aerospace

MDPI

Article

Tracking-Accuracy-Constrained Scheduling for Active–Passive Observation Constellations of Noncooperative Space-Object Clusters

Jiapeng Feng [1,2,*], Yuqing Li [1], Sensen Guo [1], Wei Zhao [3], Xuhai Wang [1] and Jun Zhou [2]

[1] Innovation Academy for Microsatellites of CAS, No. 1 Xueyang Road, Pudong District, Shanghai 201304, China; liyq@microsate.com (Y.L.)
[2] School of Astronautics, Northwestern Polytechnical University, 127 West Youyi Road, Beilin District, Xi'an 710072, China
[3] Department of Data Science and Big Data, Shanghai Ocean University, No. 999 Huchenghuan Road, Pudong District, Shanghai 201306, China
* Correspondence: fengjp@microsate.com

Abstract

Tracking noncooperative space-object clusters couples observation duration, revisit timing, and shared constellation resources. We develop a scheduler that constrains position–velocity covariance throughout a 200 s measurement-free interval after each observation. Variable-duration visits combine time-dependent common-field acquisition, serial laser ranging, and independent target posteriors. A hierarchical differential-evolution search coordinates these visits. Nonlinear truth and an extended Kalman filter evaluate all targets in eight synthetic scenes with 72 observers and 200 Monte Carlo realizations per target. Six matched methods, five paired search seeds, and duration and budget scans separate scheduling effects. Under the reference condition, initialization-priority adaptive duration (IP-AD) gives position and velocity root-mean-square errors of 18,443.8 ± 1432.6 m and 57.102 ± 6.491 m/s. Uncertainty is the sample standard deviation of five eight-scene means. Relative to fixed 20 s visits with identical initialization priority and hold screening, the means decrease by 66.00% and 46.46%. However, relaxed planning thresholds combined with lower measurement noise reverse this benefit. Adaptive duration therefore offers conditional average gains; covariance-feasible plans do not ensure that every target meets the actual-error requirements.

Keywords: noncooperative space objects; joint angle–range observations; space mission planning; observation duration; tracking accuracy; nonlinear estimation

Check for updates

Academic Editors: Marco Felice Montaruli and Mauro Massari

Received: 25 September 2026
Revised: 19 October 2026
Accepted: 24 October 2026
Published: 28 October 2026

Copyright: © 2026 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license.

1. Introduction

Sustained space-object tracking requires repeated state updates under limited observing resources. Observation constellations can extend tracking coverage, while multi-spacecraft operations more broadly require coordinated planning and control [1,2]. For noncooperative objects, a useful measurement does not by itself ensure adequate accuracy until the next observation. Nearby lines of sight may permit shared pointing, but the targets can differ in prior uncertainty, observing geometry, and subsequent tracking requirements. A longer observation reduces one target's uncertainty while consuming time that other targets need for initialization or revisits. Scheduling must therefore couple satellite assignment, service order, observation duration, and revisit timing.

Passive optical angles and active laser ranges provide complementary information. Angles constrain the line-of-sight direction, whereas range measurements constrain radial distance; their combined value also depends on geometry and sampling times. Li et al. [3] studied real-time orbit determination during the transition from optical tracking to debris laser ranging. Sang and Bennett [4] examined debris-orbit prediction using laser data from a single station. These ground-based studies support the measurement rationale, but do not establish the feasibility of the spaceborne laser link assumed here. In the present model, several targets may occupy the optical field of view, whereas the laser serves one target at a time. Shared pointing can reduce acquisition overhead, but each target retains its own state estimate and covariance.

Orbit uncertainty has already been incorporated into sensor tasking. Ravago and Jones [5] studied risk-aware scheduling for large-constellation tracking, and Li et al. [6] considered orbital uncertainty in space-target sensor scheduling. The Heimdall framework of Dhingra et al. [7] evaluates posterior covariance from current orbits and planned measurements. It supports collection windows, frequencies, durations, and accuracy requirements at specified future times. Building on this idea, we impose a separate requirement on every observation: its posterior must satisfy covariance limits throughout a prescribed interval without further measurements. The same posterior then determines when the target requests its next visit. Observation stopping and revisit demand thus remain coupled through both uncertainty propagation and shared resource calendars.

Satellite scheduling methods also address visibility windows, attitude transitions, and resource limits [8]. Wu et al. [9] combined adaptive simulated annealing with dynamic task clustering. Galloua et al. [10] studied multilevel clustering and priority-driven prescheduling for Earth imaging. For space situational awareness, Shtofenmakher and Balakrishnan [1] jointly considered constellation tasking and observation sequencing, including incidental targets in the optical field. Shimane and Ho [11] formulated ground-based multisensor tasking as a job-shop scheduling problem with prescribed exposure durations. These studies inform the organization of shared observing opportunities. In our setting, however, common-field eligibility changes with satellite and service time, while each target's duration depends on its current prior and hold requirement. Resource reuse must therefore be checked over the actual service interval together with serial laser occupancy and independent target uncertainty.

Duration and service order are also coupled in radar scheduling. Yuan et al. [12] jointly allocated tracking sequences and dwell times for phased-array radar; Li et al. [13] developed adaptive dwell scheduling using three-way decisions. Repeated service introduces an additional timing constraint. Han et al. [14] studied continuous monitoring of moving targets by Earth-observation satellites, while Xing et al. [15] addressed replanning under changes including satellite failures and urgent tasks. For active–passive observation constellations, the remaining task is to coordinate measurement-free hold requirements, shared acquisition opportunities, and serial laser service within one schedule. Time saved on one visit is useful only when the resulting calendar supports timely initialization and sustained tracking of other targets.

We study the coupled choice of observation duration, revisit release, and resource allocation for noncooperative target clusters with coarse orbit priors and known associations. Covariance-aware tasking, adaptive dwell allocation, differential evolution, and strategy adaptation provide established foundations. Our contribution concerns their scheduling-specific formulation and coupling, rather than the independent invention of each component.

First, we formulate a per-visit measurement-free hold requirement. Each observation's posterior must satisfy the planning covariance limits throughout the prescribed interval

without credit from later measurements. This specializes future-accuracy tasking [7] to a visit-level condition that determines feasible durations from the current prior and geometry. The resulting posterior also determines the release of the next revisit request. The contribution is this explicit coupling of visit sufficiency and subsequent service demand; covariance propagation and adaptive dwell allocation themselves are established techniques.

Second, we construct resource-feasible calendars that combine this posterior-dependent service demand with time-dependent common-field acquisition and serial laser ranging. Shared pointing can reduce repeated acquisition, but does not provide parallel range measurements or merge target posteriors. A separate initialization-priority heuristic favors feasible first observations that can finish before the initialization deadline. It changes resource allocation rather than the estimator, and does not guarantee timely initialization for every target.

Third, we implement the coupled formulation through scheduling-specific random-key encoding, a covariance-guided feasible-plan constructor, and complete-calendar evaluation within H-SADE. Differential evolution and strategy adaptation are adopted from established methods [16,17]. The scheduling architecture connects these search components to the state-dependent candidate construction. The experiments compare construction policies under this common optimizer; they do not establish H-SADE as superior to alternative outer-search algorithms.

Final performance is measured by position and velocity errors against nonlinear truth trajectories. The experiments use eight synthetic orbital scenes, with common truth and standardized noise across methods. Rejected, unobserved, and late-initialized targets remain in every accuracy statistic. Single-seed experiments cover four planning thresholds and four measurement-noise pairs. Matched controls separate duration, initialization priority, and hold screening; five paired search seeds, fixed-duration scans, and budget tests assess sensitivity. Throughout, planning feasibility and actual-error compliance are reported separately.

2. Observation Scenario and Tracking-Accuracy Hold Model

2.1. Target Clusters and Observation Tasks

Let $\mathcal{S} = \{1, \ldots, N_S\}$ denote the observing satellites, $\mathcal{J} = \{1, \ldots, N_I\}$ the noncooperative targets, and $[0, T_f]$ the planning horizon. Coarse orbit priors are available without navigation information supplied by the targets. Target discovery and association, observer ephemerides, and attitude information are assumed known. The problem concerns joint observation scheduling under assumed measurement availability, rather than target search or identification.

During a joint observation, one satellite measures a target's angles and range with co-boresighted optical and laser payloads. The optical field may contain several objects, but targets without laser measurements do not count as jointly observed. A target cluster is a group of independent objects with potential shared pointing opportunities; common-field eligibility depends on the satellite and time. This relationship reduces repeated acquisition overhead without merging target estimates or covariances. Figure 1 illustrates the task and serial service sequence.

An observation is written as $a = (i, s, t_a^{\text{setup}}, t_a^{\text{cap}}, u_a, e_a)$, identifying the target, satellite, setup start, acquisition start, measurement start, and measurement end. Its measurement duration is $T_a = e_a - u_a$. Slewing and settling, acquisition, and measurement occupy the resources sequentially. A schedule π contains all observations and their execution order. Geometric visibility windows follow from the orbit and pointing constraints. Candidate observation intervals are then constructed within those windows using the current prior, hold requirement, and resource calendar.

https://doi.org/10.3390/aerospace13110987

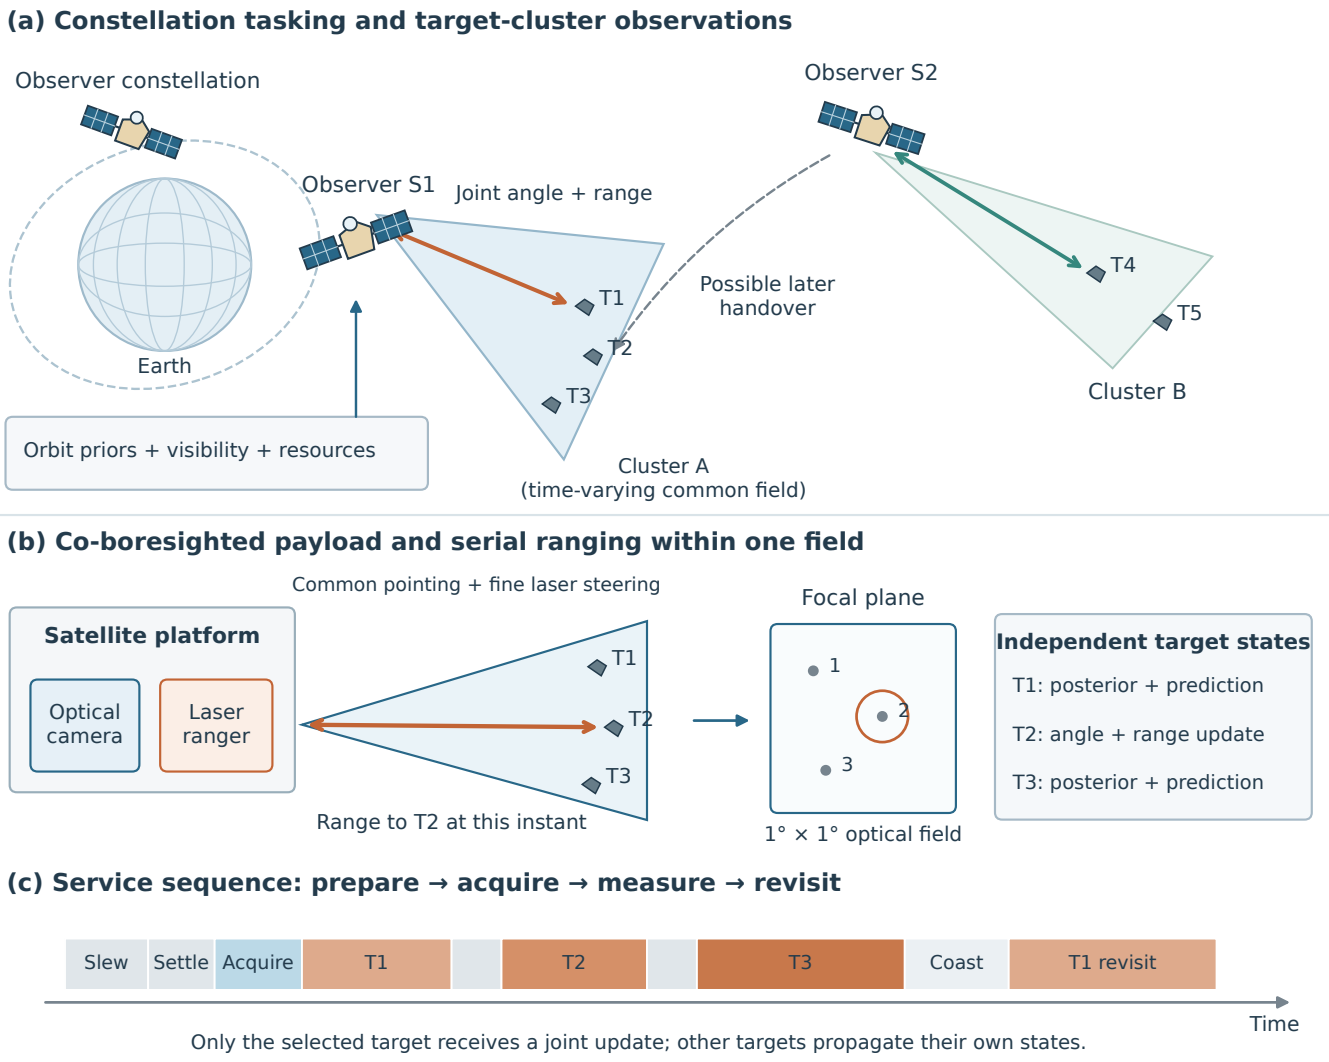

Figure 1. Joint observation and serial target service. Blue and green shading distinguishes the two cluster fields; colored bidirectional links indicate joint angle–range observations. Orange marks laser ranging, blue denotes optical acquisition, and gray arrows indicate handovers or time progression. (**a**) Satellites serve different target clusters; dashed links indicate possible later handovers. Geometry changes with time. (**b**) Co-boresighted optical and laser payloads range individual targets within a shared field. The focal-plane circle identifies the currently ranged target. Only that target receives a joint angle–range update; the others propagate independently. (**c**) Setup, acquisition, measurement, and revisits form a serial schedule. Light-gray blocks denote transition setup; redundant acquisition is omitted where common-field conditions allow. This conceptual diagram is not to scale and does not depict a simulated plan.

2.2. Orbit Propagation and Joint Angle–Range Measurements

The target state in an Earth-centered inertial frame is $\boldsymbol{x}_i = [\boldsymbol{r}_i^{\mathsf{T}}, \boldsymbol{v}_i^{\mathsf{T}}]^{\mathsf{T}}$, where $\boldsymbol{r}_i$ and $\boldsymbol{v}_i$ are position and velocity. Nominal orbits follow two-body dynamics:

$$\dot{\boldsymbol{r}}_i = \boldsymbol{v}_i, \qquad \dot{\boldsymbol{v}}_i = -\mu \frac{\boldsymbol{r}_i}{\|\boldsymbol{r}_i\|^3}. \tag{1}$$

Here, μ is Earth's gravitational parameter.

The reference model omits oblateness, atmospheric drag, and other perturbations. These can change orbit phase and orientation, range, line-of-sight rate, and visibility; the white acceleration noise is not a calibrated substitute for these forces. Appendix C separately examines J_2 propagation and fixed-calendar estimator mismatch, without claiming

validation of drag or all perturbations. Linearization along the nominal orbit gives the error-state matrix

$$A_i(t) = \begin{bmatrix} 0 & I_3 \\ G_i(t) & 0 \end{bmatrix}, \qquad G_i(t) = \mu\left(\frac{3\boldsymbol{r}_i\boldsymbol{r}_i^{\mathsf{T}}}{\|\boldsymbol{r}_i\|^5} - \frac{I_3}{\|\boldsymbol{r}_i\|^3}\right). \tag{2}$$

With time step Δt, the local gravity gradient is held constant within each step. A second-order transition approximation and white acceleration noise give

$$\begin{aligned} F_{i,k} &= I_6 + A_{i,k}\Delta t + \tfrac{1}{2}A_{i,k}^2(\Delta t)^2, \\ Q(\Delta t) &= q_a\begin{bmatrix} ((\Delta t)^3/3)I_3 & ((\Delta t)^2/2)I_3 \\ ((\Delta t)^2/2)I_3 & \Delta t I_3 \end{bmatrix}. \end{aligned} \tag{3}$$

The acceleration-noise spectral density q_a has units of $\mathrm{m}^2\,\mathrm{s}^{-3}$. The matrix Q in Equation (3) approximates a local constant-velocity error model. It is not the exact discretization of the system with a gravity gradient. For fixed A, the exact covariance is $Q_d = \int_0^{\Delta t} e^{A\tau} L q_a L^{\mathsf{T}} e^{A^{\mathsf{T}}\tau}\, d\tau$, where $L = [0_{3\times3}^{\mathsf{T}}, I_3^{\mathsf{T}}]^{\mathsf{T}}$. A block matrix exponential evaluates this integral [18]; Appendix A quantifies the local approximation error. The initial covariance is $P_{i,0} = \mathrm{diag}(\sigma_{r,i}^2 I_3, \sigma_{v,i}^2 I_3)$, with separately specified position and velocity prior standard deviations. Prediction and measurement updating use

$$\begin{aligned} P_{i,k}^- &= F_{i,k-1}P_{i,k-1}^+F_{i,k-1}^{\mathsf{T}} + Q(\Delta t), \\ K_{is,k} &= P_{i,k}^- C_{is,k}^{\mathsf{T}}\left(C_{is,k}P_{i,k}^- C_{is,k}^{\mathsf{T}} + R\right)^{-1}, \\ P_{i,k}^+ &= (I_6 - K_{is,k}C_{is,k})P_{i,k}^-(I_6 - K_{is,k}C_{is,k})^{\mathsf{T}} + K_{is,k}RK_{is,k}^{\mathsf{T}}. \end{aligned} \tag{4}$$

The superscripts $-$ and $+$ indicate states before and after a measurement update. Without measurements, only prediction is performed. The recursion follows the Kalman filtering framework [19]; Equation (4) uses the Joseph covariance form to reduce numerical loss of symmetry and positive semidefiniteness.

Define the observer–target range $\rho_{is} = \|\boldsymbol{r}_i - \boldsymbol{r}_s\|$ and unit line of sight $\boldsymbol{u}_{is} = (\boldsymbol{r}_i - \boldsymbol{r}_s)/\rho_{is}$. Let $\boldsymbol{e}_1, \boldsymbol{e}_2$ form an orthonormal basis in the plane perpendicular to that line. The two angular components and range have local measurement matrix and noise covariance

$$C_{is,k} = \begin{bmatrix} \boldsymbol{e}_1^{\mathsf{T}}/\rho_{is,k} & 0 \\ \boldsymbol{e}_2^{\mathsf{T}}/\rho_{is,k} & 0 \\ \boldsymbol{u}_{is,k}^{\mathsf{T}} & 0 \end{bmatrix}, \qquad R = \mathrm{diag}(\sigma_\theta^2, \sigma_\theta^2, \sigma_\rho^2). \tag{5}$$

Angle and range noise are independent in this model. Measurements occur from $u_a + \Delta t$ through e_a; no additional sample is taken at u_a. Each sample is preceded by prediction. After a satellite handover, the target's latest posterior is propagated to the new observation start. Neither reassignment nor common-field eligibility resets its prior.

2.3. Covariance Constraint over a Measurement-Free Interval

Observation duration depends on uncertainty growth after departure, as well as uncertainty at the last measurement. Partitioning the joint covariance by position and velocity gives the three-dimensional root-mean-square uncertainties and their normalized maximum:

$$\begin{aligned} e_{r,i}(t) &= \sqrt{\mathrm{tr}\, P_{rr,i}(t)}, \quad e_{v,i}(t) = \sqrt{\mathrm{tr}\, P_{vv,i}(t)}, \\ \kappa_i(t) &= \max\left\{\frac{e_{r,i}(t)}{\varepsilon_r}, \frac{e_{v,i}(t)}{\varepsilon_v}\right\}. \end{aligned} \tag{6}$$

The planning thresholds are ε_r for position and ε_v for velocity. The condition $\kappa_i(t) \le 1$ means that both modeled uncertainties meet their thresholds. Although the metrics use diagonal blocks, propagation retains the complete joint covariance, including position–velocity

cross terms. The normalized quantity is neither an actual-error exceedance probability nor the radius of a three-dimensional confidence sphere.

Even an unbiased and correctly calibrated covariance gives a second-moment description, not a bound on every realization. Bias and covariance inconsistency further weaken that interpretation. Accordingly, tracking accuracy in the scheduling constraint denotes a covariance-based uncertainty surrogate; actual-error compliance is evaluated independently.

For measurement start u, duration T, and measurement-free hold duration H, subsequent measurements are suppressed when propagating the posterior. The independent hold constraint is defined by

$$C_{is}(u, T; H) = \max_{\delta \in \{0, \Delta t, \ldots, H\}} \kappa_i^{\text{coast}}\big(u + T + \delta \mid P_i^+(u + T)\big). \tag{7}$$

A feasible observation satisfies $C_{is}(u, T; H) \leq 1$. Its posterior must support H seconds of propagation without credit from any later measurement, even when the schedule revisits earlier. This specializes future-accuracy tasking [7] to an independent requirement on each visit. When $u + T + H > T_f$, the check still covers the entire hold interval; actual measurements remain within the planning horizon.

