77th International Astronautical Congress (IAC), Antalya, Türkiye, 5-9 October 2026.
Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

IAC–26–D2.5.8

GEORyder System-Level Validation of Embedded Guidance and Navigation Algorithms for Autonomous Rendezvous and Docking

Annachiara Ippolito[a,b*] , **Luca Giorcelli**[a,c] , **Niccolò Faraco**[a,d] , **Michele Maestrini**[a,e] , **Mauro Massari**[a,f]

[a] *Department of Aerospace Science and Technology, Politecnico di Milano, Via La Masa, 34, 20156 Milan, Italy*

[b] *PhD Student in Aerospace Engineering, annachiara.ippolito@polimi.it*

[c] *PhD Student in Aerospace Engineering, luca.giorcelli@polimi.it*

[d] *Research Fellow, niccolo.faraco@polimi.it*

[e] *Assistant Professor, michele.maestrini@polimi.it*

[f] *Associate Professor, mauro.massari@polimi.it*

[*] *Corresponding author*

Abstract

The deployment of autonomous rendezvous and docking capabilities requires Guidance and Navigation (GN) architectures capable of operating reliably under tight computational constraints and imperfect sensing conditions. Despite extensive simulation-based investigations, structured validation approaches explicitly connecting numerical modeling and embedded real-time implementation are still limited.

Within the GEORyder project, a staged validation campaign has therefore been developed, integrating closed-loop Software-in-the-Loop (SIL) simulations with vision-based relevant environment experiments, enabling systematic comparison between simulated and hardware-executed performance.

In particular, in the SIL configuration, an operationally representative simulation environment based on Unreal Engine reproduces the relative dynamics expected during proximity operations in Geostationary Transfer Orbits (GTO) and generates synthetic image streams consistent with the on-orbit visual conditions. These images are consequently processed by the Navigation module to successfully estimate the relative 6-DoF state of the chaser with respect to the target. The estimated state is finally fed to the Guidance module, which has the role of computing a safe approach trajectory and the associated control actions, closing the loop through numerical propagation of the dynamics. To move beyond pure simulation, the same GN software is deployed without structural modifications on an embedded development platform, where real-time operability, computational load, and numerical robustness are assessed under realistic hardware constraints.

The relevant environment experimental phase concentrates on the terminal approach, starting from approximately 15 meters of relative distance from the target spacecraft and covering the critical close range segment. Scaled 3D-printed targets, designed consistently with the optical properties of the test camera, are utilized to preserve geometric consistency with the expected operational scenario. In practice, a monocular camera mounted on a robotic manipulator observes the targets, generating the images of representative approach trajectories under varying illumination conditions. Finally, the navigation algorithm processes the acquired image streams to reconstruct the full 6-DoF relative state, allowing the quantitative evaluation of pose accuracy and sensitivity to perception uncertainties, as well as the direct comparison with SIL outcomes during the most delicate phase of docking.

Overall, the SIL and close-range relevant environment activities provide a consistent view of how the GN architecture behaves when moving from simulation to hardware execution. The results suggest that the GN architecture sustains realistic sensing effects and closed-loop operation without compromising navigation performance, while the embedded characterisation quantifies the computational margin still to be recovered on the target processor, supporting its applicability to autonomous rendezvous and docking operations in GTO.

Keywords: Autonomous Rendezvous and Docking, Vision-based Navigation, Software-in-the-Loop (SIL), Hardware-in-the-Loop (HIL), Real-time Embedded Systems

1. Introduction

In recent years, the increase of launches and constellations in the crowded near-Earth orbital regimes has led to a focus on On-Orbit Servicing (OOS) research for autonomous operations. The interest of the European governing bodies in this research field was displayed in the European In-Space Operations and Services (ISOS) workshop, organized in Brussels in December 2023. In this workshop, multiple ISOS European projects were presented, aiming to achieve an organized and standardized OOS ecosystem to reduce the number of launches by prolonging the operational lifetime of the assets. One of

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

the different ISOS tasks is Autonomous Rendezvous and Docking (ARD), where a servicer, also known as chaser, approaches and mates with a customer to refuel it or perform other activities. ARD missions are divided into multiple phases in which the chaser gradually approaches the target [1, 2]. After each phase, sanity checks are performed to ensure that the next operations are performed without any issues. The closer the chaser is to the target, the higher are the required safety level and the number of difficulties for the system. The last approach is usually divided into two sub-phases, the far- and close-range phases. The first starts from some kilometers of separations targeting a holding point (HP), a stable point of zero relative velocity in which the final checks are performed, in the range of tens of meters from the target. The second starts from this HP and ends when docking is complete. This phase is characterized by a fine navigation strategy that produces precise relative poses using features of the target spacecraft or markers on the docking face, which are required to respect the strict tolerances expected by the soft docking international standard [3]. For the close-range phase, developing autonomous and robust Guidance, Navigation and Control (GNC) systems is of paramount importance. GNC systems for proximity operations must be fast and responsive to unforeseen behaviors and events, maintaining the high safety level required in these tasks. To reach this goal, the GNC framework is usually run in closed loop on the on-board computer (OBC) of the spacecraft. Given the differences in working principle and performance between the hardware of modern PCs and OBCs, it is important to test the algorithms on the relevant processors to validate their deployability, especially in terms of compatibility between the different building blocks and of computational time demands. The second main problem during testing concerns the difficult conditions in which the cameras operate during the approach. In particular, harsh and changing lighting conditions lead to a degraded image quality that challenges the navigation framework, and these effects are not well simulated in synthetic environments [2, 4]. Indeed, the algorithms shall be able to function with incomplete information; for example, when some markers or other features are shadowed and impossible to extract from the image. To solve these two major issues, during the GEORyder project, the guidance and navigation algorithms have been tested in two separate ways. First, images were obtained in a realistic environment with Politecnico di Milano's RAFFAELLO facility. A 3D-printed model of the docking face of the target was created and mounted on a robotic arm, while a camera with the same performances as GEORyder's was mounted on a second arm. The guidance algorithm and a simulation environment were used to generate trajectories, which were then recreated with the two manipulators. The navigation algorithm was then tested with these images, validating its applicability in photorealistic conditions. The second validation process consisted in a Software-in-the-Loop (SIL) test, in which the guidance and navigation algorithms were deployed on the GEORyder board and tested with an in-house proximity operations simulator based on Unreal Engine [5]. This proved the end-to-end validity and deployability of the algorithms on the real hardware and the correct functioning of the interface between the two blocks. The rest of this article is organized as follows: in Section 2 the GNC architecture, the mission scenario and the reference frames are presented; in Section 3 the set-up and results of the real images test are explained; in Section 4, the SIL test is reported; in Section 5 some final remarks are drawn.

2. GNC Architecture & Mission Scenario

2.1 Reference Frames

First, it is important to define the reference frames used throughout this work. The Earth-Centered-Inertial (ECI) frame is derived as in Figure 1. The origin is in Earth's center, x-axis is along the J2000 vernal equinox, the z-axis along the J2000 pole, and the y-axis completes the triad. The state in this frame is defined as $\mathbf{X} = (\mathbf{R}, \mathbf{V})^\top = (R_x, R_y, R_z, V_x, V_y, V_z)^\top$.

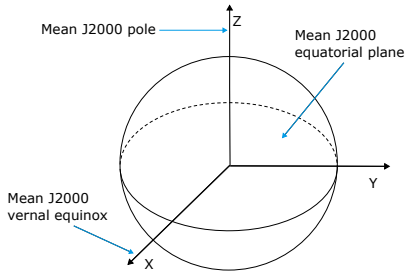


Fig. 1. ECI frame

The Local-Vertical-Local-Horizontal (LVLH) frame is defined with the origin in the target, the x-axis in its orbital radial direction, the z-axis along its orbital momentum, and the y-axis completes the triad. It is reported in Figure 2. The state in this frame is defined as $\mathbf{x} = (\mathbf{r}, \mathbf{v})^\top = (r_x, r_y, r_z, v_x, v_y, v_z)^\top$. The axes of the LVLH frame are also named R-bar, V-bar and H-bar.

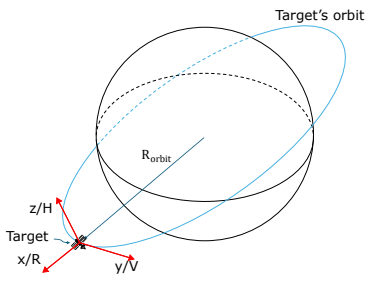


Fig. 2. LVLH frame

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

2.2 Proximity Operation Scenario

The envisioned mission scenario is a rendezvous and docking operation from the negative V-bar axis with a target on a Geostationary Transfer Orbit (GTO). In particular, this article focuses on the close-range phase, which means the final approach and the docking procedure. The mission starts from a zero-velocity holding point (HP) in $\mathbf{r} = (0, -15, 0)^\top$, with the target placed near the apogee of the GTO reported in Table 1.

Table 1. Target's GTO orbital parameters

a (km)	e (-)	i (°)	Ω_{GTO} (°)	ω_{GTO} (°)
24396.15	0.7283	6	0	178

Throughout the trajectory, the chaser must stay inside an approach cone, which is modeled to avoid possible appendages of the target, such as solar panels and antennas. The cone is designed as a semi-cubical parabola, as in Eq. 1.

$$h(\mathbf{r}) = r_x^2 + 0.08^2(r_y - 5)^3 + r_z^2 \qquad [1]$$

The chaser and the target are cubical spacecraft with edges of 0.5 meters, 780 kg of mass, and a drag coefficient of 2. The chaser thrust is limited to 22 N per LVLH axis.

