AFOSR Astrodynamics & Physics of Sensing Program Review

Mechanical, Aerospace and Nuclear Engineering
Rensselaer Polytechnic Institute
110 8th Street, Troy, NY 12180

Department of Aerospace Engineering
Texas A&M University,
400 Bizzell St, College Station, TX 77843

RCAT-CS: Reconfigurable Constellations for Adaptive Tracking in Cislunar Space

Year-1 Technical Review · September 2025 – September 2026

Sandeep K. Singh
Principal Investigator
Rensselaer Polytechnic Institute

Manoranjan Majji · **John L. Junkins**
Co-Investigators
Texas A&M University

AFOSR Program Review · **19–23 October 2026**

Supported by: Air Force Office of Scientific Research (AFOSR) Grant FA9550-25-1-0313

Outline

1 Motivation and Overview
Domain challenges, research objectives, three-year plan

2 Constellation Design and Optimization
Constraint-aware, estimator-in-the-loop architecture search over periodic orbit families

3 Maneuver Identification and Characterization
Modal innovation analysis for detecting stealthy impulsive burns

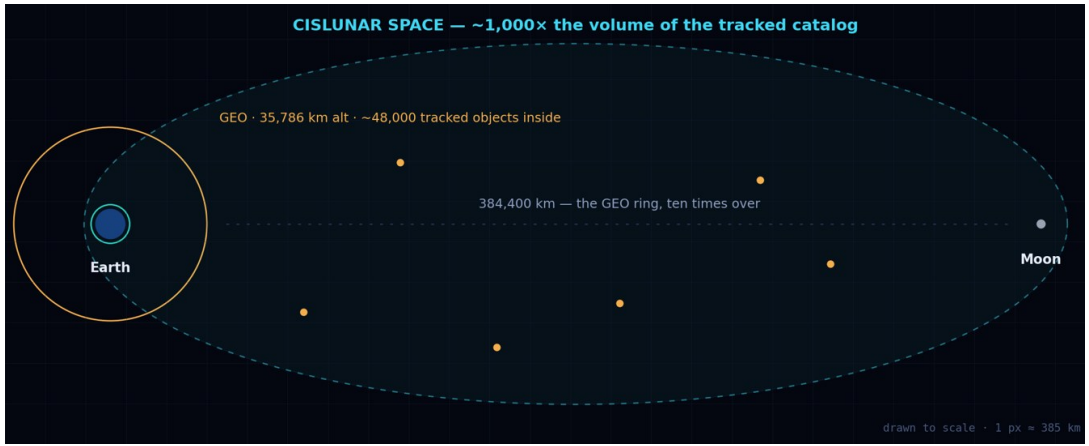
4 Year-2 Efforts
Adversarial attrition, reachability maps, lunar surface sensor networks

5 Publications and Student Involvement
Archival output and workforce development

1 Motivation and Overview

Cislunar SDA demands an agile and robust observer network.

The cislunar SDA problem and challenges.



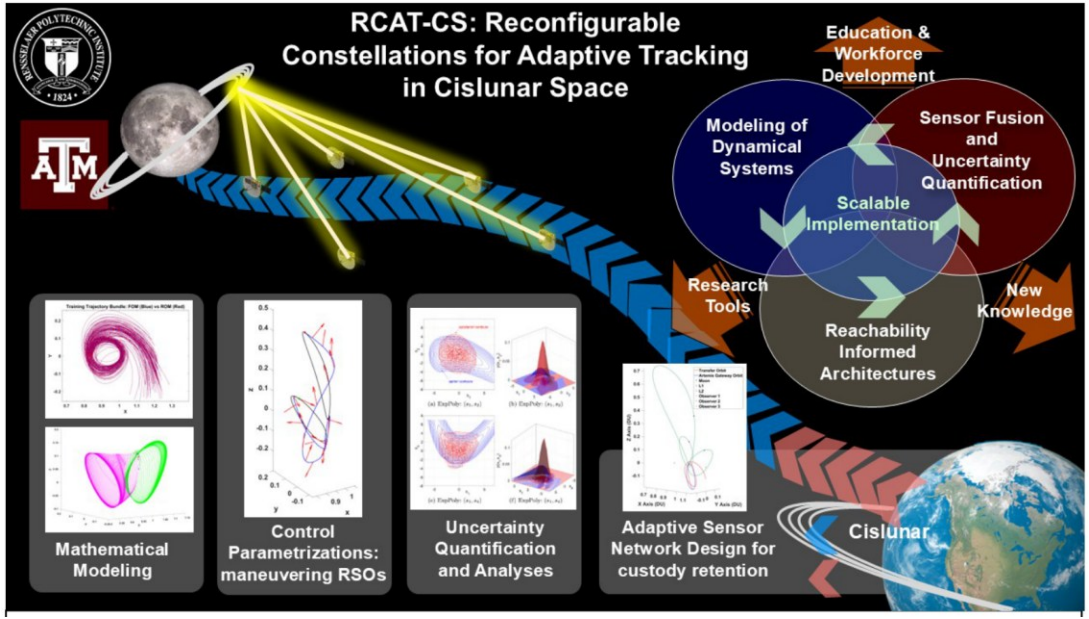
Cislunar volume relative to the GEO belt, drawn to scale. Objects shown are illustrative.

WHAT MAKES THE REGIME DIFFERENT?

- Non-Keplerian, multi-body dynamics with no near-invariant element set.
- Small **state errors diverge exponentially** along unstable manifolds.
- **Observation** opportunities **are sparse and intermittent**: lunar occultation, eclipse, bright bodies and communication dropouts.
- Range and angular resolution degrade with distance. Faster custody loss.

CONSEQUENCE Ground-based networks supporting LEO and GEO custody cannot be scaled into xGEO. At the center of the tracking problem is a sensor-architecture design problem. **1) Where observers are placed,** and **2) How that placement is adapted**, dominates achievable tracking accuracy and custody retention.

Research Objectives and Programmatic Scope



Program structure: modeling, sensor fusion and uncertainty quantification, and reachability-informed architectures.

RO1
RSO Maneuver Parameterization and Characterization

Characterize families of admissible maneuvers and RSO behaviors in the cislunar regime; derive models for low-thrust, impulsive and adversarial control.

RO2
Robust Sensor Constellation Design

Design xGEO space-based constellations for coverage, tracking and detection; maximizing performance under operational constraints.