For the current prior, satellite, start time, and resource limit, the discrete feasible durations and shortest sufficient duration are

$$\begin{aligned} \mathcal{T}_{is}(u; H) &= \{T = n\Delta t > 0 : T \leq T_{is}^{\max}(u),\ C_{is}(u, T; H) \leq 1\}, \\ T_{is}^{\min}(u; H) &= \min \mathcal{T}_{is}(u; H). \end{aligned} \tag{8}$$

The upper bound $T_{is}^{\max}(u)$ combines the continuous laser limit, remaining budget, remaining uninterrupted visibility, and remaining planning time. An empty feasible set rejects the combination. Here, "sufficient" means satisfying the modeled hold condition; it does not require the covariance derivative to approach zero. The shortest duration permits departure from one target but need not optimize the complete service sequence.

The hold duration H specifies the interval over which a completed visit must support covariance-based prediction without another measurement. It is a mission input, distinct from the revisit interval realized by the schedule. We use $H = 200$ s as the prescribed study configuration, without identifying it with a universal operational requirement or a value uniquely implied by the constellation geometry. An operational application should specify H jointly with its planning thresholds, prior uncertainty, and sensing and resource models. Changing H requires recomputing the feasible durations and optimizing new calendars; it cannot be implemented by relabeling the outcomes of an unchanged plan. For a fixed candidate and propagation grid, extending the checked interval cannot decrease the maximum in Equation (7), although actual-error performance after rescheduling need not vary monotonically. Combinations with no feasible duration are rejected.

2.4. Resource Constraints and Sustained-Tracking Objective

Range, off-nadir angle, Earth occultation, and line-of-sight rate constraints must hold throughout acquisition and measurement. Setup, acquisition, and measurement cannot overlap on one satellite. Observations of one target are also sequential, preserving an unambiguous posterior history. For a line-of-sight change $\Delta\theta$, the scheduler approximates a slew with zero initial and final angular velocities:

$$T^{\text{slew}} = \begin{cases} 2\sqrt{\Delta\theta / a_{\max}}, & \Delta\theta \leq \omega_{\max}^2 / a_{\max}, \\ \Delta\theta / \omega_{\max} + \omega_{\max} / a_{\max}, & \Delta\theta > \omega_{\max}^2 / a_{\max}. \end{cases} \tag{9}$$

The angular-speed and angular-acceleration limits are $\omega_{\max}$ and $a_{\max}$; settling time is added after each slew. A changed candidate start requires a new angle and timing

check. This scheduling approximation does not replace a full attitude-control model for moving-line-of-sight interception and dual-actuator pointing.

Let $T_L^{\max}$ be the continuous laser limit and B_s the cumulative budget of satellite s. The observation constraints are

$$0 < T_a \le T_L^{\max}, \qquad \sum_{a \in \pi : s_a = s} T_a \le B_s, \qquad C_{i_a s_a}(u_a, T_a; H) \le 1. \tag{10}$$

The planning surrogate accounts for both initialization and subsequent maintenance. Let t_{ini} be the common initialization deadline, $\mathcal{K} = \{0, \Delta t, \dots, T_f\}$ the evaluation grid, and π_i the subsequence serving target i. Define

$$\begin{aligned} A_i(\pi) &= \mathbf{1}\{\exists a \in \pi_i : e_a \le t_{\text{ini}},\ \kappa_i(e_a; \pi) \le 1\}, \\ y_i(\pi) &= A_i(\pi)\mathbf{1}\Big\{\max_{t \in [t_{\text{ini}}, T_f] \cap \mathcal{K}} \kappa_i(t; \pi) \le 1\Big\}, \\ N_{\text{plan}}(\pi) &= \sum_{i=1}^{N_I} y_i(\pi), \qquad \eta_{\text{plan}}(\pi) = \frac{N_{\text{plan}}(\pi)}{N_I} \times 100\%. \end{aligned} \tag{11}$$

Initialization requires a complete observation ending before the deadline. Rejected and uninitialized targets remain in the denominator. Any threshold violation during maintenance disqualifies sustained compliance over that interval, even if accuracy later recovers. The check uses updated covariance at measurement times and predicted covariance elsewhere. Thus, η_{plan} describes discrete model feasibility, not the final tracking accuracy or actual-error compliance. Section 4 computes those outcomes independently from nonlinear truth errors.

Search first maximizes the number of targets with sustained model feasibility, then compares covariance and resource costs through the lexicographic objective

$$\boldsymbol{J}(\pi) = \big(-N_{\text{plan}},\ \kappa_{\max},\ \bar{\kappa},\ T_{\text{occ}},\ T_L,\ J_{\text{dist}}\big). \tag{12}$$

Both $\kappa_{\max}$ and $\bar{\kappa}$ cover every target and the common maintenance grid. Total occupied time T_{occ} includes slewing, settling, acquisition, and measurement, whereas T_L includes only laser measurement. The replanning disturbance J_{dist} sums cancellations, satellite changes, and total time shifts divided by T_f; it is zero for initial planning. Lexicographic comparison prevents time savings from compensating for a higher-priority feasibility loss.

3. Hierarchical Resource Scheduling Under the Hold Constraint

3.1. Framework and Candidate Observation Intervals

Figure 2 separates upper-level encoded search from lower-level schedule construction. Covariance propagation maps target posteriors to feasible durations and revisit requests. The constructor combines these requests with common-field conditions and resource calendars; the search updates admission, order, and satellite preferences using complete-plan evaluations. Each observation changes its target's posterior and revisit timing, as well as the intervals available to other targets.

Initialization requests are released at the start of planning. The earliest setup interval follows from the target's previous observation, request release, and candidate satellite's calendar tail. Slew and acquisition times are checked iteratively. The latest target posterior is then propagated to the actual measurement start, and joint measurements are added sequentially. A hold-endpoint check provides a necessary-condition screen before the full interval in Equation (7) is evaluated. The first feasible discrete duration becomes the shortest candidate.

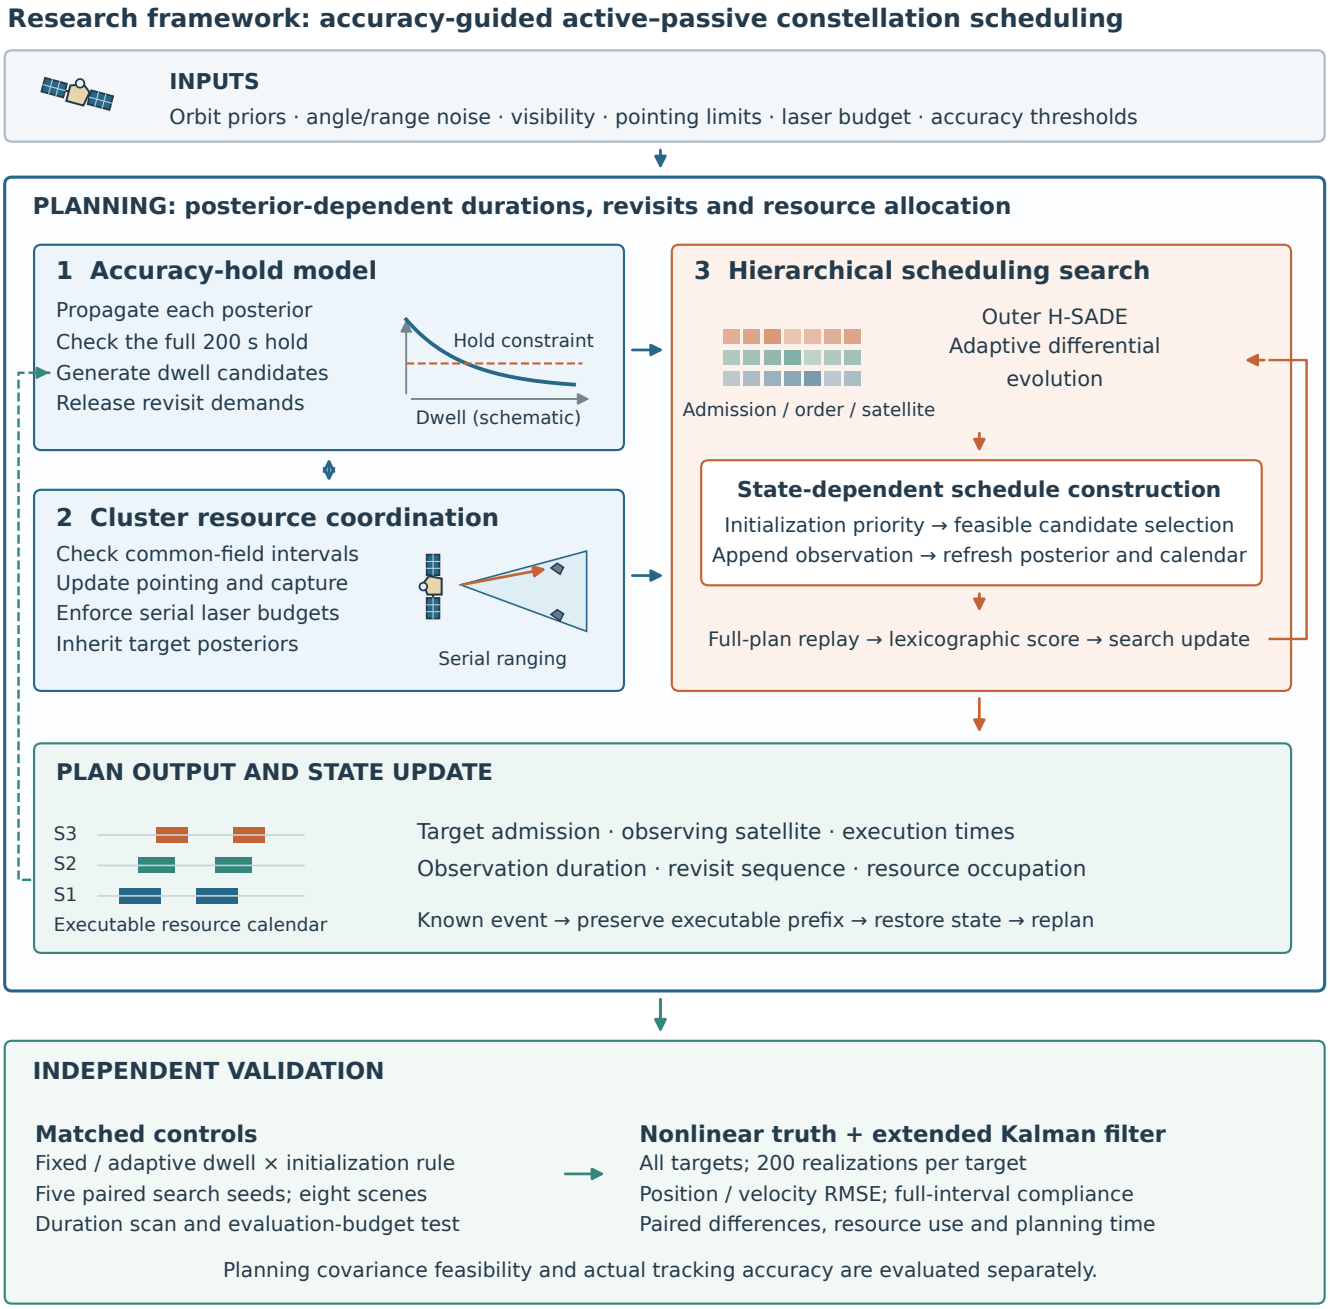

Figure 2. Scheduling framework and independent accuracy evaluation. Blue denotes the hold model and resource coordination, orange denotes scheduling search, and green denotes plan output and validation. Solid arrows indicate information flow and search updates. The planning stage combines the hold model, cluster resource coordination, and hierarchical self-adaptive differential evolution (H-SADE). Initialization priority is an optional construction rule. Nonlinear truth and an extended Kalman filter independently evaluate the output plan through root-mean-square error (RMSE), compliance, resource use, and computation time. Actual errors are not fed back into the search objective. Green dashed arrows show the interface for replanning after known events. Satellite, curve, and calendar icons are conceptual illustrations.

Besides the shortest feasible visit, the adaptive constructor retains one extended candidate. In the reference configuration its duration in grid steps is $n_{\mathrm{ext}} = \min\{n_{\max}, \lceil 1.5 n_{\min} \rceil\}$. The bound $n_{\max}$ combines the continuous-laser limit, remaining budget, and remaining horizon. This fixed heuristic gives a limited alternative between early departure and additional information; it is not an optimized multiplier. Every candidate must pass the full

hold and uninterrupted-visibility checks. A visibility boundary does not create an extra truncated candidate. Fixed-duration controls retain their prescribed durations. The local multiplier scan in Appendix C changes only this adaptive candidate setting.

After an observation, the next request is released when measurement-free propagation of its posterior first reaches the warning threshold κ_{w}:

$$\begin{aligned} t_i^{\mathrm{hit}} &= \min\{t \in [e_a, T_f] \cap \mathcal{K} : \kappa_i^{\mathrm{coast}}(t \mid P_i^{+}(e_a)) \geq \kappa_{\mathrm{w}}\}, \\ t_i^{\mathrm{rel}} &= \max\{e_a + \Delta t, t_i^{\mathrm{hit}}\}. \end{aligned} \tag{13}$$

We set $\kappa_{\mathrm{w}} = 0.8$ in the original reference configuration to release a request before the normalized uncertainty reaches the planning limit of 1. The resulting margin is in the covariance metric, not a guaranteed time allowance for queuing or transitions. For a fixed posterior propagation, a lower threshold can release the request earlier, but may increase competition for service; the effect on complete-plan performance depends on the resulting calendar. If the threshold is already reached at the observation end, the request is released at the next grid point. An empty set gives an infinite release time, so no visit is forced before T_f without a threshold crossing. Equation (11) checks the resulting complete plan. The independent hold duration H constrains visit quality, whereas the realized revisit interval emerges from scheduling. Appendix New Search Seeds for a Selected Warning Variant separately evaluates a warning-0.7 variant selected during revision; all previously reported reference results retain 0.8.

Figure 3 is a planning diagnostic using three independent two-target examples. For target 0, the unscaled initial per-axis position/velocity standard deviations are 888.205 m/96.371 m/s, 1170.569 m/100.305 m/s, and 881.129 m/83.026 m/s for seeds 4101–4103, respectively. These values are specified at initialization and propagated to the observation start. The corresponding observer–target ranges are 164.986, 169.101, and 193.055 km.

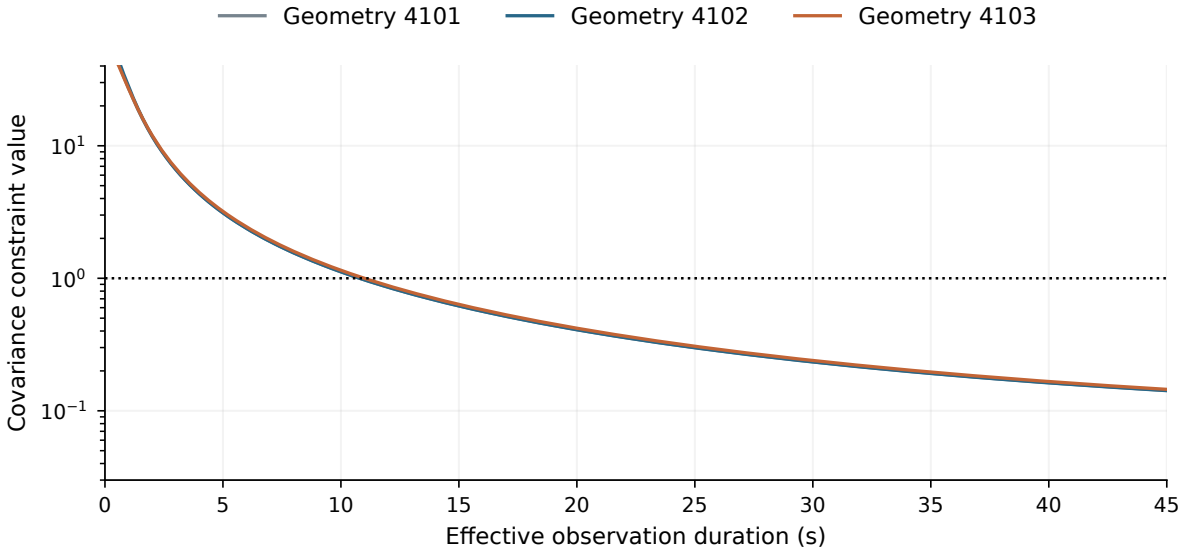

Figure 3. Hold feasibility versus observation duration. The ordinate is Equation (7), with feasibility at or below 1. Geometry 4101–4103 denotes example-generation seeds, separate from main scenes 2101–2108. Each example has two targets; curves use target index 0, observer indices 8, 23, and 26, respectively (zero-based), and measurement start 5 s. Prior scale multiplies both initial position and velocity standard deviations: scale 1 leaves them unchanged, while covariance scales with its square. P1 thresholds, 100 µrad/50 m noise, and $H = 200$ s apply.

3.2. Common-Field Eligibility and Serial Laser Service

For adjacent services of distinct targets i, j by satellite s, a shared boresight is constructed from their lines of sight. Over the checked interval $[a, b]$,

$$\boldsymbol{b}_{ij,s}(t) = \frac{\boldsymbol{u}_{is}(t) + \boldsymbol{u}_{js}(t)}{\|\boldsymbol{u}_{is}(t) + \boldsymbol{u}_{js}(t)\|}, \qquad \max_{\substack{t\in[a,b]\cap\mathcal{K} \\ \ell\in\{i,j\}}} \arccos\left(\boldsymbol{u}_{\ell s}^{\mathsf{T}}(t)\boldsymbol{b}_{ij,s}(t)\right) \le \theta_{\mathrm{FOV}}/2. \tag{14}$$

The parameter θ_{FOV} is the angular side length of the square field. Its inscribed circle provides a sufficient common-field test. Both targets must also satisfy range, occultation, and pointing constraints throughout the interval. Proximity at the midpoint alone is insufficient. A degenerate sum of unit vectors does not define a shared boresight.

Let e_i be the preceding observation end, and t_j^{cap} and e_j the next acquisition start and observation end. The next target's 5 s acquisition is waived only if $t_j^{\mathrm{cap}} - e_i \le 10$ s and Equation (14) holds over $[e_i, e_j]$. Slewing, settling, and target-specific measurements remain on the calendar. For a service sequence σ_c on one satellite,

$$T_L(c) = \sum_{a\in\sigma_c} (e_a - u_a), \qquad T_{\mathrm{occ}}(c) = \sum_{a\in\sigma_c} (e_a - t_a^{\mathrm{setup}}). \tag{15}$$

Idle waiting is excluded from T_{occ}. Shared pointing assumes that the optical axis follows the angular bisector while laser fine pointing switches between targets. Cluster service time therefore sums the individual measurement and transition times; it is not the longest member duration. Changing the service order or satellite requires recalculating start times, priors, durations, and common-field eligibility.

The implementation forms local service chains from adjacent target pairs. It does not treat transitive pairwise compatibility as proof that an entire cluster shares one field. Nor must one satellite serve every member. When acquisition cannot be reused, a member is scheduled independently; any later satellite handover preserves that target's posterior.

3.3. Random-Key Encoding and Covariance-Guided Construction

The upper level represents discrete decisions with continuous target-level random keys:

$$\boldsymbol{z} = \left[\boldsymbol{z}^{\mathrm{adm}}, \boldsymbol{z}^{\mathrm{ord}}, \boldsymbol{z}^{\mathrm{sat}}\right] \in [0,1]^{3N_I}. \tag{16}$$

An admission key $z_i^{\mathrm{adm}} \ge 0.5$ permits an attempt to serve target i. The order key z_i^{ord} enters request priority, and z_i^{sat} selects from the geometrically visible satellites sorted by identifier. Keys do not directly fix observation times. The constructor determines those times from the current state, avoiding treatment of posterior-dependent durations as static task attributes.

The constructor maintains each satellite's calendar tail and remaining laser budget, and each target's latest posterior and request release time. New observations are appended after both the target's previous service and the assigned satellite's calendar tail. Duplicate actions with the same target, satellite, and measurement interval are retained only once. A feasible candidate b receives the local score

$$S_b = \frac{[\kappa_i^0(e_b) - \kappa_i^+(e_b)]_+}{\max\{1, (e_b - e_{\mathrm{prev}})/\Delta t\}} - 0.11\{s_b \ne s_i^{\mathrm{pref}}\} - 10^{-4}\frac{T_b}{\Delta t}. \tag{17}$$

Here, $\kappa_i^0(e_b)$ and $\kappa_i^+(e_b)$ are endpoint covariance metrics without and with the candidate observation, and $[x]_+ = \max(x, 0)$. The preceding same-satellite observation ends at e_{prev}; for the first task, the current acquisition start is used instead. The preferred satellite is s_i^{pref}.

All three coefficients are fixed in advance and only rank feasible actions. Equation (7) and resource checks enforce feasibility separately.

Following the opportunity-loss idea used in insertion heuristics [20], a request with at least two distinct actions has regret $\Gamma_i = S_i^{(1)} - S_i^{(2)}$. A request with only one action receives a single-candidate priority flag. Requests are ordered first by the 30 s block containing their highest-scoring action's start, then by that flag and the combined priority

$$D_i = 0.5 z_i^{\mathrm{ord}} + 0.3 \frac{\min\{\kappa_i^0(u_i), 3\}}{3} + 0.2 \min\{\Gamma_i, 1\}. \tag{18}$$

The time u_i is the preferred action's measurement start. The fixed weights 0.5, 0.3, and 0.2 combine the order key, capped normalized uncertainty, and capped opportunity loss after the preceding priority rules. They rank feasible requests and do not relax the hold or resource constraints. These weights are retained as reference values following comparative evaluation of alternative parameter settings (Appendix C). In practical applications, the parameters should be tuned to the target population, observation geometry, measurement quality, and available resources. Each iteration commits the highest-scoring action of the highest-priority request and updates its posterior, resources, and revisit release. Cached duration calculations require identical physical and covariance inputs.