2.3 Vision-based Navigation Strategy

The close-range navigation architecture computes the relative 6-DoF state using planar fiducial markers within the dual-mode operational framework introduced in [6], tailored to handle the transition from long-range observation down to terminal docking.

In particular, this pipeline manages the operational transition across three consecutive flight regimes. In the *far-range* phase the target is still optically unresolved, and only the relative translational state is estimated from the apparent centroid of the target. In the *close-range* phase the markers are resolved, and the marker-based pose solution supports estimation of the full 6-DoF relative state. Between the two lies a short *parallel zone*, entered at approximately $16\ m$ distance from the target, in which both measurement paths run concurrently until the hand-over to the close-range filter is committed. While the closed-loop campaign presented in Section 4 evaluates the entire profile including mode transitions, the remainder of this section and the experimental HIL campaign described in Section 3 focus strictly on the close-range regime.

During close-range operations, the image processing pipeline detects visible ArUco markers on the docking interface in real time, extracting 2D image coordinates and matching them with known 3D locations on the target [7]. At each time step, these correspondences are processed via a Perspective-3-Point (P3P) algorithm [8] embedded in a RANSAC loop to yield an instantaneous measurement vector $\mathbf{z}$ consisting of the relative position vector $\tilde{\mathbf{r}}$ and orientation quaternion $\tilde{\mathbf{q}}$.

To filter visual measurement noise, mitigate transient marker dropouts, and reconstruct unmeasurable relative velocities, these instantaneous pose estimates feed an Unscented Quaternion Estimator (USQUE) [9]. The full relative state vector $\mathbf{x}$ estimated by the navigation filter is defined as:

$$\mathbf{x} = \begin{bmatrix} \mathbf{r} & \mathbf{v} & \boldsymbol{\omega} & \mathbf{q} \end{bmatrix}^\top \qquad [2]$$

where $\mathbf{r}$ and $\mathbf{v}$ represent the relative position and velocity vectors in the target LVLH frame, while $\boldsymbol{\omega}$ and $\mathbf{q}$ are the relative angular velocity and the relative orientation quaternion, respectively.

To prevent numerical inconsistencies and violation of the quaternion unit norm constraint during filter update steps, the USQUE parameterizes local attitude errors using Modified Rodrigues Parameters (MRPs) [10]. The error updates are executed in Euclidean MRP space and subsequently mapped back into unit quaternions, ensuring mathematical consistency and stable covariance propagation throughout the approach.

In the navigation filter setup, the process noise covariance $\mathbf{Q}$ accounts for dynamic unmodeled perturbations, while the measurement noise covariance $\mathbf{R}$ reflects the expected accuracy of the P3P pose solution.

Finally, a dynamic sensor-switching logic monitors relative distance and position uncertainty bounds. Beyond approximately $2.5\ m$, visual tracking is handled by the narrow angle camera. Inside the terminal guidance zone, the system transitions to the wide angle camera to maintain field of view (FoV) coverage up to physical docking.

2.4 Guidance and Control Logic

2.4.1 Translational G&C

The close-range translational guidance and control framework is based on sliding mode control (SMC). SMC is a robust control methodology that has been applied to different guidance and control tasks, such as precision landing [11] and proximity operations [12]. It consists of a two-step design: modeling a sliding surface and a control function. The sliding surface is a manifold in the state space that guides the spacecraft towards the goal state. The control function is a function of the surface that brings the system state on the manifold and keeps it sliding on top of it. If this is achieved, asymptotic stability and convergence to the goal state is ensured [13]. In the current work, taking inspiration from [14], the sliding surface is modeled by using artificial potential fields to include the path constraints in the guidance law, as reported in Eq. 3.

$$\boldsymbol{\sigma} = (\mathbf{v} - \mathbf{v}_{dock}) + \mathbf{K}_s^\top \nabla_{\mathbf{r}} U_{tot} \qquad [3]$$

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

$\mathbf{v}_{dock} = (0, 0.055, 0)^\top$ is the selected docking velocity, $\mathbf{K}_s$ is a positive 3×1 gain vector, and U_{tot} is the potential field, reported in Eq. 4, where $h(\mathbf{r})$ is the approach cone of Eq. 1.

$$U_{tot} = \frac{1}{2}\left(10^3\, \mathbf{r}^\top \mathbf{r} + \frac{\mathbf{r}^\top \mathbf{r}}{10^{-4} h(\mathbf{r})}\right) \quad [4]$$

$\mathbf{K}_s$ is adaptively tuned during the approach to achieve faster convergence when the chaser is closer to the docking axis but avoiding overshooting. Being $\boldsymbol{\zeta} = (-\tanh(2|r_x| + |v_x|) + 1, 0, -\tanh(2|r_z| + |v_z|) + 1)^\top$ and $\mathbf{K}_s^0 = \left(2 \cdot 10^{-5}, 5 \cdot 10^{-8}, 2 \cdot 10^{-5}\right)^\top$ the initial value of $\mathbf{K}_s$, the adaptive law is reported in Eq. 5.

$$\begin{cases} \dot{\mathbf{K}}_s^t = 5 \cdot 10^{-2}\, \boldsymbol{\zeta}^\top \mathbf{K}_s^{t-1} \\ \mathbf{K}_s^t = \mathbf{K}_s^{t-1} + \dot{\mathbf{K}}_s^t\, dt \end{cases} \quad [5]$$

Additionally, $\mathbf{K}_s$ is clipped to 10^{-4}. Lastly, the control function is designed as in Eq. 6.

$$\mathbf{u} = -0.8|\boldsymbol{\sigma}|^{1/1.5} \tanh(\boldsymbol{\sigma}) \quad [6]$$

2.4.2 Rotational G&C

The rotational G&C was selected as a simple proportional controller on the quaternion error $\mathbf{q}_{err}$ and the chaser's angular velocity $\boldsymbol{\omega}_C$, as reported in Eq. 7. In the first phase of the approach, the quaternion error is evaluated to align the chaser's camera axis, which in this study is the y-axis of its body frame, with the chaser-target line of sight. Then, at $|r_y| \leq 3$ meters, the quaternion error is evaluated to align the chaser's docking port to the one of the target, so aligning the y-axis of the chaser along the target's LVLH V-bar. Throughout this procedure, the continuity of the quaternion is always enforced to avoid possible ambiguities that would flip the chaser around its camera axis.

$$\boldsymbol{\tau} = -(8\mathbf{q}_{err} + 5\boldsymbol{\omega}_C) \quad [7]$$

3. Relevant Environment Test

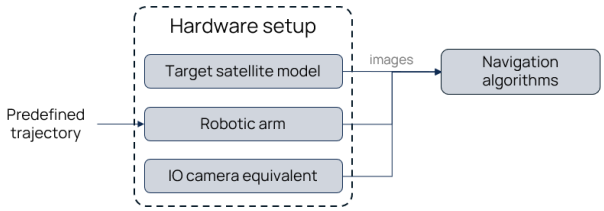


Fig. 3. Relevant environment test design

The experimental validation in a relevant environment relies on the Hardware-in-the-Loop (HIL) setup shown in Figure 3. This campaign specifically targets the perception challenges of the final approach, where the navigation module must operate under reduced geometric margins, constrained field of view, and variable illumination conditions.

Specifically, the primary objective is to evaluate the navigation module when driven by real optical measurements rather than synthetic images. By incorporating physical sensing effects, such as actual sensor noise, optical distortions, exposure dynamics, and realistic lighting variations, the HIL facility exposes the vision pipeline to true operational uncertainties. The navigation filter processes these live image streams without any modification to its internal architecture, providing a direct assessment of estimation robustness under realistic perception inputs.

3.1 Pass/Fail criteria

To quantitatively establish the success of the experimental campaign, a structured set of Pass/Fail requirements was defined. Table 2 summarizes this list of criteria, outlining the specific metrics and aspects evaluated during testing.

Table 2. Relevant environment test pass/fail criteria

Metric	Symbol	Evaluated Aspect
Navigation Module		
Visual detection robustness	P_{det}	Valid frame processing ratio
Navigation accuracy	$P_{\text{pos}}, P_{\text{rot}}$	Steady-state error boundaries
Filter statistical consistency	$\sigma_{\text{pos}}, \sigma_{\text{rot}}$	3σ covariance bounds
Overall System		
Environmental robustness	$\mathcal{R}_{\text{env}}$	Illumination sensitivity
System stability	$\exists P_{\text{limit}}$	Absence of filter divergence

Each requirement is systematically assessed through quantitative indicators computed on the testing datasets, as discussed in the following.

- **Visual detection robustness** is evaluated using the frame detection success rate, defined as:

$$P_{det} = \frac{N_{det}}{N_{valid}} \geq 90\%$$

where N_{det} is the number of frames successfully processed, while N_{valid} is the total number of valid frames in the trajectory.