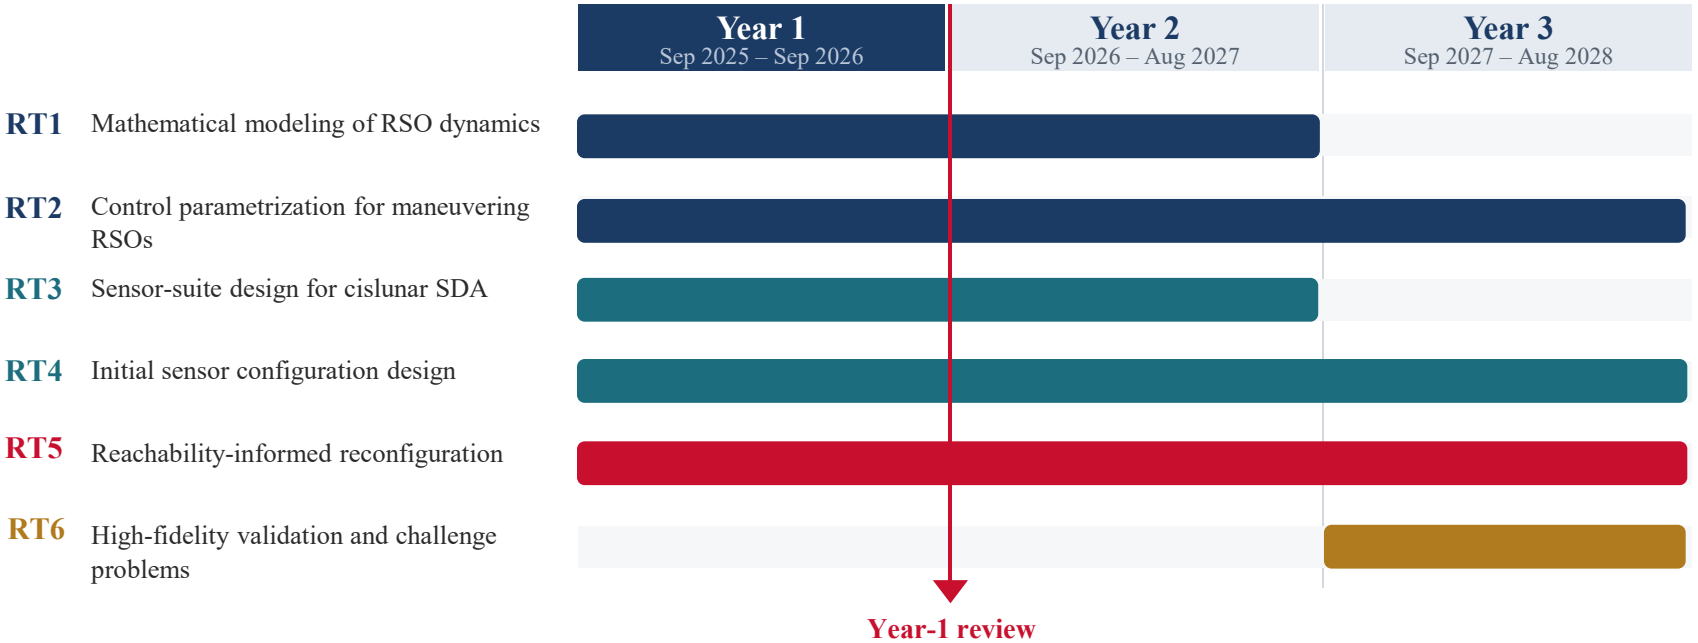
RO3
Adaptive and Rapid Reconfiguration

Models for reconfig. of constellations in response to projected RSO motion, natural perturbations and/or node loss, subject to operational constraints.

RO4
Implementation and Characterization

Demonstrate the methodology on legacy RSO ephemerides; evaluate applicability to real-time and near-real-time operation.

Project Timeline



YEAR-1 SUMMARY **Significant progress on RTs 1,2 and 4 which has** produced *multiple archival journal articles*. RTs 3 & 5 are planned to be implemented in Year 2. Outputs from RT3 will be used to refine results from RT4 and 5 . RT6 remains scheduled for Year 3.

2

Constellation Design & Optimization

Challenges:

Dealing with an infinite design space.
Solution Space littered with local extrema.
Constrained Optimization, Multi-Objective.
Robust Optima? Which Algorithm?

Problem Formulation & Design Space.

DESIGN VECTOR

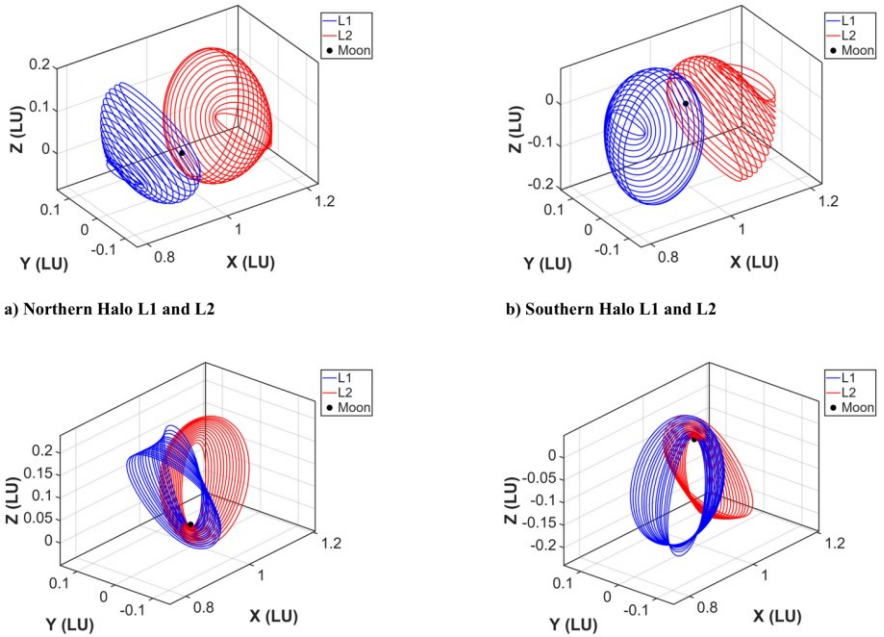
$$DP = [\, F_1, O_1, \varphi_1, \ \ldots, \ F_n, O_n, \varphi_n \,]^T$$

- F_i - discrete orbit-family label, drawn from eight libration-point families
- O_i - discrete orbit selection within that family (amplitude, period)
- $\varphi_i \in [0, 2\pi)$ - continuous initial phase along the periodic orbit

Mixed-integer, 3N-dimensional. N ∈ {3, 4, 5, 6} observers. 100 members catalogued per family.

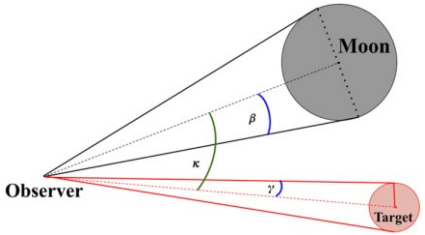
TARGET REGIMES

- **Coasting**: passive motion on a 9:2 southern L_2 NRHO (the lunar Gateway orbit), one full period.
- **Maneuvering:** time-optimal low-thrust transfer between an L_1 halo (C = 3.0767) and an L_2 halo (C = 3.1472): 21.72 days, 9.18 kg of propellant.



Orbit-family database, ten of one hundred members shown per family: northern and southern halo, and northern and southern rectilinear, at L_1 and L_2.

Measurement Model, Fidelity and Constraints.



