

Adaptive End-to-End Architecture for Autonomous Spacecraft Navigation and Control During Rendezvous and Proximity Operations

Justin Kruger*, Tommaso Guffanti†, Tae Ha Park*, Mason Murray-Cooper*, Samuel Low*, Toby Bell‡ and Simone D'Amico§

Stanford University, Stanford, CA, 94305

Christopher Roscoe¶ and Jason Westphal¶
Ten One Aerospace, Washington, DC, 20002

Onboard guidance, navigation, and control (GNC) during rendezvous and proximity operations (RPO) between spacecraft poses unique challenges for associated algorithms. Future missions will require greater on-board autonomy while maintaining in-orbit safety guarantees at varying distances, and scenarios of interest may feature multiple spacecraft which may be cooperative or non-cooperative. This paper presents the conception and development of a new GNC software payload for distributed space systems, which enables safe, autonomous RPO between multiple objects with maximum flexibility and modularity. Navigation algorithms fuse far-range camera images, close-range camera images, differential carrier-phase Global Navigation Satellite System data, and inter-satellite crosslink data to estimate absolute orbits, relative orbits, target pose, and auxiliary states across the entire regime of interest. A control algorithm suite provides optimal maneuver solutions for efficient long-term formation maintenance at far range, centimeter-level rendezvous accuracy at close range, and fast, robust collision avoidance. Prototype simulations featuring cooperative and non-cooperative targets at far, medium and close ranges demonstrate highly capable GNC performance for distributed space systems, and are a significant step towards achieving complete integration of a flexible autonomous RPO kit for spacecraft.

I. Nomenclature

a	=	orbit semi-major axis
α	=	azimuth bearing angle
$\boldsymbol{\alpha}$	=	orbit element state vector
C	=	camera image reference frame
$\delta\boldsymbol{\alpha}$	=	relative orbit element state vector
δa	=	relative semi-major axis
δe_x	=	x-component of relative eccentricity vector
δe_y	=	y-component of relative eccentricity vector
δi_x	=	x-component of relative inclination vector
δi_y	=	y-component of relative inclination vector
$\delta\lambda$	=	relative mean longitude
$\Delta\mathbf{v}$	=	maneuver delta-v
e	=	orbit eccentricity
ϵ	=	elevation bearing angle
f	=	orbit true anomaly

*Ph. D. Candidate, Stanford Space Rendezvous Laboratory.

†Postdoctoral Fellow, Stanford Space Rendezvous Laboratory.

‡Ph. D. Student, Stanford Space Rendezvous Laboratory.

§Associate Professor, Department of Aeronautics and Astronautics. AIAA Associate Fellow.

¶Co-Founder, Ten One Aerospace.

i	=	orbit inclination
\mathcal{I}	=	inertial reference frame
M	=	orbit mean anomaly
ω	=	orbit argument of periapsis
Ω	=	orbit right ascension of the ascending node
\mathbf{p}	=	Rodrigues relative orientation vector
\mathbf{q}	=	quaternion vector
r	=	orbit radius
\mathcal{R}	=	radial/along-track/cross-track (RTN) reference frame
σ	=	standard deviation
t_i	=	image integration time
t_{sk}	=	station-keeping time horizon
t_{wp}	=	waypoint time horizon

II. Introduction

IMMENSE growth in usage of the space environment renders conducting rendezvous and proximity operations (RPO) between spacecraft an increasingly necessary capability. Strong interest in applications such as on-orbit servicing, assembly and manufacturing [1], space domain awareness [2], space debris management [3], and space sustainability requires new software and hardware packages for satellites which can support these objectives in a robust, safe, flexible, autonomous fashion. Of particular interest are integrated algorithms and architectures which enable the characterization, approach, and rendezvous of one or more servicer or observer satellites with respect to one or more target resident space objects (RSO).

To enable this objective, it is necessary to characterize targets and robustly determine their relative orbits during a servicer’s approach. Furthermore, there is a need for robust knowledge of relative attitudes too. However, targets may be known or unknown, cooperative or uncooperative, which requires a flexible navigation architecture that may handle various scenarios. Guidance and control algorithms should provide ΔV -efficient maneuver planning accurately while maintaining the passive safety constraints for collision avoidance on the rendezvous trajectory. The on-board navigation and control algorithm suites are required to span the complete operational range of the system, from tens or hundreds of kilometers of inter-satellite separation at the commencement of an approach, down to meter-level separations during a rendezvous. More generally, the resulting guidance, navigation and control (GNC) software architecture should be maximally autonomous, to support responsive in-orbit use and mission scalability, and should be suitable for low size, weight and power (SWaP) and commercial off-the-shelf (COTS) hardware, for broad applicability to future missions. These pose extremely challenging requirements.

RPO has been demonstrated or attempted in orbit during several previous missions. Examples include the DARPA Orbital Express demonstration, in 2007 [4]; the ARGON experiment, in 2012 [5]; the AVANTI experiment, in 2016 [6]; the CPOD mission, in 2022 [7]; and the upcoming ClearSpace-1 mission in 2026 [8]. RPO is also regularly performed for vehicles approaching the International Space Station. Nevertheless, RPO operations are not yet routine and have required a high degree of supervision and specialized software and hardware. In response, Stanford’s Space Rendezvous Laboratory (SLAB) is developing a novel GNC software architecture for distributed space systems, intended to enable autonomous RPO between multiple cooperative or non-cooperative objects. Navigation algorithms fuse data from multiple sensors for accurate, robust system state estimation, and a control algorithm suite provides accurate, efficient maneuver planning and collision avoidance. The architecture has been developed in concert with an RPO hardware kit from Ten One Aerospace which provides the necessary sensing, actuation, and avionics hardware in a COTS and low-SWaP format. Together, the software and hardware kits provide simple, modular and flexible integration of end-to-end RPO capability into future mission proposals.

This paper describes the conception and design of the RPO kit GNC architecture, its capabilities, its constituent GNC algorithms, and their integration. Initial development of many of these algorithms by SLAB has been motivated by prior or current flight projects. These include TanDEM-X (2010), an synthetic aperture radar interferometry mission consisting of two satellites [9]; PRISMA (2010), a technology demonstrator consisting of two SmallSats [5]; Starling (2023), an autonomous swarming technology demonstration consisting of four CubeSats [10]; VISORS (2024), distributed telescope mission for solar imaging in the extreme ultra-violet spectrum [11]; and SWARM-EX (2024), a long-term distributed aeronomy experiment consisting of three CubeSats [12]. The flight heritage these missions have provided, or will provide, is valuable. However, the combination of their algorithms into a single architecture requires

an extensive integration effort and algorithms require adaptations to optimally support the RPO context.

For example, consider Figure 1, which presents simulated hardware-in-the-loop images of the same spacecraft at varying inter-satellite distances (ISD) obtained through SLAB’s Optical Stimulator [13] testbed. Optical characteristics and navigational requirements in each of these regimes are very different, producing challenges for algorithmic development and integration which must be overcome. The new navigation system must fuse data from a variety of sources, including far- and close-range imagery, Global Navigation Satellite System (GNSS) receivers, and an inter-satellite crosslink, and must intelligently transition between regimes in which each sensor is most appropriate. It must also adapt to distributed systems which may contain multiple servicers or targets of interest. Equivalently, the new maneuver planner must provide control solutions which are appropriate to all relevant mission modes and are as optimal as possible, while simultaneously ensuring passively safe operation. The resulting set of highly autonomous, flexible and transitional capabilities has not before been demonstrated by a single space system. Further benefits are provided by a modular design such that if only a particular subset of capabilities is needed, relevant components may be integrated piecemeal in a more lightweight fashion.

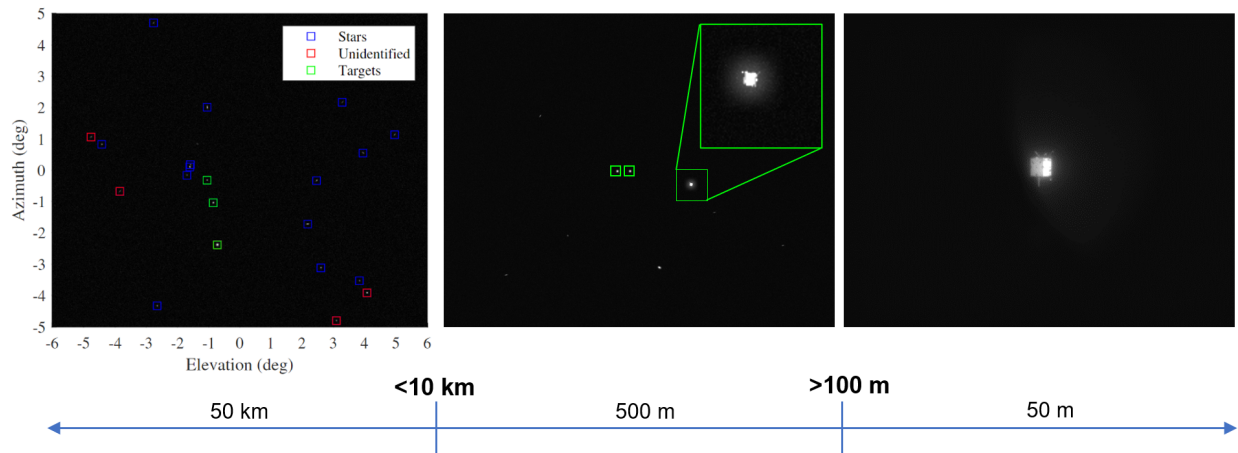


Fig. 1 Simulated images of a rendezvous with the PRISMA mission’s Tango spacecraft [5] at varying ISD using the SLAB Optical Stimulator [13] and a Blue Canyon Technologies Nano Star Tracker [14].

Development of the GNC package is divided into ‘Tier 1’ and ‘Tier 2’ stages. The subject of this paper is Tier 1, which addresses integration of four components: 1) far-range navigation for targets via optical angles-only techniques, 2) close-range navigation for cooperative targets only via carrier-phase differential GNSS (CDGNSS) techniques, 3) far-range control using efficient closed-form Δv solutions, and 4) close-range control using highly accurate optimal maneuver planning. Tier 2, which is the subject of a future publication, will address close-range navigation for non-cooperative targets via data-driven pose estimation, as well as navigation and control for n -object systems consisting of multiple servicers and targets. Tier 2 also implements intelligent, autonomous transition logic for algorithms when moving from far- to mid- to close-range regimes. This work verifies a prototype Tier 1 architecture with high-fidelity simulations at far range, medium range and close range. Results display highly capable performance for GNC of distributed space systems, and are a significant step towards achieving complete integration of a fully-featured software/hardware RPO kit for spacecraft, with the ultimate goal of enabling autonomous end-to-end RPO missions.

Following this introduction, the ‘Architecture’ section gives a high-level overview of the structure and operation of the GNC architecture. The ‘Algorithms’ section summarizes each constituent algorithm, its purpose, and its methodology. The ‘Simulation’ section presents results from testing of the Tier 1 prototype, followed by conclusions and plans for future development.

III. GNC Architecture

The new GNC software architecture is deployed on one or more observer or servicer satellites. Observers obtain measurements of targets using their on-board sensors and use their on-board actuators to perform planned maneuvers. Together, the servicers and targets form a distributed space system (DSS). Conceptually, the architecture aims to support autonomous RPO for DSS across all relevant operational stages and inter-satellite distances – from initial acquisition,

to approach, to rendezvous, to eventual docking – using a single integrated system. In the following discussions, the spacecraft running the instance of the GNC architecture being discussed is the ‘host’. Other cooperative spacecraft running instances of the GNC architecture are ‘peers’. A ‘target’ is a spacecraft whose relative state is being estimated by the host. Targets may be uncooperative or may be peers themselves.

A. Modular Structure

A top-level architecture overview is presented in Figure 2. Core software modules of the architecture are light gray blocks. Individual algorithmic blocks are shaded within each for navigation and control. External subsystems of the space and ground segments are dark gray ovals. Exchanged data is colored in blue and data flow is indicated by arrows.