- **Navigation accuracy** is evaluated over the steady-state portion of the trajectory, excluding the initial 10% transient phase associated with the filter initialization. Within the steady-state interval, performance assessment is focused on the final segment of the trajectory, corresponding to the last 10% of the

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

valid samples, which represents the most demanding portion of the close-range approach. Let

$$e_r = ||r_{est}(t) - r_{true}(t)||, \quad d(t) = ||r_{true}(t)||$$

be the position error and the corresponding ground truth distance. The position requirement is expressed as

$$P(e_r(t) \leq 0.2d(t)) \geq 80\%$$

Similarly, let $e_q(t)$ denote the angular attitude error expressed in degrees. The attitude requirement is

$$P(e_q(t) \leq 6^\circ) \geq 80\%$$

- **Filter statistical consistency** is verified by comparing the estimation error with the corresponding 3σ covariance bounds. Denoting by $\sigma(t)$ the standard deviation derived from the filter covariance matrix, the consistency metric to be satisfied is

$$P(|e(t)| \leq 3\sigma(t)) \geq 90\%$$

for both position and attitude components.

- **Environmental robustness** is assessed by evaluating the sensitivity of the marker detection performance to illumination variations. Let P^1_{det} and P^2_{det} be detection ratios obtained under two exposure settings. The robustness metric is defined as the relative degradation between the two configurations:

$$\mathcal{R}_{env} = \frac{\min(P^1_{det}, P^2_{det})}{\max(P^1_{det}, P^2_{det})}$$

The requirement is satisfied if $\mathcal{R}_{env} \geq 90\%$, meaning that the worst illumination condition shall not reduce the detection performance by more than 10% with respect to the best one.

- **System stability** is verified over the entire test sequence. The criterion is satisfied if no numerical divergence occurs, the filter covariance remains bounded, and no unbounded growth of the state estimate is observed.

3.2 Experimental Setup

As already anticipated, the experimental campaign reproduces the terminal close-range segment of the rendezvous scenario under laboratory conditions. Specifically, the operational close-range interval, spanning from $15\ m$ down to $2\ m$ from the target, has been divided into two distinct sub-phases: the first covering the 15-$5\ m$ and the second covering 5-$2\ m$.

Each section is independently replicated in the laboratory over a reduced physical distance range between $3\ m$ and $1\ m$. Operating in such manner allows to preserve visual representativeness across the entire close-range domain while remaining compatible with the spatial constraints of the test facility.

Thus, to reproduce the relative motion between chaser and target spacecraft, a robotic arm is employed, which executes a set of predefined trajectories provided by the Guidance module. Specifically, the trajectories are obtained by executing the close-range guidance algorithm with perfect navigation in the real environment, starting from different initial conditions, and are then rescaled and rotated into the robotic arms frames. Nine initial conditions are generated for both aforementioned sub-phases: a nominal one, starting along the V-bar axis of the target LVLH frame, and eight off-axis conditions along both R-bar and H-bar, with two different displacement values. The initial conditions are reported in Table 3. These have been selected to obtain different poses to change the illumination conditions of the face of the target.

Table 3. Initial relative positions (LVLH) in [m] for the 15 m (Test 1) and 5 m (Test 2) sub-phase trajectories. Initial relative velocity is zero in all test cases.

ID	Test 1	Test 2
1	$[\ 0.0, -15.0,\ \ 0.0]$	$[\ 0.0, -5.0,\ \ 0.0]$
2	$[\ 0.5, -15.0,\ \ 0.0]$	$[\ 0.2, -5.0,\ \ 0.0]$
3	$[-0.5, -15.0,\ \ 0.0]$	$[-0.2, -5.0,\ \ 0.0]$
4	$[\ 0.0, -15.0,\ \ 0.5]$	$[\ 0.0, -5.0,\ \ 0.2]$
5	$[\ 0.0, -15.0, -0.5]$	$[\ 0.0, -5.0, -0.2]$
6	$[\ 1.6, -15.0,\ \ 0.0]$	$[\ 0.5, -5.0,\ \ 0.0]$
7	$[-1.6, -15.0,\ \ 0.0]$	$[-0.5, -5.0,\ \ 0.0]$
8	$[\ 0.0, -15.0,\ \ 1.6]$	$[\ 0.0, -5.0,\ \ 0.5]$
9	$[\ 0.0, -15.0, -1.6]$	$[\ 0.0, -5.0, -0.5]$

While the robotic manipulator operates as chaser spacecraft, images are acquired using a Basler acA1300-22gm industrial CCD camera [15]. Rigidly mounted on the robotic arm, its pose is entirely dictated by the executed trajectory. This camera serves as a representative substitute for the narrow-angle unit of the operational Infinite Orbits (IO) OrbSight sensor used during the close-range approach; wide-angle representativeness below 2.5 m is not addressed in this campaign. Table 4 compares the operational and surrogate specifications. Intrinsic parameters were calibrated prior to testing and remained fixed throughout the experimental campaign.

Table 4. Specifications of operational and test cameras.

Parameters	OrbSight	Basler
Resolution	2048×2048	1296×966
Field of View	18.8°	28°
Focal length	50 mm	12 mm

Considering the test camera parameters, two scaled 3D-printed targets are employed, each associated with

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

one of the operational distance intervals. The laboratory targets were obtained by geometrically scaling the real docking interface to preserve optical similarity (i.e., sensor FoV occupancy) within the spatial constraints of the test facility. Based on the pinhole camera model, the anisotropic scale factor $\mathbf{s} = [s_x, s_y, s_z]^T$ applied to the target dimensions is computed by enforcing equal angular coverage between the operational and test configurations at a reference distance Z:

$$\mathbf{s} = \frac{Z_{test}}{Z_{orig}} \cdot \frac{\tan\left(\frac{FoV_{orig}}{2}\right)}{\tan\left(\frac{FoV_{test}}{2}\right)}$$

Since the close-range approach is divided into two operational segments, two distinct scaled targets were manufactured: Target 1, sized for the 15–5 m interval with a front face footprint of 9.12×9.12 cm, and Target 2, tailored for the 5–2 m regime with dimensions of 27.4×27.4 cm.

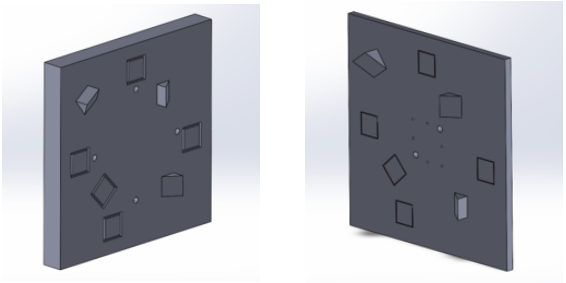


Fig. 4. Target 1 on the left, Target 2 on the right



Fig. 5. Dual-arm robotic setup at DART Lab: chaser camera (left) and scaled target mockup (right).

To complete the physical target definition, each face is equipped with a specific spatial arrangement of ArUco markers. Their locations and relative orientations were optimized to ensure robust feature extraction across the operational distance range while mitigating the inherent pose ambiguities of Perspective-n-Point (PnP) solvers. Specifically, an asymmetric pattern was implemented, incorporating three markers mounted on tilted out-of-plane ramps. This non-coplanar relief guarantees unique 2D-3D correspondences and prevents sign flips and degenerate geometric solutions during relative pose reconstruction.

Finally, the entire setup is operated in a dark laboratory environment (see Figure 5) to allow controlled variation of illumination conditions. In particular, artificial lighting is adjusted across different experimental runs to evaluate the robustness of feature detection and pose estimation against lighting variability.

3.3 Test Process Description

The testing methodology covers the complete perception-to-estimation chain. As outlined below, the evaluation process begins with image acquisition and dataset generation on the HIL setup, proceeds with visual feature extraction via ArUco detection, and concludes with the performance assessment of the pose estimation and filtering algorithms.

3.3.1 Image Acquisition and Datasets Generation

The first stage of the experimental campaign concerned the image acquisition process. For each close-range configuration, the corresponding scaled target was rigidly mounted on the robotic manipulator. Meanwhile, the reference approach trajectories generated by the Guidance module were provided to the robotic system, which executed the prescribed relative motion profiles in open loop, thereby reproducing the relative kinematics of the terminal rendezvous phase.

For each trajectory, images were acquired at two exposure settings (350 ms and 550 ms). This approach emulates orbital illumination changes, yielding diverse brightness and contrast levels to evaluate the sensitivity and robustness of marker detection and 6-DoF pose estimation. The resulting sequences (Figure 7) form the dataset for the subsequent navigation analysis.