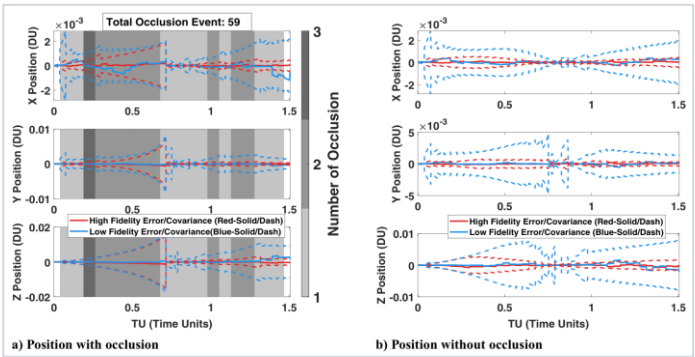
Conical occlusion test: an event is declared whenever $\kappa < \beta + \gamma$.

EMBEDDED CONSTRAINTS

- **Angles-only:** right ascension and declination of the line-of-sight vector.
- **Angles-with-range:** slant range appended, rendering the position Jacobian full rank.
- **Occlusion embedded inside the filter:** only unoccluded observers contribute measurements at each epoch.
- **Measurement cadence 4 hours,** effective target radius *100 m*, centralized update across all observers.

CALIBRATED NOISE LEVELS

Quantity	Low fidelity	High fidelity
Angular $\sigma\angle$	9.3×10^{-4} rad	1.3×10^{-4} rad
Range $\sigma\rho$	1 km	0.01 km
Process noise $\mathbf{Q}$	$10^{-4} \cdot I_{6x6}$	$10^{-4} \cdot I_{6x6}$
Initial σ_r^2 / σ_v^2	1 km² / 0.01 m²s⁻²	1 km² / 0.01 m²s⁻²

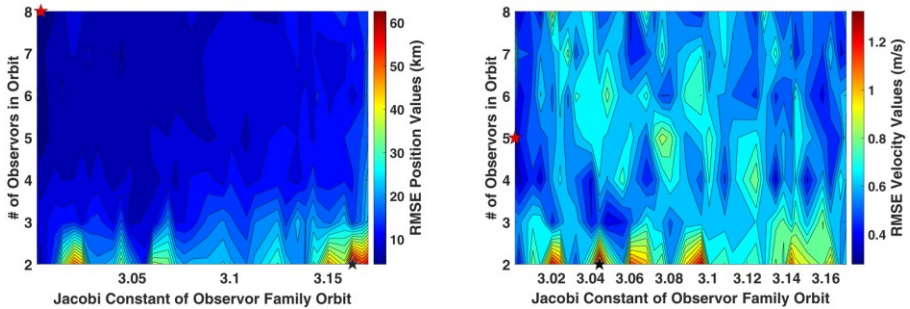


Three-sigma position envelopes with occlusion (left) and without (right); high- and low-fidelity noise overlaid. Shaded bands mark the number of observers occluded.

WHY THIS MATTERS? In the worst-case geometry studied, **a three-observer constellation experienced 59 occlusion events in a single target period.** Uncertainty grows monotonically through each blackout and will recontract only after a direct viewing angle is regained.

Stage I: Heuristic Pruning of Orbit-Family Database.

Before optimizing, **ten orbits per family were drawn by Latin hypercube sampling** and tracked with observer counts from two to eight, placed at equal angular separations. **Filter accuracy was mapped against the Jacobi constant of the observer orbit and against observer count.**



Position and velocity RMSE contours for the northern halo L_1 family, occlusion included.

- **Increasing observer count has diminishing returns beyond five to seven spacecraft**.
- The **four halo families outperform the rectilinear families**, which suffer excessive line-of-sight occlusion near the Moon.

Family	Avg. pos.	Avg. vel.
NHL1	13.98	0.63
NHL2	7.98	0.46
SHL1	9.70	0.55
SHL2	5.62	0.38
NNRL1	41.74	1.27
NNRL2	24.38	1.62
SNRL1	41.20	1.40
SNRL2	7.41	0.55

Target: lunar gateway orbit. Position in km, velocity in m/s. Rows in grey were excluded.

RESULT The database was reduced to four halo families, NHL1, NHL2, SHL1 and SHL2, before any optimization was run. **This makes the subsequent mixed-integer search tractable** without resorting to a brute-force sweep over all family and observer-count combinations.

Stage II: GA constellation Optimization.

$$\overline{\mathrm{RMSE}}_r = \frac{\mathrm{RMSE}_r}{\max(\mathrm{RMSE}_r)}, \qquad \overline{\mathrm{RMSE}}_v = \frac{\mathrm{RMSE}_v}{\max(\mathrm{RMSE}_v)} \qquad \bar{J}_{\mathrm{trace},r} = \frac{\langle \mathrm{tr}(P_r) \rangle}{\max \langle \mathrm{tr}(P_r) \rangle}, \qquad \bar{J}_{\mathrm{trace},v} = \frac{\langle \mathrm{tr}(P_v) \rangle}{\max \langle \mathrm{tr}(P_v) \rangle}$$

- **Two error metrics:** time-averaged position and velocity RMSE against the high-fidelity truth trajectory.
- **Two confidence metrics:** time-averaged trace of the posterior position and velocity covariance submatrices.
- **Normalized** by their maximum over the sampled design space, so no objective dominates by scale alone.
- Estimation **error** and posterior **uncertainty** are treated **as complementary**: they decouple under unmodeled maneuvers, strong anisotropic observability, and intermittent occlusion.

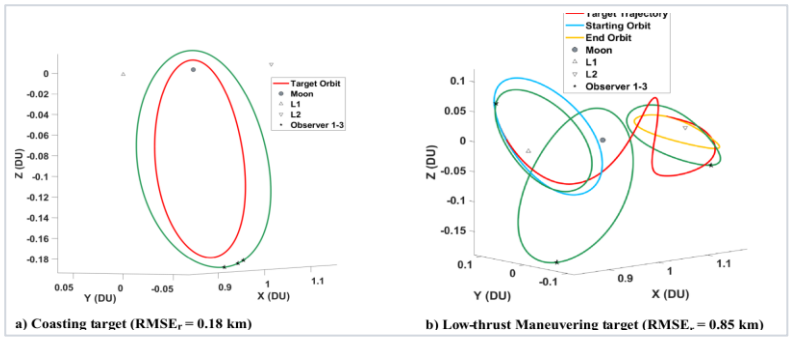
WHY A METAHEURISTIC?

Cost surface is **nonconvex** and **periodic** in phase.
The family and orbit selections are discrete.
Each function evaluation is a full closed-loop, occlusion-embedded EKF run.
Gradient information is unavailable and the evaluations are expensive and noisy.

ALGORITHM SETTINGS

Hyperparameter	Single-objective	Multi-objective
Population size	50	50
Maximum generations	**600**	**1200**
Crossover fraction	0.80	0.80
Elite count	2	2
Mutation rate	0.02	0.02
Function tolerance	10^{-6}	10^{-4}

Optimized Architectures: Single Objective



a) Coasting target ($RMSE_r = 0.18$ km) b) Low-thrust Maneuvering target ($RMSE_r = 0.85$ km)

Three-observer constellations optimized for high-fidelity AO sensors, minimizing position RMSE.