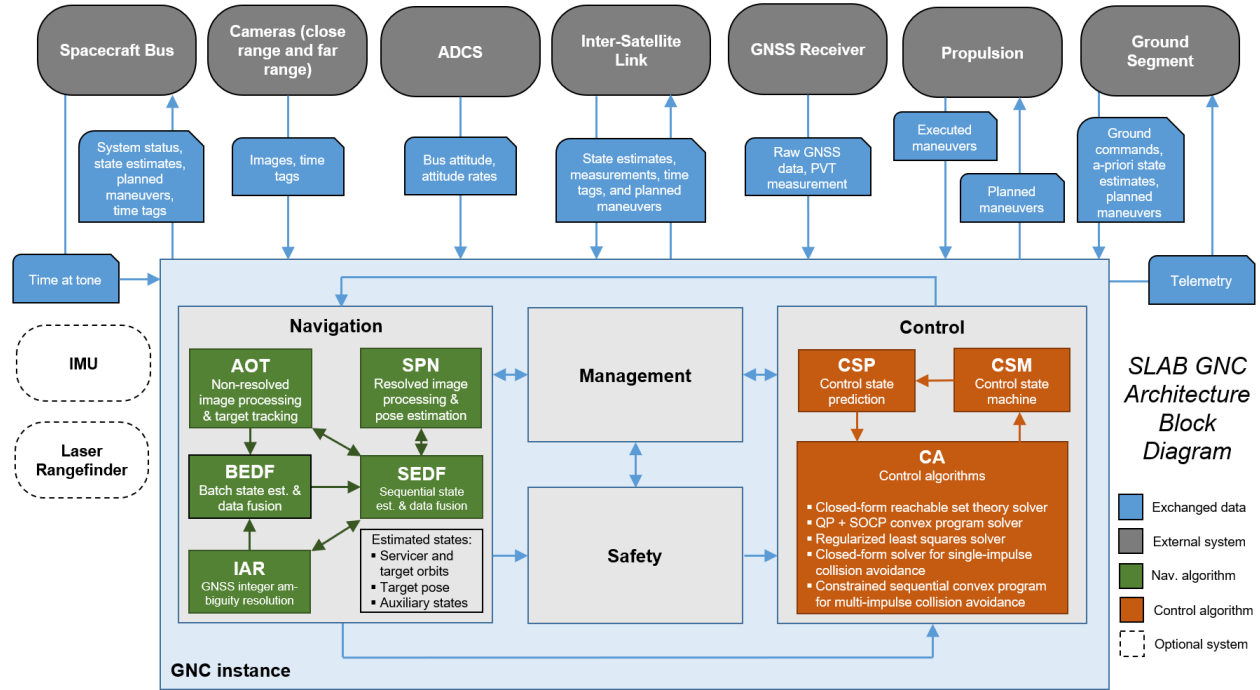


Fig. 2 The GNC architecture including external systems (dark gray), exchanged data (blue), and core algorithms (green and orange).

The GNC software consists of four core modules: management, navigation, control, and safety. The management module handles transitions between different operational modes of the navigation and control modules (e.g. transitioning from far range optical navigation to close range optical navigation during approach, or transitioning from a control approach mode to a control station-keeping mode after arriving at a waypoint) as well as fault detection, isolation and recovery (FDIR). The navigation module uses input measurements to generate output state estimates. Simultaneously, the safety module performs orbit predictions to produce minimum separation estimates for any system agents whose state estimates are available, for collision detection. The control module uses guidance, navigation and safety inputs to produce a buffer of planned system maneuvers.

To navigate, the architecture fuses data from a far range visible-wavelength camera (e.g. a star tracker), a close range long-wave infrared (LWIR) camera, an inter-satellite crosslink (ISL), pseudorange and single-difference carrier phase measurements from a GNSS receiver, and potentially other sensors (e.g. a laser rangefinder or accelerometer), as per Figure 2. The far-range camera is utilized to obtain bearing angles to targets for absolute and relative orbit determination during target approach. The close-range camera is utilized to obtain resolved target imagery for target pose estimation during rendezvous. The ISL enables sharing of measurement and state information between cooperative spacecraft for improved performance, robustness, and distribution. The GNSS receiver is used for absolute orbit determination of the servicer and precise relative orbit determination of cooperative targets via carrier-phase differential GNSS (CDGNSS). Two estimation frameworks fuse data from each sensor, a batch estimation and data fusion system, and a sequential estimation and data fusion system. The batch system is used to perform target state initializations in flight, if a-priori

state knowledge is limited or unknown. The sequential system applies an adaptive, efficient unscented Kalman filter (UKF) to refine initialized states over time.

Guidance is provided via a combination of manual commands from the ground and autonomous re-planning on board. Maneuvers are computed using a suite of solver algorithms, each with different capabilities as per their accuracy, flexibility, robustness, safety, and computational speed. These algorithms interact with a control state machine and a maneuver buffer which allows the system to perform autonomous closed-loop control and seamlessly transition between different operational scenarios. Autonomous collision avoidance for single- or multi-target systems is simultaneously implemented to ensure vehicle safety during RPO sequences.

B. Concept of Operations

Consider a typical RPO scenario in which a servicer is approaching a target. Nominal execution of the GNC architecture proceeds as follows.

First, the management module selects which algorithms are active and controls transitions between algorithms. Examples of transitions include moving from non-resolved image processing (for navigation during a far range approach) to close-range image processing (for rendezvous at close ranges); or moving from a numerical optimization control solver (for highly accurate control during an approach trajectory) to a closed-form control solver (for efficient control during station-keeping periods). Transitions can be handled autonomously or may be planned from the ground. The module also manages FDIR tasks and fault reporting. The relevant system status and error flags that are generated are reported to the spacecraft bus and ground segment.

In parallel, the navigation module takes sensor inputs and control inputs and updates the system state estimate. Processing of camera data and GNSS data is necessary before state updates are performed via the angles-only tracking (AOT), spacecraft pose network (SPN) and integer ambiguity resolution (IAR) blocks in Figure 2. The batch estimation and data fusion (BEDF) algorithm performs batch estimation tasks which may be necessary to initialize the navigation system, e.g. angles-only initial relative orbit determination for a previously unknown target. The sequential estimation and data fusion (SEDF) algorithm flexibly handles subsequent state propagation for the servicer and its targets, measurement updates from each local sensor, and remote measurement updates from the ISL. The active set of estimated states may vary depending on system status and data availability, and encompasses orbits, attitudes, and auxiliary states (e.g. differential ballistic coefficients, clock offsets, sensor biases, and float ambiguities) for the observer and its targets. The up-to-date SEDF state estimate is subsequently used to aid the AOT, SPN and IAR algorithms. State estimates are also sent to the management, safety and control modules, and to architecture outputs.

The safety module uses system state estimates to assess the probability of collision between the servicer and other system members. It aims to mitigate the risk of collision in case of critical subsystem failure, employing passive and reactive strategies in tandem to ensure safe in-orbit operation. Its collision probability assessments are sent to the control module for incorporation into upcoming maneuver plans. Collision assessments take the form of minimum predicted target separations in the servicer's radial-tangential-normal (RTN) frame, and associated confidence levels.

The control module uses updated state estimates, indicators from the collision detection system, and guidance information from the ground to generate new maneuver plans. This task is managed by an internal state machine which tracks current system status with respect to its operational mode and control objectives, performs autonomous maneuver planning and re-planning, and performs autonomous collision avoidance. Maneuvers are computed using a suite of control solver algorithms, noting that the specific control solver applied may vary depending on the current GNC task. Updated maneuver plans, in the form of ΔV sequences, are subsequently sent to navigation module and to architecture outputs.

C. Software Integration

The GNC architecture encompasses a large set of capabilities which may or may not be required for a particular scenario. It has therefore been designed to be inherently flexible with core software modules and algorithms integrated in a modular fashion. Each of the four modules and their algorithms are accessed through a unified application programming interface (API). This API is kept consistent across the software regardless of whether a user is interacting with the GNC architecture at architectural, modular, or algorithmic level. Algorithms and modules are additionally designed to be separable so that e.g. the control module can be integrated without navigation, or a particular algorithm can be called upon outside of the context of the complete module.

In support of these goals, the software applies event-driven timing and synchronization. It is not called at fixed sample times but instead, the arrival of new inputs triggers the appropriate functions, such as a new camera image

triggering image processing algorithms. Limits are placed on execution intervals such that modules do not run too often (thus monopolizing computation) or too rarely (thus becoming out of date). This allows for maximum flexibility as the presence of different inputs and outputs varies. Outputs are provided using callback functions such that higher levels of software may invoke specific callbacks to obtain the desired, most up-to-date output of interest at any time.

IV. GNC Algorithms

Tables 1 and 2 present a summary of individual algorithms applied within the architecture. In concert, the algorithms enable RPO for both cooperative and non-cooperative targets, for both binary systems and n -spacecraft systems. The only requirement is that the target must have a known 3D model if pose estimation is to be performed on-board the servicer. The TanDEM-X [9] and PRISMA [5] missions have partially demonstrated some of the applied navigation, control, and safety concepts, and the NASA Starling mission is validating the applied angles-only navigation algorithms. The upcoming VISORS [11] and SWARM-EX [12] missions (both scheduled to launch in 2024) will further verify a subset of constituent algorithms. This heritage ensures that the algorithms are robust and can feasibly be operated in orbit with respect to computation, sensing and actuation requirements.

A. Absolute and Relative Navigation

The overall navigation task is managed by a set of measurement processing algorithms and a navigation filter which performs data fusion and state estimation. If targets are non-resolved, bearing angles are extracted from images and used to perform absolute and relative orbit determination. When targets become resolvable in imagery, a convolutional neural network (CNN) is employed for target pose estimation. If a target is cooperative, CDGNSS and IAR are leveraged for improved positioning accuracy. Included algorithms are Angles-Only Tracking (AOT), the Spacecraft Pose Network (SPN), Integer Ambiguity Resolution (IAR), Batch Estimation & Data Fusion (BEDF), and Sequential Estimation & Data Fusion (SEDF).

Table 1 provides a list of possible inputs and outputs to each algorithm, noting that in many cases, the presence of a specific input is optional. For example, the system does not require access to peer GNSS information, but may utilize such information when provided; or if multiple peers are sharing bearing angle measurements of the same target, the host does not require a GNSS receiver for absolute orbit determination [10].

1. Angles-Only Tracking

AOT supports measurement acquisition at larger ISD, e.g. as part of the initial target approach during which it is approximately a point source in imagery. AOT reads sequential images from the far range camera and begins by performing a series of common image processing steps [10]: region of interest (ROI) detection, ROI centroiding, conversion of pixel centroids to unit vectors, Pyramid star identification [21], and q-method attitude determination [22]. It subsequently produces a list of unidentified bearing angles in the field of view (FOV) which must be associated with previously tracked or un-tracked targets.

If target orbit states are known or being estimated, they may be used to associate bearing angles. Otherwise, AOT applies a multi-hypothesis tracking (MHT) algorithm known as SAMUS [15] to detect and track previously unknown targets. SAMUS is part of the Absolute and Relative Trajectory Measurement System (ARTMS) software payload currently conducting flight experiments aboard the NASA Starling swarm [10]. In general, MHT is a robust tracking approach [23] that nonetheless requires costly computation as it enumerates the growing set of target track hypotheses (as formed from unidentified bearing angles over time) and propagates them forwards. To compensate, SAMUS leverages domain-specific knowledge regarding the evolution of target kinematics in the servicer's reference frame. Consider the following equation

$$\begin{bmatrix} \epsilon \\ \alpha \end{bmatrix}^{\mathcal{R}} \approx \frac{r}{a} \begin{bmatrix} x_1 - x_2(\cos(f - x_3) + \frac{\epsilon}{2} \cos(2f - x_3)) \\ x_4 + x_5 \sin(f + \omega - x_6) \end{bmatrix} \quad (1)$$

which maps bearing angles in \mathcal{R} to a set of x_1, \dots, x_6 constants analogous to the quasi-nonsingular relative orbital elements (ROE) [24], using f of the servicer as the only quickly-changing parameter. Given three bearing angle pairs and the accompanying orbital parameters of the servicer r, a, e, f, ω (or the equivalent non-singular orbit parameters), the model can be fitted to target tracks and used to develop kinematic criteria with which to assess hypotheses. This allows particularly efficient pruning of unlikely tracks and highly accurate determination of the best hypothesis. AOT

Table 1 Summary of constituent navigation algorithms.