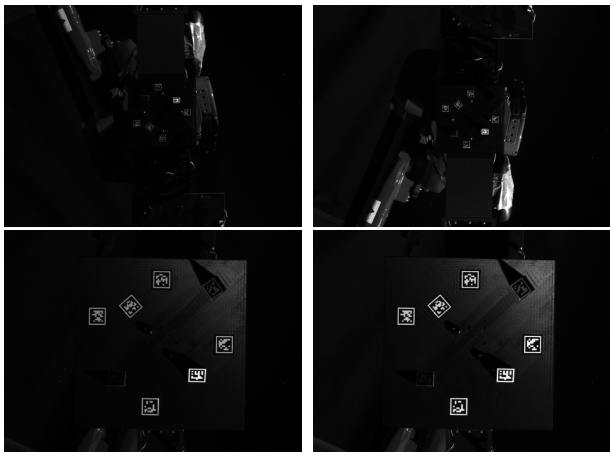
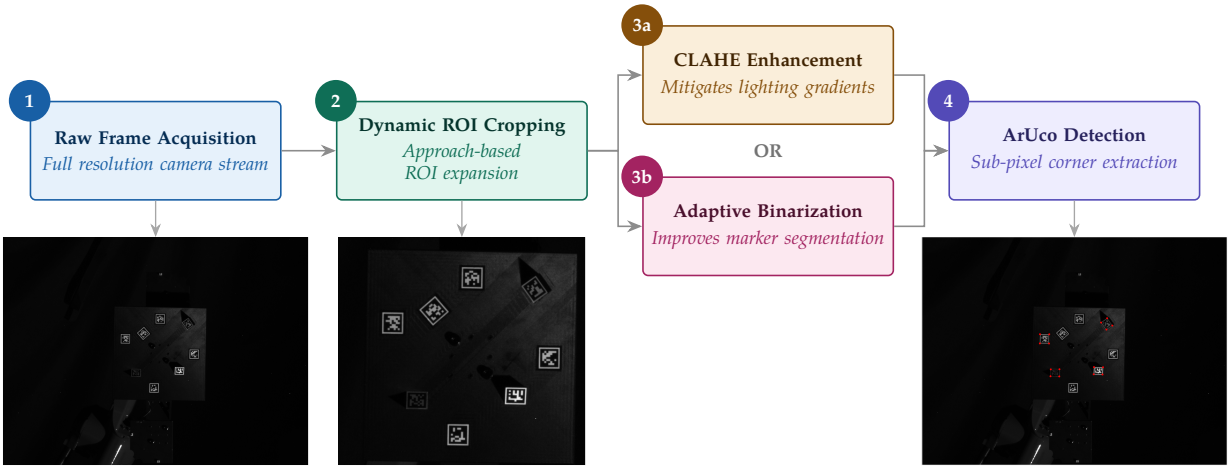


Fig. 7. Examples of acquired images of Target 1 (top row) and Target 2 (bottom row) under different exposure times: 350 ms (left) and 550 ms (right) [16]

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.



(a) Raw Frame (b) Cropped Target ROI (c) Detected Corners

Fig. 6. Overview of the image processing pipeline. Depending on the target configuration, either CLAHE enhancement or adaptive binarization is applied after dynamic cropping, prior to ArUco detection.

3.3.2 Image Processing and ArUco Detection

The acquired datasets are processed through a dedicated preprocessing and detection pipeline aimed at enhancing marker detectability while preserving geometric consistency with the camera model. A dynamic central cropping strategy is adopted to improve robustness and reduce computational load. For each trajectory, the cropping window is progressively expanded as a function of the frame index, reflecting the increase in apparent target size during the approach phase. Formally, the crop fraction along each dimension is defined as

$$f = f_0 + \alpha(f_1 - f_0) \quad \text{where } \alpha = \left(\frac{i-1}{N-1}\right)^{\gamma}$$

where i is the frame index, N the trajectory length, and $\gamma > 1$. Setting $\gamma > 1$ keeps the region of interest (ROI) compact early in the trajectory, when the target appears small, and expands it more rapidly as the chaser approaches, reducing the risk of premature marker clipping. The resulting value f is clamped between minimum and maximum limits to keep the bounding box within valid image bounds. Both the crop bounds and γ were tuned independently for each target, reflecting their different scale and apparent size evolution along the approach trajectory.

Depending on the target configuration, the cropped grayscale image undergoes one of two alternative contrast-enhancement strategies. For Target 2, contrast is enhanced through Contrast Limited Adaptive Histogram Equalization (CLAHE) [17], applied with a clip limit of 0.02 and a 8×8 tile grid. These values were empirically selected to enhance local contrast under non-uniform illumination and mitigate exposure-dependent brightness variations, while limiting noise amplification in low-signal regions. For Target 1, the smaller markers are more susceptible to local lighting gradients across the face. Therefore, an adaptive thresholding step [18] is introduced prior to detection. The local threshold sensitivity is set to 0.35, effectively isolating marker edges while rejecting background noise. When enabled, nearest-neighbor interpolation is used for resizing, preserving sharp edge transitions in the resulting binary image. In both cases, no aggressive spatial filtering is applied, in order to avoid distortion of marker contours, which are critical for accurate corner extraction.

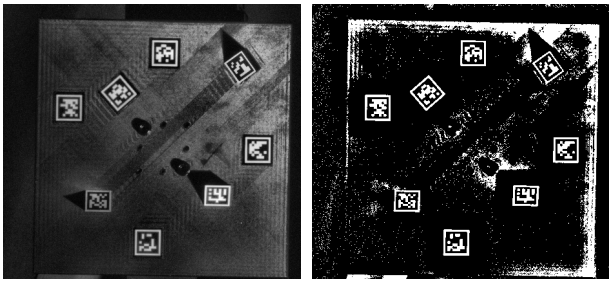


Fig. 8. Comparison of the two alternative contrast-enhancement strategies: CLAHE enhancement (left), and adaptive binarization (right)

Following preprocessing, ArUco detection is performed frame-by-frame. For each image, the algorithm extracts marker IDs and corner pixel coordinates, enabling 2D–3D correspondences with the known marker geometry in the target reference frame.

3.3.3 Pose Estimation and Filtering Performance

Once valid marker detections are obtained, the navigation pipeline performs pose estimation followed by state

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

(a) Target 1 compliance assessment

Test		Accuracy			Consistency			Overall
ID	Config.	P_{det}	P_{pos}	P_{rot}	σ_{pos}	σ_{rot}	$\mathcal{R}_{\text{env}}$	Score
Test1_1	350ms	0.02	0.00	0.00	**1.00**	0.85	0.63	0.17
	550ms	0.01	0.00	0.00	**1.00**	0.85	0.63	0.17
Test1_2	350ms	0.02	0.26	0.00	**1.00**	**0.98**	0.50	0.33
	550ms	0.01	0.05	0.00	**1.00**	**1.00**	0.50	0.33
Test1_3	350ms	0.01	0.08	0.00	**1.00**	0.89	0.86	0.17
	550ms	0.02	0.00	0.00	**1.00**	**0.90**	0.86	0.33
Test1_4	350ms	0.02	0.00	0.00	**1.00**	**0.97**	0.50	0.33
	550ms	0.01	0.00	0.00	**1.00**	**0.98**	0.50	0.33
Test1_5	350ms	0.03	0.23	0.00	**1.00**	**0.98**	0.73	0.33
	550ms	0.02	0.00	0.00	**1.00**	**0.96**	0.73	0.33
Test1_6	350ms	0.02	0.13	0.00	**1.00**	0.75	0.71	0.17
	550ms	0.01	0.00	0.00	**1.00**	**0.97**	0.71	0.33
Test1_7	350ms	0.01	0.08	0.00	**1.00**	**0.95**	0.80	0.33
	550ms	0.01	0.00	0.00	**1.00**	**0.98**	0.80	0.33
Test1_8	350ms	0.02	0.00	0.00	**1.00**	**0.97**	0.80	0.33
	550ms	0.02	0.00	0.00	**1.00**	**0.97**	0.80	0.33
Test1_9	350ms	0.03	0.00	0.00	**1.00**	**0.97**	0.17	0.33
	550ms	0.00	0.00	0.00	**1.00**	**0.97**	0.17	0.33

(b) Target 2 compliance assessment

Test		Accuracy			Consistency			Overall
ID	Config.	P_{det}	P_{pos}	P_{rot}	σ_{pos}	σ_{rot}	$\mathcal{R}_{\text{env}}$	Score
Test2_1	350ms	**0.95**	**1.00**	**0.93**	0.78	**1.00**	**0.92**	0.83
	550ms	0.88	**1.00**	**1.00**	**0.99**	**1.00**	**0.92**	0.83
Test2_2	350ms	**0.91**	**1.00**	**1.00**	**1.00**	**1.00**	**0.96**	1.00
	550ms	0.87	**1.00**	**0.86**	**1.00**	**1.00**	**0.96**	0.83
Test2_3	350ms	0.89	**0.86**	**1.00**	**1.00**	**1.00**	**0.94**	0.83
	550ms	0.84	**0.93**	0.71	**1.00**	**1.00**	**0.94**	0.67
Test2_4	350ms	0.89	0.57	**0.93**	**0.96**	**1.00**	**0.92**	0.67
	550ms	0.81	0.00	0.71	**0.95**	**1.00**	**0.92**	0.50
Test2_5	350ms	**0.99**	**1.00**	**1.00**	0.84	**0.99**	**0.95**	0.83
	550ms	**0.95**	**1.00**	**1.00**	**1.00**	**1.00**	**0.95**	1.00
Test2_6	350ms	**0.94**	**1.00**	0.57	**1.00**	**1.00**	**0.93**	0.83
	550ms	0.87	**0.86**	**1.00**	**1.00**	**1.00**	**0.93**	0.83
Test2_7	350ms	**0.96**	**1.00**	0.64	**1.00**	0.89	**0.96**	0.67
	550ms	**0.92**	0.43	0.14	**1.00**	**1.00**	**0.96**	0.67
Test2_8	350ms	**0.95**	**1.00**	**0.93**	0.89	**0.99**	**0.96**	0.83
	550ms	**0.91**	0.00	**1.00**	**0.99**	**1.00**	**0.96**	0.83
Test2_9	350ms	**0.92**	**0.93**	**0.93**	**0.92**	**1.00**	**0.99**	1.00
	550ms	**0.91**	**1.00**	0.36	**1.00**	**1.00**	**0.99**	0.83

Table 5. Compliance assessment for Target 1 and Target 2. Bold values denote compliant parameters.

filtering through the USQUE framework. The estimated relative position and attitude are compared against the ground truth trajectory provided by the robotic manipulator to enable quantitative performance assessment. As already mentioned, to avoid bias associated with filter initialization, the first 10% of the samples are excluded from the analysis. The evaluation is then conducted over the steady-state portion of each trajectory, with particular attention to the final segment corresponding to the last 10% of valid samples, which represents the most critical phase of the close-range approach. The previously defined performance and consistency criteria are applied independently to each exposure configuration. The resulting compliance levels are subsequently compared to assess the impact of illumination variations on overall navigation performance.