- Optimized constellations **implicitly yield geometries with zero or very few occlusion events.**
- **Geometric diversity** across halo families, combined with modest phase adjustment, **outperforms simply adding observers**.
- **Steady-state tracking** of a coasting target and **transient tracking** of a maneuvering target favor different architectures.

0.18 km

Position RMSE, coasting target

Gateway 9:2 southern L_2 NRHO, one full period, three high-fidelity angles-only observers.

0.85 km

Position RMSE, maneuvering target

Continuous low-thrust L_1-to-L_2 halo transfer tracked end to end by the same observer count.

GEOMETRIC CHARACTER OF THE SOLUTIONS

- Coasting optima places observers in geometrically close orbits at nearly equal Jacobi constants.
- The maneuvering case selects a different geometry, spreading observers to limit transient observability loss during the thrust arc.

DESIGN IMPLICATION **A single fixed constellation cannot serve both mission regimes at optimum.** This motivates reconfiguration, and it emerged from the optimization rather than being assumed at the outset.

Sensor Modality: AO vs. AR

Position RMSE for GA-optimal constellations with **low-fidelity sensors**, evaluated across four target scenarios and observer counts from three to six.

Target scenario	N = 3	N = 4	N = 5	N = 6
Artemis (gateway NRHO)	0.1008 | 3.1585	0.0674 | 3.8142	0.0294 | 1.9974	0.0387 | 1.5371
Distant retrograde orbit	0.1547 | 4.7247	0.0571 | 4.3688	0.0863 | 4.3370	0.0534 | 3.6580
Northern halo L_2	0.0554 | 4.1208	0.0392 | 2.6650	0.0220 | 2.5981	0.0345 | 2.4105
Low-thrust transfer	0.4626 | 8.8371	0.4388 | 8.0666	0.3735 | 7.0601	0.3593 | 5.5635

Each cell reports angles-with-range | angles-only position RMSE in km.

OBSERVATIONS

- **Appending range reduces position RMSE by 1-2 orders of magnitude in every scenario and at every observer count tested.**
- The **benefit holds even with low-fidelity sensors**, where range noise is 1 km, two orders of magnitude coarser than the high-fidelity case.
- Range also yields substantially smaller posterior covariance traces and fewer occlusion-driven filter degradations.

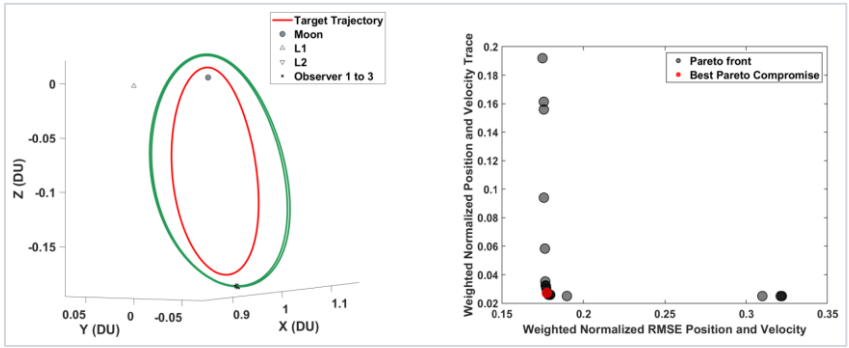
MECHANISM

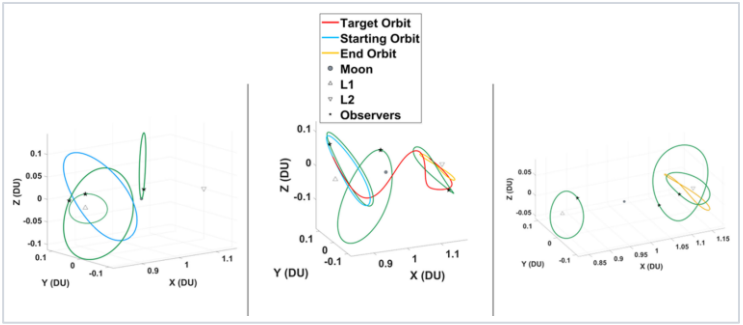
- **AO measurements** - range direction unobservable at a single epoch. The **filter must accumulate information** across widely separated observers and across time.
- The position Jacobian becomes full rank once range is included, so the measurement null space collapses to the zero vector.

RECOMMENDATION **Deploy range-capable sensors wherever mission constraints permit.** This is the single largest performance lever, more substantial than observer count and family selection.

Multi-objective Optimization and the Pareto Structure

Coasting scenario, high-fidelity angles-only model: Pareto-compromise architecture (left) and the corresponding Pareto front (right).





Optimal geometries for the low-thrust case: departure orbit, transfer arc, and arrival orbit.

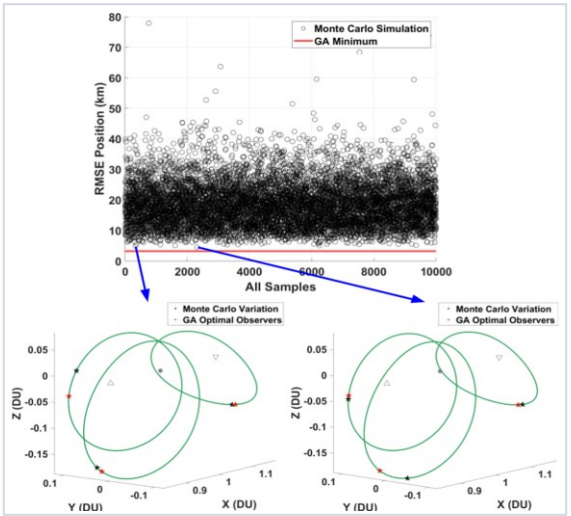
FINDINGS

- The **Pareto fronts show that meaningful gains** in one require large degradation in another.
- **Geometric adjustments between the weighted single-objective solution and the compromise are modest**. Halo-family geometry exerts the dominant influence on tracking performance.
- Estimation error and posterior uncertainty are positively correlated but not identical. Pairing the error metric with a directional robustness metric might be useful.

INTERPRETATION Where the mission tolerates a modest increase in nominal error, **the Pareto-compromise architecture is the better operational choice as** it trades nominal accuracy for confidence.

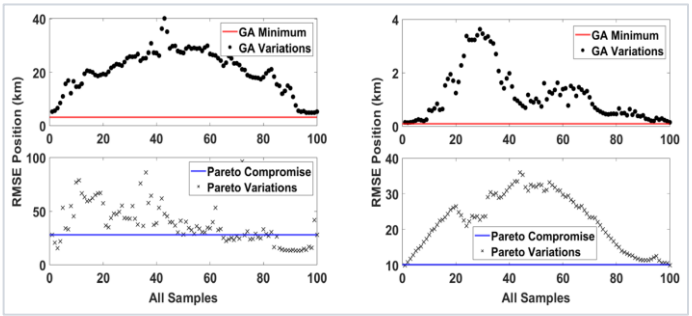
Assessing an Optimized Constellation's Robustness

Three independent stress tests were applied to the converged architectures: **local perturbation of the observer geometry**, **variation of the target's initial phase**, and **extension of the evaluation horizon well beyond the design period**.



