Resilient Navigation Architectures and Proactive Routing for Unmanned Aircraft Under Uncertain Interference

Krishna R. Pattipati, *Life Fellow, IEEE*

***Abstract*—What does it take for an unmanned aircraft to keep its navigation within a corridor when satellite signals may be jammed, and what does it cost? We treat navigation resilience as a joint property of the navigation sources the aircraft carries, the error growth of its inertial system between fixes, the route it flies, and the decisions it makes in flight, and we study each link with a model chosen to isolate one mechanism. The planning law for inertial error growth must come from flight data: in simulated flight conditions (vibration and error processes prescribed in the truth model), laboratory values under-predicted the drift by a factor of three to five, and the identification window must match how fast conditions change. When every fix resets the error alike, route risk is governed by the longest outage rather than total exposure, and routes built on the threat's statistical law generalized better than routes fitted to sampled scenarios. Sensing jammers and replanning every 5 km raised the probability of meeting the requirement from 0.70 to about 0.90, and, in the same case (LEO layer, detection at the onset of denial), replanning as soon as a jammer was detected met the requirement on all 200 test fields with the same mean number of replans; sensing as often but replanning only every 5 km did not help, so the gain came from acting on detections when they occurred. With an exact belief for the modeled threat and sensor, the tested look-ahead methods (rollout, Monte Carlo tree search, and AND-OR search) showed no statistically detectable improvement over open-loop-optimal-feedback model predictive control (OLOF-MPC), which replans the best fixed route for the current belief at about one-hundredth of their computation, under either replanning rule and including rollout with OLOF-MPC as its base policy. Among the architectures evaluated, surviving the loss of any one navigation source required three independent sources, a navigation-grade inertial unit, and a route chosen for that loss: the same hardware that failed 21 of 200 test missions after losing satellite navigation on the direct route failed only once in 200 test missions on a route 22 km longer. All results are conditional on the stated models.**

***Index Terms*—Unmanned aircraft, GNSS jamming, resilient navigation, inertial navigation, adaptive Kalman filtering, chance constraints, robust route planning, model predictive control, rollout, Monte Carlo tree search, LEO PNT.**

I. Introduction

A. Motivation

Satellite navigation signals are weak enough that inexpensive jammers deny them over kilometers and spoofers mislead receivers silently [1], [2]. Unmanned aircraft systems (UAS) in contested airspace therefore carry antenna arrays, inertial sensors, cameras, and alternative radio receivers, and new low-Earth-orbit (LEO) navigation constellations promise much stronger signals [3], [4]. Part I [5] quantified what LEO constellations buy: about 25–33 dB of jamming margin depending on size, and few-centimeter carrier-phase positioning within a minute or two of a cold start with a large constellation, with integer-fixed solutions only when the satellite orbit and clock corrections are accurate to about 2 cm (with 5 cm errors, no LEO receiver fixed its integers). It also showed that the common practice of modeling an antenna array by counting its nulls is too pessimistic when jammers outnumber the array's elements, and too optimistic when a jammer lies close in direction to a satellite.

B. The question and a running example

This paper asks the following question: *what does it take for an unmanned aircraft to keep its navigation uncertainty within a corridor under uncertain interference, and what does it cost?* The answer depends on four things that are usually designed separately: the navigation sources carried (hardware), how fast the inertial error grows between fixes (estimation), the route flown (planning), and the decisions made in flight (control).

One example, developed in Section VIII, shows why the four things cannot be separated. An aircraft carries a LEO receiver with a small antenna array, an eLoran receiver, a thermal camera, and a navigation-grade inertial measurement unit. On the direct 139 km route, it met its navigation requirement on every one of 200 test fields; but when the satellite layer was lost, it failed on 21 of them. On a 161 km route through a more textured terrain, the same hardware failed on only 1 in 200 test fields after losing the satellite layer. The extra 22 km costs about 12 minutes of flight. Resilience is a property of hardware, estimation, and route together, and this paper develops the models needed to see it.

C. Approach: two studies with different aims

We use two levels of modeling, summarized in Table I. *Study A* isolates mechanisms: with one satellite layer and a simple model in which every fix resets the navigation error to the same small value, it asks what route risk should measure, how routes should be selected, and how much in-flight replanning helps. *Study B* evaluates architectures: it propagates

K. R. Pattipati is with Qualtech Systems, Inc., Rocky Hill, CT 06067 and the Department of Electrical and Computer Engineering, University of Connecticut, Storrs, CT 06269 USA (e-mail: krishna@teamqsi.com/krishna.pattipati@uconn.edu).