3.4 Compliance Assessment

For each test case, the outcomes of the detection, accuracy, statistical consistency, and environmental robustness analyses are combined to determine the overall compliance of the navigation module. The detailed results for all trajectories and illumination configurations are summarized in Table 5a and 5b.

For Target 2, the position accuracy evaluated over the final steady-state segment satisfies the prescribed relative distance requirement in the large majority of configurations. Compliance levels remain consistently above the acceptance threshold, confirming reliable translational performance in the 5-$2\ m$ operational interval. Attitude accuracy shows slightly higher sensitivity to illumination conditions, with a limited number of configurations approaching the compliance boundary. Nevertheless, no systematic degradation is observed. The statistical consistency analysis confirms coherent filter behavior, with estimation errors predominantly contained within the predicted 3σ bounds across all trajectories. Environmental robustness evaluation indicates limited sensitivity to exposure variations, as performance differences between configurations remain within the admissible tolerance. Overall compliance values therefore validate the perception-estimation pipeline for the short range segment represented by Target 2, where marker observability is inherently favorable due to increased apparent feature size.

In contrast, Target 1 reveals a severe perception-driven limitation inherent to the 15–$5\,\text{m}$ regime under the surrogate optical setup. At this range, the apparent pixel footprint of the markers remains critically small, causing the ArUco detector to fail across nearly the entire profile, as confirmed by detection ratios $P_{\text{det}} \leq 0.03$ in Table 5a. Consequently, the visual pose solution was unavailable for the vast majority of frames, preventing the navigation pipeline from meeting the operational pose accuracy bounds (P_{pos} and P_{rot}). Crucially, however, this test validates the mathematical stability and conservatism of the filtering framework under an extreme measurement-outage scenario. Even while operating almost purely in open-loop covariance propagation, the USQUE main-

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

tained bounded uncertainties without encountering numerical singularities, divergence, or state flips. The high consistency metrics (σ_{pos} and σ_{rot}) demonstrate that the filter process noise and propagation models correctly bounded the unobserved physical drift, confirming that the performance bottleneck lies strictly in the front-end optical resolution rather than in the estimation architecture.

Overall, the experimental campaign demonstrates that the navigation framework is fully adequate in the short range regime characterized by strong visual observability (Target 2), while the more distant portion of the close-range approach (Target 1) is primarily limited by perception constraints related to feature scale and optical resolution, rather than by intrinsic weaknesses of the estimation algorithm.

4. Closed-loop Software-in-the-Loop Test

The relevant environment campaign of Section 3 exercises the navigation module in open loop: the trajectories are prescribed, the image stream is replayed, and the estimated pose is compared against the manipulator ground truth. This is the right instrument for assessing perception under realistic optics and illumination, but it is structurally unable to answer some additional concerns that need to be assessed before flight.

The first one is *error feedback*. In closed loop the navigation error enters the guidance law, modifies the commanded trajectory, and therefore changes the next image the camera acquires. Whether that feedback path is self-correcting or divergent cannot be observed on a prescribed open-loop trajectory. Then, the *mode logic correctness* needs to be checked. The transition between the far- and close-range navigation modes, and the narrow-to-wide camera switch of Section 2, are triggered by the estimated state and cannot be exercised end-to-end in an open-loop test. Third, is evaluating the *mission-level outcome*: while Section 3 reports pose error along a fixed path, the quantity of operational interest is whether the spacecraft arrives at the docking point within prescribed margins. Lastly, *real-time execution* against a declared cycle budget needs to be measured.

This section reports a closed-loop Software-in-the-Loop (SIL) campaign designed around those four questions.

4.1 Objectives and Pass/Fail Criteria

The criteria are summarised in Table 6. Where the same quantity is measured as in Section 3, the symbol and the threshold are deliberately carried over unchanged so that the two campaigns can be read against each other; the remaining criteria address the closed-loop behaviour that only this configuration can expose.

The quantitative definitions are as follows.

- **Visual detection robustness** retains the definition of Section 3, $P_{det} = N_{det}/N_{valid} \geq 90\%$, but is reported per navigation phase rather than per trajectory, since detection performance is strongly range-dependent. The denominator is the number of navigation cycles that consumed a camera frame; because the cameras run slower than the navigation task, this is not the same as the number of cycles executed.

- **Navigation accuracy** uses the position criterion of Section 3 with an absolute floor,
$$P\left(e_r(t) \leq \max\left(0.2\, d(t),\ 0.05\ m\right)\right) \geq 80\%$$
A floor is here introduced to maintain the criterion meaningful at short range. The attitude criterion is unchanged, $P(e_q(t) \leq 6^\circ) \geq 80\%$, evaluated over the close-range segment where the rotational filter is active.

- **Filter statistical consistency** is $P(|e(t)| \leq 3\sigma(t)) \geq 90\%$ as before. For the rotational filter the containment is evaluated on the relative angular rate, for which the filter propagates an explicit covariance, rather than on the attitude angles.

- **Path constraint compliance** is the fraction of samples satisfying $h(\mathbf{r}) < 0$ with h the approach cone of Eq. 1, evaluated on the *true* state. The requirement is 100%; the minimum margin is reported alongside.

- **Terminal accuracy** is assessed against a docking box of eight criteria, reproduced in Table 7, with e_{dock} the margin left on each. The box is evaluated on the *true* relative state at the instant the vehicle declares arrival. That distinction matters because the on-board check can only test the *estimated* state: a positive declaration means that navigation believes the vehicle has docked, but does not guarantee that corresponds to reality. Comparing the two gives the rate of false arrival declarations,
$$P_{fp} = \frac{N_{fp}}{N_{runs}} = 0$$
where N_{fp} is the number of false positives, i.e. runs declaring arrival with the true state outside the box, and N_{runs} the total number of runs. No other test configuration in this work can detect such a false positive.

- **Computational load** is the worst-case cycle time normalised by the task period, $C = \max(t_{\text{cycle}})/T \leq 1$, with $T = 500\ ms$ for both the navigation and the guidance task. P_{dl} is the fraction of cycles completing before the next release.

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

Table 6. Closed-loop SIL test pass/fail criteria

Metric	Symbol	Evaluated Aspect
Navigation Module		
Visual detection robustness	P_{det}	Valid frame processing ratio, per phase
Navigation accuracy	$P_{\text{pos}}, P_{\text{rot}}$	Error bounds along the closed-loop trajectory
Filter statistical consistency	$\sigma_{\text{pos}}, \sigma_{\text{rot}}$	3σ covariance bounds
Estimator continuity	N_{reinit}	Filter reinitialisations
Guidance and Control		
Path constraint compliance	P_{cone}	Approach-cone containment
Propellant expenditure	ΔV_{tot}	Total commanded ΔV
Actuator compliance	P_{sat}	Commands within control authority
Mode Transition Logic		
Phase sequencing	$\exists T_{\text{phase}}$	Correct far $\rightarrow$ close ordering
Hand-over integrity	ε_{ho}	State discontinuity at hand-over
Sensor switching	$\exists T_{\text{cam}}$	Camera switch at the design threshold
Docking Performance		
Terminal accuracy	e_{dock}	True state inside the docking box
False arrival declaration	P_{fp}	Arrival declared with truth outside the box
Real-Time Execution		
Computational load	$C_{\text{nav}}, C_{\text{gui}}$	Worst-case cycle time vs. period
Timing determinism	P_{dl}	Deadline compliance

4.2 Experimental Setup

The setup couples the guidance and navigation implementation of Section 2, used unchanged, to the PROX-SIMA simulation environment that plays two roles at once: it integrates the vehicle dynamics on which the resulting commands act, and it renders the camera stream the navigation module consumes. The first role also supplies the truth reference for every error reported below. It is available at every step and for the full relative state, translational and rotational alike.

Nine closed-loop runs were executed unattended. Each starts from a holding point at $20\ m$ on the negative V-bar with a closing velocity of $55\ mm/s$ and terminates at the docking box. The dispersion set mirrors the structure of Table 3: a nominal on-axis case plus eight off-axis initial conditions, displaced by $\pm 1\ m$ and $\pm 2\ m$ along R-bar and H-bar. The starting range is extended from the $15\ m$ of the relevant environment campaign to $20\ m$ so that the complete far-to-close mode sequence, and not only the terminal segment, is exercised. Each run spans approximately 5.5 minutes of simulated time and the same in wall clock.