10,000-trial Monte Carlo exploration around the GA optimum, low-fidelity angles-only, low-thrust scenario.

Sensitivity to the target's initial phase, 100 temporally equidistant samples along the Artemis orbit.



- **Local perturbation:** across 10,000 runs, **GA solution remained at or near** the lowest sample mean and median for both metrics.
- **Target phase:** Peak position RMSE reached 40.189 km for the GA-optimal architecture and 96.215 km across Pareto variations, against an optimal value of 3.159 km. **Degradation peaks when the target is initialized near periapsis.**
- **Long horizon:** over 20 revolutions, peak RMSE reached 33.380 km for angles-only but only 2.427 km with range. **Extended-duration sensitivity is less damaging than initial-phase sensitivity.**

CONSEQUENCE FOR OPERATIONS A constellation optimized for one target phase may degrade substantially at another. Candidate architectures must be **validated before adoption**, and **a fixed geometry cannot be assumed to hold custody indefinitely**.

Design Recommendations and Research Limitations

I Prefer range-capable sensors

Augmenting angles-only measurements with range reduces errors by 1-2 orders of magnitude and yields substantially smaller posterior covariance traces. The benefit persists at low sensor fidelity.

II Prefer geometry over raw count

Combining observers drawn from multiple halo families and applying modest phase adjustments improves geometric observability more effectively than simply increasing the number of spacecraft. Returns from additional observers diminish..

III Embed constraints inside the estimator

Occlusion handling should be a part of the filter loop. In-the-loop suppresses spurious innovations and produces more reliable filter behaviour during and after line-of-sight blackouts.

RESEARCH LIMITATIONS

Single-target scope. · Geometric **occlusion model** that omits stray-light and background-brightness effects. · **Mission-level constraints** such as communication windows, station-keeping cost and retasking latency are not modelled. · **Resilient Network design -** Under Attrition. · **Hybrid Networks -** Both Space based and Terrestrial sensors.

3

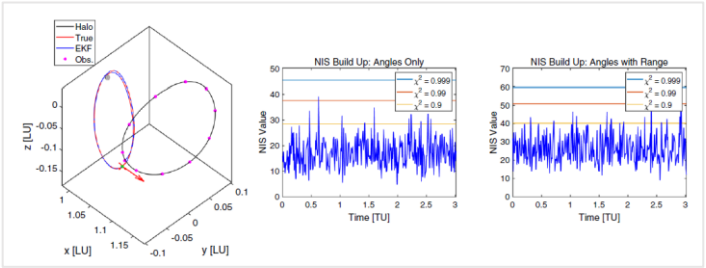
Maneuver Identification and Characterization

Challenges:

Small and frequent maneuver detection through measurements?

Sophisticated adversarial intent? Maneuver designed to evade.

Limitations of the Full-Spectrum Innovation Testing

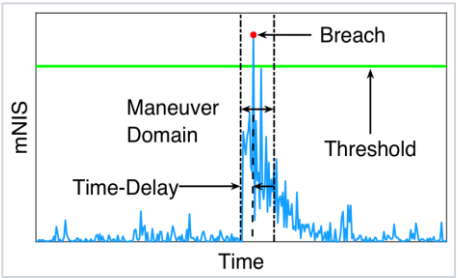


Innovation test orbit/maneuver configuration and the corresponding NIS spectra using a uniform 10 sat. observer constellation

THE CONVENTIONAL TEST

$$NIS_k = \nu_k^T S_k^{-1} \nu_k , \qquad NIS_k \sim \chi^2(m \cdot N_{obs})$$

- NIS-test flags any deviation across the full measurement space, so a maneuver signature competes against every other source of residual.
- False positives possible when filter is over-confident, initial covariance is tight, or nonlinearity in the dynamics introduces residual leakage.
- **An adversary aware of the constellation geometry can plan a burn so that the full-space statistic never breaches its threshold or introduces several false detections.**



***Detection nomenclature**: threshold, breach, maneuver domain and time-delay.*

DETECTION METRICS

- **True positive** - a breach inside the maneuver domain, defined as the first 10% of total flight time after the burn.
- **False positive** - any breach outside that domain.
- **Time-delay** - interval between the true burn epoch and the first true-positive breach. Undefined when no true positive occurs.

PREMISE If an adversary can place a burn where the estimator is least sensitive, detection could be performed in a reduced "subspace" that amplifies the signature.

The Modal Innovation Analysis

TWO REDUCED-ORDER BASES

Proper orthogonal decomposition (POD)

A snapshot matrix of predicted measurement histories is decomposed via SVD. The leading ***r left singular vectors*** form an orthonormal modal basis spanning the most energetic measurement directions under nominal dynamics.

Extended dynamic mode decomposition (EDMD)

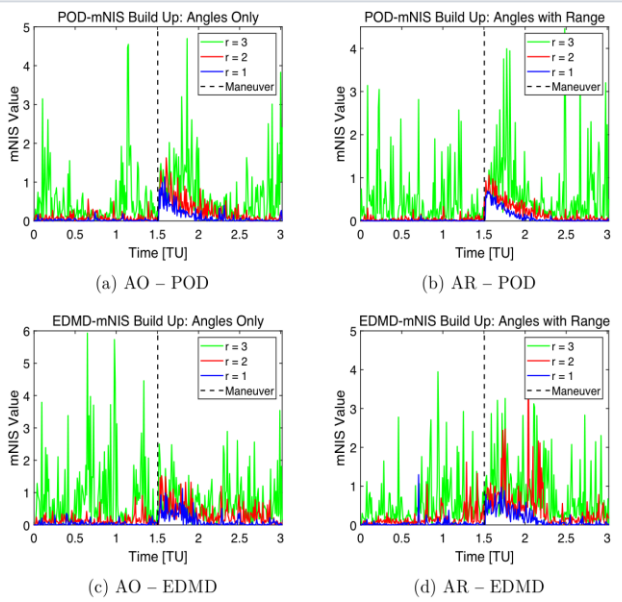
The measurement is lifted through a ***dictionary of nonlinear observables*** and a finite-dimensional Koopman operator is fitted by least squares. Its dominant eigenvectors form the basis. The lifting separates nonlinear signatures that a purely linear projection folds together.