Working draft (revision 32, Part II of two), November 2026. The models and computations were developed with the assistance of an AI system; all tables are generated directly from the accompanying code and result files, and all results should be independently verified before publication.

TABLE I
THE TWO STUDIES.

	Study A (Sections V–VII)	Study B (Section VIII)
Purpose	Isolate outage, route-selection, and replanning effects	Compare navigation architectures and their routes
Sources	One satellite layer	LEO or GNSS, eLoran, camera, IMU grades
Navigation model	Every fix resets error to the same value; error then grows by the in-flight coasting law	Kalman filter with fixes of different quality and correlated source errors
Outcome measure	Horizontal RMS position uncertainty	Cross-track position uncertainty
Interpretation	Mechanism under controlled assumptions	Conditional system-level performance

the navigation error through fixes of different quality, with correlated source errors, explicit outages of whole source domains, and the array model of Part I. Study A's conclusions concern mechanisms under controlled assumptions; Study B's concern system-level performance under its assumptions. The simplified reset model of Study A should not be carried into the architecture results.

D. Paper organization and contributions

Section II provides the background and terminology used in this paper. The remaining sections make the following contributions.

1) A problem statement in which navigation error is propagated through fixes of different quality, explicit failure domains, and routes, flight policies, and architectures are designed jointly (Section III).
2) A comparison of ways to obtain the inertial coasting law under five kinds of mismatch, showing that measurement and process noise must be identified jointly and re-identified continually, as a streaming identifier does on board, over a window matched to the rate of change (Section IV).
3) Evidence that, under a common reset model, the longest outage rather than total unaided distance governs failure, an approximate Poisson model of outage risk with measured calibration, and an audit of chance-constrained routing (Section V).
4) A comparison of offline route-selection criteria on maritime robustness metrics, under the endurance limit and under misspecified threats (Section VI).
5) A partially observed formulation of proactive routing with an exact Bayesian belief, compared across seven policies, with imperfect sensing and an adaptive attacker (Section VII).
6) Pareto fronts over combinations of hardware and route (cost, nominal failure, failure after losing any one source, availability of centimeter-level accuracy), selected on validation data and reported on separate test data, with sensitivity to the array model and to costs (Section VIII).

II. Background and Terminology

Navigation sources. GNSS (GPS, Galileo, GLONASS, BeiDou) broadcasts from medium Earth orbit (*MEO*). A *LEO* navigation layer broadcasts from a few hundred kilometers up with much stronger signals; the satellites of a broadband constellation, such as Starlink, can also be used for positioning as *signals of opportunity*, with meter-level accuracy so far (Part I). *eLoran* uses ground transmitters at 100 kHz, accurate to about 10 m and hard to jam over wide areas [6]. An *inertial measurement unit* (IMU) integrates accelerometers and gyroscopes and cannot be jammed, but its error grows during *coasting* (navigation without fixes), roughly as the square or cube of time [7]; IMUs are graded industrial, tactical, and navigation (in increasing cost, mass, and accuracy). A *camera* gives absolute *map fixes* by matching the ground to a map over textured terrain (by day, or day and night with a thermal camera) and *visual odometry* (velocity) elsewhere.

Filters and identification. A *Kalman filter* fuses the IMU with fixes and reports an error covariance, its own estimate of its uncertainty. Its *innovations* (measurement minus prediction) are uncorrelated over time when its noise model is correct; *innovation-based identification* uses this to estimate the noise model from flight data. *Allan variance* characterizes sensors from static laboratory recordings.

Failure domains. A *failure domain* is a set of sources that can be removed together with a single cause (all satellite signals in a band, a terrestrial network, the camera in bad weather). The *N−1 criterion* requires that the mission still succeeds after losing any one domain.