This construction places most geometry and covariance checks below the outer search, allowing complete schedules to be compared within a small evaluation budget. Its search space is nevertheless restricted by causal appending and local priority rules. It does not fill arbitrary calendar gaps, and a high-scoring later action can displace an early initialization opportunity. The score is therefore a construction heuristic, not a capacity bound or a guarantee of initialization.

3.4. Hierarchical Self-Adaptive Differential Evolution

The three random-key groups and covariance-guided decoder are combined with differential evolution [16] and strategy adaptation [17]. We refer to the resulting implementation as hierarchical self-adaptive differential evolution (H-SADE). Its scheduling-specific elements are the encoding, feasible-plan constructor, and complete-calendar evaluation. The upper level uses three established mutation strategies:

$$\begin{aligned} \boldsymbol{v}_j^{(1)} &= \boldsymbol{z}_{r_1} + F(\boldsymbol{z}_{r_2} - \boldsymbol{z}_{r_3}), \\ \boldsymbol{v}_j^{(2)} &= \boldsymbol{z}_j + F(\boldsymbol{z}_{\mathrm{best}} - \boldsymbol{z}_j) + F(\boldsymbol{z}_{r_1} - \boldsymbol{z}_{r_2}), \\ \boldsymbol{v}_j^{(3)} &= \boldsymbol{z}_{\mathrm{best}} + F(\boldsymbol{z}_{r_1} - \boldsymbol{z}_{r_2} + \boldsymbol{z}_{r_3} - \boldsymbol{z}_{r_4}). \end{aligned} \tag{19}$$

The current individual is j; r_1 to r_4 are distinct randomly selected indices. The best encoding is $\boldsymbol{z}_{\mathrm{best}}$, and F is the scale factor. Binomial crossover takes at least one component from the mutant, and values outside $[0, 1]$ are clipped. Construction and complete-plan replay evaluate each trial with Equation (12). Replacement requires a strict lexicographic improvement.

For strategy q, success and attempt counts n_q^{s} and n_q^{t} both start at 1. The selection probability is

$$p_q = \frac{n_q^{\mathrm{s}} / n_q^{\mathrm{t}}}{\sum_{h=1}^{3} n_h^{\mathrm{s}} / n_h^{\mathrm{t}}}. \tag{20}$$

Parameter memories start at $\bar{F} = 0.55$ and $\overline{CR} = 0.75$. Normal draws with standard deviations 0.12 and 0.10 generate F and the crossover probability CR, respectively. They are clipped to $[0.15, 0.95]$ and $[0.10, 0.98]$. Successful replacement updates the memories by $\bar{F} \leftarrow 0.9\bar{F} + 0.1F$ and $\overline{CR} \leftarrow 0.9\overline{CR} + 0.1CR$. The population contains 12 individuals. Initial admission keys lie in $[0.5, 1]$, including one guided individual admitting every target; evolution may subsequently lower admission keys.

A complete-plan evaluation includes decoding and independent replay from the common initial state. Replay recomputes covariance metrics and resource use rather than accepting cached construction labels as final evidence of compliance. Initial and evolved individuals both count toward the budget. The best available plan is returned when that budget is exhausted. Adaptation in Equations (19) and (20) changes neither feasibility constraints nor the need to compare alternative outer search algorithms.

3.5. Replanning Interface and Execution Procedure

At a replanning event, a frozen interval is a short commitment window, distinct from the covariance hold duration. The constructor accepts a retained prefix selected by its caller. In the event-replanning rule, an observation is retained when its setup starts strictly before the event time plus the freeze duration, provided it remains executable. Its complete setup, acquisition, and measurement interval is retained even if it extends beyond that boundary. Completed visits are preserved; a failed satellite's visits ending after failure are removed. Replaying the retained prefix reconstructs planning posteriors, occupied time, and remaining laser budgets before appending actions. A committed future observation is a prediction, not an acquired measurement. Historical 0 s and 20 s windows are examples of caller policies, not a universal parameter of the reference experiments. A warm start retains the previous encoding as one initial individual while rebuilding the remaining calendar.

Algorithm 1 summarizes the execution procedure.

Algorithm 1. Hierarchical observation scheduling with a covariance hold constraint.

1. Read observer and target orbits, target priors, accuracy thresholds, hold duration, resource budgets, and the executed prefix. Compute nominal trajectories and geometric visibility windows.
2. Initialize 12 individuals containing the three random-key groups. If a retained prefix exists, restore target posteriors and resource states, and include a warm-start individual when needed.
3. Release initialization or revisit requests for the individual being evaluated. Construct shortest and extended candidates at their actual start times; check geometry, common-field eligibility, laser budgets, and hold feasibility.
4. If initialization priority is enabled, screen candidates using Equation (21). Select an observation by time block, single-candidate flag, and covariance–regret priority. Update posteriors and calendars until no further request can be appended.
5. Replay the complete schedule independently. Evaluate sustained model compliance, covariance metrics, and resource costs, then update the best plan.
6. Generate trial encodings using Equations (19) and (20), and repeat Steps 3–5 within the remaining budget. Return the best plan; return to Step 2 after a new known resource event.

This procedure provides a prefix-restoration and replanning interface for known resource events. Section 4 evaluates posterior, observation-window, and calendar updates in fixed scenarios. It does not test repeated stochastic disruptions. Detection of unknown tracking loss and closed-loop intersatellite communication remain outside the validation scope.

3.6. Initialization-Priority Rule

The base constructor chooses each target's preferred action by information gain before ordering targets by that action's time block. A high-scoring later action can suppress an earlier candidate that would finish before the initialization deadline. An additional priority

level addresses this mismatch. Let $\mathcal{B}_i$ contain the current feasible actions for target i, and let n_i be its number of scheduled observations. Define

$$\mathcal{B}_i^{\mathrm{ini}} = \{b \in \mathcal{B}_i : e_b \le t_{\mathrm{ini}}\}, \qquad \widetilde{\mathcal{B}}_i = \begin{cases} \mathcal{B}_i^{\mathrm{ini}}, & n_i = 0,\ \mathcal{B}_i^{\mathrm{ini}} \neq \emptyset, \\ \mathcal{B}_i, & \text{otherwise.} \end{cases} \tag{21}$$

Scores and regret are computed within $\widetilde{\mathcal{B}}_i$. Uninitialized targets with an on-time candidate receive priority; within each level, the 30 s blocks and Equation (18) still apply. Later observations remain available when no candidate meets the deadline, allowing overdue targets to reduce subsequent errors. The rule changes neither the measurement model nor the hold constraint or laser budget. It cannot ensure on-time initialization of every target.

Section 4 denotes the base constructor by AD and the initialization-priority variant by IP-AD. Both use the same H-SADE outer search. The reference comparison uses 48 complete-plan evaluations; Sections 4.9 and 4.10 specify the additional seeds and budgets. Because initialization priority changes later calendars and revisit opportunities, its benefit must be evaluated with actual-error simulations rather than inferred from planning feasibility.

4. Nonlinear Tracking-Accuracy Evaluation

4.1. Scenarios and Paired Evaluation Protocol

The reference architecture is Walker-Delta 72/8/1 at 800 km and 60° inclination. The eight original scenes contain 138 distinct targets in total, with 12, 18, or 24 in each independently scheduled scene, not 138 simultaneous targets. These illustrative populations and the constellation were not selected from an operational demand model or optimized as a constellation design. Two synthetic families near selected orbital planes deliberately exercise persistent local service and shared acquisition. They are not a representative debris-catalog sample. The horizon is 0–900 s, initialization ends at 120 s, and actual errors are evaluated every 0.2 s over 120–900 s. Appendix B specifies generation and geometry; Appendix C adds separate nested workloads and an overload boundary. The reference hold is 200 s, while actual revisit intervals depend on posterior propagation and resources.

Tables 1 and 2 summarize the common parameters and scene-specific settings.

The reference experiment compares three constructors. F20 requires exactly 20 s of measurement and rejects a candidate if its window or resources are insufficient. It does not use Equation (7) to accept or reject candidates. AD uses Equations (7) and (8) to generate shortest and extended durations. IP-AD adds Equation (21). All three share orbits, geometry, budgets, encoding, initial populations, outer search, and revisit warnings. Each method is optimized independently with 48 complete-plan evaluations before nonlinear tracking simulation. The 200 s independent hold, the 0–900 s planning horizon, and the 120–900 s error-evaluation interval have distinct roles. The comparison concerns complete construction policies and their resulting calendars, rather than attributing every difference to one stopping condition.

Truth states use fourth-order Runge–Kutta integration of Cartesian two-body dynamics, with process increments drawn from the discrete Q in Equation (3). Measurements use nonlinear tangent-plane angles and true range. A relinearized extended Kalman filter (EKF) applies Joseph updates and wraps angular innovations to $[-\pi, \pi)$. Observer ephemerides and target associations are known. Estimates and covariances remain continuous across handovers. Separate random streams generate each target's initial errors, process noise, and measurement noise. Measurement-noise streams advance even when no observation is scheduled, preventing schedule-dependent changes to the truth. Corresponding truth-array hashes at 120 s and 900 s are checked across methods and parameter settings. Whenever noise standard deviations change, the planner, measurement generator, and EKF

use the same updated covariance. Standardized draws are shared, while physical noise amplitudes differ.

Table 1. Common model and computation parameters.

Item	Setting
Dynamics and Earth	Two-body; $\mu = 3.986004418 \times 10^{14}$ m^3/s^2; $R_E = 6{,}378{,}137$ m
Target orbits	Altitude 610–690 km; inclination and node offsets $\pm 0.025^\circ$ from the selected observer; initial phase offset -1.1° to -0.55°
Prior standard deviations	Position 800–1200 m and velocity 70–110 m/s, independently uniform by target; equal axis variances and zero initial cross terms
Joint measurements	Reference: 100 µrad, 50 m; noise pairs: 50 µrad/30 m, 100 µrad/60 m, 150 µrad/90 m, 200 µrad/120 m; period 0.2 s; independent components
Geometric availability	Range $\leq$ 2000 km; off-nadir angle $\leq 60^\circ$; Earth-occultation check; optical field $1^\circ \times 1^\circ$
Attitude and acquisition Laser limits	Maximum rate 5°/s; acceleration 2.5°/s^2; acquisition 5 s; settling 2 s Continuous measurement $\leq$ 60 s; satellite budgets in Table 2
Planning thresholds	P1: 500 m/25 m/s; P2: 1000 m/50 m/s; P3: 1500 m/75 m/s; P4: 2000 m/100 m/s
Hold and revisit	Independent measurement-free hold 200 s at every level; normalized warning threshold 0.8
Outer search	H-SADE; population 12; reference experiments: seed 3101, 48 evaluations; matched seeds 3101–3105; budgets 24/48/96 in two scenes (Sections 4.8–4.10)
Actual-error simulation	200 realizations per target; shared initial errors, process noise, and standardized measurement noise across methods, seeds, and budgets; all targets retained

Table 2. Target counts and resource settings in the eight scenes.

Scene	Targets	Laser Budget per Satellite (s)	q_a (m^2 s^{-3})
2101	12	75	10^{-4}
2102	12	120	10^{-4}
2103	18	120	10^{-4}
2104	18	180	10^{-4}
2105	24	180	10^{-4}
2106	24	240	10^{-4}
2107	12	120	5×10^{-4}
2108	18	180	5×10^{-4}

The P1–P4 planning sweep uses reference noise of 100 µrad and 50 m. Crossing four additional noise pairs, M1–M4, with P1–P4 produces 16 further settings. Across all 20 settings, every one of the 138 targets has 200 realizations. Three methods produce 1,656,000 estimated trajectories sharing 27,600 truth trajectories. All 480 schedules and 8280 target–method result chains are recomputed. Sections 4.3–4.5 use reference P1; Sections 4.6 and 4.7 report the planning and crossed sweeps. The 200 repetitions represent random initial errors and noise within each scene, not independent orbital scenarios. Rejected, unobserved, and late-initialized targets propagate to 900 s and remain in the statistics.

The 2000 km bound is a geometric filter, and 5 Hz is an assumed measurement rate. Illumination, reflectance, laser-return link budgets, observer orbit errors, and acquisition failures are not modeled. Scheduled measurements are assumed available. The results therefore evaluate scheduling and estimation conditional on these availability assumptions.

4.2. Metrics Computed from Actual Tracking Errors

For realization m of target i, define the position error $\boldsymbol{e}_{r,i}^{(m)}(t) = \boldsymbol{r}_{i,\text{true}}^{(m)}(t) - \hat{\boldsymbol{r}}_i^{(m)}(t)$ and the corresponding velocity error $\boldsymbol{e}_{v,i}^{(m)}(t)$. At each evaluation time, the all-target root-mean-square errors (RMSEs) are

$$\begin{aligned} E_r(t) &= \sqrt{\frac{1}{N_I M}\sum_{i=1}^{N_I}\sum_{m=1}^{M}\|\boldsymbol{e}_{r,i}^{(m)}(t)\|^2}, \\ E_v(t) &= \sqrt{\frac{1}{N_I M}\sum_{i=1}^{N_I}\sum_{m=1}^{M}\|\boldsymbol{e}_{v,i}^{(m)}(t)\|^2}, \qquad M = 200. \end{aligned} \tag{22}$$

Let $\mathcal{K}_{\text{eval}} = [120, 900] \cap \mathcal{K}$. Averaging squared errors over time before taking the square root gives the primary scene-level metrics:

$$\overline{E}_r = \sqrt{\frac{1}{|\mathcal{K}_{\text{eval}}|}\sum_{t\in\mathcal{K}_{\text{eval}}} E_r^2(t)}, \qquad \overline{E}_v = \sqrt{\frac{1}{|\mathcal{K}_{\text{eval}}|}\sum_{t\in\mathcal{K}_{\text{eval}}} E_v^2(t)}. \tag{23}$$

These errors are reported in m and m/s, without conversion from covariance metrics. A scene mean in the tables is the equally weighted average of eight values from Equation (23). It is neither an RMSE pooled over all scenes nor an average of target-level RMSEs. Terminal errors are calculated separately from Equation (22) at 900 s. A lower instantaneous or terminal error does not imply compliance throughout the interval.

We also report full-interval actual-error compliance at position and velocity acceptance thresholds:

$$\eta_e(d_r, d_v) = \frac{100\%}{N_I M}\sum_{i=1}^{N_I}\sum_{m=1}^{M}\mathbf{1}\left\{\begin{array}{l}\max\limits_{t\in\mathcal{K}_{\text{eval}}}\|\boldsymbol{e}_{r,i}^{(m)}(t)\| \le d_r, \\ \max\limits_{t\in\mathcal{K}_{\text{eval}}}\|\boldsymbol{e}_{v,i}^{(m)}(t)\| \le d_v\end{array}\right\}. \tag{24}$$

A target realization is compliant only if both error limits hold at every evaluation grid point. This statistic uses neither the initialization labels nor the covariance test in Equation (11), and excludes no unscheduled target. Acceptance thresholds (d_r, d_v) differ in purpose from planning thresholds $(\varepsilon_r, \varepsilon_v)$. The common acceptance baseline is 500 m and 25 m/s, while planning thresholds vary across P1–P4. Section 4.5 changes only acceptance thresholds for fixed P1 schedules. Sections 4.6 and 4.7 change planning thresholds or noise and optimize new schedules.

Table 3 defines the planning and evaluation terms.

Table 3. Definitions used in planning and evaluation.

Term	Meaning
Planning thresholds	Limits on $\sqrt{\text{tr}\,P_{rr}}$ and $\sqrt{\text{tr}\,P_{vv}}$, not realized errors (Equation (6)).
Hold feasibility	Each posterior passes the full measurement-free grid without credit from later measurements (Equation (7)).
Revisit warning	A request trigger, 0.8 in the reference configuration, not guaranteed service before violation (Equation (13)).
Sustained planning feasibility	Targets initialized on time and within modeled covariance limits throughout maintenance (Equation (11)).
Actual-error thresholds	Joint limits on nonlinear truth–estimate error norms, reference 500 m/25 m/s (Equation (24)).
Full-interval compliance	Target–realization pairs meeting both limits at every 120–900 s evaluation point; denominator $N_I \times 200$ in each scene.
RMSE	Squared errors pooled within each scene across targets, realizations, and evaluation times; scene RMSEs are then averaged equally.

An improvement in RMSE does not establish full-interval compliance. Changing planning thresholds requires new schedules; changing acceptance thresholds alone reclassifies fixed-plan errors.

4.3. Reference Accuracy and Resource Use Across All Targets

The reference condition uses planning thresholds of 500 m and 25 m/s and measurement standard deviations of 100 µrad and 50 m. Table 4 reports eight-scene means; Table 5 and Figure 4 show scene-level variation. Every RMSE includes all targets and 200 realizations. Large errors combine insufficient observations with nonlinear estimation bias; their separate contributions have not been quantified.

Table 4. Actual tracking errors and laser use, equally weighted over eight scenes.

Method	$\bar{E}_r$ (m)	$\bar{E}_v$ (m/s)	$\eta_e(500, 25)$ (%)	T_L (s)
F20	51,655.0	108.855	19.16	315.0
AD	20,844.6	80.268	29.79	177.2
IP-AD	16,847.7	57.709	31.19	183.5

Table 5. All-target accuracy by scene. Each method has position RMSE in m and velocity RMSE in m/s.

Scene	F20 pos.	F20 vel.	AD pos.	AD vel.	IP-AD pos.	IP-AD vel.
2101	61,095.5	118.51	24,214.8	51.76	621.0	4.06
2102	53,849.3	101.49	54,157.4	318.72	929.3	4.02
2103	52,350.1	98.21	4295.5	22.42	19,865.0	54.57
2104	50,813.6	108.29	10,896.0	43.38	17,726.2	59.39
2105	49,695.6	96.21	18,406.1	59.61	17,701.5	57.48
2106	56,758.8	180.75	50,769.0	122.23	73,658.3	257.79
2107	38,053.7	71.37	369.5	3.04	369.5	3.04
2108	50,623.3	96.01	3648.1	20.98	3910.9	21.31

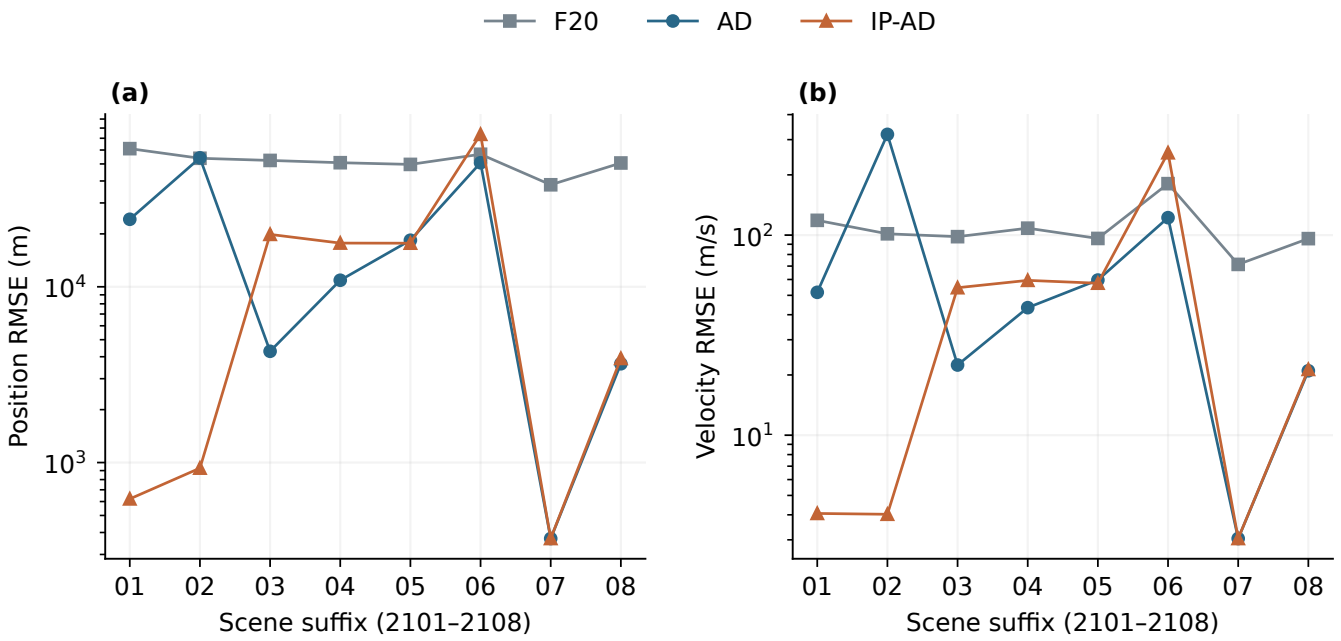

Figure 4. All-target RMSE by scene: (**a**) position and (**b**) velocity. Both ordinates are logarithmic. Suffixes 01–08 denote scenes 2101–2108. Equation (23) includes every 0.2 s grid point over 120–900 s and 200 realizations per target, without covariance-based eligibility filtering.

Relative to F20, IP-AD reduced the scene-mean position and velocity RMSEs by 67.38% and 46.99%, respectively. Mean laser use fell from 315.0 s to 183.5 s. AD gave

20,844.6 m and 80.268 m/s. These seed-3101 results compare complete construction policies; Sections 4.8 and 4.9 separate duration and initialization effects with matched controls. Because all methods use H-SADE, this experiment does not compare differential evolution against other outer search algorithms.