The simulation environment and the dynamics model it embeds are described in turn below.

4.2.1 Simulation Environment

The campaign is run against PROX-SIMA [5], an in-house modular simulator for close proximity operations built on Unreal Engine 5.5. Images are produced by a physically based rendering pipeline with dynamic global illumination and hardware ray tracing, so that specular response, shadowing and the strong illumination contrast typical of the orbital environment are reproduced. The rendered target is a textured model of the GEORyder docking interface carrying the same ArUco marker arrangement described in Section 2. Two cameras are simulated with the operational OrbSight parameters of Table 4: a narrow-angle unit with an $18.8°$ field of view and a wide-angle unit with a $117°$ field of view, both 2048×2048 and monochrome, sampled at $1\ Hz$.

What distinguishes this environment from a rendering tool, and what the present campaign exploits, is its closed-loop capability. The architecture is sketched in Figure 9. The vehicle dynamics are integrated as a Functional Mock-up Unit (FMU) under the FMI 2.0 co-simulation standard [19]; the sensor framework derives the camera stream from the resulting true state; and both the sensor data and the returning control actions are exchanged with an external GN module over a WebSocket interface. The GN module is a separate process running the same guidance and navigation implementation described in Section 2, driving its navigation and guidance tasks at $2\ Hz$ each under a real-time scheduling policy.

The simulator advances under real-time pacing: the dynamics are stepped against wall clock and the GN module participates as an independent asynchronous client, holding no ability to stall the physics. The FMU applies the last control action received until a new one arrives. This reproduces the operational situation, in which the onboard computer and the vehicle dynamics share no clock, and it means that transport delay and control staleness are part of what is being measured rather than artifacts removed by construction. This, it must be noted, comes at the cost of slightly decreased reproducibility of the results, as the communication latency is not deterministic. The results below are reported as a nine-run set rather than as a single trace to mitigate this problem.

4.2.2 Dynamics Model

The dynamics model, shown in Figure 10, is written in Modelica and exported as a co-simulation FMU. Chaser and target are modeled as full six-degree-of-freedom rigid

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

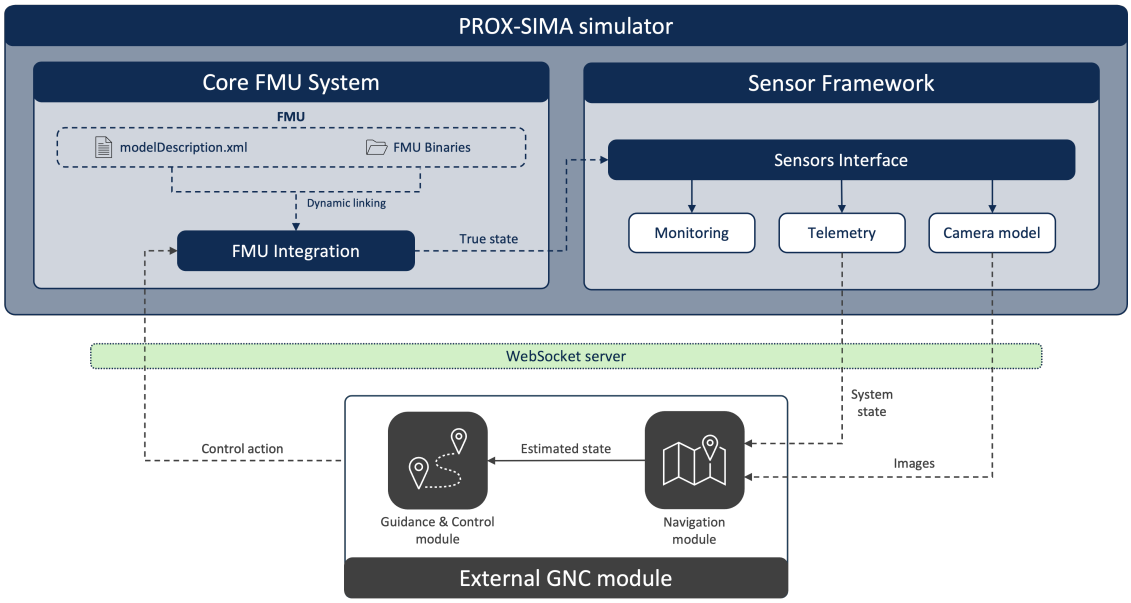


Fig. 9. Closed-loop SIL architecture. The simulator integrates the vehicle dynamics as an FMI co-simulation unit, derives the rendered camera stream from the resulting true state, and exchanges sensor data and control actions with the GN module over a WebSocket interface. The interface is identical whether the GN module executes on a workstation or on the target processor.

bodies, with the physical characteristics already described in Section 2. The model accepts a commanded acceleration in the target LVLH frame and a commanded body torque for the chaser, and exposes the true inertial states and attitude quaternions of both vehicles together with the actually applied force and torque.

The relative translational dynamics are those of Franzini and Innocenti [20]: the nonlinear equations of relative motion of a deputy with respect to a chief, written directly in the rotating target-centred LVLH frame, retaining the exact differential gravity term and propagating the angular velocity and angular acceleration of the LVLH frame itself as states driven by the chief orbit.

This is deliberately *not* the model used inside the navigation filter, which propagates its state and covariance using linearised eccentric relative motion. Adopting the same formulation on both sides would make the estimator process model exact by construction, and the resulting navigation accuracy would reflect that coincidence rather than the performance to be expected in flight. Keeping the two distinct means the filter is evaluated against dynamics it does not know, which is the purpose of the exercise. The same reasoning applies to the attitude channel, where the simulator integrates the full rigid-body equations with translation-attitude coupling through the body-frame rotation of the commanded thrust, while the navigation module propagates translation and attitude in two decoupled filters.

4.3 Results

All nine runs reached the docking box. There were no filter reinitialisations, no approach-cone violations, and no missed task deadlines across the set. Below, the most interesting results are discussed in more detail.

4.3.1 Guidance and Trajectory

Figure 11 shows the nine trajectories in the two LVLH projections. From every dispersed start the guidance law recovers the docking axis well before the terminal segment and converges without overshoot; the lateral excursion visible between 15 and 12 m follows the transition into the close-range mode and is common to all nine. The reason is to be found in the Parallel Zone hand-off logic from the far-range EKF filter to the close-range USQUE. A more refined switch logic, with a stricter convergence check on the USQUE performance, could easily mitigate the already minor problem even further.

Approach-cone compliance is $P_{cone} = 100\%$ in every run, with a worst-case margin of $h = -2.17$, i.e. the trajectories remain well inside the admissible corridor throughout. Total commanded ΔV ranges from 0.137 to 0.184 m/s across the set, and no command reached the control authority limit, $P_{sat} = 0\%$.

4.3.2 Navigation

Figure 12 plots the closed-loop position error against true range for all nine runs, with the accuracy requirement overlaid. The error remains below the bound over the entire approach and by a widening margin as the target is

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

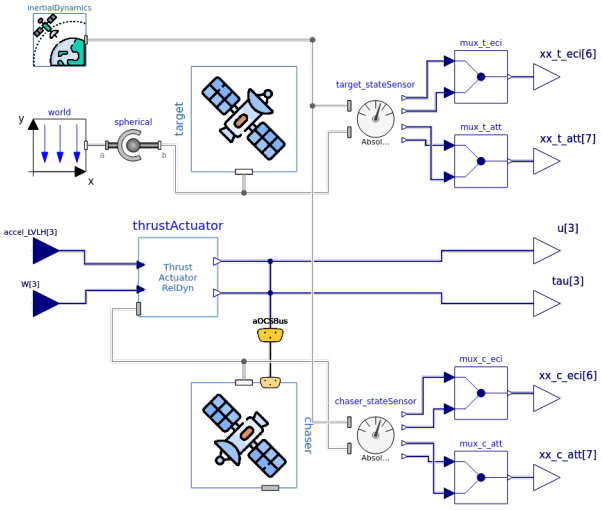


Fig. 10. Modelica model exported as the co-simulation FMU. Chaser and target are full six-degree-of-freedom rigid bodies; the actuator converts the commanded LVLH acceleration into the applied body-frame force, and the state sensors expose the true ECI states and attitudes of both vehicles.

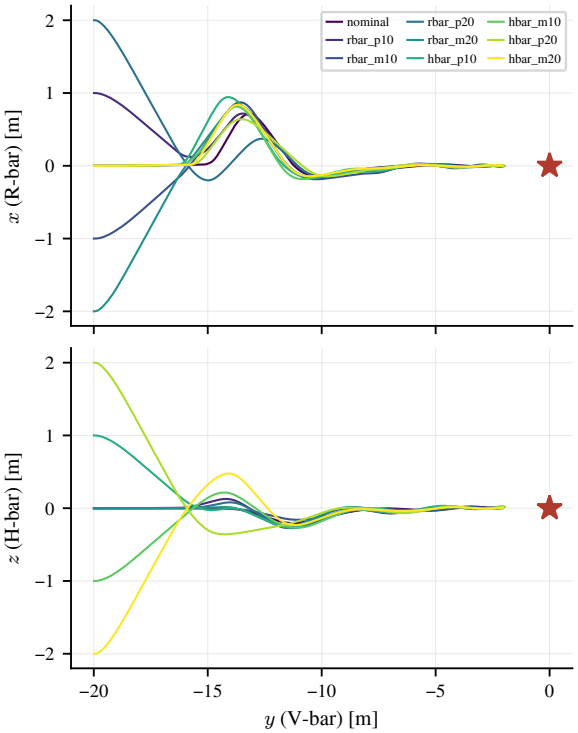


Fig. 11. Closed-loop trajectories of the nine-run dispersion set in the target LVLH frame, R-bar (top) and H-bar (bottom) against V-bar. Lateral axes are magnified with respect to the approach axis. The target is at the origin.

neared, giving $P_{pos} = 100\%$ in every run. The attitude criterion is met with P_{rot} between 89% and 100%, against the 80% requirement. Filter consistency is likewise satisfied, $\sigma_{pos} \geq 94\%$ and $\sigma_{rot} = 100\%$, and no run required a filter reinitialisation, $N_{reinit} = 0$.