Decision terms. A *Pareto front* is the set of non-dominated solutions. A *chance constraint* must hold with a given probability. *Minmax* minimizes the worst case over scenarios; *regret* subtracts the best cost achievable in each scenario in hindsight; $CVaR_{0.9}$ averages the worst 10% of scenarios [8]. A *Copeland score* ranks methods by pairwise wins across metrics. In a partially observed problem, the *belief* is the probability distribution over what cannot be seen; it is a sufficient statistic for optimal Bayesian decision-making. *Certainty equivalence* (CE) plans as if estimates were exact; *open-loop-optimal feedback* (OLOF) replans the best fixed plan for the current belief at every step; *model predictive control* (MPC) applies only the first step of each re-planned solution; *rollout* scores candidate actions by simulating a base policy, which amounts to one policy-improvement step of policy iteration in dynamic programming; *Monte Carlo tree search* (MCTS) grows a search tree by sampled simulations; an *AND-OR tree* alternates between the decision (OR) nodes, where the best action is chosen, and the chance (AND) nodes, where outcomes are averaged over samples. We call rollout, MCTS, and AND-OR search *look-ahead methods*: they simulate the consequences of each candidate action in sampled futures.

III. Problem Statement

A. Mission, threat, and environment

A fixed-wing UAS flies at 30 m/s and 300 m altitude across a 150×100 km theater, from $(5, 50)$ to $(144, 50)$ km (139 km in a straight line), with a 90-minute endurance

(162 km). Jammers form a Poisson field: a background density of one per 400 km^2 plus a dense band near $x = 75$ km; jammer strengths, given as equivalent (also called effective) isotropically radiated power (EIRP, the transmitter power times the antenna gain toward the receiver), are spread evenly on a logarithmic scale from 10 W to 100 kW. A Poisson field is the natural model when jammer positions are unknown and independent: the number of jammers in any region is random with a known mean, and what happens in one region tells nothing about another. 30% of the ground is featureless for a camera. Whether a jammer denies a satellite layer at a cell follows the physics of Part I: the denial radius for one jammer in Study A, and aggregate jamming with an element-level adaptive array model in Study B. In Study A, receiver protection is the jamming tolerance relative to a single-antenna GPS receiver, and it adds a signal advantage and an antenna-array contribution (Part I). The main case, 59 dB, combines a LEO layer's signal advantage of 29 dB, below Part I's 33 dB ceiling and corresponding to roughly 2,000 satellites at 42°N, with 30 dB of array nulling; the dense case, 45 dB, has the same array with a much weaker layer (about 15 dB); 69 dB pairs the 29 dB layer with a 40 dB array.

B. The navigation requirement and what "failure" means

The aircraft must stay within a 50 m corridor half-width. We translate this into a requirement on the navigation filter's own uncertainty: three standard deviations of position must stay within 50 m throughout the flight, $\max_t \sigma(t) \le 16.7$ m. We call a violation of this requirement a *mission failure*. It is a navigation-uncertainty event: it is neither the probability that the aircraft actually leaves the corridor (a first-passage probability that also involves the realized error and the flight controller) nor a certified integrity statement. Study A measures σ as the horizontal RMS position uncertainty (both axes, $\sqrt{2}$ times one axis), which is conservative for a corridor; Study B measures the cross-track uncertainty, the quantity a corridor constrains, and reports the horizontal RMS as a sensitivity.

C. Objectives and failure domains

An architecture x (satellite layer, terrestrial tier, user antenna, IMU grade, camera, flight policy) is evaluated by

$$\begin{aligned} f_1(x) &= \text{annualized cost},\\ f_2(x) &= \Pr\{\text{failure by day } \cup \text{ failure by night}\},\\ f_3(x) &= \max_{d\in\mathcal{D}(x)} \Pr\{\text{failure} \mid \text{domain } d \text{ lost}\},\\ f_4(x) &= 1 - \mathbb{E}[\text{flight fraction with a formal 5 cm}\\ &\qquad \text{LEO solution}], \end{aligned} \tag{1}$$

where $\mathcal{D}(x)$ contains the domains present in x (the satellite layer, eLoran, the camera). The union in f_2 is evaluated field by field: a field counts as a failure if the route fails by day or by night on it, which is in general more than the larger of the day and night failure probabilities. Each domain also fails on its own: eLoran is unavailable for a whole mission with probability 0.1 (a regional outage), and the camera with probability 0.1 (weather or obscurants). We call f_4 *LEO 5 cm formal-solution unavailability*. It is computed from Part I's convergence curve of the LEO receiver (Section III-E) and includes the initial (cold-start) convergence and re-convergence after every outage. Because the cold start is a fixed interval, a fraction favors longer flights; the representative designs therefore also report unavailable seconds. We report failure counts out of 200 fields with 95% Wilson intervals;[1] zero failures in 200 trials bound the failure rate only below about 1.9%.