Absolute performance remained limited. At acceptance thresholds of 500 m and 25 m/s, the equally weighted scene-mean full-interval actual-error compliance rates were 19.16%, 29.79%, and 31.19% for F20, AD, and IP-AD, respectively, for search seed 3101. Equation (24) requires both limits at every 0.2 s evaluation point over 120–900 s for the same target realization. These percentages are neither the fraction of physical targets that are always successfully tracked nor the fraction of compliant time points. The relative RMSE improvements therefore do not establish adequate tracking of the entire population. Rejected, unobserved, and late-initialized targets remain in all accuracy statistics, but their inclusion is not the only source of large errors: inconsistent estimates after measurement updates can also contribute, as discussed in Section 4.11.

4.4. Benefits and Regressions from Initialization Priority

Under the reference condition, AD gave position and velocity RMSEs of 20,844.6 m and 80.268 m/s; IP-AD gave 16,847.7 m and 57.709 m/s. Mean laser use was 177.2 s and 183.5 s, respectively. Across eight scenes, IP-AD improved position RMSE in three, worsened it in four, and tied in one (Table 5). A local candidate's gain cannot predict the full-interval outcome because initialization priority changes durations and subsequent revisit calendars.

Earlier measurements shorten unobserved propagation for some targets, but redirecting resources can delay other targets' revisits. Initialization priority is therefore a heuristic for using early opportunities, with no demonstrated improvement in every scene. Figure 5 shows equally weighted scene RMSE histories. Figure 6 reports full-interval compliance alongside laser use so that the comparison does not rest on one mean alone.

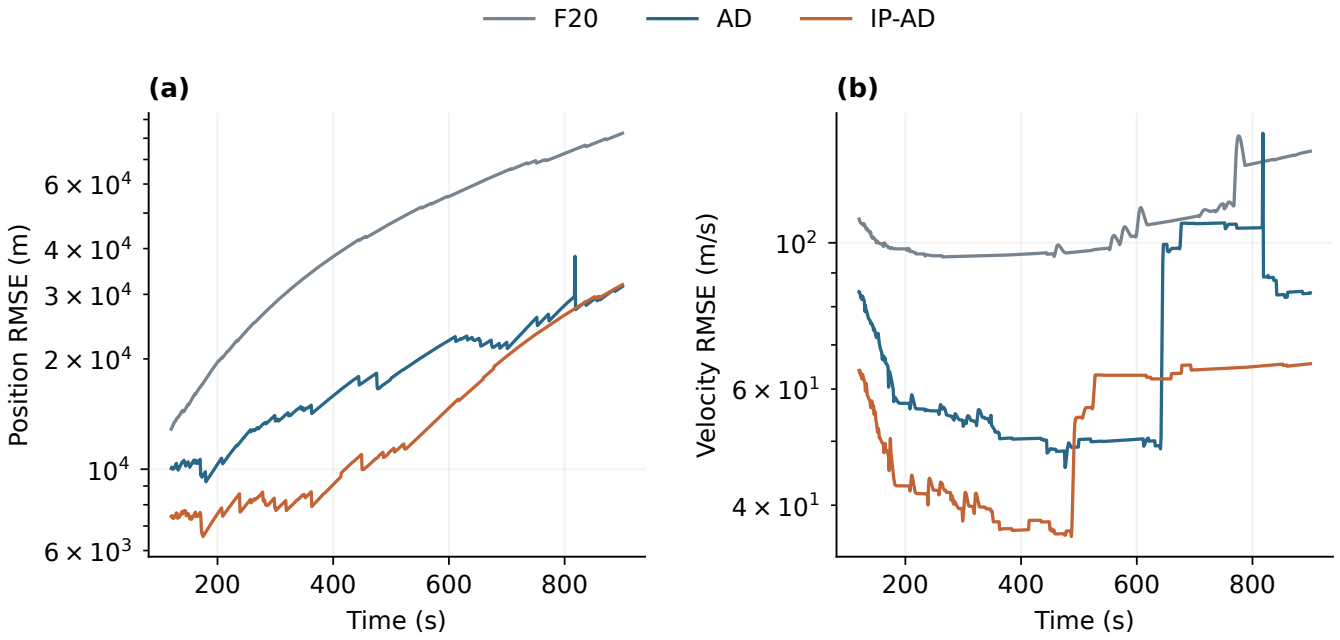

Figure 5. Actual tracking-error histories during maintenance: (**a**) position and (**b**) velocity RMSE. Equation (22) is first evaluated within each scene using all targets and 200 realizations, then averaged equally across eight scenes. Curves display one point per 1 s; summary metrics use the complete 0.2 s grid. Ordinates are logarithmic.

Representative scenes illustrate this conditional behavior. In Table 5, IP-AD reduces position RMSE relative to AD from 24,214.8 m to 621.0 m in scene 2101, but increases it from 50,769.0 m to 73,658.3 m in scene 2106. Appendix B shows that these scenes have 12 and 24 targets and per-satellite laser budgets of 75 s and 240 s, respectively. In the single-seed IP-AD diagnostic described in Section 4.11, all 12 targets in scene 2101 complete their first visit by 120 s. Only 11 of 24 do so in scene 2106, where four targets remain unobserved and the illustrated inconsistent posterior also occurs. More nominal laser time does not by itself secure timely coverage or reliable nonlinear estimates; the availability windows and allocation of early service remain relevant. The evaluation-budget comparison in Section 4.10 also retains these contrasting scenes.

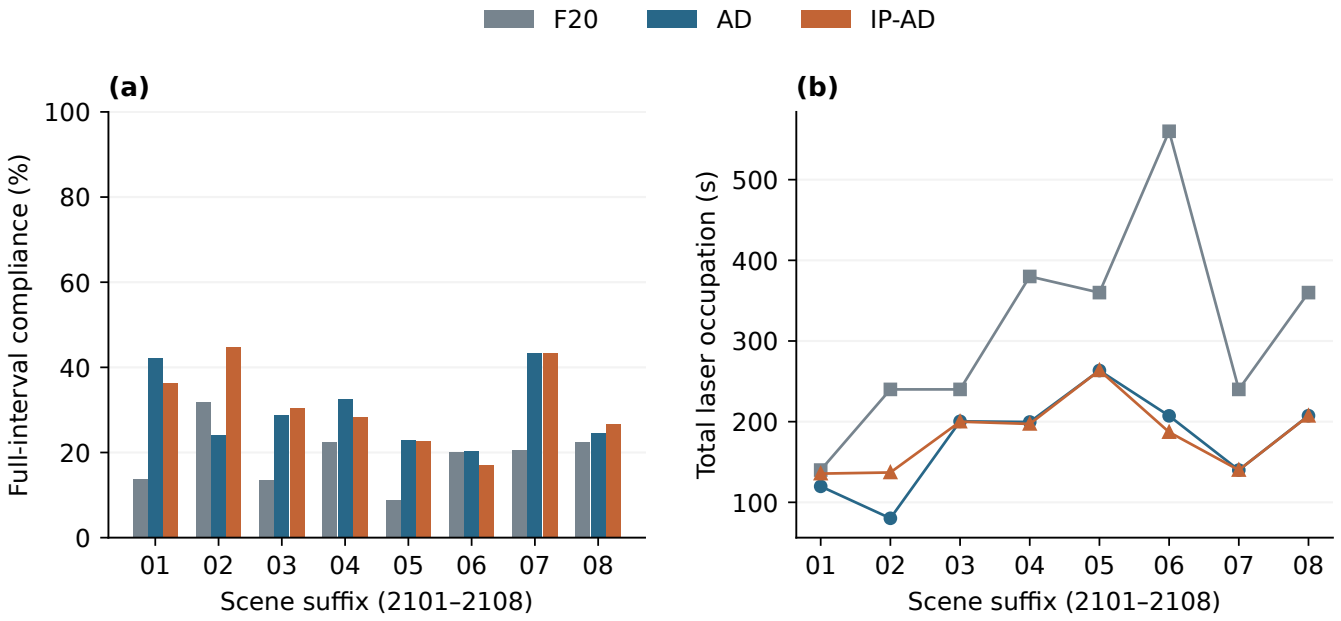

Figure 6. Actual-error compliance and laser use: (**a**) full-interval compliance at 500 m and 25 m/s, and (**b**) total laser measurement time. Compliance requires both limits at every grid point over 120–900 s for the same target realization. Each scene's denominator is its target count multiplied by 200.

Scene 2107 provides a different boundary case. AD and IP-AD have the same rounded position and velocity RMSEs, 369.5 m and 3.04 m/s. The scene has 12 targets, a 120 s per-satellite budget, higher process noise, and only one geometrically available observer for each target, yet all 12 targets complete their first IP-AD visit by 120 s in the diagnostic. Favorable performance in this case cannot be explained by observer redundancy or low process noise. These comparisons describe the joint effects of scene configuration and the resulting calendars. Since geometry, target count, budget, and other inputs vary together, they do not isolate individual causal effects; the IP-AD-only diagnostic also does not fully explain the difference from AD.

4.5. Sensitivity to Acceptance Thresholds

Equation (24) was recomputed at four position–velocity threshold pairs using unchanged truth trajectories, estimates, and schedules (Table 6, Figure 7). The baseline is 500 m and 25 m/s; the other pairs are 1000 m/50 m/s, 1500 m/75 m/s, and 2000 m/100 m/s. These scenario settings illustrate different acceptance requirements. They are not presented as standard thresholds for cataloging, collision warning, or other operational missions.

Table 6. Full-interval actual-error compliance (%) under different acceptance thresholds. Observation plans and planning thresholds are unchanged.

Position (m)/Velocity (m/s) Limits	F20	AD	IP-AD
500/25	19.16	29.79	31.19
1000/50	32.64	52.16	56.49
1500/75	38.54	57.77	63.54
2000/100	41.81	60.36	67.23

Relaxing acceptance limits reclassifies trajectories without improving their realized errors. IP-AD compliance rose to 56.49% at 1000 m/50 m/s and 67.23% at 2000 m/100 m/s. The primary RMSE values remain those in Table 4. Section 4.6 instead supplies these threshold levels to the scheduler and recomputes durations and revisits.

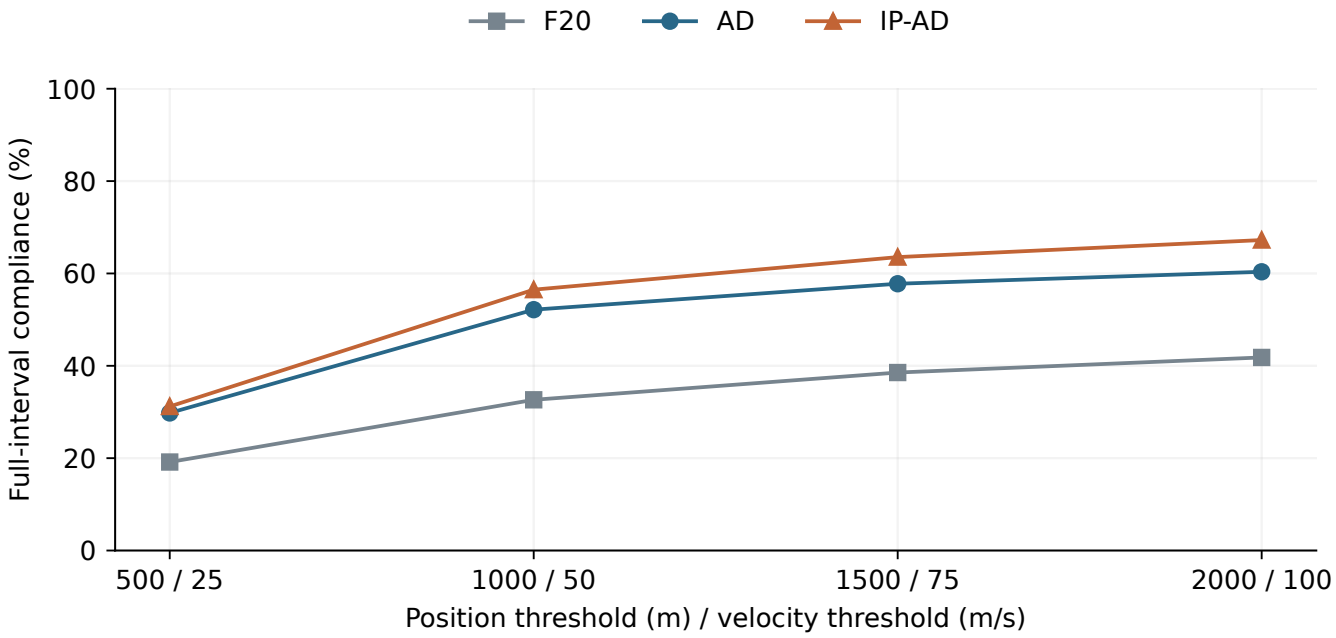

Figure 7. Sensitivity of full-interval actual-error compliance to acceptance thresholds. Each horizontal-axis label gives position (m) and velocity (m/s) limits, both of which must hold for the same trajectory. The three curves use the same scenes, realizations, and fixed schedules; planning thresholds remain at P1 (500 m and 25 m/s).

4.6. Replanning at Four Planning-Accuracy Levels

Planning thresholds were P1: 500 m/25 m/s, P2: 1000 m/50 m/s, P3: 1500 m/75 m/s, and P4: 2000 m/100 m/s. Substitution into Equations (6)–(8) changes stopping conditions, candidate durations, and normalized warning-release times. Each setting was optimized independently with 48 evaluations. The hold duration, warning threshold, and measurement noise remained 200 s, 0.8, and 100 µrad/50 m. Because warning metrics are normalized by planning thresholds, their absolute covariance levels also change.

Table 7 reports actual errors and laser use under a common acceptance criterion of 500 m and 25 m/s. The same criterion applies to P2–P4 despite their looser planning limits. The result files also retain compliance at all four acceptance pairs, including the pair matching each planning level. Because those matched acceptance limits vary, that statistic alone cannot establish an accuracy improvement.

Table 7. Actual errors and resource use after independently optimizing each planning-accuracy level. Noise is fixed at 100 µrad/50 m; compliance is evaluated uniformly at 500 m/25 m/s.

Level	Method	$\overline{E}_r$ (m)	$\overline{E}_v$ (m/s)	$\eta_e(500,25)$ (%)	T_L (s)
P1	F20	51,655.0	108.855	19.16	315.0
P1	AD	20,844.6	80.268	29.79	177.2
P1	IP-AD	16,847.7	57.709	31.19	183.5
P2	F20	46,744.1	90.750	9.79	310.0
P2	AD	28,345.6	70.895	8.99	95.8
P2	IP-AD	10,606.1	29.748	10.53	110.1
P3	F20	24,932.4	58.283	8.21	302.5
P3	AD	23,390.1	55.255	3.88	70.4
P3	IP-AD	12,685.6	31.910	4.16	83.3
P4	F20	24,954.0	58.267	8.16	302.5
P4	AD	24,535.4	58.723	1.59	56.5
P4	IP-AD	11,565.8	33.268	1.93	67.2

Across P1–P4, IP-AD position RMSEs were 16,847.7, 10,606.1, 12,685.6, and 11,565.8 m; velocity RMSEs were 57.709, 29.748, 31.910, and 33.268 m/s. Figure 8 also shows common-threshold compliance and laser use. Looser planning limits reduce the information required from one visit, but can shorten measurements and alter revisit timing. Actual errors reflect all these changes, so easier planning feasibility does not imply better tracking accuracy.

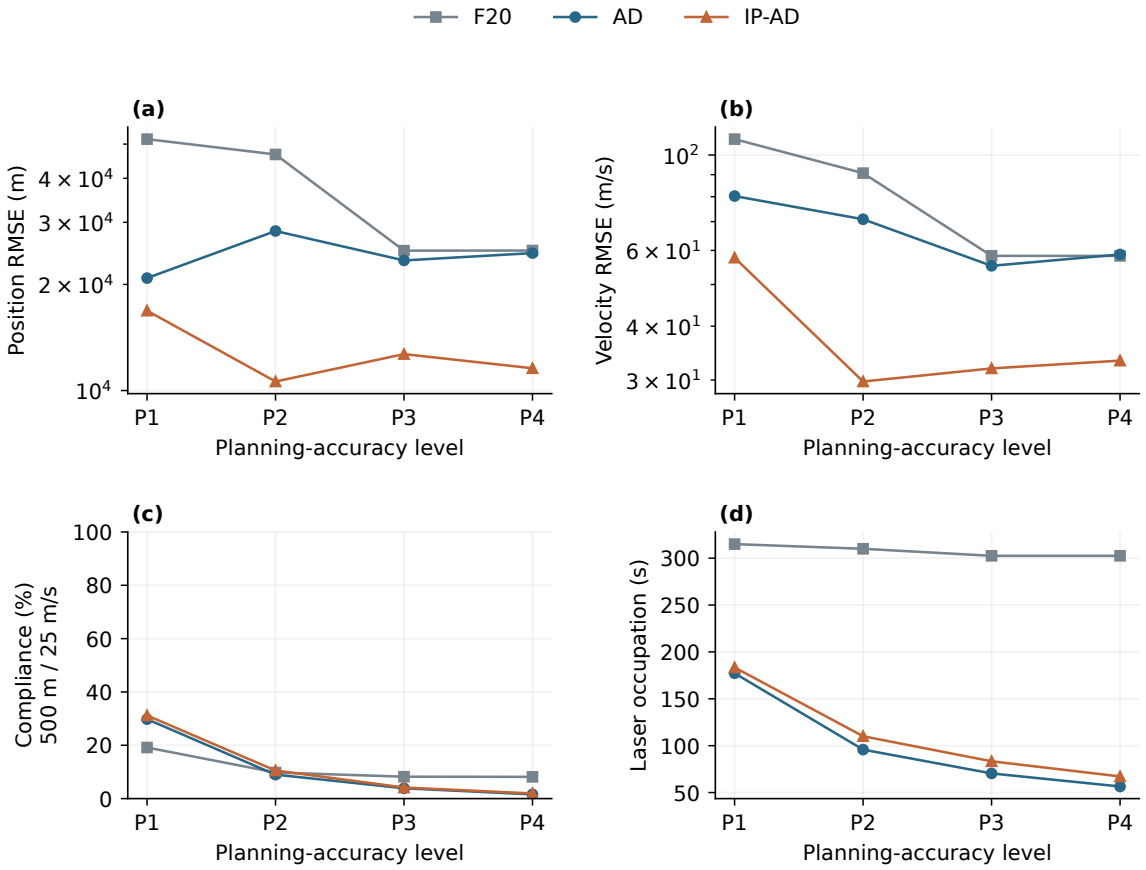

Figure 8. Planning-threshold sensitivity with independently optimized schedules. Noise is fixed at 100 µrad/50 m. Panels show (**a**) position RMSE, (**b**) velocity RMSE, (**c**) full-interval compliance at the common 500 m/25 m/s limits, and (**d**) laser use. Each point is an equally weighted mean over eight scenes.

4.7. Crossed Measurement-Noise and Planning-Threshold Experiment

The additional measurement-noise pairs were M1: 50 µrad/30 m, M2: 100 µrad/60 m, M3: 150 µrad/90 m, and M4: 200 µrad/120 m. Each was crossed with P1–P4, yielding 16 combinations. M2 has a 60 m range standard deviation, distinct from the 50 m reference value. All combinations use a 0.2 s sampling period and independent noise in the two angular components and range. As angular and range standard deviations vary together, their separate effects cannot be identified from this experiment.

For every combination, noise matrices were changed consistently in planning, measurement generation, and the EKF before schedules were optimized and evaluated for all targets. The results measure the combined response to sensor accuracy and schedule adaptation, not filter sensitivity under a fixed schedule. Orbits, priors, process noise, 200 s hold duration, budgets, and initialization deadline were held constant. Corresponding targets shared truth and standardized noise streams.

Table 8 and Figures 9 and 10 report position and velocity RMSEs. Table 9 gives IP-AD compliance at four acceptance pairs for all 16 combinations. Within a row, schedules and simulated trajectories are fixed, and only acceptance limits change. Across rows, planning thresholds and measurement noise may both change. Complete compliance matrices and scene-level values are retained in the result CSV files, including unobserved and late-initialized targets.

Table 8. All-target RMSE across the measurement-noise and planning-threshold grid. Each method has position error (m) followed by velocity error (m/s). Position values are rounded only for compact display.

Plan	Noise	F20 pos.	F20 vel.	AD pos.	AD vel.	IP-AD pos.	IP-AD vel.
P1	M1	52,208	118.15	54,244	232.27	45,823	116.10
P1	M2	52,523	112.29	32,016	101.36	27,707	82.52
P1	M3	54,488	110.02	33,696	95.37	30,144	68.16
P1	M4	54,928	105.25	28,943	69.30	26,845	69.94
P2	M1	23,876	57.66	47,603	113.49	23,542	53.09
P2	M2	53,151	114.02	30,731	78.62	11,534	35.72
P2	M3	54,607	109.78	26,865	65.76	12,688	36.92
P2	M4	52,830	106.17	20,876	53.66	13,233	37.18
P3	M1	25,642	62.44	40,474	91.48	21,726	54.71
P3	M2	25,676	60.15	24,706	57.21	16,790	40.14
P3	M3	52,682	106.12	23,261	54.38	13,401	34.34
P3	M4	53,989	108.57	19,461	49.37	9063	26.95
P4	M1	25,681	62.31	45,128	100.36	62,067	132.39
P4	M2	25,577	59.79	22,305	54.61	10,977	32.62
P4	M3	30,058	66.86	21,514	51.87	8883	25.93
P4	M4	53,103	106.88	19,022	47.46	7622	23.19

Across these 16 settings, IP-AD position RMSE ranged from 7622.1 m at P4/M4 to 62,066.7 m at P4/M1. This range reflects both the chosen conditions and heuristic scheduling; its minimum is not a general performance estimate. Noise or threshold changes can select different candidates, so reduced information per measurement need not change all-target RMSE monotonically. Figures 9 and 10, Tables 8 and 9, and scene-level results jointly describe the accuracy–coverage tradeoff.