Marker detection succeeded on every processed frame in all three phases, with an average of 7.53 of the 8 available markers contributing to each close-range solution. This is a markedly better figure than the relevant environment campaign obtained on Target 1, and the comparison is instructive: the difference is not in the estimation architecture, which is identical, but in the fidelity of the illumination and of the sensor response. The synthetic imagery reproduces the geometry of the approach faithfully while remaining optically benign, whereas the laboratory campaign of Section 3 exposes the detector to the exposure-dependent contrast and residual illumination gradients that dominate at small apparent marker size. The two campaigns are therefore complementary rather than redundant: the closed loop validates the feedback architecture under nominal perception, and the relevant environment test bounds the perception performance that architecture can be given.

4.3.3 Mode Transition Logic

The phase sequence was correct and monotone in all nine runs, with no oscillation between modes. Entry into the proximity zone occurred between 16.0 and 17.2 m against a 16 m design threshold, and the close-range filter took over between 15.3 and 16.5 m. The hand-over introduced no state discontinuity in any run, $\varepsilon_{ho} = 0$; this is a property of the hand-over path that was actually taken, which commits the translational state unchanged, rather than a general statement about the transition. The narrow-to-wide camera switch fired between 2.14 and 2.26 m true range against its 2.5 m design threshold, the offset being the confirmation count required before the switch is committed.

4.3.4 Real-Time Execution

Figure 13 reports the navigation cycle time against the 500 ms budget, with the test running on a consumer laptop. Across 5715 navigation cycles and 5837 guidance cycles there were *no* deadline overruns, and the scheduling jitter remained below 1.34 ms. The worst-case computational load is $C_{nav} = 0.364$ and $C_{gui} \leq 0.0013$, for the navigation and guidance processes, respectively. The guidance task is almost computationally free and the navigation task retains roughly a factor of three of margin.

Looking at the figure, it is clearly visible how the average computational time for frame-processing cycles (the upper measurements cluster) is higher in the Parallel Zone mode (red dots): this operative mode runs both measurement pipelines concurrently, the centroid-based far-range measurement and the marker-based pose solution, in order to arbitrate between them before committing to the

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

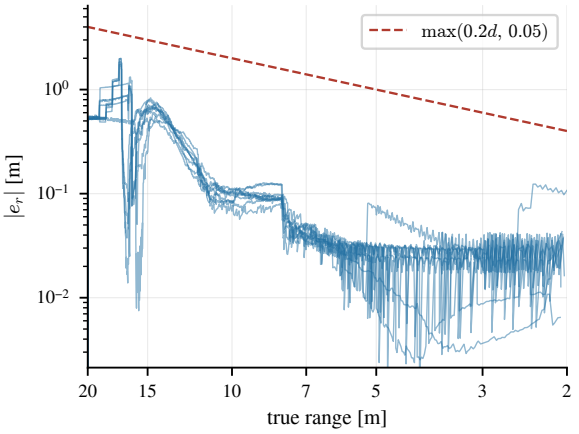


Fig. 12. Closed-loop navigation position error against true range for the nine-run dispersion set, with the accuracy requirement $\max(0.2d, 0.05\ m)$ overlaid.

close-range filter.

Cycles in which no new frame is available complete in under 2 ms, the navigation output being dead-reckoned between camera samples.

Because the simulator and the GN module share no clock, the measurement latency is itself a result. A camera frame is on average 724 ms old, and at worst 1.28 s old, at the moment the navigation task consumes it, this interval covering rendering, encoding, transport and buffering. Adding the navigation computation and up to one guidance period, the control action derived from a given image is applied approximately 1.1 s after that image was formed. Nevertheless, the approach converges from every dispersed initial condition: the guidance loop is tolerant to a transport delay much larger than what could be obtained on properly engineered flight hardware.

4.3.5 Terminal Accuracy

Table 7 reports the true terminal state against the eight docking criteria. The four translational criteria are satisfied in all nine runs and with substantial margin: the tightest is the axial position, at 82% of its limit in the worst run, while lateral position, lateral rate and closing rate reach at most 25%, 8% and 56% respectively. The closed-loop navigation position error at the instant of arrival lies between 6.5 and 96 mm.

The attitude criterion behaves differently. In five of the nine runs the combined yaw and pitch tilt at arrival lies between 2.12° and 3.89°, against a 2° limit, while the on-board estimate reports between 0.27° and 0.86° in *every* run. The arrival check therefore passes and the vehicle declares a successful docking with a true tilt of up to twice the admissible value, giving $P_{fp} = 5/9$ against a requirement of zero.

The cause is geometric rather than algorithmic. The marker set spans approximately 30 mm of depth across a

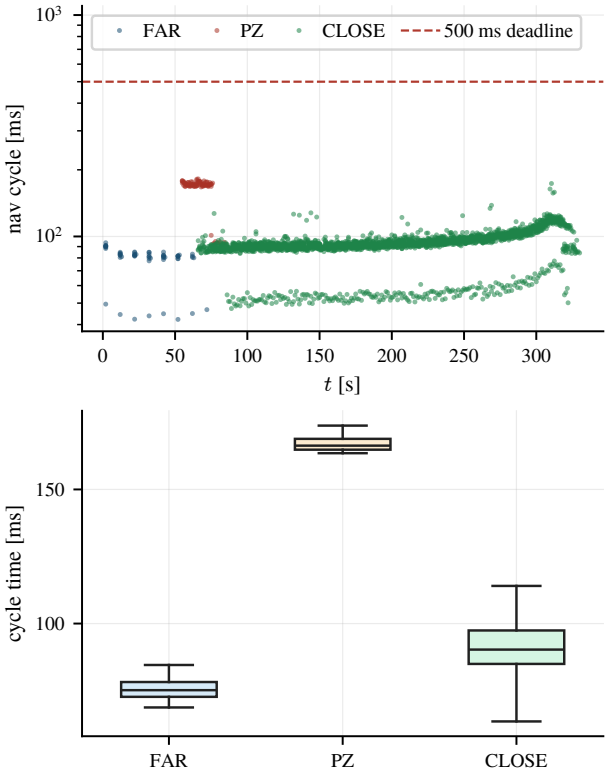


Fig. 13. Navigation cycle time against the 500 ms deadline. Top: frame-processing cycles of the nominal run, coloured by phase. Bottom: distribution per phase over the nine-run set.

Table 7. True terminal state against the docking box over the nine-run dispersion set

Quantity	Limit	Range	Worst
Lateral position [m]	0.10	0.005–0.025	25%
Axial position [m]	0.08	0.026–0.066	82%
Lateral rate [m/s]	0.04	0.001–0.003	8%
Closing rate [m/s]	0.10	0.055–0.056	56%
|yaw| + |pitch| [°]	2.0	0.90–3.89	**194**%
|roll| [°]	4.0	0.01–0.74	18%
Lateral ang. rate [°/s]	0.20	0.007–0.164	82%
Axial ang. rate [°/s]	0.20	0.013–0.037	19%

0.42 m extent, so the two rotations about axes normal to the approach direction are weakly observable: a 3° error in either rotation displaces the marker corners by well under one pixel at 13 m, which is below the residual noise of the corner extraction. The rotation about the target face normal, mostly coincident with the Y axis, is five to seventeen times better determined. Notably, the navigation module remains compliant with its own accuracy requirement throughout, $P_{rot} \geq 89\%$ against a 6° bound, and its covariance remains consistent. The navigation module meets the accuracy it was specified to meet, but such accuracy is not sufficient to support a 2° arrival criterion

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

under the current geometrical design of the target.

Two remedies follow directly from the diagnosis. Introducing genuine depth into the marker layout on the docking interface would make the out-of-plane tilts observable at the corner-extraction precision already achieved; alternatively, an attitude source independent of the marker plane could be fused into the rotational filter. It is worth observing that both campaigns reported in this work terminate on a perception or geometry limit rather than on an estimation limit: the performance of the relevant environment test in Section 3 is bounded by apparent marker size at range, and the closed-loop performance discussed here by marker-plane depth at contact.

Table 8 summarises the compliance assessment over the whole criteria set.