Algorithm	Angles-Only Tracking (AOT) [15]	Spacecraft Pose Network (SPN) [16]	Integer Ambiguity Resolution (IAR) [17]	Batch Estimation and Data Fusion (BEDF) [18]	Sequential Estimation and Data Fusion (SEDF) [19, 20]
Objective	Process sequential images from on-board cameras to produce target bearing angles	Process sequential images from on-board cameras to estimate target pose	Process DDCP GNSS measurements to resolve integer wavelength ambiguities	Use batches of inertial target bearing angles to produce target relative orbit estimates	Fuse multi-sensor and multi-observer measurements to refine servicer and target state estimates
Relevant Techniques	Multi-hypothesis tracking Kinematic target modeling	Convolutional neural networks	mLAMBDA algorithm	Iterative least squares Target range sampling	Unscented Kalman filter Adaptive process noise Statistical data association
Possible Inputs	Host camera images Host orbit estimate	Host camera images Host orbit estimate Host attitude estimate	Host GNSS Peer GNSS	Host bearing angles Host camera attitude Host orbit estimate	Host & peer bearing angles Host & peer GNSS Host & peer camera attitude Host SPN heatmaps Host & peer state estimates Host & peer maneuver plan
Possible Outputs	Target bearing angles Sensor attitude	Target pose w.r.t. servicer	Resolved DDCP integers	Target relative orbit est. and covariance	Host absolute orbit estimate Peer relative orbit estimates Host absolute attitude estimate Peer relative attitude estimates Host & peer auxiliary state est.
Nominal Sample Time (LEO)	60 sec	10 sec	20 sec	1 orbit	5-120 sec
Representative Accuracy ($1-\sigma$)	<35'' angle error <20'' attitude error	<10 cm relative pos. <1° orientation	<1 cm relative position error	<20% target range error	<5 m abs. pos. (GNSS w/out IAR) <5 cm rel. pos. (GNSS w/out IAR) <1 cm rel. pos. (GNSS with IAR) <1 km abs. pos (angles-only) <2% $\delta\lambda$ rel. pos. (angles-only) <10 cm rel. pos. (SPN) <1° rel. att. (SPN)
Applicability	Non-resolved targets	Resolved targets with known 3D model	Cooperative targets	Non-resolved targets	-
Flight Heritage	PRISMA Starling	-	VISORS	Starling	Starling VISORS
System Tier	1	2	1	1	1

outputs target bearing angle tracks and corresponding camera attitude quaternions which are used by BEDF and SEDF. AOT expects a sample rate of 25-100 images per orbit and has been tested for up to 5 persistent targets in the FOV. Targets may be cooperative or noncooperative but are assumed to have consistent optical presence (with allowances for regular camera blinding or eclipse periods). For nominal cases in low Earth orbit (LEO), >99.9% tracking precision is achieved.

2. Spacecraft Pose Network

The SPN algorithm [16] supports measurement acquisition at smaller ISD, e.g. during a final rendezvous with the target when it is resolved in imagery. SPN reads sequential images from a close range monocular camera, detects the target, extracts its ROI, and provides outputs that enable estimation of the target object’s translation and orientation, i.e. its pose with respect to the servicer.

SPN applies a multi-scale, multi-task CNN with a shared featured encoder for this task. First, SPN crops the ROI around the target, using either the state estimate from SEDF (if available) or a separate CNN for object detection. Then, three prediction heads provide predictions of the following information from the input image:

- 1) The target bounding box, with direct regression of target translation and orientation vectors.
- 2) A set of k heatmaps, whose peaks are associated with the 2D projected locations of a set of k pre-designated target keypoints.
- 3) Pixel-wise binary classification of the spacecraft foreground.

The heatmap measurements, which include the 2D locations and dispersions of the target’s features, are provided to the SEDF module for target pose estimation.

SPN, in concert with SEDF, expects a sample rate of approximately 5-10 seconds between images. It assumes

that only one resolved target is present in imagery. Targets may be cooperative or noncooperative but their 3D model must be known a-priori in order to train the CNN. In representative cases, <10 cm of position error and <1 degree of orientation error is achieved, even when SPN is trained exclusively using synthetically rendered images and is tested on substantially different hardware-in-the-loop (HIL) imagery.

The next-generation SPNv3, as developed for the RPO kit, is designed to work over a larger range of ISD in an effort to support a seamless transition from far-range to close-range vision-based navigation. It also proposes several improvements to support robust operation in flight, such as on-line retraining using on-orbit imagery [25], as well as substantial efficiency improvements to support operation on computationally-constrained spacecraft processors. Figure 3 presents a schematic of SPNv3 in which the CNN operates on camera images \mathcal{I}_t to generate heatmaps $\hat{\mathcal{H}}_t$; heatmaps are fed as inputs into a filter; and the filter is used to create ‘pseudo-labelled’ images $\mathcal{H}_t^{\text{pl}}$ for on-line retraining.

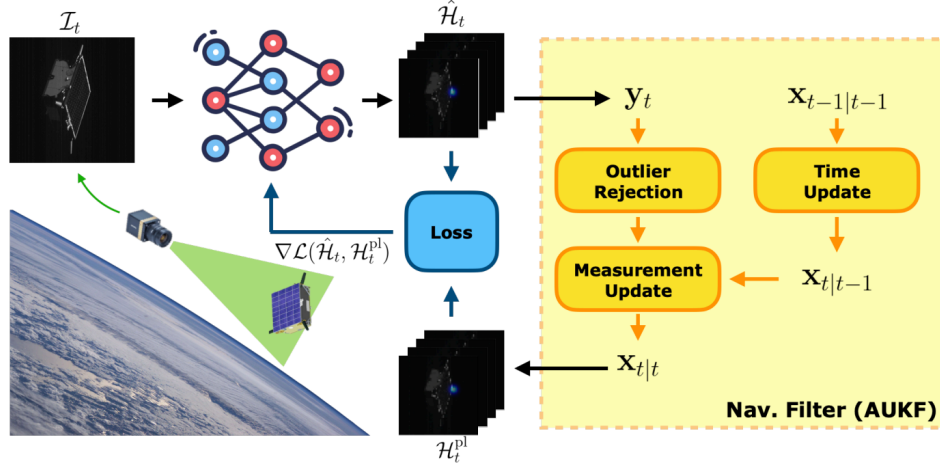


Fig. 3 A schematic of SPNv3 operating with on-line retraining.

3. Integer Ambiguity Resolution

Highly precise relative navigation can be achieved by exploiting the powerful range-error cancelling techniques of CDGNSS on range measurements between two communicating and cooperative spacecraft. Key to this process is to the estimation of the ambiguous number of differenced carrier wave cycles occurring between a GNSS transmitter and receiver antenna. If resolved into their true integer values by the IAR algorithm, a differential measurement with very low thermal noise is produced [17].

IAR operates in concert with the CDGNSS measurement updates performed in the main SEDF filter. If differential GNSS measurements are available, SEDF initially estimates the numbers of carrier wave cycles as floats. The floats are provided as inputs to IAR, which leverages the modified Least Squares Ambiguity Decorrelation Adjustment (mLAMBDA) algorithm [26]. Incoming single-differenced carrier phase (SDCP) measurements are double-differenced and decorrelated and an integer weighted least squares search is used to solve for the exact integer outputs. Statistical acceptance tests ensure robustness of selected integers. The presence of additional metrologies (such as angles-only measurements) in the GNC architecture also offers the opportunity to apply sensor coupling to the IAR process, with benefits for its overall robustness and efficiency [27].

GNSS measurements are expected to be processed by the architecture with an approximate frequency of 10-20 seconds. If double-differenced carrier phase (DDCP) measurements are available and IAR is active, sub-centimeter relative position accuracy can be achieved. In-orbit CDGNSS with IAR will be validated during the upcoming VISORS formation-flying mission [11].

4. Batch Estimation and Data Fusion

BEDF performs angles-only orbit determination at far range, in which batches of target bearing angles from AOT (plus their corresponding time tags and sensor attitudes) are used to produce absolute and relative orbit estimates and associated covariances. If the target state is not known a-priori, this relative state estimate may be used to initialize a

navigation filter. Alternately, it may be employed as a consistency check for the current SEDF estimate.

In single-observer systems, BEDF applies a sampling approach [18] that is made possible by use of an ROE state representation. Target range is weakly observable with bearing angles but is primarily contained within a single ROE parameter: for example, $\delta\lambda$ in the set of quasi-nonsingular ROE [28]. Consequently, BEDF generates set of $\delta\lambda$ samples across the expected ISD interval. For each sample, the remaining ROE are solved for via iterative batch least squares refinement, in which the residuals between measured angles and modeled angles (using the current ROE estimate) are minimized. In multi-observer systems, BEDF may instead solve for $\delta\lambda$ directly (without sampling) thanks to improved observability and simultaneously refine the servicer’s absolute orbit estimate if it is coarse. BEDF also produces an associated state covariance estimate using the converged measurement residuals.

BEDF performs best when provided with batches of 50-200 bearing angles across a period of 1-2 orbits. For nominal cases in LEO, initialization errors are less than 20% of target range when using a single observer. Single-observer angles-only BEDF is currently being validated as part of the NASA Starling mission [10].

5. Sequential Estimation and Data Fusion

The SEDF algorithm uses new measurements to continually refine the system state estimate. The algorithm takes in local bearing angles, local GNSS data, local pose outputs from the CNN, crosslink bearing angles and crosslink GNSS data, and seamlessly fuses these sources to update onboard state estimates, in concert with a high-fidelity orbit propagator. The SEDF algorithm accomplishes this via an adaptive and efficient UKF [19]. A UKF is chosen due to its fully nonlinear dynamics and measurement models that allow higher-order information to be retained, enhancing state observability.

Several additional features are included to maximize performance. First, adaptive process noise estimation techniques such as covariance matching [29] and adaptive state noise compensation [30] are used to improve convergence speed and robustness to errors in the dynamics model. Second, the state definition is organized in a way that exploits the structure of the Cholesky factorization to reduce the number of calls to the orbit propagator by almost a factor of two [31], increasing computational efficiency. Third, received measurements from remote observers are matched to locally-tracked targets using data association algorithms, for improved multi-observer estimation performance [10]. Finally, optional estimation of auxiliary states such as differential ballistic coefficients and inter-satellite clock offsets within the filter contributes to long-term robustness and autonomy in domains where external knowledge of these parameters is stale or coarse.

SEDF measurement updates operate at typical sample rates of 60-120 seconds for far range angles-only navigation, 20-30 seconds for GNSS measurements, and 5-10 seconds for close-range pose estimation. Representative steady-state estimation accuracies are on the order of 2% of target range (angles-only); sub-centimeter position knowledge (differential GNSS); and sub-decimeter position error and sub-degree orientation error (target pose), depending on the scenario and quality of the sensor. It is assumed that multiple targets may be involved in angles-only or differential GNSS navigation but only one target is involved in pose estimation via SPN.

B. Relative Orbit Control

The control module takes a desired set of ROE and the current system state estimate and computes a ΔV sequence to achieve the desired ROE. Inputs are typically provided in the form of waypoints, consisting of a starting epoch, desired ROE with respect to the target, allowable control error, and time horizon. In the Tier 1 system, waypoints are pre-planned from the ground.

A variety of control solvers are applied depending on the current scenario. Each attempts to solve a fuel-optimal (minimum Δv) two-point boundary value problem

$$\underset{\Delta v(t)}{\text{minimize}} \quad \int_{t_e}^{t_f} \|\Delta v(t)\|_2 dt \quad (2)$$

$$\text{subject to} \quad \Delta\delta\alpha = \delta\alpha(t_f) - \Phi(t_e, t_f)\delta\alpha(t_e) = \int_{t_e}^{t_f} \Gamma(t)\Delta v(t) dt \quad (3)$$

where $\delta\alpha$ is the ROE state of the maneuvering spacecraft with respect to the passive spacecraft, t_e is the time of the state estimate, and t_f is the final time, such that $\delta\alpha(t_f) \equiv$ guidance and $\delta\alpha(t_e) \equiv$ state estimate. $\Delta v(t) \in \mathbb{R}^3$ is a velocity change applied by the maneuvering spacecraft expressed in the RTN frame of the passive spacecraft and Φ is the analytic dynamics constraint. Solvers act to minimize the p -norm of the time-varying cost function. In Equation 3, to minimize fuel consumption, the minimum 2-norm of the applied Δv is sought.