D. Navigation error

In *Study A*, every fix resets the uncertainty to the same small value (the LEO fix), and during an outage it grows by the in-flight coasting law of the IMU: the uncertainty predicted by a Kalman filter with the true in-flight error model, including vibration (Section IV). The coasting law is increasing, so failure occurs exactly when an outage lasts longer than the time t^* the law takes to reach the requirement, that is, when an uninterrupted outage is longer than the coasting length $\ell^* = V\,t^*$, with $V = 30$ m/s the airspeed: 3.24 km for the industrial IMU and 5.93 km for the tactical one.

In *Study B*, the cross-track navigation error is propagated by a Kalman filter with a five-state inertial error model (position, velocity, accelerometer bias, tilt, gyroscope bias) whose noise reflects in-flight vibration, augmented with one first-order Gauss–Markov state for the correlated error of each absolute source (Appendix B gives the matrices and their discretization). In each 1 km cell (33 s of flight), the filter receives the updates available there:

- a *LEO fix* if the layer is not denied;
- otherwise an *eLoran fix* if eLoran is available: 2 m white noise error (zero mean, standard deviation of 2 m) plus a 4.6 m bias (the propagation correction error) correlated over 20 km of flight;
- if there is no LEO fix and the aircraft has a camera, also a *map fix* over textured ground (5 m white noise error (zero mean and standard deviation of 5 m) plus an 8.7 m bias correlated over 2 km) and a *visual-odometry* velocity update (0.3 m/s).

Fixes of different quality therefore reset the error to different levels, and repeated coarse fixes taken within the correlation length of their bias reduce the error little below that bias; only fixes far apart carry independent information. How the LEO fix is derived from Part I is set out next.

E. The interface with Part I

The LEO fix uses the carrier-phase study of Part I (1 Hz, industrial IMU, two frequencies, 10 cm ionosphere residual,

[1]For k failures in n trials, with $\hat{p} = k/n$ and $z = 1.96$ for 95%, the Wilson score interval is the set of rates p with $|\hat{p} - p| \le z\sqrt{p(1-p)/n}$, that is, $\big(\hat{p} + \frac{z^2}{2n} \pm z\sqrt{\hat{p}(1-\hat{p})/n + z^2/(4n^2)}\big)/(1 + z^2/n)$ [9]. Unlike the textbook (Wald) interval $\hat{p} \pm z\sqrt{\hat{p}(1-\hat{p})/n}$, it stays within $[0, 1]$ and keeps close to its nominal coverage for rare events; the Wald interval collapses to the single point 0 when $k = 0$ [10]. With $n = 200$, the 95% upper bound is 1.9% for no failures and 2.8% for one. A statement about three domains at once uses the simultaneous (Bonferroni) version, with each interval at confidence $1 - 0.05/3$ ($z = 2.39$); one failure in 200 then gives 3.7%.

2 cm orbit-clock residual). Part I reports the median formal 3-D uncertainty $\sigma_{3D}(t)$ of the receiver's own solution as a function of the time t since it acquired the signals. Let $s(t) = \sigma_{3D}(t)/\sqrt{3}$ be the per-axis value and s_b its converged level (8–13 mm, the median over the last five minutes of Part I's 30-minute runs). A LEO fix on the cross-track axis at time t_k after (re)acquisition is modeled as

$$\begin{aligned} z_k &= x_k + b_k + v_k, \\ \mathrm{Var}(b_k) = s_b^2, \quad \mathrm{Var}(v_k) &= \max\{s^2(t_k) - s_b^2,\, (1\,\mathrm{mm})^2\}, \end{aligned} \tag{2}$$

where x_k is the aircraft's position, b_k a first-order Gauss–Markov error with 600 s correlation time, and v_k white. Except where the variance floor is active, the variances add to $s^2(t_k)$, so each fix matches the imported uncertainty of the receiver, and the correlated part is not counted twice. After convergence, the imported curves fluctuate about their converged level and fall up to 9% below s_b at times; there the floor of $(1\,\mathrm{mm})^2$ applies, and the fix variance slightly exceeds $s^2(t_k)$. Four further points define the interface.