Table 8. Closed-loop SIL compliance assessment over the nine-run dispersion set

Metric	Result	Req.	
P_{pos}, P_{rot}	100%, 89%	$\geq 80\%$	✓
$\sigma_{pos}, \sigma_{rot}$	94%, 100%	$\geq 90\%$	✓
P_{det}	100%	$\geq 90\%$	✓
N_{reinit}	0	—	✓
P_{cone}	100%	100%	✓
ΔV_{tot}	0.14–0.18 m/s	—	✓
P_{sat}	0%	—	✓
$\exists T_{phase}$	monotone, $9/9$	monotone	✓
ε_{ho}	0	—	✓
$\exists T_{cam}$	2.14–2.26 m	≈ 2.5 m	✓
e_{dock}	$4/9$ inside box	inside	×
P_{fp}	$5/9$	0	×
C_{nav}, C_{gui}	0.364, 0.0013	≤ 1	✓
P_{dl}	0 overruns	$\geq 99.9\%$	✓

4.4 Deployment on the Target Processor

The guidance and navigation implementation validated above is not tied to the workstation on which those runs were executed. The same source, without structural modification, was cross-compiled for the RISC-V system-on-chip (SoC) selected as the platform on which the OrbSight camera is being developed, and executes there against the identical simulator interface of Figure 9. Numerical equivalence was verified by replaying a recorded stream of 1034 frames through both builds: the phase transitions and every discrete decision of the supervisor are identical on all steps, and the continuous states agree to within 10^{-8} in relative terms.

Real-time behaviour, however, does not transfer. On the SoC processor the navigation update requires between 3.5 and $5.8\ s$ against its $500\ ms$ budget, a huge decrease in performance with respect to the workstation figure. The cost is almost entirely image handling: marker detection and image decoding together account for more than 99% of it, while the estimation filters themselves complete in a comparable amount of time to the one characterized on the workstation. The gap is due to a shortcoming of the processor, since its cores implement no vector extension and every image kernel therefore executes scalar over 4.2-megapixel frames. The resulting measurement latency is large enough to prevent the closed loop from converging on such platform today.

This is a hardware margin problem that can be solved without requiring a different platform, as the same SoC carries a Field Programmable Gate Array (FPGA) fabric to which the marker-detection front-end can be offloaded. Alternatively or in addition to that, reducing the working resolution of the detection stage could be evaluated as an option as well, as such precaution alone recovers more than an order of magnitude.

5. Conclusions

This work presented a system-level validation framework for the GEORyder autonomous rendezvous and docking GN architecture, bridging numerical simulation and embedded hardware execution.

The dual validation campaign confirmed the operational viability of the design. Hardware-in-the-Loop tests at the RAFFAELLO facility proved the vision pipeline robustness against real illumination gradients and scale variations, ensuring consistent covariance propagation, while closed-loop SIL simulations in the PROX-SIMA environment demonstrated 100% approach-cone compliance and successful sensor/mode transitions across all dispersion runs.

However, two key constraints emerged from testing. The shallow out-of-plane relief of the marker layout left the attitude filter unable to prevent five of the nine runs from declaring arrival outside the 2° tilt box, even though the navigation module met its own 6° accuracy requirement throughout: a geometric limit at touch-down rather than an algorithmic one. On the target RISC-V processor, the navigation cycle time reached 3.5–5.8 s against its 500 ms budget due to scalar image processing, a computational bottleneck rather than a design flaw.

Future developments will explore augmenting the out-of-plane depth of the marker arrangement or fusing complementary attitude measurements to tighten terminal tilt estimation, alongside offloading the ArUco detection front-end to the on-chip FPGA fabric to recover the timing margins required for flight deployment.

Acknowledgements

This work was carried out within the GEORyder project, led by Infinite Orbits and funded by the European Commission under Horizon Europe 2023. The authors gratefully acknowledge this support, as well as the DART Laboratory (Politecnico di Milano) for granting access to the RAFFAELLO facility [21] and providing technical support during the experimental dataset acquisition.

Copyright © 2026 by Annachiara Ippolito. Published by the IAF, with permission and released to the IAF to publish in all forms.

References

[1] W. H. Gerstenmaier, D. Parker, G. Leclerc, and R. Shirama, "International rendezvous system interoperability standards (irsis)," NASA, Tech. Rep., 2019. [Online]. Available: https://internationaldeepspacestandards.com/wp-content/uploads/2024/02/rendezvous_baseline_final_3-2019.pdf.

[2] A. Comellini, J. Vasconcelos, and J. McHugh, "Verification and validation of rendezvous and proximity operations safety," ESA, Tech. Rep., 2024. [Online]. Available: https://nebula.esa.int/content/verification-and-validation-rendezvous-and-proximity-operations-safety.

[3] *International docking system standard interface definition document (idss idd), revision g*, 2026. [Online]. Available: https://internationaldeepspacestandards.com/wp-content/uploads/2026/01/M2M-IDSS-IDD-Rev-G-1-23-2026.pdf.

[4] T. H. Park, J. Bosse, and S. D'Amico, *Robotic testbed for rendezvous and optical navigation: Multi-source calibration and machine learning use cases*, 2021. arXiv: 2108.05529 [cs.RO]. [Online]. Available: https://arxiv.org/abs/2108.05529.

[5] N. Faraco, M. Massari, P. Di Lizia, *et al.*, "Prox-sima: A modular simulator for the validation of in-orbit servicing and close proximity missions gnc techniques," in *IAF Space Operations Symposium (75th International Astronautical Congress)*, International Astronautical Federation (IAF), 2024, pp. 259–265.

[6] A. Ippolito, N. Faraco, M. Maestrini, and M. Massari, "Dual-mode vision based navigation framework for autonomous rendezvous and docking operations in geostationary transfer orbits," in *Proceedings of the 2026 CEAS EuroGNC conference*, Madrid, Spain, May 2026. DOI: 10.82124/CEAS-GNC-2026-067.

[7] C. Vela, G. Fasano, and R. Opromolla, "Pose determination of passively cooperative spacecraft in close proximity using a monocular camera and ArUco markers," *Acta Astronautica*, vol. 201, pp. 22–38, 2022.

[8] X.-S. Gao, X.-R. Hou, J. Tang, and H.-F. Cheng, "Complete solution classification for the perspective-three-point problem," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 25, no. 8, pp. 930–943, 2003.

[9] J. L. Crassidis and F. L. Markley, "Unscented filtering for spacecraft attitude estimation," *Journal of Guidance, Control, and Dynamics*, vol. 26, no. 4, pp. 536–542, 2003.

[10] F. L. Markley and J. L. Crassidis, "Attitude estimation using modified rodrigues parameters," in *Flight Mechanics/Estimation Theory Symposium*, 1996.

[11] R. Furfaro, D. Cersosimo, and D. R. Wibben, "Asteroid precision landing via multiple sliding surfaces guidance techniques," *Journal of Guidance, Control, and Dynamics*, vol. 36, no. 4, pp. 1075–1092, 2013, ISSN: 15333884. DOI: 10.2514/1.58246.

[12] Q. Li, B. Zhang, J. Yuan, and H. Wang, "Potential function based robust safety control for spacecraft rendezvous and proximity operations under path constraint," *Advances in Space Research*, vol. 62, no. 9, pp. 2586–2598, Nov. 2018, ISSN: 18791948. DOI: 10.1016/j.asr.2018.08.003.

[13] Y. Shtessel, C. Edwards, L. Fridman, A. Levant, *et al.*, *Sliding Mode Control and Observation*. Springer, 2014, vol. 10, pp. 1–41.

[14] X. Zhao and S. Zhang, "Adaptive saturated control for spacecraft rendezvous and docking under motion constraints," *Aerospace Science and Technology*, vol. 114, Jul. 2021, ISSN: 12709638. DOI: 10.1016/j.ast.2021.106739.

[15] Basler AG. "Basler ace aca1300-22gm – EMVA data." (2026), [Online]. Available: https://www2.baslerweb.com/en/downloads/document-downloads/basler-ace-aca1300-22gm-emva-data/.

[16] A. Ippolito, L. Giorcelli, N. Faraco, M. Maestrini, and M. Massari, *GEORyder Close Range Spacecraft Rendezvous Dataset for Vision-Based Navigation in Geostationary Transfer Orbits (v1.0)*, version 1.0, Zenodo, Mar. 2026. DOI: 10.5281/zenodo.10849311. [Online]. Available: https://doi.org/10.5281/zenodo.10849311.

[17] K. Zuiderveld, "Contrast limited adaptive histogram equalization," in *Graphics Gems IV*, 1994, pp. 474–485.

[18] D. Bradley and G. Roth, "Adaptive thresholding using the integral image," *Journal of Graphics Tools*, vol. 12, no. 2, pp. 13–21, 2007.

[19] Modelica Association, "Functional mock-up interface for model exchange and co-simulation, version 2.0.4," Modelica Association Project FMI, Tech. Rep., 2022. [Online]. Available: https://fmi-standard.org/downloads/.

[20] G. Franzini and M. Innocenti, "Relative motion dynamics with arbitrary perturbations in the local-vertical local-horizon reference frame," *The Journal of the Astronautical Sciences*, vol. 67, no. 1, pp. 98–112, 2020, ISSN: 0021-9142, 2195-0571. DOI: 10.1007/s40295-019-00185-0. [Online]. Available: http://link.springer.com/10.1007/s40295-019-00185-0.

[21] DART Laboratory, *RAFFAELLO: Robotic HIL testing for autonomous navigation and control in close-proximity operations*, https://dart.polimi.it/facilities/raffaello/, Accessed: 2026-09-08, Politecnico di Milano, Department of Aerospace Science and Technology.