P4/M1 provides a counterexample. IP-AD position and velocity RMSEs were 62,066.7 m and 132.391 m/s, compared with F20 values of 25,680.6 m and 62.305 m/s. Laser use was 49.4 s and 302.5 s, respectively. Lower resource use did not yield better all-target accuracy. This result suggests a mismatch between covariance feasibility and the information needed

by nonlinear estimation. The experiment does not fully separate that mismatch from resource allocation and search variability.

Table 9. IP-AD actual-error compliance (%) for all planning/noise combinations. Column headings are position (m) and velocity (m/s) acceptance limits.

Plan	Noise	500/25	1000/50	1500/75	2000/100
P1	M1	30.83	54.15	62.74	67.68
P1	M2	32.48	58.31	64.56	67.40
P1	M3	31.25	55.14	58.61	59.81
P1	M4	28.68	50.92	54.07	55.18
P2	M1	10.77	35.05	48.89	56.98
P2	M2	11.56	40.98	56.63	64.75
P2	M3	11.12	43.39	60.31	67.35
P2	M4	9.75	40.81	58.44	65.13
P3	M1	4.27	20.85	35.06	45.28
P3	M2	4.55	25.16	43.19	55.19
P3	M3	4.31	27.14	46.92	58.97
P3	M4	4.09	26.32	46.61	59.26
P4	M1	2.23	12.40	24.74	35.51
P4	M2	2.19	15.36	30.67	43.46
P4	M3	1.73	15.35	32.81	46.64
P4	M4	1.49	15.65	34.53	48.99

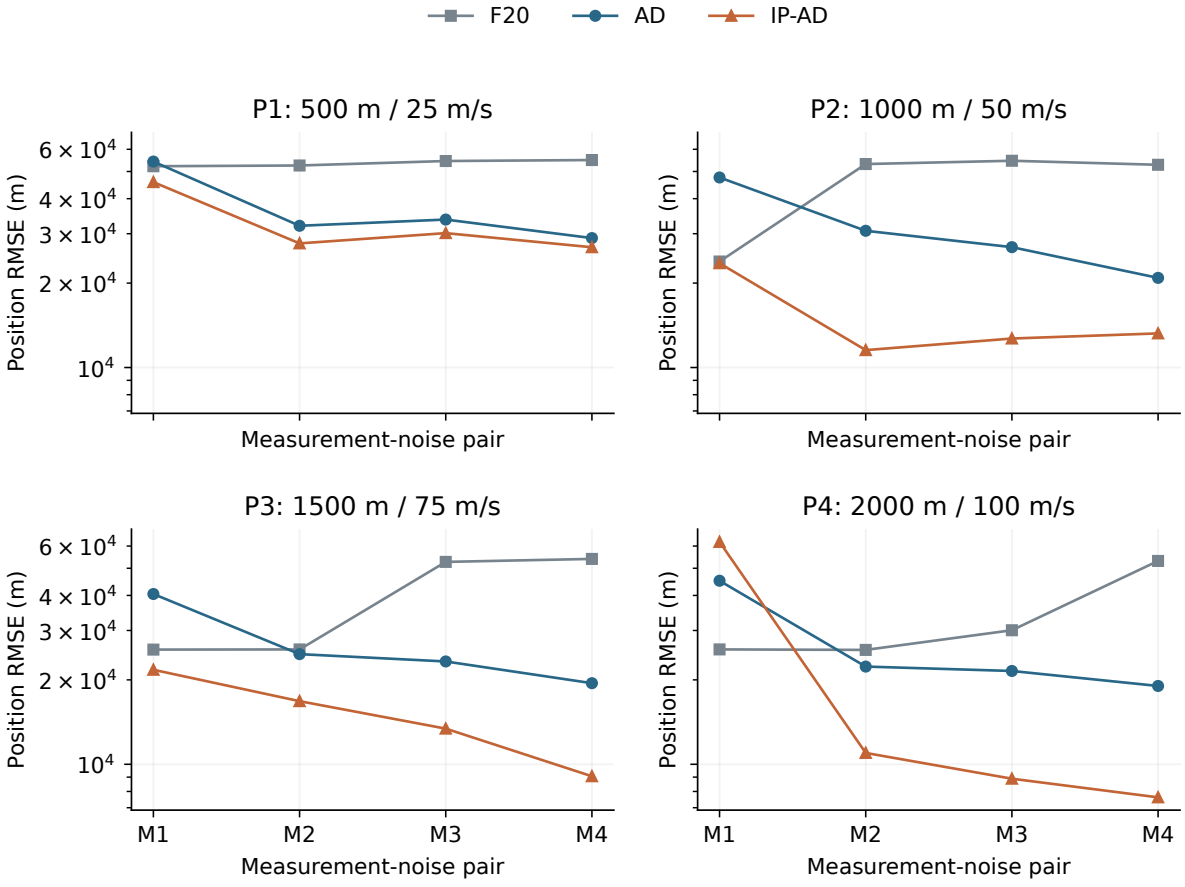

Figure 9. Position RMSE across the planning-threshold and measurement-noise grid. Panels correspond to P1–P4, and M1–M4 denote the noise pairs defined in the text. Each point and method uses an independently optimized schedule. Ordinates are logarithmic.

All settings reuse the same eight orbital scenes. The 20 parameter settings are not 20 independent orbital samples. We report descriptive averages without claims of statistical significance or global optimality.

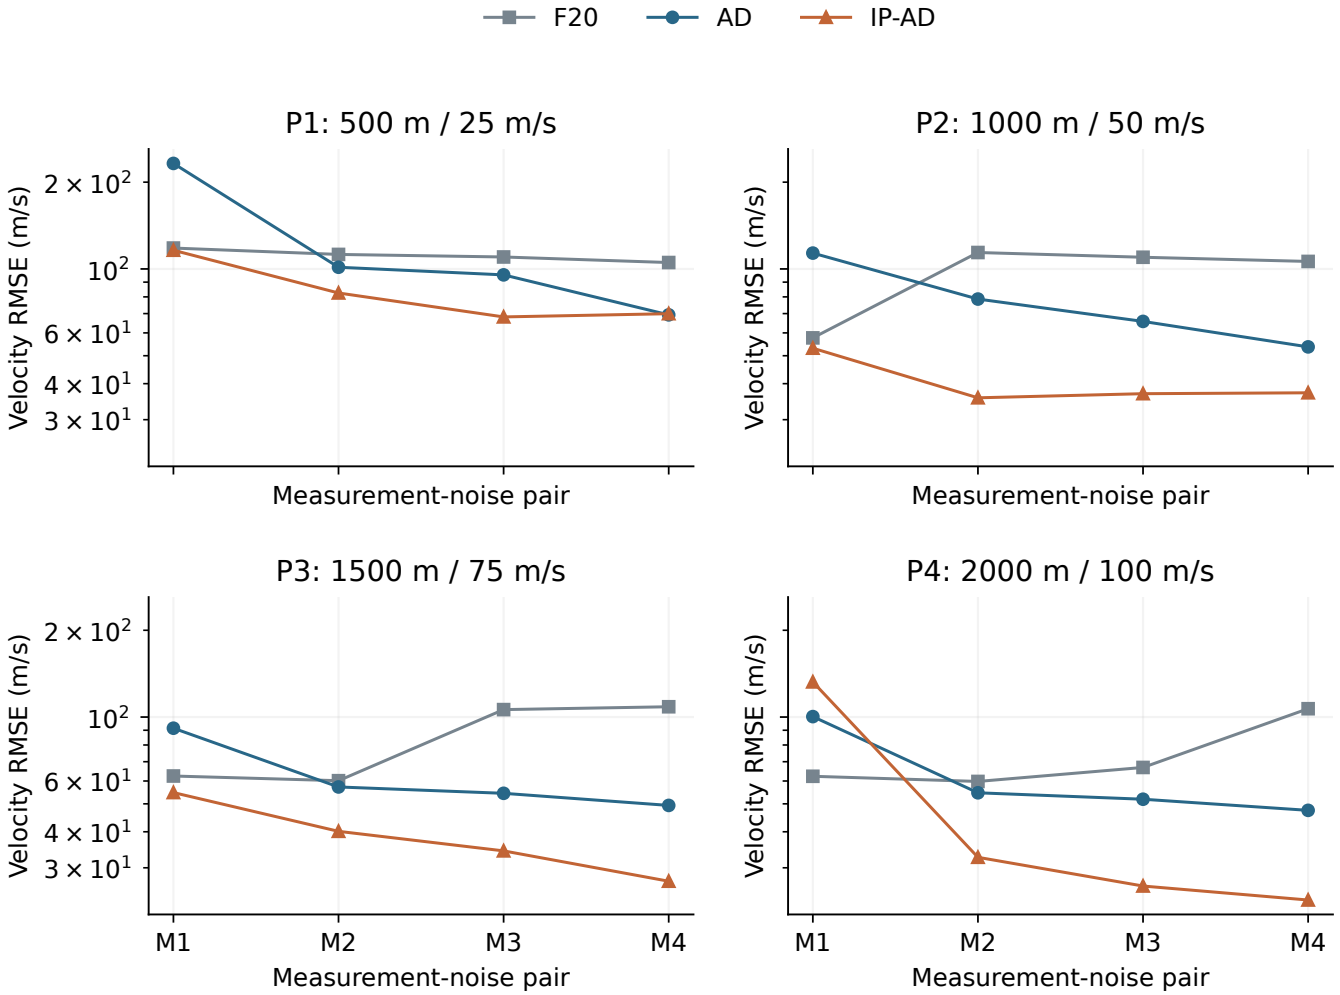

Figure 10. Velocity RMSE for the same grid as Figure 9. Each point includes all targets and 200 realizations over the complete 0.2 s grid from 120 to 900 s, followed by equal averaging over eight scenes.

4.8. Matched Controls for Duration and Initialization Priority

Table 10 separates duration policy, initialization priority, and hold screening. In F20, Equation (7) is calculated and recorded but does not reject the prescribed 20 s observation. The hold constraint in Equation (10) therefore applies to AD, IP-AD, and fixed-duration controls with the H suffix, but not to F20 or IP-F20. Every method uses the same geometry, attitude, serial-occupancy, and laser-budget checks.

IP-F20 adds only Equation (21) to F20. F20-H and IP-F20-H additionally require the prescribed 20 s observation to pass the full 200 s hold check. A failed candidate is rejected without shortening or extension. Thus, F20-H, IP-F20-H, AD, and IP-AD form a fixed/adaptive-duration by base/initialization-priority comparison under the same hold constraint. AD versus F20-H and IP-AD versus IP-F20-H assess duration under each initialization rule. IP-F20-H versus F20-H and IP-AD versus AD assess initialization priority. Every contrast uses newly constructed complete calendars, rather than local information gains as final outcomes.

Table 10. Matched controls separating duration, initialization priority, and hold screening.

Method	Measurement Duration	Initialization Priority	Equation (7) Screening
F20	Fixed 20 s	No	No
IP-F20	Fixed 20 s	Yes	No
F20-H	Fixed 20 s	No	Yes
IP-F20-H	Fixed 20 s	Yes	Yes
AD	Shortest and 1.5× candidates	No	Yes
IP-AD	Shortest and 1.5× candidates	Yes	Yes

The 20 s duration is retained as the original reference, not claimed to be optimal. A predefined scan uses 10, 20, 30, and 40 s, with initialization priority and complete hold screening enabled throughout; these controls are denoted IP-F*d*-H. The scan uses reference P1, eight scenes, seed 3101, and 48 evaluations per plan, followed by 200 realizations per target. Every duration is optimized separately. The four values span shorter and longer visits below the 60 s continuous laser limit. They describe performance–resource tradeoffs without treating the best observed fixed value as an unbiased performance estimate.

Changing duration changes the measurement count, opportunities for initialization, and subsequent revisit calendar. Performance need not improve monotonically with longer fixed visits. Hold screening tests the planning covariance and cannot ensure that a short feasible visit is sufficient for nonlinear estimation. Table 11 and Figure 11 retain all four durations; the next section compares six methods across multiple search seeds.

Table 11. Fixed-duration sensitivity under identical initialization and hold rules. Reference P1, seed 3101; values are equally weighted eight-scene means. Compliance uses 500 m/25 m/s acceptance limits.

Method	$\bar{E}_r$ (m)	$\bar{E}_v$ (m/s)	η_e (%)	T_L (s)
IP-F10-H	90,103.0	326.500	0.00	58.8
IP-F20-H	49,483.2	94.639	19.85	317.5
IP-F30-H	57,895.2	109.817	16.14	315.0
IP-F40-H	63,424.4	120.287	17.74	300.0
IP-AD	16,847.7	57.709	31.19	183.5

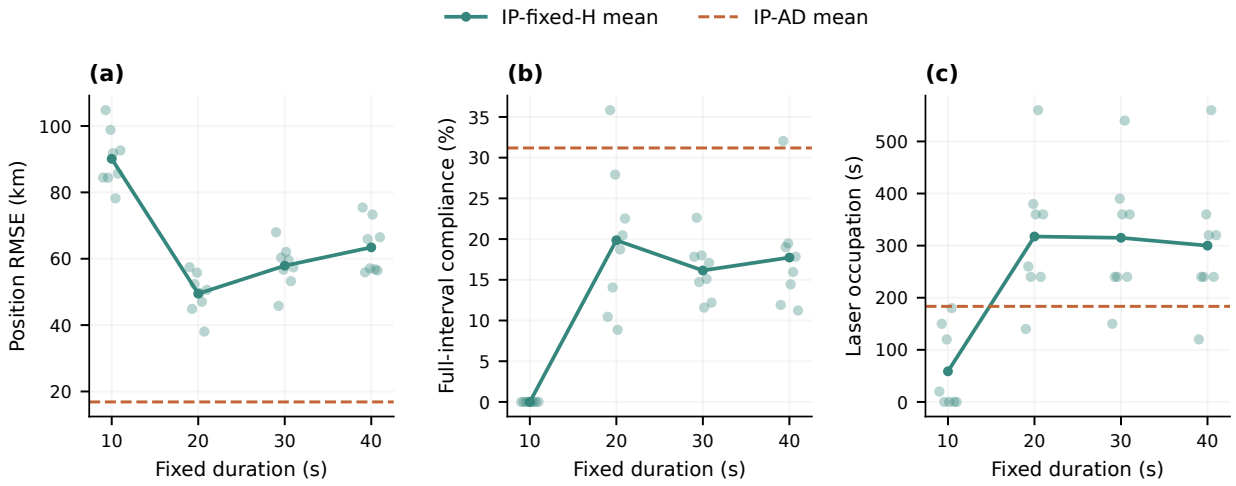

Figure 11. Fixed-duration scan under the same initialization-priority and hold rules: (**a**) position RMSE, (**b**) full-interval actual-error compliance, and (**c**) laser use. Small teal points show all eight scene values; solid lines show their equal-weight means. Dashed lines show IP-AD means for the same search seed. These single-seed data do not quantify optimization variability.

4.9. Paired Search Seeds and Method-Difference Distributions

The predefined search seeds were 3101, 3102, 3103, 3104, and 3105. For reference P1 and the adverse P4/M1 condition, six methods were optimized over eight scenes with 48 evaluations, producing 480 plans. Within a scene, condition, and seed, all methods share the initial population and search seed. Different decoder outcomes can alter successful replacements, so final search paths need not coincide. Initial errors, process noise, and standardized measurement noise remain shared by target, with 200 realizations each. Search seeds are separate from estimation-error random streams.

For each seed, Table 12 first averages the eight scene metrics equally, then reports the mean and sample standard deviation of those five averages. The standard deviation describes optimization variability on this fixed set of scenes. It is neither a confidence interval for an orbital population nor the dispersion of 200 Monte Carlo trajectories. Tables 4–9 retain the descriptive seed-3101 results and are not interchangeable with these multiseed averages.

Table 12. Six-method results across five search seeds. Values are the mean ± sample standard deviation of five eight-scene averages. Every seed uses the same scenes; compliance is evaluated at 500 m/25 m/s.

Condition	Method	$\bar{E}_r$ (m)	$\bar{E}_v$ (m/s)	η_e (%)	T_L (s)
P1	F20	54,260 ± 2426	114.1 ± 4.6	21.18 ± 1.44	316.5 ± 2.2
	IP-F20	54,239 ± 3416	106.6 ± 8.3	21.20 ± 1.27	313.5 ± 5.2
	F20-H	54,260 ± 2426	114.1 ± 4.6	21.18 ± 1.44	316.5 ± 2.2
	IP-F20-H	54,239 ± 3416	106.6 ± 8.3	21.20 ± 1.27	313.5 ± 5.2
	AD	29,547 ± 5142	112.1 ± 18.6	29.55 ± 0.38	169.7 ± 6.4
	IP-AD	18,444 ± 1433	57.1 ± 6.5	31.43 ± 1.06	186.6 ± 2.6
P4M1	F20	25,114 ± 690	60.7 ± 2.3	15.17 ± 0.20	302.0 ± 1.1
	IP-F20	24,642 ± 613	60.0 ± 1.6	15.76 ± 0.24	302.5 ± 0.0
	F20-H	25,114 ± 690	60.7 ± 2.3	15.17 ± 0.20	302.0 ± 1.1
	IP-F20-H	24,642 ± 613	60.0 ± 1.6	15.76 ± 0.24	302.5 ± 0.0
	AD	44,753 ± 6416	99.1 ± 12.1	1.80 ± 0.09	40.5 ± 0.4
	IP-AD	54,451 ± 31,179	117.4 ± 61.5	2.10 ± 0.13	49.1 ± 0.5

Under reference P1, IP-AD reduced mean position RMSE by 35,795.4 m and velocity RMSE by 49.541 m/s relative to IP-F20-H. Both methods use the same initialization priority and hold screening. Full-interval compliance increased by 10.22 percentage points. Across 40 scene–seed pairs, position-error improvement/tie/regression counts were 35/0/5; velocity counts were 33/0/7.

Under P4/M1, IP-AD instead increased mean position RMSE by 29,809.1 m and velocity RMSE by 57.407 m/s relative to the matched IP-F20-H control. Full-interval compliance decreased by 13.66 percentage points. Position improvement/tie/regression counts were 13/0/27 across the 40 pairs; velocity counts were 14/0/26.

For all 40 pairs in each condition, F20-H produced the same final plans as F20, and IP-F20-H the same plans as IP-F20. Hold screening therefore did not change the final 20 s plans in these experiments, despite changing the candidate-acceptance rule. Duration comparisons still use methods with identical hold rules.

Paired differences subtract the second method from the first within the same condition, scene, and seed. Negative error differences favor the first method; positive compliance differences favor it. Table 13 reports position-difference means, medians, interquartile ranges, and direction counts. Figure 12 shows every position and velocity pair. The 40 pairs comprise eight scenes by five seeds, not 40 independent orbital samples. Time points

and shared-scene parameter settings are likewise not treated as independent replicates for significance testing.

Table 13. Paired position-RMSE differences in km. Intervals are the 25th–75th percentiles of 40 scene–seed differences, not confidence intervals. Counts denote improvement/tie/regression for the first method.