1) *Correlation.* Because b_k decorrelates over 600 s, fixes taken within a few minutes of each other cannot reduce the error much below s_b, while fixes separated by much more than 600 s carry new information. The 600 s value is the correlation of Part I's orbit-clock residual; the receiver's position error need not have the same correlation, because the filter, the changing geometry, and satellite turnover transform it. We checked the surrogate against Part I's own per-axis error sequences (30 start times each for 1,000 and 5,000 satellites). Its marginal standard deviation matched within about 10% (for example 1.29 versus 1.40 cm and 0.81 versus 0.82 cm at 10–20 min), but the actual error decorrelated much faster: its autocorrelation was 0.48–0.50 at 33 s (the spacing of fixes here) and 0.17–0.22 at 2 min, against 0.92–0.94 and 0.79–0.81 for the surrogate. The surrogate therefore lets successive fixes average down less than the actual errors would. For estimating a nearly constant quantity, this makes it pessimistic. However, stronger correlation does not in general give a larger covariance in a dynamic estimator, so we do not claim it is conservative. We checked its effect directly: with the LEO correlation time reduced from 600 s to 33 s, close to the measured decorrelation, the test failure counts of all 39 front pairs, nominal and after losing each domain, were unchanged, and the median peak cross-track uncertainty changed by 0 to $+1.8\%$; the faster-de-correlating errors gave slightly *larger* peaks, which illustrates that the surrogate is not uniformly conservative. The effect is confined to the centimeter-level part of the uncertainty, far below the 16.7 m requirement; f_4 (Section III-C) uses the receiver curve directly and is unaffected.
2) *Update interval.* Part I runs at 1 Hz; the navigation filter here takes one LEO fix per 1 km cell (33 s). The fix is the receiver's current solution, not a raw measurement, so the update interval does not change its uncertainty, only how many correlated samples are used.
3) *Shared IMU.* Part I's receiver uses the same industrial-grade inertial model as an aid, so its solution and the aircraft's IMU are not strictly independent; the navigation filter neglects this cross-correlation. Part I found that in-flight inertial aiding gives the same convergence as no motion knowledge, which suggests that the shared information, and hence the neglected correlation, is small; we have not computed the cross-correlation itself.
4) *Reacquisition.* The cold-start curve applies only at take-off. After an outage, the receiver restarts warm: it keeps its troposphere and orbit-clock states and an inertial position prior, and only its ambiguities are new. Part I simulated warm restarts after outages of 15, 60, and 180 s; with 5,000 satellites, the float solution regained 5 cm in 15–18 s instead of 25 s, and with 1,000 satellites in 42–61 s instead of 55 s. These warm start curves were simulated for a receiver that had tracked for 900 s before the outage. Part I also found, for 60 s outages, that two minutes of continuous prior tracking already gave the full float benefit, while an interruption after only 30 s recovered almost like a cold start. We therefore use the cold-start curve at take-off and after any outage that follows less than 120 s of continuous tracking immediately before it, and otherwise the warm-restart curve for the shortest simulated outage at least as long as the actual one. The 120 s threshold is supported by the tested 60 s outages; applying it to other outage lengths and to repeated interruptions is an approximation, not a general convergence threshold. After outages longer than 180 s, whose inertial position prior is worse than any simulated, we use whichever of the 180 s and cold-start curves is slower; this is an extrapolation rule, not an established bound. After 120–300 s of prior tracking, the 900 s curves somewhat overstate how fast the fixed solution returns, which affects only the centimeter-level part of the fix noise; the time to a 5 cm float solution, which defines f_4, is the same. The same curves set re-convergence in f_4.

The objective f_4 (Section III-C) uses the same curve, but it measures the LEO receiver's solution, not the fused navigation uncertainty: it is the share of flight time during which the LEO receiver does *not* have a formal 5 cm solution, counting cold start and re-convergence after every outage.

F. Route and policy

An offline route is a path on the 1 km grid chosen before flight. A proactive policy replans in flight using what the aircraft has sensed (Section VII). The policies are first compared with replanning every 5 km; because that interval is longer than the industrial IMU's coasting length (3.24 km), Section VII-D then tests intervals of 2 and 1 km, and sensing every 1 km cell while replanning as soon as a new jammer is detected (otherwise every 5 or 2 km). We recommend this detection-driven rule: in every case it matched replanning every 1 km to within one field in 200, with far fewer replans (29–45 per flight, against about 140).