Table 2 Summary of constituent control algorithms.

Algorithm	Closed-Form Control Solver [32, 33]	SOCP/QP Control Solver [34]	Least Squares Control Solver [11]	Single-Impulse Col. Avoidance[28, 35]	Multi-Impulse Col. Avoidance [36]
Objective	Compute impulsive control solutions for formation maintenance and reconfiguration	Compute optimized impulsive control solutions for approach and rendezvous	Compute time-fixed impulsive control solutions for approach and rendezvous	Compute impulsive control solutions for collision avoidance	Compute impulsive control solutions for collision avoidance
Relevant Techniques	Reachable set theory	Reachable set theory Convex optimization	Least squares Tikhonov regularization	Analytical solutions Passive safety	Constrained SCP Passive safety
Maneuver Type	3 T, 1-2 N	4-6 RTN	2-3 RTN	1 RTN	1+ RTN
Maneuver Frequency	0.5 orbits	Optimized	Time-fixed	-	-
Planning Horizon	>1.5 orbits	<1.5 orbits	<1.5 orbits	ASAP	ASAP
Maneuver Plan Update Frequency	2 orbits	60 sec	60 sec	-	-
Nominal Accuracy (1σ)	1-10 m	1-10 cm	1-10 cm	1-10 m	1-10 m
Applicability	Large inter-satellite sep. Station-keeping and reconfiguration man.	Small inter-satellite sep. Trajectory tracking	Small inter-satellite sep. Trajectory tracking	Escape maneuvers	Escape maneuvers
Flight Heritage	PRISMA TanDEM-X VISORS	VISORS	VISORS	VISORS	-
System Tier	1	1	1	1	2

1. Closed-Form Solver

Optimal maneuver locations and magnitudes may be computed analytically in closed form [32] by leveraging an ROE state. Triplets of in-plane maneuvers and one or two out-of-plane maneuvers are implemented every $n/2$ orbits to achieve meter-level control accuracy over horizons of ≥ 1.5 orbits. This algorithm is used for efficient station-keeping or for transfers between waypoints at far range. It has partial flight heritage from the PRISMA [5], TanDEM-X [9] and BIROS [37] missions.

2. Convex Optimization and Reachable Set Theory Solver

Reachable set theory may be used to compute optimal maneuver locations and magnitudes by solving a pair of convex optimization problems [33, 34]. A second-order cone program (SOCP) is formalized to solve for optimal times, whereas a quadratic program (QP) is formalized to solve for optimal magnitudes. A key benefit of this optimization-based algorithm is that it is able to solve for optimal times directly through reachable set theory, without requiring a fine discretization of the control window, thus becoming much more computationally efficient. Furthermore, optimal maneuver sequences of (typically) 4-6 maneuvers can be found over horizons down to a fraction of the orbit, allowing fine tracking of an approach trajectory over a progressively shrinking control horizon.

The algorithm is able to achieve centimeter-level accuracy in closed-loop as part of a stochastic model predictive control (MPC) pipeline. At each new state estimate, relative motion is propagated on-board, including planned maneuvers, in open-loop. The plan is updated and the loop is closed whenever the difference between the predicted and desired ROE state strays above a user-specified threshold, with nominal loop closure times as per Table 2. This algorithm will be validated in flight during the VISORS mission [11].

3. Regularized Least Squares Solver

Reduced controllability may be observed in terminal phases of a rendezvous as the control horizon continually shrinks. To ensure terminal satisfaction of constraints, it is effective to compute a 2-3 impulse sequence at fixed times for final trajectory updates. This is achieved by solving a Tikhonov-regularized least squares (LSQ) problem, which directly solves for Δv components at specified time instants while penalizing fuel consumption [11]. This algorithm may also be applied if the SOCP/QP-based solver fails to converge. It will be validated in orbit during the VISORS mission [11].

4. State Machine

A state machine within the control module is used to guide and manage the current control objectives of the system and the resulting actions of the control module. It consists of a set of states, transition conditions, and actions, each of which may be adapted for the scenario of interest to provide appropriate decision-making on board. For example, a typical RPO scenario may consist of a series of ROE waypoints that describe an approach towards a target from far range to docking. The state machine implements logic for autonomous transfers between waypoints and station-keeping at waypoints.

Whenever a specific waypoint is prescribed, control propagates the current state estimate to the control horizon, determines whether the current maneuver plan needs to be updated, and computes new maneuvers if so, using the appropriate solver. The module includes a buffer to spread and store planned maneuvers which allows the system to meet maneuver quantization constraints imposed by actuation, or alter previously computed maneuvers for closed-loop corrections. At any time, if commanded by the safety module, the control module can plan one or more escape maneuvers.

5. Safety and Collision Detection

Passive safety conditions are frequently checked to mitigate risks of collisions in case of critical subsystem failure. Every time a new state estimate is available, the system propagates the state up to a specified time horizon (including and excluding planned maneuvers) and computes minimum RTN separations between the servicer and targets. An escape is triggered if the predicted minimum separation is lower than a specified threshold with a specified confidence level.

Safety must also be maintained at the guidance level such that any provided waypoints describe passively safe geometry and maintain eccentricity/inclination vector separation between spacecraft [38]. This delegation of safety to both guidance (in the sense of pre-planned safety) and the collision detection system (to guarantee safety at every instant of the transfer between waypoints) allows for more efficient, deterministic maneuver plans which maintain safe operations as the system proceeds between control objectives.

6. Single-Impulse Collision Avoidance Solver

In the case of single-target safety violations, a quickly executable escape maneuver is computed to enforce a minimum RN separation between spacecraft while increasing along-track separation. The escape is computed as a single impulse that achieves a minimum RN separation as per analytical conditions on the ROE [35] and an along-track drift rate prescribed by a tuned relative semi-major axis. Finding appropriate ROE is inherently a guidance problem which can be solved exactly in closed form via Gauss' Variational Equations (GVE) [28]. The time of the maneuver is selected to be as soon as possible after the contingency is detected while permitting time for the execution of buffered maneuvers. The algorithm will be validated in orbit during the VISORS mission.

7. Multi-Impulse Collision Avoidance Solver

During more complex collision avoidance scenarios, a constrained sequential convex program (SCP) is used to generate multiple-impulse solutions which remain passively safe for multiple spacecraft in the system [36]. Typically, safety may be enforced in this type of control problem via point-evaluated constraints on system trajectories. The proposed algorithm instead replaces these constraints with single functions of the integration constants of the equations of motion (e.g. the ROE) at the contingency instants. Uncertainty effects are accounted for through variation of parameters. This reduces the number of constraints by one polynomial degree in the number of discrete time samples, allowing computationally efficient enforcement of passive safety within a multi-agent optimal control problem, solvable with direct methods.

C. Algorithm Management

The management module provides on-board decision-making regarding which algorithms shall be active at the current time. For example, consider the handoff between far-range optical navigation (supported by AOT) and close-range optical navigation (supported by SPN). Navigation in the mid-range region must remain robust, even when both algorithms are operating on semi-resolved images in non-optimal conditions (see Figure 1). Similarly, multiple control solvers may be suitable for a specific guidance problem and the system must choose which to apply. The management logic is outlined below.

1. Navigation

The RPO kit hardware provides two cameras for optical navigation with respect to the target: 1) a far-range visible wavelength camera with narrower FOV, for imaging at large ISD, and 2) a close-range LWIR camera with wider FOV, for imaging at small ISD. The integration time of each camera is independently modifiable.

Figure 4 indicates the conditions for transitioning between five distinct navigation modes during a target approach. Each mode uses images from one or both cameras and sends the images to AOT and/or SPN. Examples of applicable ISD for each mode are given, noting that such values will vary depending on camera FOV, target dimensions, and so on. The example values assume an approximate 10° FOV for Camera 1, a 35° FOV for Camera 2, and target dimensions of ~ 1 m. t_i refers to integration time of the camera image.

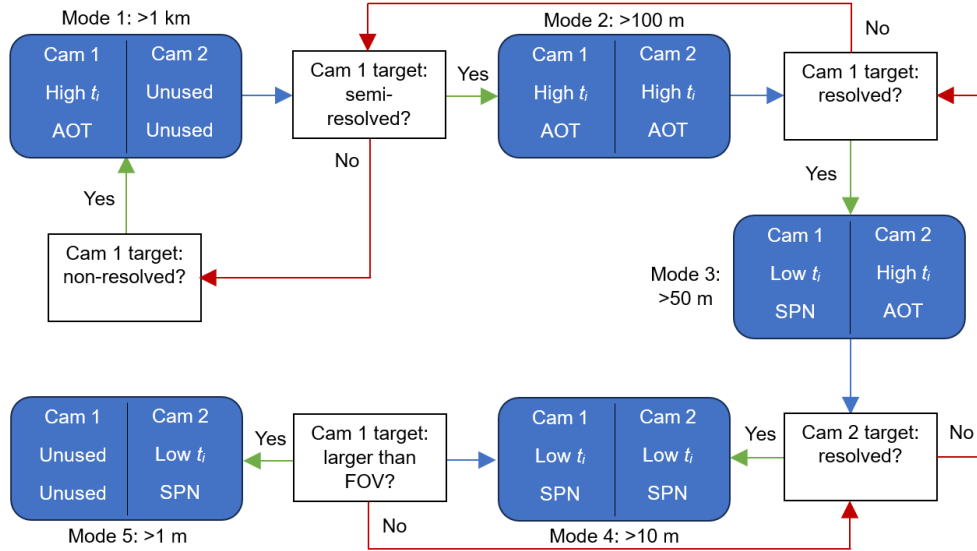


Fig. 4 Transitions between modes for optical navigation during a target approach.

In Mode 1, at far range, AOT is performed using Camera 1 only. Targets are assumed to be Gaussian point spread functions (PSF) and may be completely unresolvable with Camera 2. When the target becomes semi-resolved in Camera 1, it is no longer a Gaussian PSF. Larger pixel centroiding errors and corresponding bearing angle errors are subsequently encountered. Mode 2 allows Camera 2 to potentially assist by providing AOT images in which the target remains primarily non-resolved. Mode 3, which first engages SPN, is applied once the target becomes resolved in Camera 1. AOT may be continued with Camera 2 to provide added robustness. Mode 4 commences once the target is resolved in both camera images and is driven entirely by SPN. Eventually, the target becomes too large to be contained within the FOV of Camera 1 and only Camera 2 is used. The sequence may also be performed in reverse during egress from the target. The 3D model of the target is assumed known in Tier 1 such that the system may autonomously determine whether the target is non-resolved, semi-resolved or resolved with a reasonable degree of certainty using available measurements, state estimates, and covariances. Otherwise, mode switches can be managed from the ground.

2. Control

Figure 5 presents selection of an appropriate control solver during an approach towards a waypoint. The solver is nominally selected based on the current horizon, with adjustments to parameters based on the current quality of the target relative position estimate. t_{sk} and t_{wp} refer to the station-keeping horizon and waypoint time respectively.

The system commences tracking when a waypoint is commanded. Initially, the closed-form solver is applied. The closed-form solver is most suitable when horizons are long, when very high accuracy is not yet necessary, and when computational efficiency is paramount. The control loop is closed less regularly (e.g. every two orbits) and maneuvers may be re-planned if predicted ROE control errors are larger than a user-specified bounds. Bounds may be specified as either a fixed ROE error threshold, or as a Mahalanobis distance [23] threshold proportional to the current ROE uncertainty from navigation, if the estimate is coarse. When the horizon reaches 1.5 orbits or less, the SOCP/QP-based solver is engaged. The control loop is closed much faster (e.g. every few minutes) to allow finer tracking of the waypoint

approach, with more varied maneuver solutions that focus on achieving high accuracy. When the horizon reaches 0.1 orbits or less, and if ROE estimation errors are small, the least squares-based solver is engaged. This solver generates a final series of fixed-time maneuvers with which to fine-tune the rendezvous and minimize final errors.