MODAL NIS

$$\tilde{\nu}_k = \Phi_m^T \nu_k , \qquad \tilde{S}_k = \Phi_m^T S_k \Phi_m$$

$$\mathrm{mNIS}_k = \tilde{\nu}_k^T \tilde{S}_k^{-1} \tilde{\nu}_k \sim \chi^2(r)$$

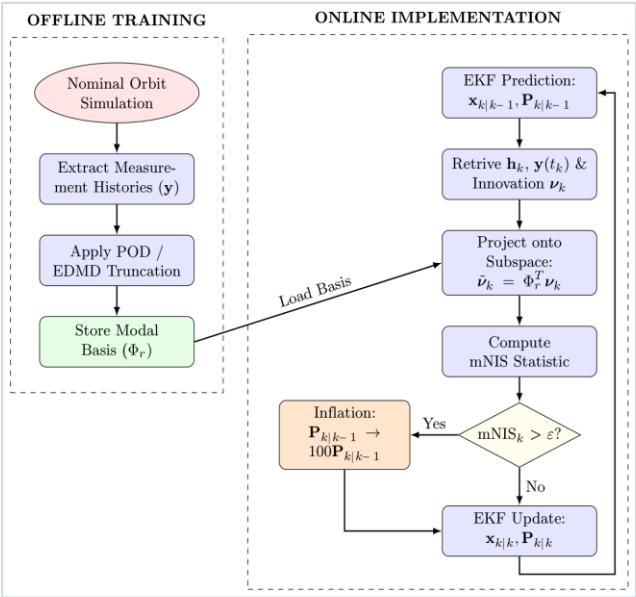
Retaining all $2N_{obs}$ modes recovers the standard NIS exactly, so the method degrades gracefully to the conventional test.



Reduced-order mNIS spectra as a function of retained mode count, POD (top) and EDMD (bottom), angles-only and angles-with-range.

KEY PROPERTY **Projecting** the innovation onto a small number of dominant modes **suppresses the noise floor**. The maneuver signature is concentrated rather than distributed, which raises detection probability and reduces false alarms at the same threshold.

The Integrated Detection & Estimation Framework



Offline training establishes the modal basis and the online loop filters innovations and triggers adaptive covariance inflation.

OFFLINE TRAINING

- **Nominal trajectories are simulated** in the CR3BP for the mission profile of interest and measurement histories are generated with calibrated noise.
- POD or EDMD truncation extracts the dominant **modal basis representing the expected innovation structure** under nominal conditions.
- Requires only the measurement time history, so the **basis can be regenerated for any profile at low cost.**

ONLINE IMPLEMENTATION

- The **innovation** from each update **is projected onto the stored basis** and the mNIS statistic is evaluated against a percentile threshold.
- **On a breach** the **prior covariance is inflated** before the measurement update, for that time step only.
- Inflation lets the update pull the estimate toward the post-maneuver trajectory without divergence, while the breach records the detection.
- Cost scales as O(m · Nobs) for POD and O(M · m · Nobs) for EDMD.

ARCHITECTURE NOTE **Detection and mitigation are coupled.** The mNIS statistic flags the maneuver and triggers the covariance response. No separate decision stage before recovery.

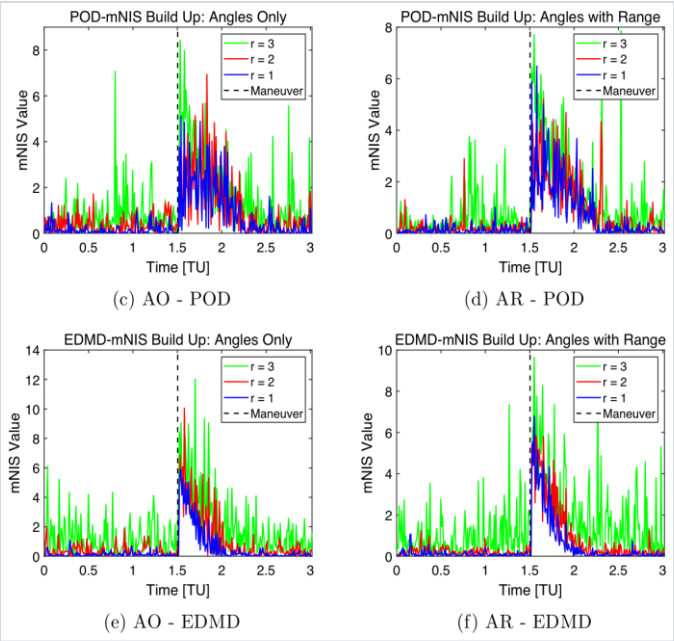
Strategically designed evasive maneuvers

Four impulsive burn scenarios were analyzed to test the methodological limits of the detector. Each applies **25 m/s in total near apolune of the gateway NRHO**.

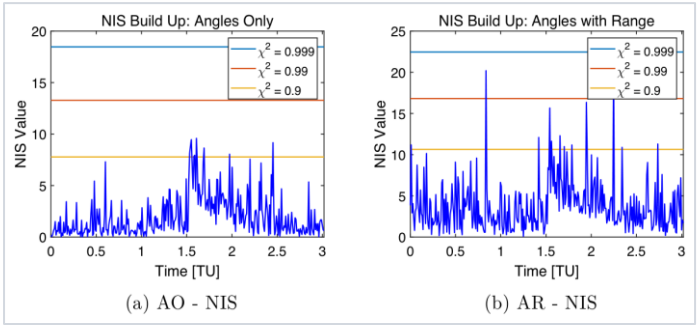
S1	**Null-space burn**	• Burn direction lies in the LOS Jacobian's null space, causing no first-order angle change. • Under angles-only, the maneuver is invisible at execution; detection depends on higher-order deviation over time. • Adding range collapses the null space to zero.
S2	**Covariance-aligned burn**	• Direction aligns with the dominant covariance eigenvector i.e., the filter's least-certain axis. • The resulting innovation can be masked by normal filter drift and uncertainty growth.
S3	**Unstable eigendirection burn**	• Burn aligns with the periodic orbit's unstable-manifold direction i.e., the dominant Monodromy eigenvector. • The resulting departure mimics natural dynamical divergence, making it hardest to isolate.
S4	**Sub-noise micro-impulses**	• 2.5 m/s impulses every 3 h, each below the sensor noise floor. • No impulse is immediately detectable; only accumulated residuals reveal the sequence.

DESIGN PRINCIPLE The scenarios target: **measurement geometry, filter confidence, dynamical structure, and sensor resolution**.

Detecting the “null-space” burn



Scenario S1, reduced-order mNIS: POD (top row) and EDMD (bottom row), angles-only and angles-with-range, for one, two and three retained modes.



The same scenario under the conventional full-spectrum NIS test.

READING THE SPECTRA

- Standard NIS registers anomalies in both measurement models, but spikes are not pronounced and cause several false IDs, especially for AR.
- EDMD-based mNIS resolves the timing distinctly. Nonlinear lifting is largely unaffected by a burn constructed to lie in the linear measurement null space.
- The POD-based mNIS also exposes the reduced-order modes, though its signature is harder to separate from background.