IV. The Coasting Law: Datasheet or Flight Data?

Question. During outages the aircraft coasts on its IMU, and every route and policy depends on how fast its error grows. Can this law be taken from a datasheet, or must it be identified from flight data, and with which method?

A. Why the coasting law matters, and how it enters planning

The planning studies need one input from the inertial system: the predicted growth of position uncertainty during an outage, $\sigma(t)$ for t seconds since the last fix. Its use is a chain:

flight data $\rightarrow$ identified noise levels $\rightarrow$ predicted $\sigma(t)$ during an outage $\rightarrow$ route and architecture evaluation.

The identified quantities are the process noise (velocity random walk and total acceleration random walk) and the measurement noise of the aiding source; together with the filter model, they give $\sigma(t)$. The planning studies (Sections V–VIII) use the law an ideal identifier would deliver: the true in-flight model, including vibration. In particular, Study B's eight-state filter (Appendix B) uses the true in-flight noise intensities; the single-axis polynomial (α-β-γ) model that gain-first identifies in this section is not run inside it. This section is a separate supporting study: it shows how far a laboratory law is from that, and which identification methods recover it from flight data. A planner given a law that is too optimistic accepts routes with outages it cannot bridge; one too pessimistic rejects good routes.

B. Setting

The truth is a single-axis inertial error model with velocity and angle random walk inflated by in-flight vibration, accelerometer-bias random walk, a constant gyroscope bias, and (in one scenario) Gauss–Markov colored accelerometer and gyroscope errors, aided at 1 Hz by position fixes. In three stationary scenarios, each method identifies a coasting law from one aided record and is then deployed as its own filter for 30 min of aided flight followed by an outage; its predicted error envelope is compared with its actual error (100 Monte Carlo runs per scenario). In two nonstationary scenarios, the vibration level doubles, either in a step 20 min before the outage or as a ramp over the last 60 min, and each method identifies on a sliding window of 5, 15, 30, or 60 min ending at the outage; the window length plays the role of the forgetting horizon of a streaming identifier [11].

C. Methods

- *Datasheet:* a Kalman filter with the laboratory (Allan-variance) noise levels, without vibration.
- *Inflation:* the datasheet process noise scaled by a factor chosen so that the normalized innovation squared averages one on the record. It is the simplest adaptive fix, but it attributes every excess of innovation variance to the IMU.
- *ALS (batch reference):* autocovariance least squares [12] over 20 lags of the innovations of the datasheet filter. It fits the noise variances of a given model so that the predicted innovation autocovariances match the observed ones. We give it the correct model in an identifiable form (velocity random walk, total acceleration random walk, measurement noise). In one position-aided axis, the accelerometer-bias and tilt random walks enter identically and cannot be separated; fitting them separately made one of the two negative in every record while their sum was estimated without bias, so only the sum is estimated. The unconstrained solution is used when it is nonnegative and the non-negativity is imposed otherwise (it was negative in 5–42% of records, most often on the noisier and shorter records). A deployed system would not know the true model structure, however.
- *Gain-first:* the finite-record recipe of [13] for a polynomial (α-β-γ) filter. Instead of guessing noise levels and computing the gain from them, it chooses the gain that makes the innovations white, $\hat{W} = \arg\min_W \sum_{\ell=1}^{20} [\hat{r}_\nu(\ell; W)/\hat{r}_\nu(0; W)]^2$ over stabilizing gains, where $\hat{r}_\nu(\ell; W)$ is the ℓ–lag cross-correlation. Once the gain is optimized, it then recovers the noise levels, $\hat{R} = (1-\hat{\alpha})\hat{S}$, $\hat{Q}_C = 4\hat{\gamma}^2\hat{S}/T^5$, positive by construction. It needs the polynomial filter order but no IMU error model. The order itself need not be fixed in advance: filters of several orders (α, α-β, α-β-γ [13]) can run in parallel as a bank, and the order selected by its maximum a posteriori probability given the innovations; MAP selection over a bank of filters with different state dimensions picked the correct model in 100 of 100 trials per truth in simulation, a stress test rather than a proved result [11]. On board, the same principle runs in a streaming form, with piecewise-constant gains updated online over a sliding window [11]; the nonstationary scenarios below apply the finite-record recipe to sliding windows as a stand-in.
- *Oracle:* a Kalman filter with the true model at outage conditions.