Station-keeping commences after the waypoint is reached. It may employ either the closed-form or SOCP/QP algorithm, depending on current navigation accuracy and the desired error bounds, and applies a re-planning horizon of several orbits.

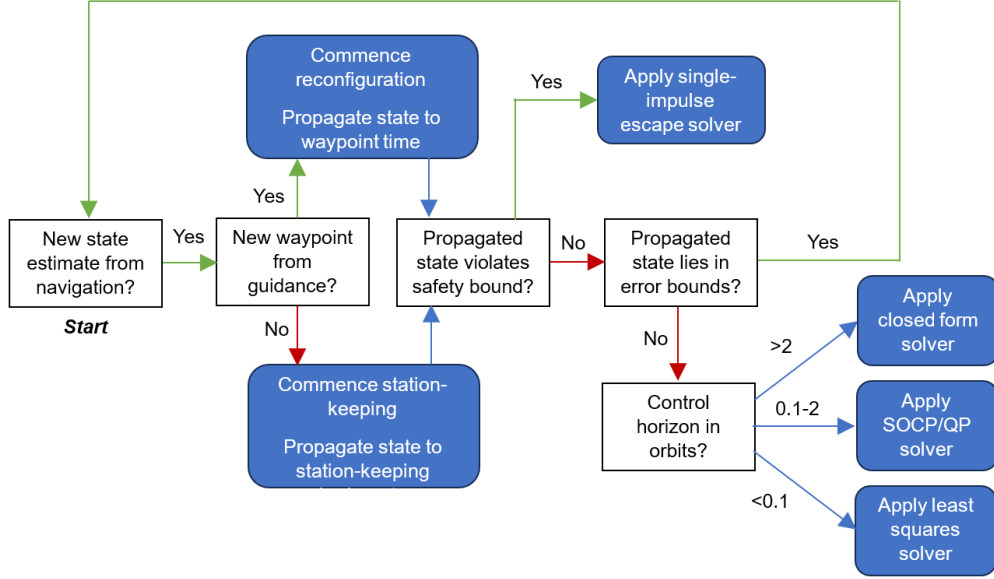


Fig. 5 Selection logic for control solvers during waypoint following.

3. State Parametrization

The system employs several state parametrizations in concert to take advantage of specific benefits of each. Applied parametrizations are listed in Table 3, noting that the choice of appropriate auxiliary states is flexible and dependent on the domain.

In the navigation module, the primary orbit state is the Earth-Centered inertial (ECI) position & velocity of each system member. This provides simplicity, avoids singularities, and is especially suited to computing measurement updates in SEDF. It also remains valid for arbitrary inter-satellite separations. For dynamics updates and orbit propagation in SEDF, the equinoctial absolute orbit elements (OE) are employed. OE enable more efficient and accurate propagation via numerical integration of GVE [39], and nonsingular OE are necessary for operating in geosynchronous orbits.

In the control module, the primary orbit states are servicer OE ($\alpha \in \mathbb{R}^6$) and target ROE ($\delta\alpha \in \mathbb{R}^6$). The ROE are especially useful for control as they allow the derivation of optimal analytic solutions as well as computationally-efficient framing of the optimal control problem. Control outputs are provided as RTN Δv sequences for the servicer. Servicer OE and target ROE are similarly used for providing inputs and outputs (I/O) to the architecture. The ROE are preferred as they allow simple visualization of system geometry, passive safety, and relative motion over time [28].

At close range, when estimating target attitude, navigation applies a modified Rodrigues parameter (MRP) state [20]. A reference quaternion vector $\mathbf{q} \in \mathbb{R}^4$ tracks relative orientation of the target, and the MRP vector $\mathbf{p} \in \mathbb{R}^3$ tracks step changes in relative orientation of the target with respect to \mathbf{q} . \mathbf{q} is updated by applying \mathbf{p} via the multiplicative approach of the Unscented Quaternion Estimator [40]. A vector $\boldsymbol{\omega} \in \mathbb{R}^3$ tracks relative angular velocity with respect to the target's principal axes.

V. Simulation Results

The intermediate Tier 1 architecture is validated across three simulation scenarios. Each exercises a different domain of the system: 1) angles-only navigation and formation control at far range, 2) CDGNSS navigation and formation

Table 3 Summary of applied state parametrizations.

Estimated State	Parametrization	Application	Auxiliary State	Purpose
Servicer orbit	ECI position/velocity	Measurement modeling	Empirical accelerations	Robustness to dynamics errors
	Equinoctial OE	Orbit propagation and architecture I/O	Differential clock offsets and drift rates	Long-term synchronization of multi-observer measurements
Target orbit	ECI position/velocity	Measurement modeling	Ballistic coefficients	Robustness to modeling errors
	Equinoctial OE	Orbit propagation	Sensor biases	Robustness to calibration errors
	Nonsingular ROE	Control solutions and architecture I/O	DDCP float/integer ambiguities	Precise relative positioning with GNSS
	RTN position/velocity	Collision detection		
Target attitude	Modified Rodrigues parameters and relative angular velocity	Pose estimation		

control at close range, and 3) handling of semi-resolved imagery at mid-range by AOT. Verification of the fully-integrated system, including multiple servicers and targets and SPN-assisted navigation at close range, is left for Tier 2 work.

A. Scenarios

1. Far-Range Approach (Scenario #1)

The far range approach consists of a servicer approaching a noncooperative target in LEO following a series of fixed waypoints, presented in Table 4. The servicer begins by drifting towards the target, gradually slowing its rate of drift while reducing the inter-satellite along-track separation to zero (Waypoints 1-5). The servicer then transitions to a passive safety ellipse (Waypoint 5); then into a position to potentially begin a final approach (Waypoint 6); then retreats to another passive safety ellipse (Waypoint 7). The final waypoints re-introduce a drift to depart from the target. The simulated trajectory of the target in the RTN frame of the servicer is plotted in Figure 6. The servicer has access to camera images for target relative orbit determination, as well as PVT solutions from its GNSS receiver for servicer absolute orbit determination.

2. Close-Range Approach (Scenario #2)

Table 5 presents servicer OE and target ROE over a 36-hour approach period. The servicer is initialized in a ‘standby’ relative orbit, with a passive safety margin of 200 m by virtue of the eccentricity-inclination vector separation between itself and the target. It then enters an ‘approach’ phase to achieve a ~45 m inter-satellite separation after 12 hours, while maintaining passive safety. The final waypoint places it ~20 m from the target after 24 hours, in position to begin a final v -bar rendezvous. The servicer maintains at this inspection point until 36 hours have passed. The target’s trajectory in the servicer’s RTN frame is plotted in Figure 7. The target is cooperative in this scenario and is sharing its own received GNSS measurements, GNSS broadcast orbits, clock information and attitude information with the servicer.

3. Semi-Resolved Approach (Scenario #3)

Table 6 presents servicer OE and target ROE for the semi-resolved test. Semi-resolved imagery occurs in the intermediate domain where a target is no longer a Gaussian point source in an image but specific spacecraft features can not yet be resolved (see Figure 1). This scenario is not a continuous simulation with control-in-the-loop; instead, the system is simulated at each waypoint for a 24-hour period, to investigate the ability of the AOT + SEDF pipeline to support a robust handoff to SPN when operating in the challenging semi-resolved regime.

B. Data Generation

Data generation for test cases employs high-fidelity dynamics and measurement models. Truth trajectories with control in the loop are generated using SLAB’s high-fidelity S^3 astrodynamics library. Dynamics models used for ground truth simulation and in the flight software are specified in Table 7. Spacecraft are physically modeled with a mass of 200 kg, drag coefficient of 2.2, and constant cross-sectional area of 2 m² (servicer) and 4 m² (target). The target visual model used for Scenario 3 is that of the PRISMA mission’s Tango spacecraft [5]. Target attitude is random in

Table 4 Far range scenario absolute and relative orbits.

Servicer OE	Time (hr)	a (km)	e	i ($^\circ$)	Ω ($^\circ$)	ω ($^\circ$)	M ($^\circ$)
1	-	6878	0.001	97.4	0	0	0
Waypoint	Time (hr)	δa (m)	$\delta \lambda$ (m)	δe_x (m)	δe_y (m)	δi_x (m)	δi_y (m)
1	0	-2000	-75000	0	0	0	0
2	6	-87	-50000	0	1000	0	1000
3	18	-70	-25000	0	500	0	500
4	30	-17	-5000	0	500	0	500
5	42	0	0	0	500	0	500
6	54	0	0	0	500	0	0
7	66	0	0	0	500	0	-500
8	78	159	0	0	1000	0	1000
9	90	0	-75000	0	1000	0	1000

Table 5 Close range scenario absolute and relative orbits.

Servicer OE	Time (hr)	a (km)	e	i ($^\circ$)	Ω ($^\circ$)	ω ($^\circ$)	M ($^\circ$)
1	-	6878	0.001	97.4	0	0	0
Waypoint	Time (hr)	δa (m)	$\delta \lambda$ (m)	δe_x (m)	δe_y (m)	δi_x (m)	δi_y (m)
1	0	0	0	0	200	0	200
2	12	-2.62	45.21	-34.51	4.78	-18.72	2.72
3	24	0	20	0	0	0	0

Table 6 Semi-resolved scenario absolute and relative orbits.

Servicer OE	a (km)	e	i ($^\circ$)	Ω ($^\circ$)	ω ($^\circ$)	M ($^\circ$)
1	6878	0.001	97.4	0	0	0
Waypoint	δa (m)	$\delta \lambda$ (m)	δe_x (m)	δe_y (m)	δi_x (m)	δi_y (m)
1	0	1000	0	30	0	30
2	0	600	0	18	0	18
3	0	200	0	6	0	6
4	0	100	0	3	0	3
5	0	50	0	1.5	0	1.5

each image, whereas the servicer’s attitude consistently points the camera boresight at the target’s center of mass.

Five synthetic measurement types are generated from the ground truth: camera attitude quaternions, raw GNSS pseudorange and carrier phase messages, processed position/velocity/time (PVT) solutions from the GNSS receiver, a-priori system orbit estimates, and camera images. The a-prioris are used for filter initialization. Gaussian zero-mean measurement noise and Gaussian errors in executed maneuvers are included as specified in Table 8, where \mathcal{I} , \mathcal{R} , \mathcal{C} refer to the inertial, RTN, and camera reference frames respectively.

For modelling GNSS, GPS-only and L1-only signals are emulated. The ionospheric error is modelled via the Klobuchar model with code phase delay equal and opposite to carrier phase advance. Receiver clock errors are modelled as a random walk with a step size of 1m per step, with step cadence every second. GPS satellite ephemerides are corrupted in the emulated broadcast messages by injecting undesired ROEs. This introduces realistic periodic variations of the Cartesian state error of the GPS satellites, as predicted by the mapping between ROE space and the Cartesian inertial coordinates. Thermal noise for pseudorange and carrier phase are simulated using an elevation-dependent carrier-to-noise (C/N_0) model [27, 28], based on the antenna gain pattern of an actual patch antenna.

Camera images are generated using 3D vector graphics in OpenGL via SLAB’s Optical Stimulator software [13], which provides physically realistic visual modeling of both non-resolved point sources and resolved objects. In addition to the noise sources in Table 8, a small amount of glare and blur are added to the target 3D model to emulate camera




Fig. 6 Trajectory of the target in the servicer’s RTN frame for the far range scenario, with waypoints numbered.