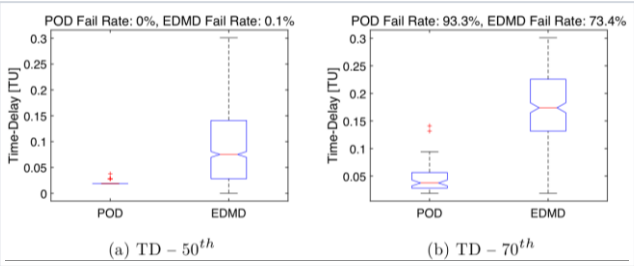
RESULT A **maneuver engineered to be first-order invisible to an AO measurement model** is detected in a 1D modal subspace, with the burn epoch recovered to within a small number of measurement intervals.

Performance across scenarios

1000 Monte Carlo trials per config., iterating over measurement-noise realization, at a 90th-percentile detection threshold. AO measurements, 1 retained mode.

Scenario	Method	FP trigger rate	TP hit rate	Time-delay
S1 Null-space	POD	100%	100%	3 hr
	EDMD	**1.9%**	**100%**	**2 hr**
S2 Covariance-aligned	POD	100%	99.9%	6 hr
	EDMD	**8.9%**	**100%**	**2 hr**
S3 Unstable eigendirection	POD	64.1%	14.4%	26 hr
	EDMD	**18.0%**	**100%**	**2 hr**
S4 Sub-noise micro-impulses	POD	87.5%	100%	26 hr
	EDMD	**0%**	**100%**	**23 hr**

Averages over 1000 trials. Trigger and hit rates report how often the respective metric fired at least once.



Time-delay distributions at the 50th and 70th percentile thresholds.

WHAT THE STATISTICS SHOW

- **EDMD eliminates false positives in every scenario at this threshold**, while retaining a perfect true-positive hit rate.
- Unstable-eigendirection burn is still obscured for POD with a single retained mode.
- Retaining two modes was the most robust truncation on average across scenarios but a single mode is faster when it works.
- **Threshold selection matters**. At 70th percentile, POD fails to breach within the maneuver domain in 93.3% of trials and EDMD in 73.4%.

4

Year-2 Efforts

Importance:

Adaptability can alleviate network fragility.

Reachability information for fast/autonomous decision making.

Problem is RESOURCE-CONSTRAINED!

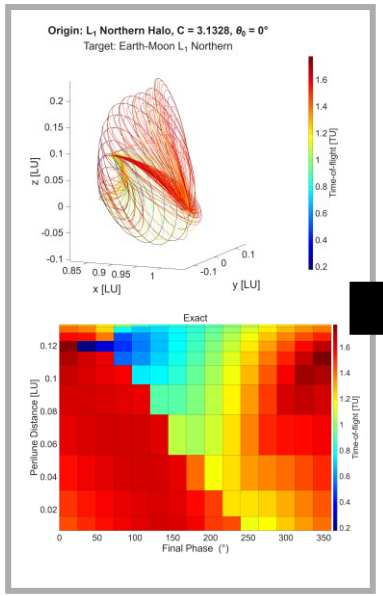
Thrust I: Custody Retention under Adversarial Attrition

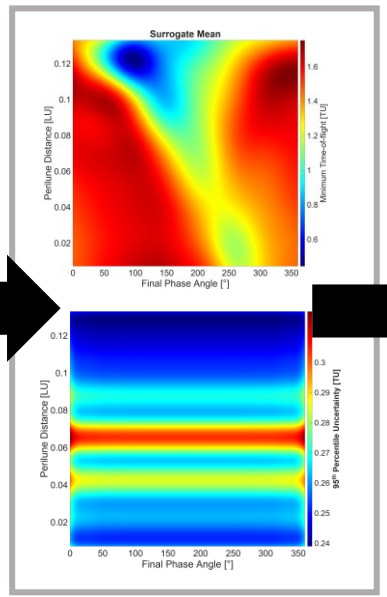
Feedback to constellation optimization by formalizing metrics that evaluate resilience to attrition, including loss, jamming and degradation of nodes.

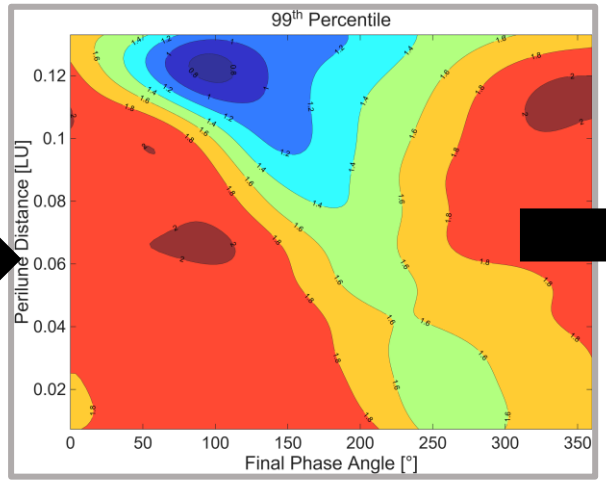
S1	**Threat Model**	• **Keystone observer:** Removal/degradation of the node that creates the largest low-observability window along the target's future trajectory. • **Adversarial maneuver:** The adversary then maneuvers in that direction. • **Availability models:** Kinetic loss, persistent jamming, and intermittent link denial.
S2	**Observability Analysis**	• **Information loss:** Track the surviving network from available sensor bearing vectors. • **Weak direction:** The smallest eigenvalue will identify the least-constrained direction. • **Exploitable direction:** Its eigenvector will define the direction most vulnerable to exploitation. • **Geometric interpretation:** Directly quantify what information was lost due to node denial.
S3	**Custody Metrics**	• **Custody loss:** Declared when the true state exits the filter's confidence ellipsoid. • **No fixed threshold:** Avoids relying on an arbitrary error bound. • **Bounded error vs. divergence:** Separates large but credible errors from silent filter divergence. • **Silent failure:** Captures cases where the filter remains confident despite an incorrect estimate.
S4	**Feedback**	• Feed the resilience metrics into constellation design and optimization. • Design becomes resilient to node loss while maintaining acceptable tracking performance.

DELIVERABLE An informative **map relating node loss to time-to-custody-loss** for a candidate architecture, and a redundancy specification stating how many and which observers are key to surviving targeted denials without losing efficacy.

Thrust II: Reachability Informed Reconfiguration







Autonomous Reconfiguration

Quick Trade Studies

OPEN QUESTIONS

How much tracking fidelity may be surrendered before custody is at risk?

What reconfiguration cadence is sustainable over a multi-year mission lifetime?

Does reconfiguration beat redundant, stationary observers on a per-kg basis?

Thrust III: Lunar Surface Sensor Networks

Why surface sites?

- No station-keeping cost and no orbital phase to drift. A surface node holds its geometry indefinitely.
- Persistent near-Moon viewing geometry, precisely where orbital observers are weakest.
- Polar sites offer favourable thermal and illumination conditions and a direct view of the polar regions where lunar surface activity is concentrated.

Disadvantages?

- Two surface sites subtend a small mutual baseline compared with widely separated orbital observers, angular parallax is poor.
- Horizon masking limits each site to roughly a hemisphere of sky.
- Power, thermal cycling and the lunar night impose duty-cycle constraints different from orbital observers.