Condition	First Method − Second Method	Mean	Median [Q1,Q3]	Counts
P1	IP-F20-H − F20-H	−0.02	0.00 [−3.78, 0.00]	15/16/9
	IP-AD − AD	−11.10	−0.70 [−22.50, 0.66]	21/6/13
	AD − F20-H	−24.71	−34.01 [−48.11, −5.69]	31/0/9
	IP-AD − IP-F20-H	−35.80	−44.16 [−54.83, −30.41]	35/0/5
P4M1	IP-F20-H − F20-H	−0.47	0.00 [−0.19, 0.48]	15/9/16
	IP-AD − AD	9.70	0.00 [−25.65, 11.03]	13/15/12
	AD − F20-H	19.64	12.04 [−0.64, 32.66]	11/0/29
	IP-AD − IP-F20-H	29.81	7.50 [−0.64, 35.73]	13/0/27

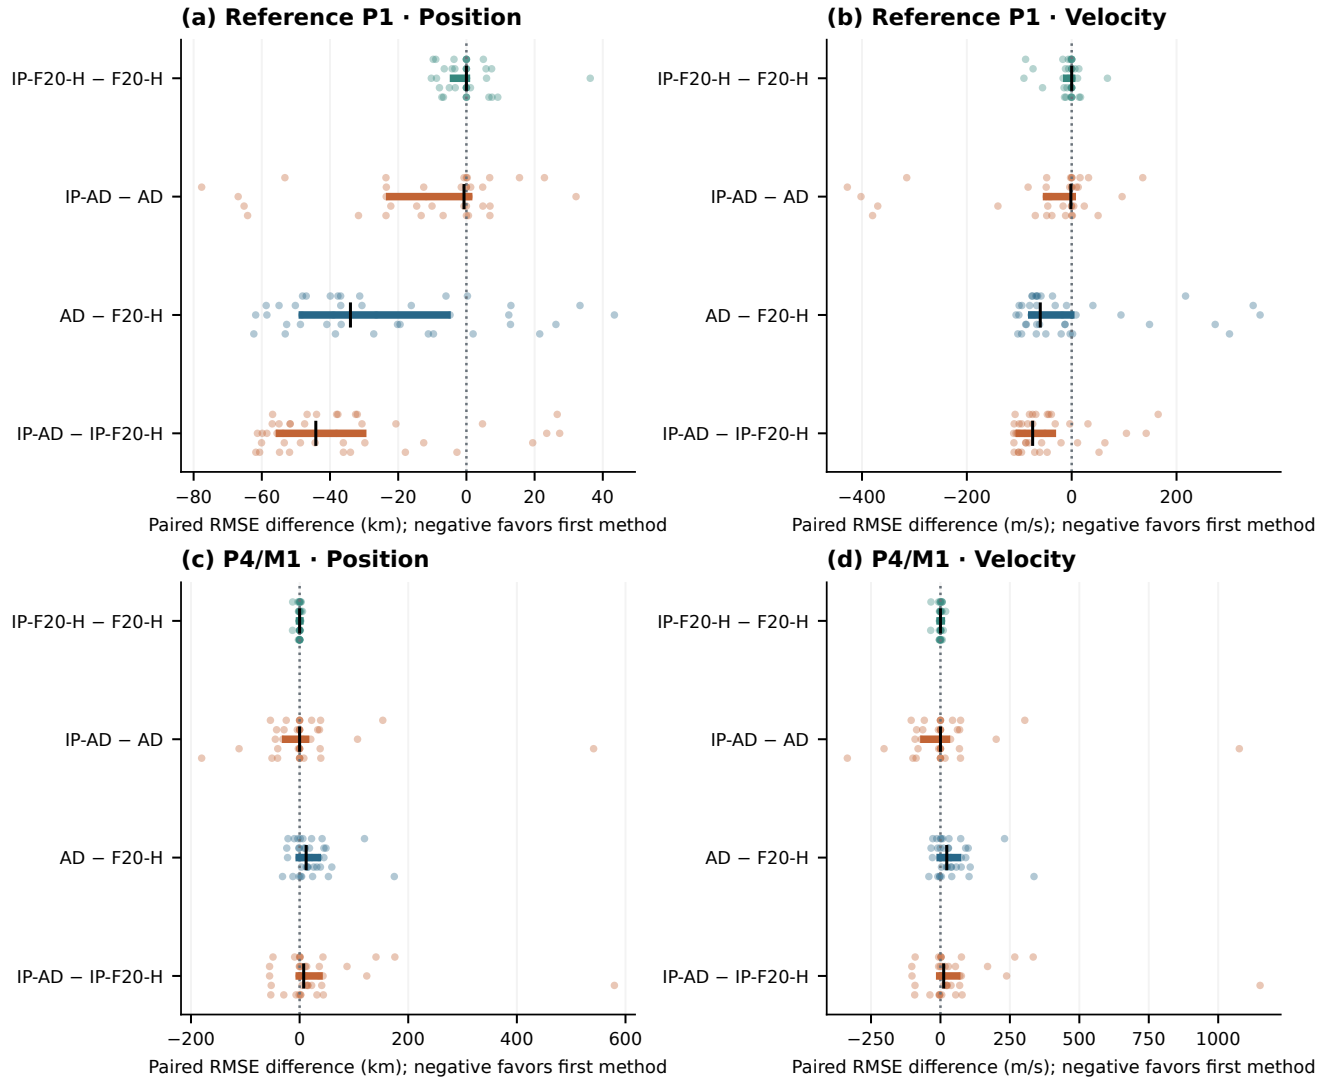

Figure 12. Paired effects of initialization priority and duration under identical hold constraints. Rows show reference P1 and P4/M1; columns show position and velocity differences. Every contrast retains all 40 scene–seed points. Thick lines indicate interquartile ranges and black marks indicate medians; vertical offsets only separate overlapping points. Negative values favor the first method.

At P4/M1 with seed 3101, IP-AD retained the 62,066.7 m position mean reported in Section 4.7. Its five seed-specific means ranged from 22,970.2 to 104,171.2 m. Search repetition reveals calendar variability but does not remove the mismatch between the covariance surrogate and nonlinear estimation.

4.10. Evaluation Budget and Measured Solution Time

The budget experiment selected reference-P1 scenes 2101 and 2106 to represent a large initialization-priority improvement and a clear regression in the original results. Six methods and five seeds used budgets of 24, 48, and 96 complete-plan evaluations. The 48-evaluation plans were reused from identical inputs; budgets 24 and 96 were solved anew, adding 120 plans. With population size 12 included in each budget, the three budgets permit 12, 36, and 84 subsequent trial evaluations. Figure 13 highlights the four methods with identical hold constraints; run-level records retain all six.

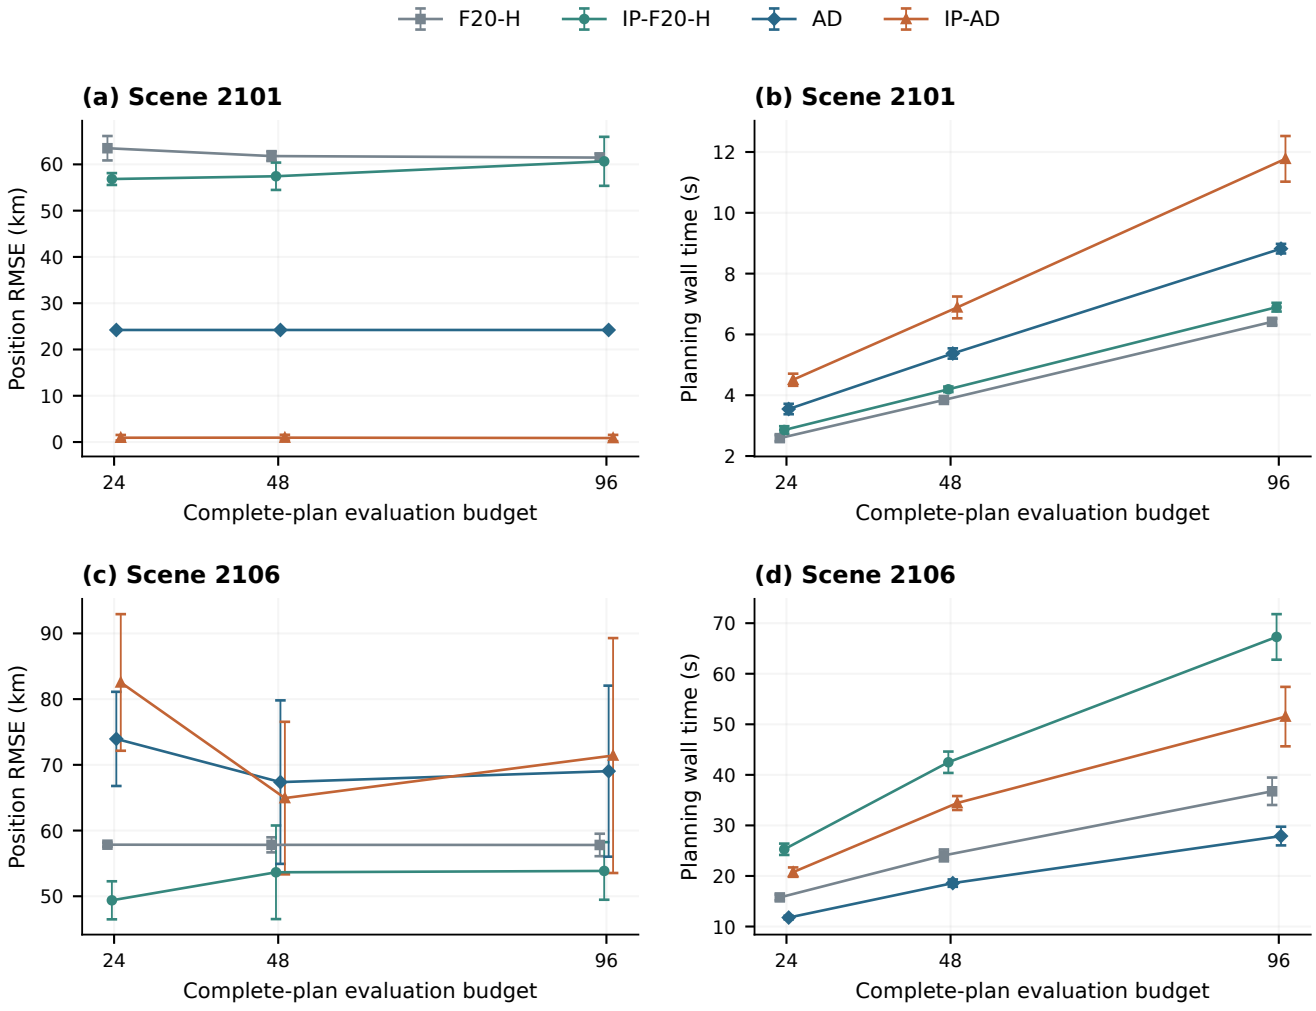

Figure 13. Accuracy and planning time versus evaluation budget under reference P1. Rows show scenes 2101 and 2106; columns show actual position RMSE and planning wall time. Points are five-seed means, and error bars are sample standard deviations. Horizontal offsets aid identification. Planning time excludes geometry preprocessing and Monte Carlo accuracy evaluation.

In scene 2101, IP-AD five-seed position means were 930.8, 946.1, and 861.0 m at budgets 24, 48, and 96. Mean planning wall times were 4.51, 6.89, and 11.78 s, respectively.

In scene 2106, the corresponding position means were 82,537.8, 64,941.9, and 71,415.5 m, with mean wall times of 20.76, 34.43, and 51.52 s.

More evaluations cannot worsen the best retained lexicographic planning objective. That surrogate differs from actual error, however, so RMSE need not decrease monotonically. Budget effects were measured only in these two diagnostic scenes. They neither establish a general convergence pattern nor show that 48 evaluations are sufficient for convergence.

The original matched-control experiments ran on an AMD RYZEN AI MAX+ 395 with Radeon 8060S, using Python 3.12.14 and NumPy 2.3.5. Eight independent worker processes ran with one BLAS/OpenMP thread each. Every plan began with an empty duration cache. A monotonic timer enclosed the full planning call, including population initialization, candidate construction, internal replay, and all trial evaluations. Geometry preprocessing, subsequent independent audits, and EKF Monte Carlo evaluation were excluded. Process CPU time was recorded separately from wall time to distinguish computation from concurrency-related variation. Table 14 spans 40 runs per method under this workstation setup; it does not measure onboard real-time capability.

Table 14. Measured planning cost at 48 evaluations, in s. Wall time is the median [minimum, maximum] of 40 runs. CPU time is the mean ± seed standard deviation of five eight-scene means. Offline accuracy evaluation is excluded.

Condition	Method	Wall-Time Median [Range]	Process CPU Time
P1	F20	15.99 [3.73, 32.49]	16.23 ± 0.29
	IP-F20	17.60 [4.22, 44.21]	19.09 ± 0.53
	F20-H	16.07 [3.76, 32.61]	16.34 ± 0.31
	IP-F20-H	17.67 [4.12, 44.66]	19.19 ± 0.52
	AD	14.60 [4.10, 37.35]	14.87 ± 0.33
	IP-AD	17.72 [6.02, 37.69]	18.14 ± 0.50
P4M1	F20	11.97 [3.20, 23.66]	11.80 ± 0.13
	IP-F20	13.09 [3.41, 26.47]	13.03 ± 0.28
	F20-H	11.99 [3.21, 24.53]	11.82 ± 0.13
	IP-F20-H	13.43 [3.43, 27.22]	13.13 ± 0.34
	AD	4.55 [0.97, 10.31]	4.51 ± 0.07
	IP-AD	5.73 [2.35, 10.96]	5.75 ± 0.03

The original matched-control extension comprised 624 independently optimized plans and 10,854 target–plan evaluation chains, corresponding to 2,170,800 target-realization evaluations. When two plans produced identical measurement satellites and times for a target, the deterministic EKF result was reused only for identical physical inputs. Scheduling searches were never reused on that basis. The cache contained 4383 distinct target-evaluation inputs. All required hold and resource checks passed. The 48 seed-3101 plans and RMSEs for the original three methods at reference P1 and P4/M1 were reproduced. Corresponding truth-checkpoint hashes matched across methods, budgets, and seeds.

4.11. Limits of Covariance-Based Accuracy Prediction

Covariance propagation predicts local measurement information, but linearization can become inaccurate after large initial velocity errors propagate without measurements. Actual nonlinear errors may therefore exceed thresholds even when each posterior passes the hold test. Evaluating every target avoids representing a method only by tracks initialized early or classified as covariance-feasible. The results retain limited velocity improvements and the redistribution of accuracy caused by initialization priority.

The experiments hold H-SADE fixed while comparing duration, initialization, and hold-screening policies. Genetic algorithms, other differential-evolution variants, and local search have not been compared using these all-target nonlinear errors. No ranking of outer optimizers follows from the results. All 20 original planning/noise parameter settings were

replanned and evaluated; the 24 reference plans and 414 target-result chains underpin Tables 4–6.

An additional diagnostic replay of the existing IP-AD reference schedules examines this limitation at P1, measurement noise of 100 µrad/50 m, $H = 200$ s, and search seed 3101. It retains the original 0.2 s grid and 200 realizations per target for all 138 targets. Of these targets, 96 complete their first visit by 120 s, 38 complete it later, and four receive no measurements. These counts describe observation timing, not compliance with the actual-error thresholds. Delayed or missing observations allow the large initial velocity uncertainty to propagate, while short local covariance-based stopping decisions need not establish a reliable nonlinear state estimate.

For example, target 10 in scene 2106 completes its first visit at 494.2 s. Its modeled posterior position and velocity uncertainties are 12.89 m and 2.430 m/s, whereas its Monte Carlo RMSEs at that same time are 892.66 m and 1153.30 m/s. The modeled hold metric is approximately 0.998, yet the posterior errors are already much larger than predicted. Subsequent measurement-free propagation can amplify this disagreement; the large errors cannot all be attributed to normal prediction growth from a well-calibrated posterior. This example demonstrates covariance–error inconsistency, but does not isolate the causal benefit of a different estimator or stopping rule. The diagnostic covers one reference search seed and is not interchangeable with the five-seed aggregates in Section 4.9. Nor does it decompose the noncompliant fraction into independent causal contributions.

4.12. Hold, Heuristic, Dynamics, and Workload Sensitivity

The supplementary experiments retain reference P1. With independent replanning at H = 100, 200, 300, and 500 s, IP-AD's mean position RMSE ranges from 18,443.8 to 41,545.6 m and full-interval compliance from 29.40% to 33.36%. Changing a visit-level hold requirement therefore needs an actual-error evaluation of the resulting calendar, rather than an inference from covariance feasibility alone.

Across nine local heuristic settings, mean IP-AD position RMSE ranges from 16,794.5 to 24,726.0 m and compliance from 28.42% to 36.18%. These one-factor comparisons describe local sensitivity; they do not establish the original selection rationale. The original baseline is retained while a selected warning variant is assessed separately.

The nested 24–96-target study exposes overload under fixed local observing resources. Admission keys permit attempts to serve a subset, but all targets remain in the accuracy denominator. The 96-target boundary, including the opposing RMSE and compliance rankings of the two methods, is retained in Appendix C; it does not establish a constellation capacity limit.

For the original fixed calendars, nominal J_2 execution checks cancel 19 of 2122 visits. With J_2 truth, IP-AD mean position RMSE is 18,770.3 m using the two-body estimator and 18,491.4 m using the J_2 estimator, compared with 18,443.8 m in the two-body control. Full-interval compliance is 0.00% with the mismatched two-body estimator and 31.41% with the matched estimator, compared with 31.43% in the original control. The two estimators under J_2 truth share the execution mask; comparison with the original control also includes geometry and execution changes. No J_2 scheduling optimization is performed.

These supplementary comparisons use reference P1 and retain unfavorable outcomes. Appendix C gives complete protocols, both methods where applicable, uncertainty, and service/resource outcomes. The tests do not establish universal robustness, a historically optimized parameter choice, or operational validity of the sensing assumptions.

A warning-0.7 variant, selected from the local sensitivity study, was assessed separately on new search seeds 3106–3110. IP-AD-W07, original IP-AD, and IP-F20-H attained mean position RMSEs of 15,753.2, 18,293.9, and 54,172.8 m and compliance of 36.49%, 32.17%,

and 21.81%, respectively. The original baseline is retained; this post-review comparison shares the original geometries and Monte Carlo inputs. Appendix New Search Seeds for a Selected Warning Variant reports the protocol and complete comparisons.

5. Discussion

Adaptive duration links the information acquired during a visit to uncertainty growth after departure. It also determines how much time remains for other targets. The reference and matched multiseed comparisons support conditional benefits, which must be interpreted with compliance, scene-level differences, and search variability. They also expose a gap between the planning surrogate and final accuracy. Maximizing covariance-feasible targets does not necessarily minimize all-target position or velocity RMSE. Error predictions calibrated against nonlinear realizations could address this gap, but their computational cost and performance on new scenes require separate evaluation.

Acceptance thresholds should follow the intended mission. Position and velocity errors concern current localization and subsequent prediction, respectively; relaxing one limit does not establish adequate overall tracking. We therefore report physical-unit RMSEs alongside full-interval compliance at joint thresholds. Mean errors, sustained compliance, and laser use can rank methods differently. Keeping these outcomes separate makes the tradeoffs visible.

Alternative observers are scarce in the simulated local geometries. Although the conceptual constellation contains 72 satellites, each plan uses only a few, and some targets have one geometrically available observer throughout 900 s (Appendix B). A larger nominal constellation does not automatically eliminate competition within local visibility windows. Including unobserved targets makes all-target RMSE sensitive to incomplete coverage and to a small number of large errors. Scene-level results and full-interval compliance help expose this structure.

Payload availability remains unvalidated. The 2000 km cutoff is geometric, 50 m is an assumed range-noise standard deviation, and 5 Hz is an assumed rate of usable joint updates, not a demonstrated return rate for debris near 2000 km. These values do not determine unique laser power, aperture, or detector requirements. A link budget would require pulse energy and duration, wavelength, divergence, pointing jitter, efficiencies, detector timing/background noise, and target optical cross-section. Size, reflectivity, orientation, and tumbling can cause range-dependent availability and intermittent returns. Missed or degraded samples would change posteriors, feasible durations, and calendars; observer ephemeris and acquisition errors are also omitted.

The attitude-dynamics effects of satellite pointing maneuvers in engineering applications are not modeled here [21]. Equation (9) supplies angular-speed and acceleration bounds plus fixed settling, without a spacecraft inertia, actuator model, flexible response, or moving-endpoint interception solution. Common-field geometry does not certify closed-loop boresight and laser pointing. The reference concerns the omitted attitude-planning and control problem; its controller is not implemented here and does not validate the adopted slew parameters.

The original comparisons share two-body truth and estimator dynamics. The separate J_2 study quantifies propagation changes and fixed-calendar mismatch, with the same observation-cancellation mask for both perturbed-truth estimators. It does not optimize calendars under J_2, validate drag or all perturbations, or remove the sensing assumptions. The 200 realizations per target describe specified random errors, and the synthetic families do not represent the full orbital population. These are conditional numerical results for a conceptual system.

A probabilistic operational requirement would need an explicit allowed violation probability and calibrated predictive errors. For example, conditional on information at visit end e, a per-target chance constraint could require probability at least $1-\alpha$ that both error norms remain below d_r, d_v at every point of the measurement-free grid $[e, e+H]$. Joint-over-time and mission-wide requirements differ from separate pointwise bounds. A Gaussian multiplier applied to an inconsistent covariance would not repair the demonstrated error gap. No chance-constrained optimizer or certified rare-event bound is established here.

6. Conclusions

The proposed constellation scheduler couples an independent post-observation hold requirement with variable durations, time-dependent common-field acquisition, and serial laser service. Its constructor derives durations and revisits from individual posteriors. Initialization priority and random-key search coordinate the resulting resource calendars, while nonlinear truth–estimate differences provide an independent accuracy assessment.

Across five paired search seeds under the reference condition, IP-AD achieved scene-mean position and velocity RMSEs of 18,443.8 m and 57.102 m/s. The matched IP-F20-H control gave 54,239.2 m and 106.643 m/s. Relative reductions of 66.00% and 46.46% support an average benefit from adaptive duration under these conditions. Full-interval compliance at 500 m/25 m/s was nevertheless only 31.43%, so the result does not establish adequate accuracy for every target.

The scheduler is applicable to repeated active–passive tracking when targets have usable orbit priors and known measurement associations, observation availability can be predicted, and the sensing and uncertainty models are sufficiently representative for planning. Its adaptive-duration mechanism is most useful when feasible visits of different lengths offer a meaningful tradeoff between information gain and the time available for initialization and revisits. Planning thresholds and the hold requirement must match the intended task, and local visibility windows and resource budgets must permit useful service. The demonstrated benefits are conditional on the tested geometry, measurement noise, and planning settings; they are not a guarantee for arbitrary target loads or sensing conditions.

The principal limitation is that covariance feasibility is a planning surrogate rather than a guarantee of actual tracking accuracy. The benefit reversed at P4/M1. IP-AD's five-seed position mean was 54,451.0 m with a seed standard deviation of 31,178.8 m, compared with 24,641.9 m for the matched fixed-duration control. The crossed planning/noise experiment, fixed-duration scan, and two-scene budget test show that these factors alter accuracy through the observation calendar. Their effects need not be monotonic, and initialization priority does not improve every scene.

