems LPVS 2021. doi: 10.1016/j.ifacol.2021.08.583.

[12] Max D. Morris. Factorial sampling plans for preliminary computational experiments. *Technometrics*, 33(2):161–174, 1991. doi: 10.1080/00401706.1991.10484804.

[13] Joost Veenman, Carsten W. Scherer, and Hakan Köroğlu. Robust stability and performance analysis based on integral quadratic constraints. *European Journal of Control*, 31:1–32, 2016. ISSN: 0947-3580. doi: 10.1016/j.ejcon.2016.04.004.

Except where otherwise noted, content of this paper is licensed under a Creative Commons Attribution 4.0 International License.

The reproduction and distribution with attribution of the entire paper or of individual pages, in electronic or printed form, including any material under non-CC-BY 4.0 licenses is hereby granted by the authors and respective copyright owners.