The constraints.

- Body exclusion: an optical sensor cannot observe a target within a few degrees of the bright lunar limb without saturating.
- Targets in LLO are angularly inseparable from the lunar disk when viewed from a distant observer.
- Surface sensors do not suffer this constraint for near-Moon targets.

OUTPUT **Terrain and bright-body occlusion informed terrestrial network design** and optimization for **near-selenocentric tracking**. Sequential placement strategies to explore. **CLPS landers provide a unique "ride-share" opportunity.**

Publications and Student Involvement

SELECTED ARCHIVAL PUBLICATIONS (Year - 1)

PUBLISHED

Multi-Objective Constrained Constellation Design and Characterization for Cislunar Space Object Tracking

Journal of Spacecraft and Rockets, 2026. DOI: 10.2514/1.A36646

PUBLISHED

Resilient Maneuver Detection through Modal Innovation Analysis in the Circular Restricted Three-Body Problem

The Journal of the Astronautical Sciences, Vol. 73, 2026, Paper 59. DOI: 10.1007/s40295-026-00604-z

UNDER REVIEW

Terrain-Horizon and Exclusion-Informed Greedy Optimization of Lunar Surface Sensor Network for Near-Moon Space Object Tracking

Astrodynamics

ACCEPTED

Benchmarking Gradient-Free Optimization Methods for Estimation-Driven Design of Cislunar Surveillance Constellations in the CR3BP

Discover Space

STUDENTS SUPPORTED

Calvin Chan, Masters	RPI · MANE
Kevin I. Alvarado, Doctoral	RPI · MANE
Hadley Earl, UG	RPI · MANE
Scott Blender, Doctoral	RPI · MANE
Quintin Nelson, Masters	TAMU · Aero
Ben Ingalls, Masters	TAMU · Aero

Thank you

Questions?



asclabrpi.com

Sandeep K. Singh

Assistant Professor, Mechanical, Aerospace and Nuclear Engineering (MANE)
Director, Advanced Space Concepts Laboratory (ASCLab)
Rensselaer Polytechnic Institute
sandes5@rpi.edu · asclabrpi.com

RCAT-CS: Reconfigurable Constellations for Adaptive Tracking in Cislunar Space

Goal:

Reconfigurable sensor networks that hold custody of maneuvering objects in cislunar space

Research Objective(s):

- ***Model and detect maneuvers***
- ***Robust constellation design***
- ***Adaptive reconfiguration***
- ***Validate on real ephemerides***

Impact to DAF:

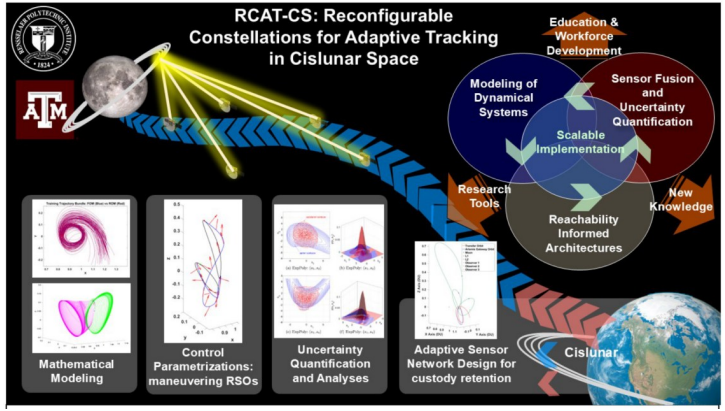
- ***Architectures informed by resource and sensor constraints.***
- ***Stealthy maneuver detection, resilient to attrition.***
- ***Lightweight models for autonomous reconfiguration.***

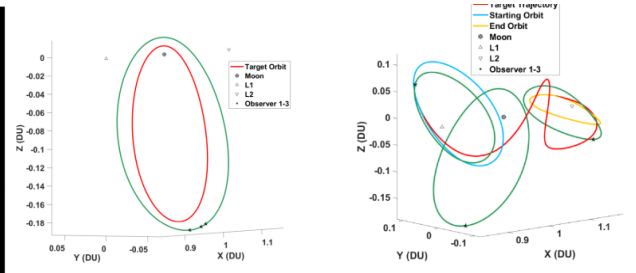
Collaborator(s):

Texas A&M: M. Majji, J. L. Junkins

Program Year(s):

PI: Sandeep K. Singh, RPI



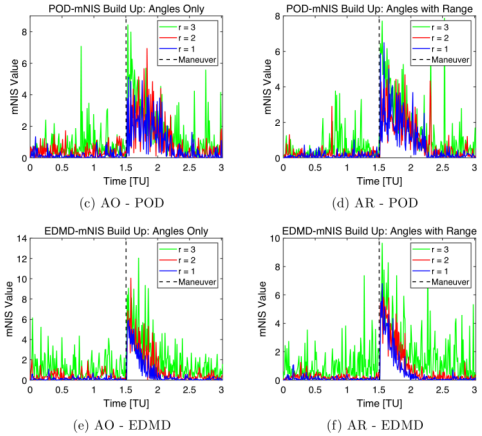


a) Coasting target ($RMSE_r = 0.18$ km)

b) Low-thrust Maneuvering target ($RMSE_r = 0.85$ km)

Optimized three-observer constellations with **occlusion in the loop**: 0.18 km position error on a coasting Gateway-orbit target (a), 0.85 km on a low-thrust maneuvering target (b).

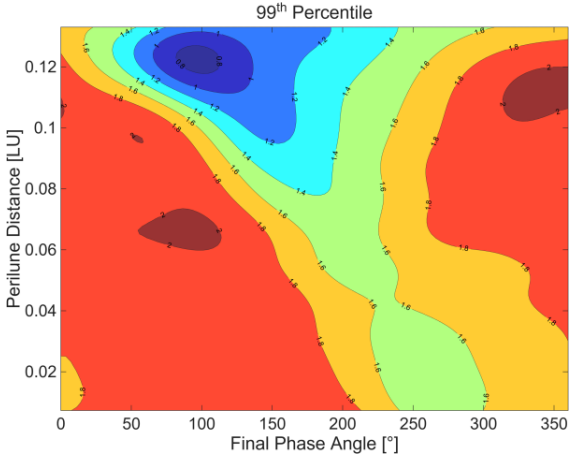
Program structure: dynamical modeling, sensor fusion with uncertainty quantification, and reachability-informed architectures.



(c) AO - POD (d) AR - POD

(e) AO - EDMD (f) AR - EDMD

A burn designed to be invisible to angles-only sensors can be detected by **modal innovation analysis** (EDMD, bottom row). Spikes mark the detected burn epoch, and the dotted line depicts true burn epoch.



Autonomous Reconfiguration: 99th-percentile minimum time-of-flight map for low-thrust observer reconfiguration within the Northern L1 Halo Family.

DISTRIBUTION STATEMENT A. Approved for public release: distribution unlimited.