These findings compare construction policies under a common H-SADE outer search. The paired-seed and evaluation-budget tests characterize variability and computational effort within that implementation; they do not establish superiority over other metaheuristics or mathematical-programming approaches, or guarantee global optimality. Replanning after repeated stochastic disruptions likewise remains outside the experimental validation.

The supplementary hold and workload tests show that visit-level covariance feasibility does not secure full-population service or monotonic accuracy gains. The J_2 replay demonstrates the importance of estimator-model consistency for full-interval compliance. The selected warning variant is assessed on shared geometries and Monte Carlo inputs, so its results do not establish transfer to new operational scenes. Further work should address admission and early-service allocation, nonlinear estimation consistency, realistic laser returns, pointing errors, and broader orbital geometries.

https://doi.org/10.3390/aerospace13110987

Author Contributions: Conceptualization, J.F.; methodology, J.F.; software, J.F. and Y.L.; validation, J.F. and Y.L.; investigation (literature review), S.G.; data curation, X.W.; writing—original draft preparation, J.F.; writing—review and editing, Y.L. and W.Z.; supervision, J.Z. All authors have read and agreed to the published version of the manuscript.

Funding: This research was funded by the Science and Technology Commission of Shanghai Municipality, grant number 25DZ3100700, for the project "Research on Models for Onboard Intelligent Mission Decision-Making and Health Management".

Institutional Review Board Statement: Not applicable. This study uses synthetic orbital scenes and does not involve humans or animals.

Informed Consent Statement: Not applicable.

Data Availability Statement: The original synthetic scene inputs, experiment configurations, and scheduling and estimation source code are available at https://github.com/terry-fjp/Tracking-Accuracy-Constrained_Scheduling (accessed on 19 October 2026). Additional numerical checks and sensitivity results are provided in Appendices A–C.

Acknowledgments: The authors acknowledge the support of the Science and Technology Commission of Shanghai Municipality. During the preparation of this work, the authors used OpenAI Codex (version 26.901.51231) to assist with translation and code development. The authors reviewed the code, simulation results, and entire manuscript and take full responsibility for the content.

Conflicts of Interest: The authors declare no conflicts of interest.

Appendix A. Numerical Checks of the Implemented Equations

Table A1 summarizes the independent numerical checks.

Equations (1)–(5) specify two-body dynamics, gravity gradients, transition and process-noise approximations, Joseph updates, and the angle–range Jacobian. Equation (7) covers the full hold interval without credit from later visits. Equation (13) includes the case where the warning threshold is already reached at observation end. Dividing times by Δt in Equation (17) matches the grid-count scoring in the implementation. The denominator of Equation (20) sums the three strategy success rates, ensuring normalization.

Independent checks use scene 2101, target 2, at $t = 20$ s. A 1 m central difference checks the gravity gradient and nonlinear angle–range Jacobian. Batch Joseph updates are compared with sequential updates of independent measurement components. Second-order transition and process-noise approximations are compared with a block matrix exponential at the frozen gravity gradient [18]. Fast 200 s covariance propagation is compared with stepwise recursion. Relative errors use Frobenius norms, except Jacobian rows, which use Euclidean norms. These checks establish local numerical agreement, not error bounds over the entire state space.

Table A1. Independent numerical checks of the implemented equations.

Check	Relative Error or Value
Gravity gradient versus central difference	4.166×10^{-10}
Maximum measurement-Jacobian relative error	7.154×10^{-10}
Batch versus sequential Joseph update	2.100×10^{-16}
Second-order transition approximation	1.499×10^{-9}
Local process-noise approximation	2.141×10^{-8}
Fast versus stepwise 200 s propagation	1.439×10^{-14}

Joseph updates were also applied to 100 random positive-definite priors. The smallest posterior eigenvalue was 6.609×10^{-6}, with no negative eigenvalues. Regression checks

passed for a 0 s hold and for the full 200 s propagation beyond the planning horizon. The reproducibility project retains the scripts, results, and equation revision record.

Appendix B. Observation Scenarios

Appendix B.1. Observer Constellation and Reference Frame

Every scene uses a Walker-Delta 72/8/1 constellation: eight planes, nine observers per plane, circular altitude 800 km, and inclination 60°. Adjacent planes differ by 45° in ascending-node longitude; within-plane phase spacing is 40°, with a 5° interplane phase offset. For plane index $p = 0, \ldots, 7$ and satellite index $j = 0, \ldots, 8$, $\Omega_p = \Omega_0 + 45°p$ and $u_{p,j} = u_0 + 40°j + 5°p$. Observer identifiers are $9p + j$, starting at 0. Scene-specific seeds generate Ω_0 and u_0 (Table A2). Displayed angles are rounded to two decimals; calculations use full stored precision.

The common epoch is 8 September 2026, 00:00:00 UTC, with planning start $t = 0$. Targets and observers follow two-body motion in a conceptual Earth-centered inertial frame. States use m and m/s; stored orbital angles use degrees. Planning covers 0–900 s, with initialization deadline 120 s and error evaluation at every 0.2 s point over 120–900 s. The independent 200 s hold check may extend beyond 900 s and is never truncated there.

Appendix B.2. Target Families and Initial Uncertainty

For each scene, two reference observers are sampled without replacement from the 72 satellites. An independent target family is generated around each reference orbit and phase. Alternating target identifiers split 12, 18, or 24 targets into families of 6, 9, or 12. Reference satellites only define initial target orbits; they are not predetermined service assignments. They need not occupy different orbital planes. Targets have neither rigid-formation constraints nor a shared orbital state.

Each target altitude is drawn independently and uniformly from 610–690 km. Relative to its reference observer, inclination and ascending-node offsets are independently uniform over $[-0.025°, 0.025°]$, and phase offsets over $[-1.1°, -0.55°]$. These local families intentionally create competition among common-field acquisition, serial ranging, and revisits. The two generated families are not fixed common-field clusters: Equation (14) is checked for each assigned satellite and actual service interval.

Per-axis position and velocity prior standard deviations are independently uniform over 800–1200 m and 70–110 m/s. Initial covariances are diagonal, with equal variances across the three axes and zero position–velocity cross terms. Targets share neither states nor posteriors. These are per-axis standard deviations, not three-dimensional RMSEs. Scene files store nominal orbits and prior parameters; separate random streams generate Monte Carlo initial errors and subsequent noise.

Appendix B.3. Settings of the Eight Scenes

Scene identifiers 2101–2108 also serve as orbit-generation seeds. Table A2 lists target counts, satellite laser budgets, process-noise spectral densities, and initial constellation orientation. Scenes 2107 and 2108 use higher process noise to cover faster uncertainty growth. Since initial orbits also vary, differences between scenes cannot isolate target count, budget, or process noise as a single causal factor. Methods within a scene use identical inputs.

Table A2. Initial configurations of the observation scenarios.

Scene	Targets	Targets/Family	B_s (s)	q_a (m^2/s^3)	Ω_0 (°)	u_0 (°)
2101	12	6	75	10^{-4}	164.00	38.94
2102	12	6	120	10^{-4}	10.30	52.35
2103	18	9	120	10^{-4}	228.63	341.77
2104	18	9	180	10^{-4}	41.87	237.15
2105	24	12	180	10^{-4}	49.69	335.57
2106	24	12	240	10^{-4}	125.84	166.44
2107	12	6	120	5×10^{-4}	64.07	127.25
2108	18	9	180	5×10^{-4}	280.88	164.96

Appendix B.4. Observing Opportunities and Resource Limits

Geometric availability is checked at 0.2 s intervals using range $\leq$ 2000 km, off-nadir angle $\leq 60°$, an unobstructed line of sight, and tracking rate $\leq 5°$/s. In Table A3, an observer is geometrically available if any target meets these conditions at any grid point in the horizon. This union count does not include complete measurement duration, setup, acquisition, or calendar conflicts. It therefore differs from the number of executable tasks.

The optical field is a square with 1° angular sides; the angular-bisector test uses its inscribed circle. The laser ranges targets serially at 5 Hz. Reference angle and range standard deviations are 100 µrad and 50 m. M1–M4 are 50 µrad/30 m, 100 µrad/60 m, 150 µrad/90 m, and 200 µrad/120 m, crossed with P1–P4 (Section 4.7). Slews assume zero initial and final rates, with limits 5°/s and 2.5°/s^2, plus 2 s settling and 5 s acquisition. Continuous laser measurements last at most 60 s, and cumulative budgets apply per satellite. Common-field service can remove repeated acquisition but adds no parallel laser channel and supplies no joint measurement to unselected targets.

Table A3. Geometrically available observing resources and satellites used by the reference plans.

Scene	Reference IDs	Available Observers	Observers per Target	Targets with One Observer	Observers Used
2101	69/5	4	1–3	6	3/3/3
2102	60/20	4	2–3	0	2/2/3
2103	1/61	2	1–1	18	2/2/2
2104	6/69	4	1–2	8	3/3/3
2105	11/48	3	1–2	22	2/3/3
2106	40/8	6	2–3	0	6/6/6
2107	35/44	2	1–1	12	2/2/2
2108	37/41	2	1–1	18	2/2/2

The final column of Table A3 gives the number of distinct observers used by F20, AD, and IP-AD, respectively, under P1 and reference noise. Reference identifiers follow the zero-based convention of the data files. Every target in scenes 2103, 2107, and 2108 has only one geometrically available observer during the horizon. This limits resource substitution: a 72-satellite constellation does not imply that every satellite can serve any target at any time.

Appendix B.5. Data Mapping and Scope

The reproducibility project stores scenes in `data/scenes/2101.json` through `2108.json`, using format version 1.0. Each observer and target has a locally unique identifier. Orbit fields distinguish geocentric radius, inclination, ascending-node longitude, and initial phase; target records also contain both prior standard deviations. Fields and units are documented in `docs/DATA_FORMAT.md`. Regenerating scenes from their identifiers,

counts, and complete parameters reproduces every stored numerical input. The shared script `tools/describe_scenes.py` checks Tables A2 and A3. Targets were neither selected nor excluded using the reported outcomes.

Position and velocity RMSE remain the primary accuracy measures. Acceptance pairs are 500 m/25 m/s, 1000 m/50 m/s, 1500 m/75 m/s, and 2000 m/100 m/s. Scene files store the P1 planning thresholds of 500 m and 25 m/s; `config/experiments.json` overrides them for P2–P4. Changing planning thresholds requires new schedules, whereas changing acceptance thresholds only reclassifies actual errors. All scenes are synthetic, with no real catalog or flight measurements. Ephemerides, target associations, and availability of scheduled measurements are assumed known. Illumination, reflectance, laser returns, and actual acquisition failures are outside the model.

Appendix C. Supplementary Sensitivity Experiments

These experiments retain reference P1: planning and actual-error thresholds 500 m/25 m/s, noise 100 µrad/50 m, step 0.2 s, horizon 0–900 s, and evaluation 120–900 s. Each scheduling run uses 48 complete-plan evaluations and search seeds 3101–3105; every target has 200 nonlinear realizations. Means and sample standard deviations use five equal eight-scene means, not 200 independent scene replicates. All targets, failed visits, and poor outcomes remain included. The archived default plans and errors are reused only after exact plan reproduction for all 80 default runs and independent nonlinear spot checks. Default planning times are remeasured in this revision. The original baseline is retained; a post-review warning variant is assessed separately in Appendix New Search Seeds for a Selected Warning Variant.

Tables A4–A9 report the sensitivity, resource, and geometry results; Figures A1–A4 visualize the four sensitivity studies.

The hold study replans IP-AD and IP-F20-H separately at $H = 100, 200, 300, 500$ s. The heuristic study varies one factor at a time in IP-AD: warning 0.7/0.8/0.9; duration multiplier 1.25/1.5/1.75; weights (0.5,0.3,0.2), (0.6,0.2,0.2), (0.4,0.4,0.2), (0.6,0.3,0.1), and (0.4,0.3,0.3). The shared baseline gives nine distinct heuristic settings. These comparisons assess local sensitivity while retaining the reference weights and duration multiplier; the separately selected warning variant is evaluated on new search seeds below.

For population scaling, each seed 2101–2108 generates a 96-target mother scene before taking prefixes of 24/48/72/96. Shared targets preserve initial states, priors, and standardized random inputs; observers, resources, and noise remain fixed within each family. New identifiers 5101–5108 distinguish these families from the original scenes. The same total constellation has 72 observers, but local geometric service availability is much smaller. This is a workload study within synthetic orbital families, not a representative catalog or a capacity bound. The 96-target case is an overload boundary, not a nominal service requirement. Admission keys allow attempts to serve a subset, but rejection or lack of observations does not remove a target from the evaluation denominator. At this boundary, IP-AD has higher compliance but worse position RMSE than the fixed-duration control (Table A6); all results are retained. IP-AD admits a mean of 87.925 targets but observes only 35.425; 8.075 targets are excluded by admission and 52.500 are admitted without receiving an observation. Admission therefore does not certify sufficient service.

The J_2 experiment applies $J_2 = 1.08262668 \times 10^{-3}$ to targets and observers with a fixed inertial z pole, the same μ and R_E, and RK4 integration. Its perturbing acceleration is

$$\boldsymbol{a}_{J_2} = \frac{3J_2\mu R_E^2}{2r^5}\begin{bmatrix} x(5z^2/r^2 - 1) \\ y(5z^2/r^2 - 1) \\ z(5z^2/r^2 - 3) \end{bmatrix}.$$

The analytic acceleration Jacobian updates the matched estimator transition; process noise remains unchanged. The potential-gradient, finite-difference Jacobian, zero-J_2, and step-halving checks are saved with the reproducibility materials. Observer ephemerides are treated as exact. Revalidation uses nominal perturbed trajectories, not realization-dependent acquisition or return detection. Both propagations start from the same Cartesian states obtained from the original circular-orbit elements. The original $H = 200$ s calendars are frozen, including setup times; a new slew must fit the reserved setup interval. Entire visits failing acquisition, geometric visibility, shared-field, or transition checks are canceled without repair. Both estimators under J_2 truth use the same execution mask; changing only the estimator dynamics then isolates that comparison. The unperturbed mask check retains all visits in all 80 original calendars. This does not optimize new J_2 schedules or validate atmospheric drag.

Physical conditions share random initial deviations, process-noise draws, and standardized measurement-noise draws. Random draws occur on every grid step regardless of observations. Different dynamics share random inputs but need not share trajectory hashes. Cache keys include physical inputs, actual measurement calendars, dynamics/evaluator versions, protocol, and realization count. No incomplete run enters these summaries. Four processes use one BLAS thread each; planning CPU and wall times are recorded separately. Replay timings include cache reuse and are not standalone estimator benchmarks.

Table A4. Hold duration. Means ± sample standard deviations of five equally weighted eight-scene means. Compliance uses 500 m/25 m/s throughout 120–900 s.

Setting	Method	Position (m)	Velocity (m/s)	Compliance (%)
$H = 100$ s	IP-AD	21,679.9 ± 2371.4	47.56 ± 3.37	30.73 ± 0.59
$H = 100$ s	IP-F20-H	54,239.2 ± 3416.2	106.64 ± 8.26	21.20 ± 1.27
$H = 200$ s	IP-AD	18,443.8 ± 1432.6	57.10 ± 6.49	31.43 ± 1.06
$H = 200$ s	IP-F20-H	54,239.2 ± 3416.2	106.64 ± 8.26	21.20 ± 1.27
$H = 300$ s	IP-AD	21,018.0 ± 3822.5	64.68 ± 15.11	33.36 ± 0.20
$H = 300$ s	IP-F20-H	54,239.2 ± 3416.2	106.64 ± 8.26	21.20 ± 1.27
$H = 500$ s	IP-AD	41,545.6 ± 927.9	95.31 ± 4.17	29.40 ± 0.43
$H = 500$ s	IP-F20-H	85,607.1 ± 309.9	228.78 ± 1.17	0.00 ± 0.00

Table A5. Local heuristics. Means ± sample standard deviations of five equally weighted eight-scene means. Compliance uses 500 m/25 m/s throughout 120–900 s.

Setting	Method	Position (m)	Velocity (m/s)	Compliance (%)
$H = 200$ s	IP-AD	18,443.8 ± 1432.6	57.10 ± 6.49	31.43 ± 1.06
w = 0.7	IP-AD	16,794.5 ± 3368.1	48.51 ± 7.04	36.18 ± 0.77
w = 0.9	IP-AD	24,726.0 ± 4932.8	75.21 ± 18.88	28.42 ± 0.94
m = 1.25	IP-AD	20,115.2 ± 1881.2	68.92 ± 11.36	31.52 ± 1.15
m = 1.75	IP-AD	18,340.8 ± 4017.3	56.38 ± 12.54	31.36 ± 0.89
W622	IP-AD	18,183.0 ± 3857.9	55.61 ± 9.64	31.73 ± 1.11
W442	IP-AD	17,165.0 ± 1854.7	55.44 ± 8.34	31.80 ± 1.25
W631	IP-AD	20,366.0 ± 3804.3	62.38 ± 11.79	31.62 ± 0.92
W433	IP-AD	16,966.2 ± 4215.5	51.16 ± 12.28	32.01 ± 1.30

Table A6. Nested populations. Means ± sample standard deviations of five equally weighted eight-scene means. Compliance uses 500 m/25 m/s throughout 120–900 s.

Setting	Method	Position (m)	Velocity (m/s)	Compliance (%)
N = 24	IP-AD	46,930.4 ± 5071.8	121.07 ± 12.04	20.95 ± 0.23
N = 24	IP-F20-H	58,127.0 ± 3914.4	115.19 ± 8.24	10.11 ± 1.96
N = 48	IP-AD	86,590.6 ± 5120.0	233.77 ± 24.99	10.11 ± 0.22
N = 48	IP-F20-H	72,889.8 ± 586.5	138.47 ± 1.47	3.68 ± 0.24
N = 72	IP-AD	85,894.2 ± 1443.0	211.65 ± 4.39	6.28 ± 0.29
N = 72	IP-F20-H	77,891.5 ± 220.6	147.54 ± 0.41	2.00 ± 0.13
N = 96	IP-AD	86,605.8 ± 2120.5	206.95 ± 8.32	4.52 ± 0.11
N = 96	IP-F20-H	81,214.3 ± 178.8	153.52 ± 0.47	1.44 ± 0.06

Table A7. Fixed-calendar dynamics. Means ± sample standard deviations of five equally weighted eight-scene means. Compliance uses 500 m/25 m/s throughout 120–900 s.

Setting	Method	Position (m)	Velocity (m/s)	Compliance (%)
H = 200 s	IP-AD	18,443.8 ± 1432.6	57.10 ± 6.49	31.43 ± 1.06
H = 200 s	IP-F20-H	54,239.2 ± 3416.2	106.64 ± 8.26	21.20 ± 1.27
J2-two-body	IP-AD	18,770.3 ± 1397.9	57.78 ± 6.16	0.00 ± 0.00
J2-two-body	IP-F20-H	54,294.7 ± 3470.4	106.38 ± 7.74	0.00 ± 0.00
J2-J2	IP-AD	18,491.4 ± 1405.4	56.92 ± 6.16	31.41 ± 1.05
J2-J2	IP-F20-H	54,288.4 ± 3476.5	106.31 ± 7.76	21.02 ± 1.35

Table A8. Service and computational outcomes. Values are five-seed means; full seed standard deviations, occupied time, wall time, evaluation time, and scene–seed paired differences are retained in the source tables. Planning CPU excludes preprocessing and EKF replay. $T/L/U$ count timely, late, and unobserved targets per scene. Laser time is summed over observers. A dash denotes no scheduling optimization in fixed-calendar replay.

Setting	Method	T	L	U	Laser (s)	CPU (s)
H = 100 s	IP-AD	13.10	2.05	2.10	107.69	10.99
H = 100 s	IP-F20-H	7.40	3.38	6.47	313.50	17.38
H = 200 s	IP-AD	11.90	5.03	0.33	186.61	16.30
H = 200 s	IP-F20-H	7.40	3.38	6.47	313.50	17.13
H = 300 s	IP-AD	10.32	6.45	0.47	245.10	20.55
H = 300 s	IP-F20-H	7.40	3.38	6.47	313.50	17.50
H = 500 s	IP-AD	7.65	5.67	3.92	287.75	28.90
H = 500 s	IP-F20-H	0.00	3.00	14.25	65.00	2.41
H = 200 s	IP-AD	11.90	5.03	0.33	186.61	16.30
w = 0.7	IP-AD	11.90	5.10	0.25	185.62	16.94
w = 0.9	IP-AD	11.70	5.28	0.28	184.08	16.30
m = 1.25	IP-AD	11.90	5.22	0.12	187.51	14.15
m = 1.75	IP-AD	11.85	5.20	0.20	188.12	18.93
W622	IP-AD	11.90	5.00	0.35	186.34	16.58
W442	IP-AD	11.88	5.10	0.28	186.57	16.23
W631	IP-AD	11.82	5.05	0.38	184.76	16.58
W433	IP-AD	11.90	5.05	0.30	184.76	16.19
N = 24	IP-AD	12.25	9.43	2.33	231.89	32.48
N = 24	IP-F20-H	7.92	5.25	10.82	317.50	23.42
N = 48	IP-AD	12.22	17.23	18.55	287.56	106.54
N = 48	IP-F20-H	7.97	7.00	33.02	322.00	47.40
N = 72	IP-AD	12.43	20.40	39.17	306.10	188.54
N = 72	IP-F20-H	8.00	7.97	56.02	324.50	71.24
N = 96	IP-AD	12.55	22.88	60.58	314.95	270.94
N = 96	IP-F20-H	8.00	8.12	79.88	325.00	94.49
H = 200 s	IP-AD	11.90	5.03	0.33	186.61	16.30
H = 200 s	IP-F20-H	7.40	3.38	6.47	313.50	17.13
J2-two-body	IP-AD	11.90	5.03	0.33	185.48	–
J2-two-body	IP-F20-H	7.40	3.33	6.53	311.00	–
J2-J2	IP-AD	11.90	5.03	0.33	185.48	–
J2-J2	IP-F20-H	7.40	3.33	6.53	311.00	–

Across the eight scenes, the maximum deterministic target position differences are 309.75 m at 200 s and 5867.65 m at 900 s; corresponding velocity differences are 3.041 and 13.363 m/s. The visibility grid has 95 lost and 378 gained target–observer–time samples out of 44,721,936 samples. Observer differences and per-scene values are also reported.