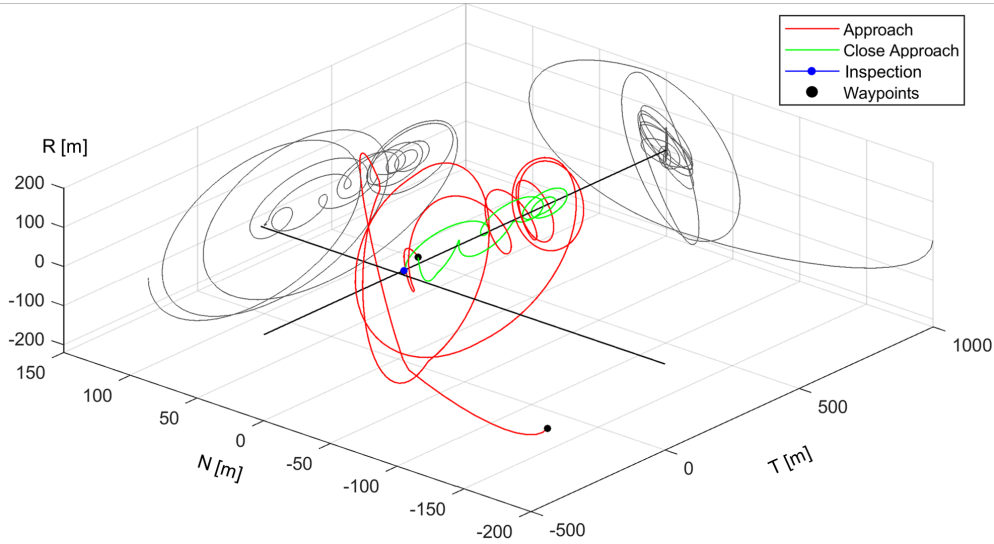


Fig. 7 Trajectory of the target in the servicer’s RTN frame for the close range scenario.

artifacts at close range. Optical measurement gaps caused by eclipse periods and sun-blinding of the camera are modeled. The simulated camera applies the physical properties of a Blue Canyon Technologies Nano Star Tracker (NST). The camera has a resolution of 1280×1024 pixels and an approximate FOV of 12×10 degrees.

C. Results

1. Far-Range Approach (Scenario #1)

The far-range scenario employs angles-only optical navigation. Control and navigation performance are plotted in Figures 8-10 with numerical values provided in Table 9. Scenario waypoints follow Table 4.

First, examine relative navigation errors in Figure 8. $\delta\lambda$ is the most weakly-observable ROE with angles-only measurements, and sees correspondingly larger uncertainties. Uncertainty in $\delta\lambda$ is also proportional to target range, because the same uncertainty in bearing angles corresponds to a larger geometric uncertainty at larger ISD. Thus, uncertainties in $\delta\lambda$ are largest at the start and end of the waypoint sequence. Estimation of $\delta\lambda$ is nonetheless robust. Errors are quickly reduced from large a-priori uncertainties of 500 m, and are kept to < 20 m when the servicer begins its close-range rendezvous in Figure 9. $\delta\lambda$ uncertainties increase as the servicer departs and the target grows more distant.

Estimation performance is also robust for the other ROE and uncertainties are consistently less than 10 m post-

Table 7 Dynamics models in simulations.

Model	Perturbations	Propagation	Context
On-board analytic [41, 42]	J_2 gravity	STM Euler integrator 5-30 sec timestep	Control solvers, AOT, BEDF
On-board numeric	20x20 GGM01S gravity model [43] Harris-Priester atmosphere with cannonball drag model	GVE RK4 integrator and Cartesian RK4 5-30 sec timestep	Collision avoidance, SEDF
Ground truth	60x60 GGM01S gravity model NRMLSISE00 atmosphere [44] with cannonball drag model SRP with conical Earth shadow and cannonball drag model Third-body lunisolar gravity from analytic ephemeris	GVE RK4 integrator 1-10 sec timestep	Simulation truth

Table 8 Gaussian measurement and control noise in simulations.

Measurement	Noise (1σ)	Axes	Sample Time
A-priori servicer orbit	[10, 10, 10, 10, 10, 50] m	$[a, ae_x, ae_y, ai, a\Omega, au]$	Single epoch
A-priori target orbit	[50, 500, 50, 50, 50, 50] m	$[a\delta a, a\delta\lambda, a\delta e_x, a\delta e_y, a\delta i_x, a\delta i_y]$	Single epoch
GNSS PVT position	[10, 10, 10] m	$[\hat{x}^I, \hat{y}^I, \hat{z}^I]$	20 sec
GNSS PVT velocity	[0.01, 0.01, 0.01] m/s	$[\dot{\hat{x}}^I, \dot{\hat{y}}^I, \dot{\hat{z}}^I]$	20 sec
Image attitude	[5, 5, 20]''	$[\hat{x}^C, \hat{y}^C, \hat{z}^C]$	60 sec
Image PSF position	[5, 5]''	$[\alpha, \epsilon]$	60 sec
Image pixel brightness	1.5%	-	60 sec
Maneuver magnitude	5%	-	-
Maneuver direction	[5, 5, 5] $^\circ$	$[\hat{x}^R, \hat{y}^R, \hat{z}^R]$	-

initialization. Spikes in error and uncertainty are concurrent with larger maneuvers, due to the effects of maneuver execution errors. However, maneuvers typically act to improve angles-only observability which results in very fast recovery of nominal estimation performance. There are also more gradual, temporary uncertainty increases such as that for δe_y between Waypoints 5-6. This is because the relative inclination vector between the servicer and target is almost zero during this period, which results in very little out of plane motion and correspondingly poorer angles-only observability. Finally, note that slight periodicity is evident due to the presence of regular eclipse periods during which angles-only measurements are unavailable.

Table 9 Control and navigation performance at the instance of each waypoint in the far range scenario. Values are expressed as the 2-norm of the associated six-dimensional ROE vector.

Waypoint	Far Range									Close Range	
	2	3	4	5	6	7	8	9	2	3	
ROE Estimation Error (m)	34.8	12.7	2.9	11.7	2.6	9.3	1.5	200.4	0.29	0.14	
ROE Est. Uncertainty ($1-\sigma$) (m)	130.0	73.7	17.2	8.2	11.5	8.7	7.8	215.3	0.60	0.57	
ROE Control Error (m)	35.9	12.2	2.8	10.7	8.1	11.0	4.6	198.2	0.31	0.25	

Figure 9 presents the trajectory of the servicer in ROE space versus the provided waypoint set. The guidance and control system is able to reach each waypoint at the desired time with the desired accuracy by using the closed-loop replanning logic described in Figure 5. Farther from each waypoint the control loop is closed more slowly and the closed-form solver is used; whereas closer to each waypoint, the control loop is closed more quickly and the SOCP/QP-based controller is used. This re-planning logic is suitably robust drives control errors down to the accuracy level of the navigation system. This is shown in Table 9, where control errors are consistently very similar to angles-only relative navigation errors and ~ 10 m accuracy is achieved for Waypoints 4-8. A further example is given by the δi_x

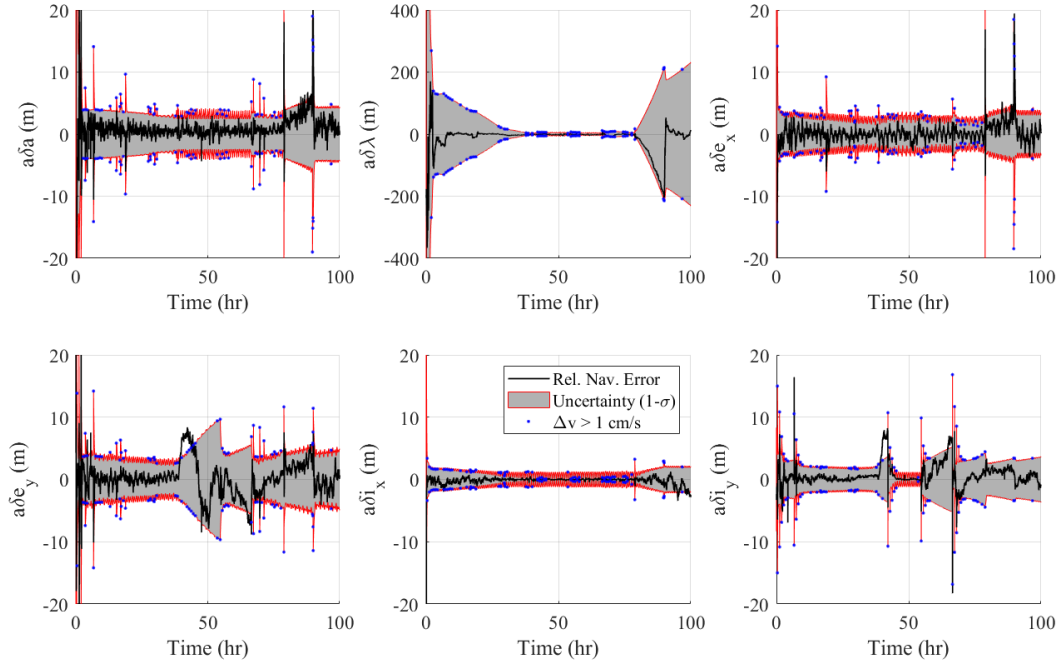


Fig. 8 ROE estimation errors and formal $1-\sigma$ covariances throughout the far range scenario. Larger maneuvers are marked with blue dots.

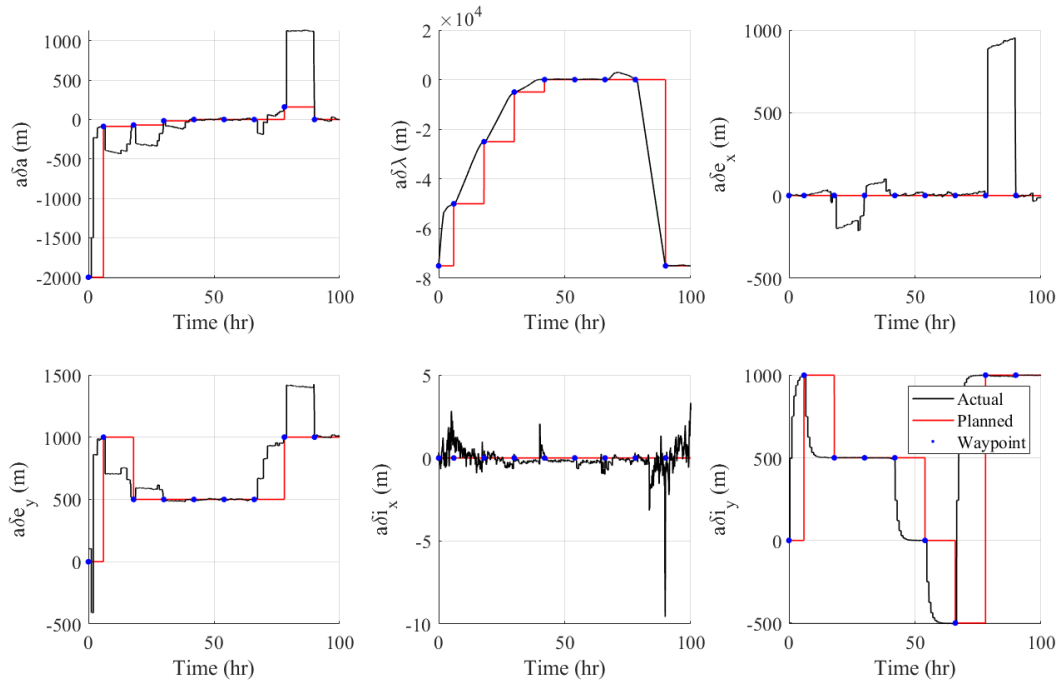


Fig. 9 ROE of the target and desired control waypoints throughout the far range scenario.




Fig. 10 Executed maneuver Δv s throughout the far range scenario in the RTN frame of the servicer.

component in Figure 9. Nominally, δi_x should be kept at zero and the slight random variations around this point reflect the accuracy of the underlying state estimate. The large variations of the relative eccentricity vector components at around 80 orbits are expected and caused by the large tangential maneuvers required to increase the relative semi-major axis and thus reduce mean along track separation $\delta \lambda$ over time. These variations are then suitably corrected before the subsequent waypoint. Executed maneuvers are plotted in Figure 10. Total Δv usage is 8.2 m/s, which is close to the theoretical lower bound [32] for this sequence of 7.2 m/s. The maneuver planner is therefore relatively efficient even in the presence of navigation and execution errors.

2. Close-Range Approach (Scenario #2)

The close-range scenario employs CDGNSS with IAR for higher precision navigation. Here, it is assumed that the target is cooperative, time-synced with GPST, and sharing with the servicer its own received GNSS measurements, GNSS broadcast orbits, clock and attitude information. Scenario waypoints follow Table 5.




Fig. 11 Relative position error in RTN.




Fig. 12 Relative velocity error in RTN.

An analysis of relative navigation performance using CDGNSS is provided in the target-centered RTN frame in Figures 11 and 12. The RTN position error norm is 1.2609 cm, while the velocity error norm is 0.1961 mm/s, successfully achieving centimeter-level positioning accuracy. Under GNSS coverage in LEO, radial-axis (R) errors are typically higher than local horizontal (TN plane) errors due to a larger local vertical dilution of precision (DOP) as compared to the local horizontal DOP.

The velocity state is not directly observable through CDGNSS measurements but is updated with high accuracy during the SEDF time update using a Cartesian numerical propagator (as per Table 7), and during the SEDF measurement

update via cross-correlations with other states. As such, the velocity estimate is more susceptible to the effects of maneuvers and maneuver execution errors. During the ‘approach’ and ‘close approach’ segments between Waypoints 2 and 3, higher Δv magnitudes are interspersed with an approximate cadence of half an orbit period. Upon arriving at the Waypoint 3, however, the control planner begins to plan much more frequent, finer maneuvers for station-keeping at the waypoint (as per the maneuver plan in Figure 15). These maneuvers produce a slight though observable degradation in RTN velocity errors in Figure 12.

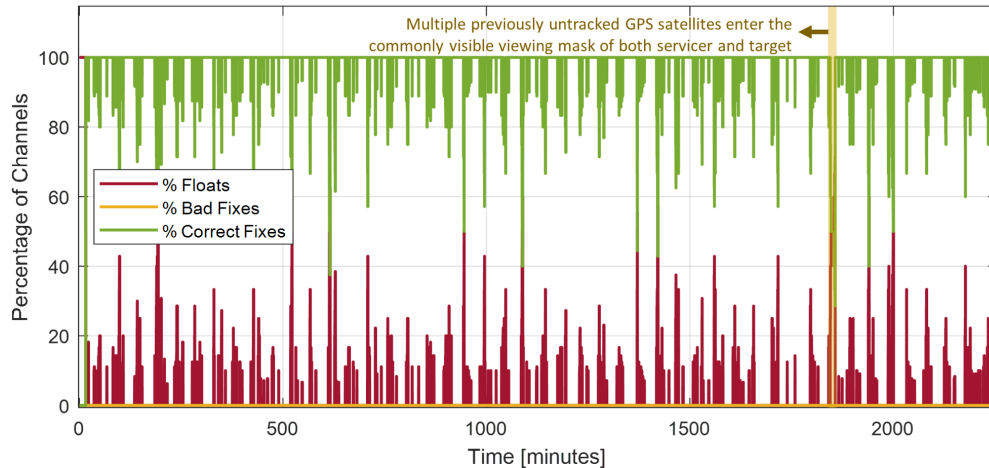


Fig. 13 IAR success rate as a percentage of tracked ambiguities.

IAR successfully resolves its first batch of integers after ~ 20 minutes, as seen in Figure 13, with a 100% success rate. An increase in the percentage of tracked ambiguities treated as floats will occur only when new commonly-visible satellites (and their corresponding new single-difference ambiguities) enter the elevation mask of the receiver. In specific instances this causes a momentary performance degradation, if a sufficient number of new, commonly-visible satellites enters the mask while previously tracked integer-resolved satellites leave the mask. An example of such an isolated event happens at 1840 minutes into the simulation; afterwards, SEDF requires approximately ~ 20 minutes of measurement updates before the float ambiguity estimates converge sufficiently to regain successful IAR.

Conformance of the trajectory to the guidance profile is presented in ROE space in Figure 14 with executed maneuvers is in Figure 15. As with the far-range scenario, the guidance and control system is able to reach each waypoint at the desired time with the desired accuracy by using the closed-loop replanning logic described in Figure 5. Control errors, listed on the right of Table 9, are driven down to decimeter levels which enable the close-range inspection. Effective station-keeping at the final waypoint is demonstrated with small oscillations about the desired ROE. Figure 14 also displays how slower control loop closure when farther away from each waypoint allows the unforced relative orbit dynamics (in this case, perturbed by differential drag) to act, causing large variations in $\delta\lambda$ up to ~ 600 meters. Variations are promptly corrected during the terminal approach to the waypoint. If required by operational constraints, this behavior can be mitigated by closing the control loop more frequently throughout the approach, but leads to a less fuel-optimal trajectory and higher Δv expenditure. Total Δv usage in this simulation is approximately 0.5 m/s.

3. Semi-Resolved Approach (Scenario #3)

Table 10 presents state estimation performance in the form of estimation errors and $1-\sigma$ uncertainties from SEDF, averaged across the final simulated orbit. The table also provides mean AOT measurement errors and SEDF bearing angle noise.

Two challenges are faced by the angles-only system when performing navigation for a semi-resolved target. First, significantly larger Gaussian centroiding errors are produced as target size in the FOV increases, due to its non-Gaussian shape. Angles-only systems are very sensitive to measurement errors which can cause swift divergence of the filter. To compensate, filter measurement noise was adaptively adjusted, as per Table 10, to be approximately 10 times the target’s diameter in pixels in the image. This was found to be a suitable empirical value in initial tests, in absence of a more sophisticated autonomous criteria based on e.g. a-priori knowledge of the target’s 3D model.

Furthermore, the observability of angles-only navigation improves as the nonlinearity of the system increases. More

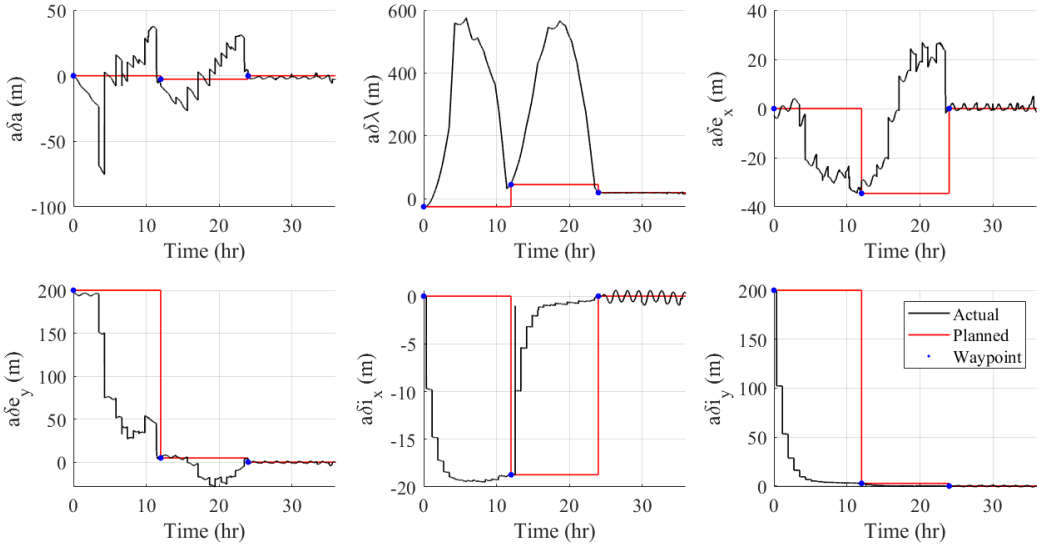


Fig. 14 ROE of the target and desired control waypoints throughout the close range scenario. Station-keeping mode is activated at the final target inspection waypoint.




Fig. 15 Executed maneuver Δv s throughout the close range scenario in the RTN frame of the servicer.

Table 10 Estimation errors and $1\text{-}\sigma$ uncertainties from AOT in the RTN frame of the servicer.

$a\delta\lambda$ (m)	Pos. Estimation Performance (m)			Vel. Estimation Performance (mm/s)			Meas. Error (")	Meas. Noise (")	Target Size (px)
	δr_R	δr_T	δr_N	δv_R	δv_T	δv_N			
1000	0.0 ± 0.9	11.2 ± 43.1	0.0 ± 0.8	0.01 ± 0.88	0.08 ± 1.90	0.02 ± 0.94	40 ± 20	50	5
600	0.0 ± 0.8	16.5 ± 39.7	0.0 ± 0.7	-0.01 ± 0.81	0.04 ± 1.75	-0.03 ± 0.86	62 ± 29	100	10
200	0.0 ± 0.4	14.7 ± 18.8	0.0 ± 0.4	-0.01 ± 0.40	0.04 ± 0.88	0.02 ± 0.41	160 ± 67	250	25
100	0.0 ± 0.3	9.9 ± 10.5	0.0 ± 0.2	0.00 ± 0.25	0.00 ± 0.55	0.02 ± 0.23	323 ± 127	500	50
50	0.0 ± 0.3	22.7 ± 8.5	0.0 ± 0.2	-0.01 ± 0.25	0.06 ± 0.56	0.04 ± 0.19	648 ± 308	1000	100

nonlinear behaviour is encountered when targets are farther apart in more distinct orbits, whereas the system is almost linear when targets are in close proximity. Robust estimation of the $\delta\lambda$ ROE is especially reliant on nonlinearities (or system maneuvers to disambiguate range, which are not present in this scenario). Proportionally larger uncertainties in $\delta\lambda$ or the along-track direction are therefore expected.

Despite these challenges, however, good performance is observed down to $a\delta\lambda = 100$ m. As expected, the largest state uncertainties are observed in the along-track direction, such that along-track position uncertainties are 4% of $\delta\lambda$ for $a\delta\lambda = 1000$ m, growing to 11% of $\delta\lambda$ for $a\delta\lambda = 100$ m (as the system becomes more and more linear). Nevertheless,

state errors remain within $1\text{-}\sigma$ bounds and uncertainties remain small enough to provide safety during an approach. Estimation errors grow outside of covariance bounds at an ISD of ~ 50 m because of very large measurement error outliers; the image processing system occasionally forms several ROI from the single target due to large brightness variations across surfaces and features. However, at such distances, the target is already resolved and SPN is expected to be active. The far-range AOT + SEDF system is therefore expected to be able to support a robust handover to the close-range SPN + SEDF system.

VI. Conclusions and Future Work

This paper describes the design, integration and validation of a comprehensive guidance, navigation and control (GNC) software architecture for end-to-end rendezvous and proximity operations (RPO) in orbit. The architecture applies a modular, adaptive framework in concert with state-of-the-art GNC algorithms with flexible handling of single or multiple servicers, single or multiple targets, and cooperative and non-cooperative targets. GNC algorithms are able to meet the differing accuracy and efficiency requirements of far-range, mid-range and close-range regimes and seamlessly transition between each in a safe and autonomous manner.

Specifically, for navigation with respect to non-cooperative targets, the architecture fuses far range angles-only navigation algorithms with a close range neural-network-based target pose estimation pipeline. For highly-precise navigation with respect to cooperative spacecraft, the architecture adopts carrier phase differential GNSS techniques with integer ambiguity resolution. For control, a suite of efficient optimal control solvers is leveraged, managed by an autonomous waypoint-following guidance system. Control solvers employ closed-form, least squares, reachable set theory and convex optimization techniques to meet a range of control precision requirements. The integrated architecture supports robust transitions from far- to close-range GNC regimes via adaptive sensor fusion, state estimation, and maneuver planning and re-planning. Passive and reactive safety is guaranteed throughout via adherence to eccentricity- and inclination-vector separation, paired with an active collision detection module which plans escape maneuvers whenever necessary.

A prototype version of the GNC software architecture is verified with high-fidelity simulations of several RPO scenarios: a far-range approach to an angles-only target, a mid-range approach to a semi-resolved target, and a close-range approach applying differential GNSS navigation. Results display meter-level navigation and control accuracy in the far range scenario and centimeter-level navigation and control accuracy in the close range scenario, with an overall level of performance and autonomy suitable for supporting a variety of future RPO missions.