Table A9. Maximum deterministic differences from two-body propagation. Columns report maxima over all eight scenes; visibility changes are reported in the text.

Bodies	Time	Position (m)	Velocity (m/s)
Targets	200 s	309.75	3.041
Targets	900 s	5867.65	13.363
Targets	Full horizon	5867.65	13.363
Observers	200 s	300.92	3.011
Observers	900 s	5755.44	12.617
Observers	Full horizon	5755.44	12.617

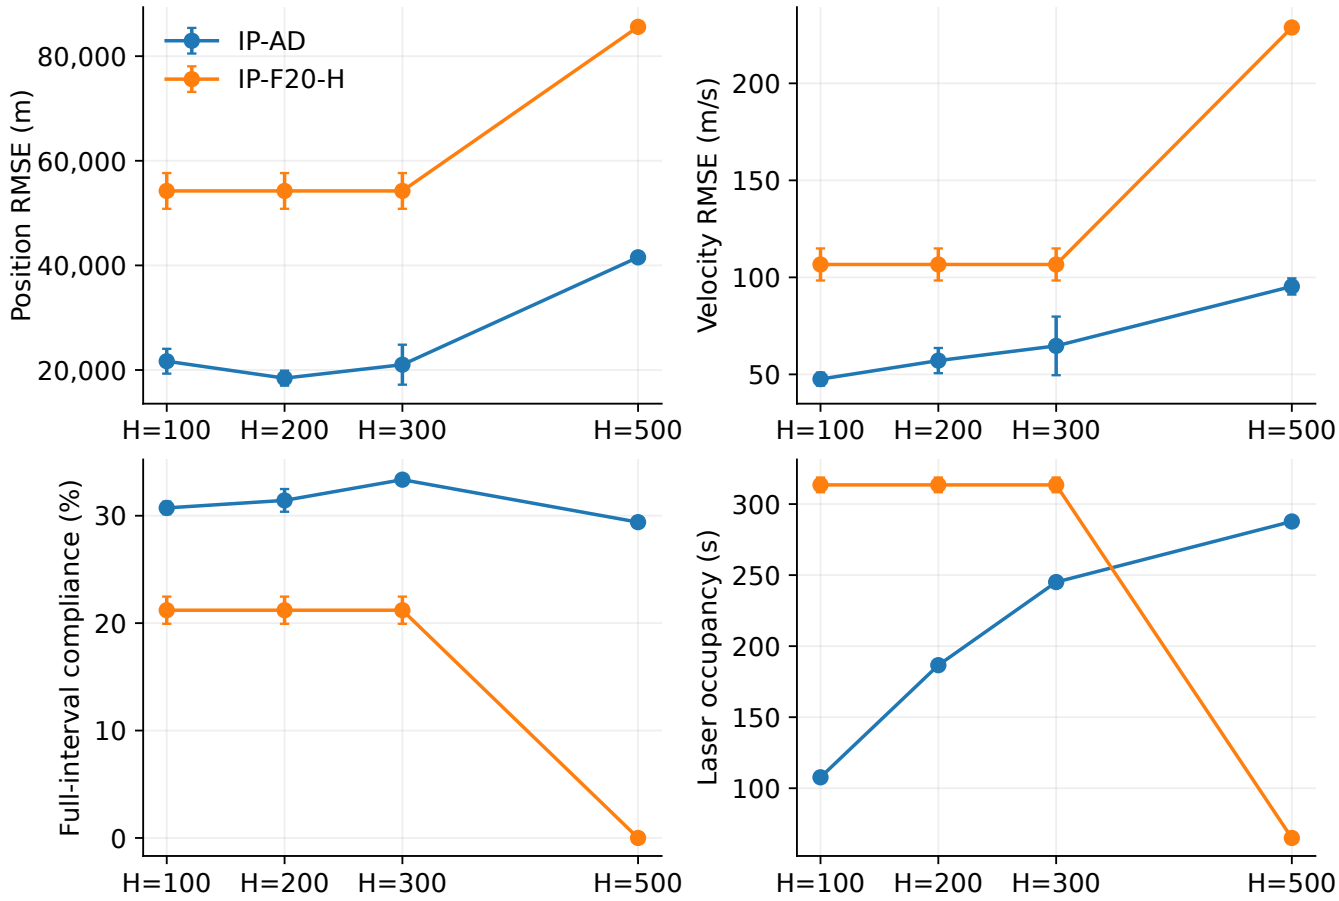

Figure A1. Hold duration. Points show five-seed means and bars show sample standard deviations of eight-scene means. All targets remain included; error bars are not confidence intervals.

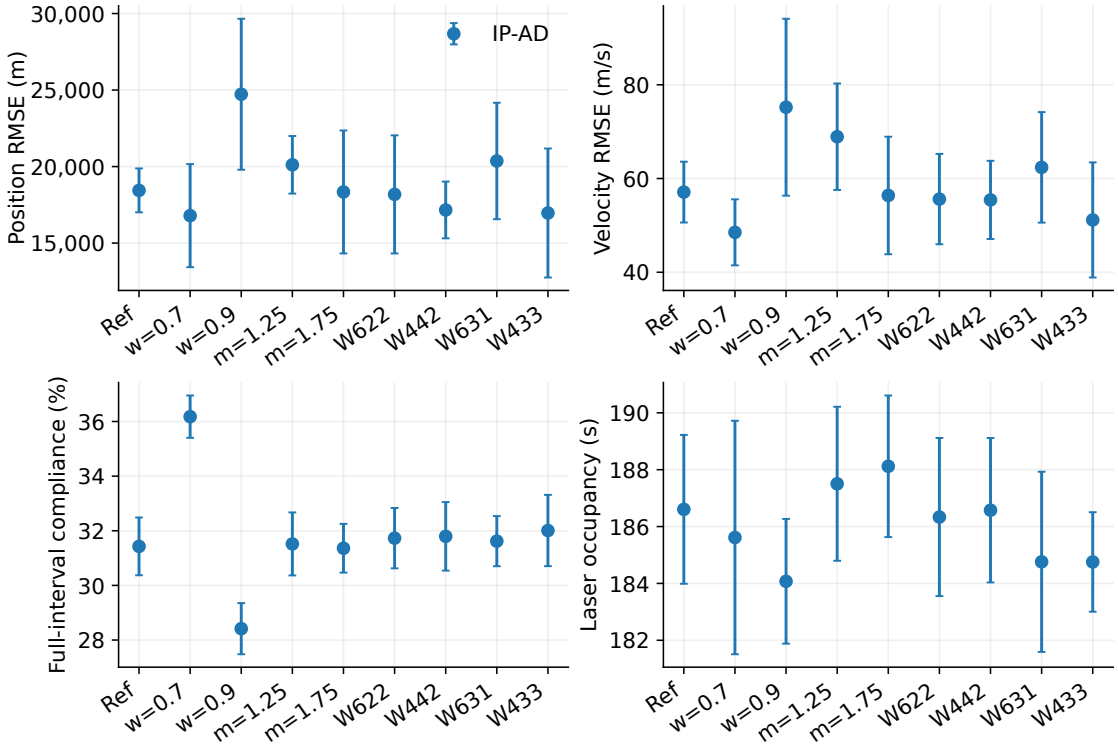

Figure A2. Local heuristics. Ref is the default setting; w is the warning threshold, m the duration multiplier, and Wabc denotes weights $(a/10, b/10, c/10)$. Points show five-seed means and bars show sample standard deviations of eight-scene means. All targets remain included; error bars are not confidence intervals.

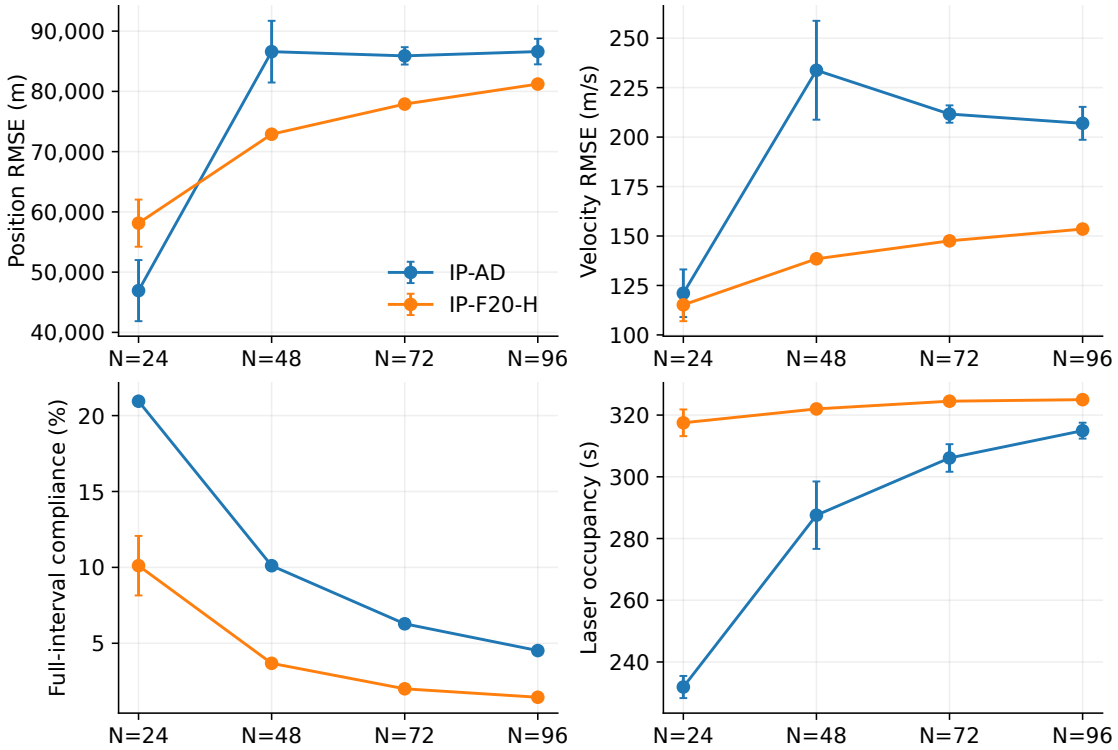

Figure A3. Nested populations. Points show five-seed means and bars show sample standard deviations of eight-scene means. All targets remain included; error bars are not confidence intervals.

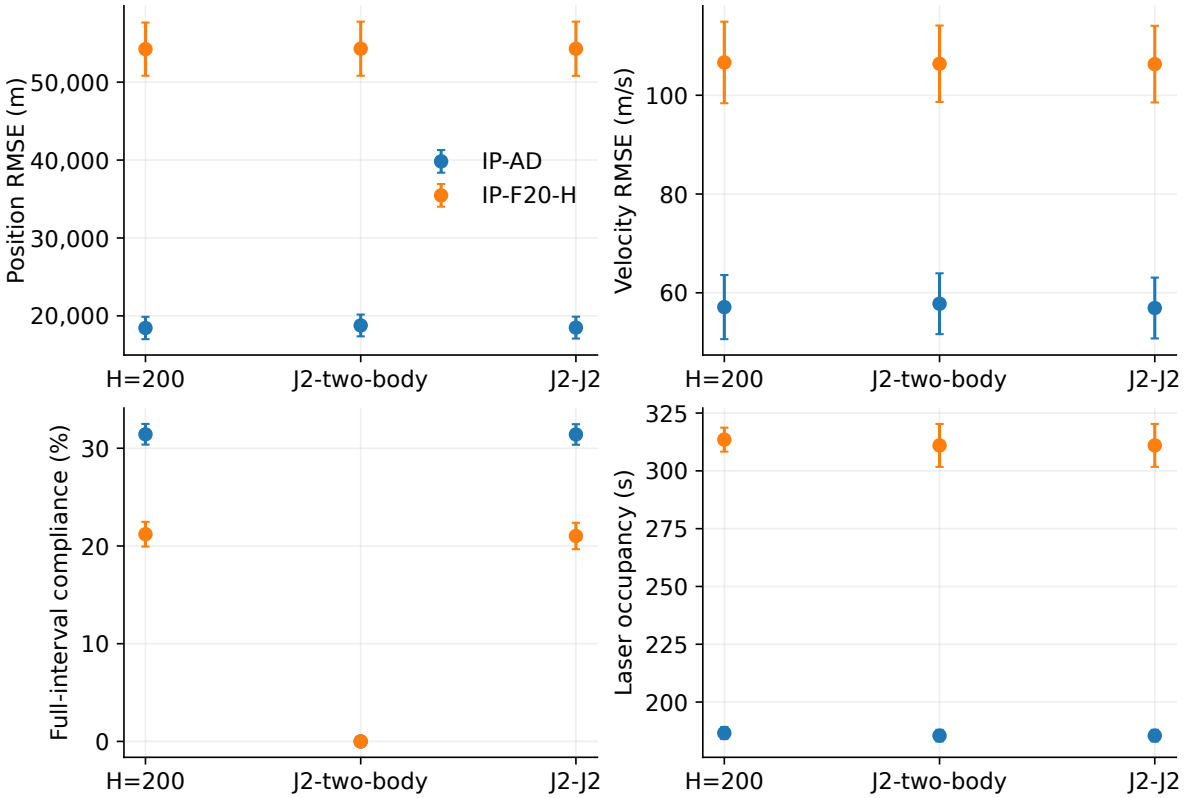

Figure A4. Fixed-calendar dynamics. Points show five-seed means and bars show sample standard deviations of eight-scene means. All targets remain included; error bars are not confidence intervals.

New Search Seeds for a Selected Warning Variant

Tables A10 and A11 report the accuracy and resource results. The nine-setting sensitivity study was used as selection data. Its warning 0.7 setting reduced mean position and velocity RMSE and increased compliance at similar laser use relative to warning 0.8. Only this warning setting was selected for IP-AD-W07; the hold remains 200 s, the multiplier 1.5, and the weights (0.5,0.3,0.2). The candidate and reporting protocol were frozen before evaluating search seeds 3106–3110. The original IP-AD and IP-F20-H controls were independently replanned with these same seeds, yielding 120 runs at 48 complete-plan evaluations each and 200 realizations per target. All eight original scenes and Monte Carlo inputs are shared with the selection study. Thus this check assesses variation across new optimization seeds; it does not validate transfer to unseen geometries or independent stochastic truth samples. All outcomes are retained. The original warning 0.8 remains the reference for previously reported experiments.

Table A10. Selected warning variant and matched controls on search seeds 3106–3110. Means ± sample standard deviations of five eight-scene means; the acceptance criterion remains 500 m/25 m/s throughout 120–900 s.

Method	Position (m)	Velocity (m/s)	Compliance (%)
IP-AD	18,293.9 ± 1658.1	52.85 ± 4.88	32.17 ± 0.62
IP-F20-H	54,172.8 ± 3341.0	107.06 ± 6.46	21.81 ± 1.36
IP-AD-W07	15,753.2 ± 2228.5	46.09 ± 6.29	36.49 ± 0.45

Table A11. Service and resource outcomes of the new-seed comparison. Means over all 40 scene–seed pairs; source tables also report sample standard deviations. $T/L/U$ denote timely, late, and unobserved targets.

Method	T	L	U	Laser (s)	CPU (s)
IP-AD	11.82	4.97	0.45	185.37	16.11
IP-F20-H	7.55	3.25	6.45	318.50	17.18
IP-AD-W07	11.95	4.75	0.55	185.65	16.68

Against original IP-AD, IP-AD-W07 has lower position RMSE in 27 of 40 pairs and higher compliance in 36 of 40 pairs. Mean paired differences are −2540.6 m in position RMSE and 4.32 percentage points in compliance. Pair counts are descriptive and are not independent statistical tests across geometries.

References

1. Shtofenmakher, A.; Balakrishnan, H. Optimal Tasking and Scheduling of Satellite Constellations for Space Situational Awareness. In *Proceedings of the AMOS 2025 Technical Papers*; Maui Economic Development Board: Kihei, HI, USA, 2025. [CrossRef]
2. Wei, Z.; Liu, S.; Feng, C.; Pan, X.; Chen, T. Distributed cooperative control for on-orbit assembly of heterogeneous spacecraft cluster. *IEEE Trans. Aerosp. Electron. Syst.* **2026**, 1–11. [CrossRef]
3. Li, B.; Sang, J.; Zhang, Z. A Real-Time Orbit Determination Method for Smooth Transition from Optical Tracking to Laser Ranging of Debris. *Sensors* **2016**, *16*, 962. [CrossRef] [PubMed]
4. Sang, J.; Bennett, J.C. Achievable debris orbit prediction accuracy using laser ranging data from a single station. *Adv. Space Res.* **2014**, *54*, 119–124. [CrossRef]
5. Ravago, N.; Jones, B.A. Risk-aware sensor scheduling and tracking of large constellations. *Adv. Space Res.* **2021**, *68*, 2530–2550. [CrossRef]
6. Li, Z.; Wang, Y.; Zheng, W.; Song, X. Sensor scheduling with orbital uncertainty for space target tracking. *Acta Astronaut.* **2023**, *205*, 23–32. [CrossRef]
7. Dhingra, N.K.; DeJac, C.; Neel, J.; Herz, A.; Wolf, T.; Jones, B. The Impact of Orbit Accuracy-Based Tasking on Sensor Network Efficiency. In Proceedings of the Advanced Maui Optical and Space Surveillance Technologies Conference, Maui, HI, USA, 27–30 September 2022.
8. Wang, X.; Wu, G.; Xing, L.; Pedrycz, W. Agile Earth Observation Satellite Scheduling Over 20 Years: Formulations, Methods, and Future Directions. *IEEE Syst. J.* **2021**, *15*, 3881–3892. [CrossRef]
9. Wu, G.; Wang, H.; Pedrycz, W.; Li, H.; Wang, L. Satellite observation scheduling with a novel adaptive simulated annealing algorithm and a dynamic task clustering strategy. *Comput. Ind. Eng.* **2017**, *113*, 576–588. [CrossRef]
10. Galloua, M.E.; Li, S.; Cui, J. Earth observation satellite imaging task scheduling with metaheuristics: Multi-level clustering and priority-driven pre-scheduling. *Adv. Space Res.* **2025**, *75*, 2929–2953. [CrossRef]
11. Shimane, Y.; Ho, K. Multi-Sensor Tasking for Ground-Based Space Situational Awareness via Job-Shop Scheduling Problem. *J. Astronaut. Sci.* **2025**, *72*, 61. [CrossRef]
12. Yuan, Y.; Yi, W.; Kong, L. Joint tracking sequence and dwell time allocation for multi-target tracking with phased array radar. *Signal Process.* **2022**, *192*, 108374. [CrossRef]
13. Li, B.; Tian, L.; Chen, D.; Liang, S. An adaptive dwell time scheduling model for phased array radar based on three-way decision. *J. Syst. Eng. Electron.* **2020**, *31*, 500–509. [CrossRef]
14. Han, X.; Yang, M.; Wang, S.; Chao, T. Continuous monitoring scheduling for moving targets by Earth observation satellites. *Aerosp. Sci. Technol.* **2023**, *140*, 108422. [CrossRef]
15. Xing, X.; Wang, S.; Liu, W.; Liu, C. A Multi-Satellite Multi-Target Observation Task Planning and Replanning Method Based on DQN. *Sensors* **2025**, *25*, 1856. [CrossRef] [PubMed]
16. Storn, R.; Price, K. Differential Evolution—A Simple and Efficient Heuristic for global Optimization over Continuous Spaces. *J. Glob. Optim.* **1997**, *11*, 341–359. [CrossRef]
17. Qin, A.K.; Huang, V.L.; Suganthan, P.N. Differential Evolution Algorithm With Strategy Adaptation for Global Numerical Optimization. *IEEE Trans. Evol. Comput.* **2009**, *13*, 398–417. [CrossRef]
18. Van Loan, C. Computing integrals involving the matrix exponential. *IEEE Trans. Autom. Control* **1978**, *23*, 395–404. [CrossRef]
19. Kalman, R.E. A New Approach to Linear Filtering and Prediction Problems. *J. Basic Eng.* **1960**, *82*, 35–45. [CrossRef]

20. Ropke, S.; Pisinger, D. An Adaptive Large Neighborhood Search Heuristic for the Pickup and Delivery Problem with Time Windows. *Transp. Sci.* **2006**, *40*, 455–472. [CrossRef]
21. Li, X.; Wei, Z.; Chen, T.; Zhang, A.; Kong, L. Path planning and control of attitude maneuver of flexible spacecraft on SO(3). *Aerosp. Sci. Technol.* **2026**, *178*, 113057. [CrossRef]

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.