Future development will focus on the integration and verification of the proposed vision-based pose estimation pipeline for close-range navigation, the extension of control and navigation algorithms to multi-observer multi-target systems, and hardware-in-the-loop testing of the complete flight architecture.

VII. Acknowledgements

This work has been supported by SpaceWERX Orbital Prime – Direct to Phase II Contract No. FA864923P0560. The authors would also like to acknowledge the contributions of Antonio Rizza and Guillem Rueda Oller towards the RPO kit project and the assistance of Ten One Aerospace when formulating the RPO kit software and scenarios.

References

- [1] Arney, D., Sutherland, R., Mulvaney, J., Steinkoenig, D., Stockdale, C., and Farley, M., “On-orbit Servicing, Assembly, and Manufacturing (OSAM) State of Play,” Tech. rep., NASA, 2021.
- [2] Holzinger, M. J., and Jah, M. K., “Challenges and Potential in Space Domain Awareness,” *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 1, 2018, pp. 15–18. <https://doi.org/10.2514/1.G003483>.
- [3] Mark, C. P., and Kamath, S., “Review of Active Space Debris Removal Methods,” *Space Policy*, Vol. 47, 2019, pp. 194–206. <https://doi.org/10.1016/j.spacepol.2018.12.005>.
- [4] Friend, R. B., “Orbital Express program summary and mission overview,” *Sensors and Systems for Space Applications II*, Vol. 6958, edited by R. T. Howard and P. Motaghedi, International Society for Optics and Photonics, SPIE, 2008, p. 695803. <https://doi.org/10.1117/12.783792>.
- [5] D’Amico, S., Ardaens, J.-S., Gaias, G., Benninghoff, H., Schlepp, B., and Jørgensen, J. L., “Noncooperative Rendezvous Using

- Angles-Only Optical Navigation: System Design and Flight Results,” *Journal of Guidance, Control, and Dynamics*, Vol. 36, No. 6, 2013, pp. 1576–1595.
- [6] Ardaens, J.-S., and Gaias, G., “Angles-Only Relative Orbit determination in Low Earth Orbit,” *Advances in Space Research*, Vol. 61, No. 11, 2018, pp. 2740–2760.
- [7] Roscoe, C. W., Westphal, J. J., and Mosleh, E., “Overview and GNC design of the CubeSat Proximity Operations Demonstration (CPOD) mission,” *Acta Astronautica*, Vol. 153, 2018, pp. 410–421. <https://doi.org/https://doi.org/10.1016/j.actaastro.2018.03.033>.
- [8] Biesbroek, R., Aziz, S., Wolahan, A., Cipolla, S., Richard-Noca, M., and Piguet, L., “The ClearSpace-1 Mission: ESA and ClearSpace Team Up to Remove Debris,” *8th European Conference on Space Debris*, Darmstadt, Germany, 2021.
- [9] Krieger, G., Moreira, A., Fiedler, H., Hajnsek, I., Werner, M., Younis, M., and Zink, M., “TanDEM-X: A Satellite Formation for High-Resolution SAR Interferometry,” *IEEE Transactions on Geoscience and Remote Sensing*, Vol. 45, No. 11, 2007, pp. 3317–3341.
- [10] Kruger, J., Koenig, A. W., and D’Amico, S., “The Starling Formation-Flying Optical Experiment (StarFOX): System Design and Pre-flight Verification,” *Journal of Spacecraft and Rockets*, 2023.
- [11] Guffanti, T., Bell, T., Low, S., Murray-Cooper, M., and D’Amico, S., “Autonomous Guidance Navigation and Control of the VISORS Formation-Flying Mission,” *2023 AAS/AIAA Astrodynamics Specialist Conference*, Big Sky, Montana, 2023.
- [12] Lowe, S., Fitzpatrick, D., Buynovskiy, A., Shoemaker, L., Palo, S., and D’Amico, S., “Concept of Operations for SWARM-EX: a Three CubeSat Formation-Flying Mission,” *IEEE Aerospace Conference*, Big Sky, Montana, 2024 (accepted).
- [13] Beierle, C., and D’Amico, S., “Variable Magnification Optical Stimulator for Training and Validation of Spaceborne Vision-Based Navigation,” *Journal of Spacecraft and Rockets*, Vol. 56, No. 4, 2019, pp. 1060–1072.
- [14] Palo, S., Stafford, G., and Hoskins, A., “An Agile Multi-Use Nano Star Camera for Constellation Applications,” *Proceedings of the Small Satellite Conference, Advanced Technologies II, SSC13-III-5*, Logan, UT, 2013.
- [15] Kruger, J., and D’Amico, S., “Autonomous angles-only multitarget tracking for spacecraft swarms,” *Acta Astronautica*, Vol. 189, 2021, pp. 514–529.
- [16] Park, T. H., and D’Amico, S., “Robust multi-task learning and online refinement for spacecraft pose estimation across domain gap,” *Advances in Space Research*, 2023. <https://doi.org/10.1016/j.asr.2023.03.036>.
- [17] Giraldo, V., and D’Amico, S., “Distributed multi-GNSS timing and localization for nanosatellites,” *NAVIGATION*, Vol. 66, No. 4, 2019, pp. 729–746. <https://doi.org/https://doi.org/10.1002/navi.337>.
- [18] Koenig, A. W., and D’Amico, S., “Observability-Aware Numerical Algorithm for Angles-Only Initial Relative Orbit Determination,” *2020 AAS/AIAA Astrodynamics Specialist Conference*, Lake Tahoe, CA, 2020.
- [19] Sullivan, J., Koenig, A. W., Kruger, J., and D’Amico, S., “Generalized Angles-Only Navigation Architecture for Autonomous Distributed Space Systems,” *Journal of Guidance, Control, and Dynamics*, Vol. 44, No. 6, 2021, pp. 1087–1105.
- [20] Park, T. H., and D’Amico, S., “Adaptive Neural-Network-Based Unscented Kalman Filter for Robust Pose Tracking of Noncooperative Spacecraft,” *Journal of Guidance, Control, and Dynamics*, Vol. 46, No. 9, 2023, pp. 1671–1688. <https://doi.org/10.2514/1.G007387>.
- [21] Mortari, D., Samaan, M. A., Bruccoleri, C., and Junkins, J. L., “The Pyramid Star Identification Technique,” *Navigation*, Vol. 51, No. 3, 2004, pp. 171–183.
- [22] Wertz, J. R., *Spacecraft Attitude Determination and Control*, Springer Science & Business Media, 2012.
- [23] Blackman, S., “Multiple Hypothesis Tracking For Multiple Target Tracking,” *IEEE Aerospace and Electronic Systems Magazine*, Vol. 19, No. 1, 2004, pp. 5–18.
- [24] Kruger, J., and D’Amico, S., “Autonomous Angles-Only Multitarget Tracking for Spacecraft Swarms,” *2020 AAS/AIAA Astrodynamics Specialist Conference*, Lake Tahoe, CA, 2020.
- [25] Park, T. H., and D’Amico, S., “Online Supervised Training of Spaceborne Vision during Proximity Operations using Adaptive Kalman Filtering,” , 2023.

- [26] Teunissen, P., “A new method for fast carrier phase ambiguity estimation,” *Proceedings of 1994 IEEE Position, Location and Navigation Symposium - PLANS'94*, 1994, pp. 562–573. <https://doi.org/10.1109/PLANS.1994.303362>.
- [27] Low, S. Y. W., and D’Amico, S., “Precise Distributed Satellite Navigation: Differential GPS with Sensor-Coupling for Integer Ambiguity Resolution,” *2024 IEEE Aerospace Conference (AERO)*, IEEE, 2024, p. 0.
- [28] D’Amico, S., “Autonomous Formation Flying in Low Earth Orbit,” Ph.D. thesis, Delft University, 2010.
- [29] Sullivan, J., and D’Amico, S., “Adaptive Filtering for Maneuver-Free Angles-Only Navigation in Eccentric Orbits,” *27th AAS/AIAA Space Flight Mechanics Conference*, San Antonio, TX, 2017.
- [30] Stacey, N., and D’Amico, S., “Analytical process noise covariance modeling for absolute and relative orbits,” *Acta Astronautica*, Vol. 194, 2022, pp. 34–47. <https://doi.org/10.1016/j.actaastro.2022.01.020>.
- [31] Dennison, K., Stacey, N., and D’Amico, S., “Autonomous Asteroid Characterization Through Nanosatellite Swarming,” *IEEE Transactions on Aerospace and Electronic Systems*, Vol. 59, No. 4, 2023, pp. 4604–4624. <https://doi.org/10.1109/TAES.2023.3245997>.
- [32] Chernick, M., and D’Amico, S., “New Closed-Form Solutions for Optimal Impulsive Control of Spacecraft Relative Motion,” *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 2, 2018, pp. 301–319. <https://doi.org/10.2514/1.G002848>.
- [33] Chernick, M., and D’Amico, S., “Closed-Form Optimal Impulsive Control of Spacecraft Formations Using Reachable Set Theory,” *Journal of Guidance, Control, and Dynamics*, Vol. 44, No. 1, 2021, pp. 25–44. <https://doi.org/10.2514/1.G005218>.
- [34] Koenig, A. W., and D’Amico, S., “Fast Algorithm for Fuel-Optimal Impulsive Control of Linear Systems With Time-Varying Cost,” *IEEE Transactions on Automatic Control*, Vol. 66, No. 9, 2021, pp. 4029–4042. <https://doi.org/10.1109/TAC.2020.3027804>.
- [35] Koenig, A. W., and D’Amico, S., “Robust and Safe N-Spacecraft Swarming in Perturbed Near-Circular Orbits,” *Journal of Guidance, Control, and Dynamics*, Vol. 41, No. 8, 2018, pp. 1643–1662. <https://doi.org/10.2514/1.G003249>.
- [36] Guffanti, T., and D’Amico, S., “Passively Safe and Robust Multi-Agent Optimal Control with Application to Distributed Space Systems,” *Journal of Guidance, Control, and Dynamics*, Vol. 46, No. 8, 2023, pp. 1448–1469. <https://doi.org/10.2514/1.G007207>.
- [37] Gaias, G., Ardaens, J.-S., and Schultz, C., “The AVANTI Experiment: Flight Results,” *10th International ESA Conference on Guidance, Navigation, & Control Systems*, Salzburg, Austria, 2017.
- [38] D’Amico, S., and Montenbruck, O., “Proximity Operations of Formation-Flying Spacecraft Using an Eccentricity/Inclination Vector Separation,” *Journal of Guidance, Control, and Dynamics*, Vol. 29, No. 3, 2006, pp. 554–563.
- [39] Sullivan, J., and D’Amico, S., “Nonlinear Kalman Filtering for Improved Angles-Only Navigation Using Relative Orbital Elements,” *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 9, 2017, pp. 2183–2200.
- [40] Crassidis, J. L., and Markley, F. L., “Unscented Filtering for Spacecraft Attitude Estimation,” *Journal of Guidance, Control, and Dynamics*, Vol. 26, No. 4, 2003, pp. 536–542. <https://doi.org/10.2514/2.5102>.
- [41] Koenig, A. W., Guffanti, T., and D’Amico, S., “New State Transition Matrices for Relative Motion of Spacecraft Formations in Perturbed Orbits,” *Journal of Guidance, Control, and Dynamics*, Vol. 40, No. 7, 2017, pp. 1749–1768.
- [42] Guffanti, T., D’Amico, S., and Lavagna, M., “Long-Term Analytical Propagation of Satellite Relative Motion in Perturbed Orbits,” *27th AAS/AIAA Space Flight Mechanics Meeting*, San Antonio, TX, 2017.
- [43] Tapley, B. D., Bettadpur, S., Watkins, M., and Reigber, C., “The Gravity Recovery and Climate Experiment: Mission Overview and Early Results,” *Geophysical Research Letters*, Vol. 31, No. 9, 2004.
- [44] Picone, J., Hedin, A., Drob, D. P., and Aikin, A., “NRLMSISE-00 Empirical Model of the Atmosphere: Statistical Comparisons and Scientific Issues,” *Journal of Geophysical Research: Space Physics*, Vol. 107, No. A12, 2002, pp. SIA–15.