D. What to expect

The datasheet omits vibration, so it should be optimistic. Inflation should be calibrated when the aiding noise is as assumed, but should over-predict drift when the aiding is noisier than assumed, because it blames the aiding's excess on the IMU. ALS and gain-first estimate measurement and process noise together, so they should not make that error. After a change of conditions, any method that averages over too long a window should under-predict, and one that averages over too short a window should be noisy.

E. Results

What it gives. Table II and Fig. 1 show the results at 60 s of coasting. The datasheet model under-predicted drift 2.9–5.5 times and its 2σ envelope covered only 29–47% of outcomes. In stationary conditions, all adaptive methods matched the oracle (coverage 0.95–0.96). When the aiding noise during calibration was four times larger than assumed, inflation over-predicted drift fourfold (34.9 m predicted versus 8.5 m actual),

whereas gain-first and ALS stayed calibrated (coverage 0.97 and 0.96). After a change of conditions (Table III), the window length decided calibration. With a 15 min window, which after the step contains only post-change data, gain-first and ALS covered 0.94–0.96 (oracle 0.97). A 60 min window mixed old and new conditions and under-predicted drift by 11–20% (coverage 0.88–0.91), and a 5 min window was too noisy for gain-first (0.88–0.89). The same ordering held for the ramp.

Insight. A laboratory coasting law is three to five times too optimistic in flight, enough to turn a route the planner believes safe into one that fails. The law must come from flight data, and the identification must estimate measurement and process noise jointly: inflation fails exactly when the aiding noise is misjudged. Given the true model structure, batch ALS and gain-first were equally accurate here; gain-first reached that accuracy from the polynomial filter order alone, with covariances valid by construction. When conditions change, calibration is restored by continuing identification whose window (or forgetting horizon) matches the rate of change: here about 15 min, long enough to average the noise and short enough to exclude stale data. Gain-first's streaming form [11] provides exactly this tuning parameter; ALS would have to be rerun on windows.

V. What Should Route Risk Measure?

A. Longest outage versus total exposure

Question. Route planners commonly minimize or constrain the total unaided distance X (the length of route without fixes). Does X rank routes the way mission failure does?

What to expect. In Study A failure is, by construction, the event that one uninterrupted outage exceeds ℓ^*. Ten separate 1 km gaps give $X = 10$ km and no failure; one 4 km gap gives $X = 4$ km and a failure. So X and failure need not agree, and the question is how much they disagree in practice. We confirmed the construction in the code (it held in all 4,000 pairs of 40 routes and 100 fields examined); this is a verification, not a finding. The informative test is across routes.

We evaluated 150 endurance-feasible routes on 200 held-out fields (LEO layer with 59 dB of protection, no camera, industrial IMU).

What it gives. Across routes, the Spearman correlation of failure with mean total unaided distance was 0.62 and with mean longest outage 0.82 (Fig. 2). With a camera, whose fixes break outages into short pieces, total unaided distance ranked routes almost perfectly (0.985). Despite the moderate correlation without a camera, the routes selected by X and by outage risk coincided or showed no statistically detectable difference in our instances (Table IV).

B. A Poisson model of outage risk

Question. Can the risk that some outage exceeds the coasting length be computed along a route, so that route search can minimize it directly, and how accurate is such a model?

What to expect. A model that counts only outages caused by a single jammer should be accurate when the effective jammer field is sparse, because long outages are then rarely formed by several overlapping jammers, and should underpredict when it is dense.

To plan on outage duration, the planner needs a risk it can add up along a route. Define for each jammer the event that it *alone* denies a continuous stretch of the route longer than ℓ^*. Because this event depends only on that jammer's own position and power, the jammers for which it occurs are themselves a Poisson field. This is the *thinning* property of Poisson fields [14]: if each point of a Poisson field with intensity $\lambda(\mathbf{y})$ is kept independently with a probability $q(\mathbf{y})$ that depends only on its own attributes $\mathbf{y}$ (here position and power), the kept points again form a Poisson field with intensity $\lambda(\mathbf{y})q(\mathbf{y})$, and the removed points another independent Poisson field. The number of jammers that alone deny a long stretch is therefore Poisson. If its mean is $\Lambda_1(p)$ for route p, the probability that at least one occurs is $1-e^{-\Lambda_1(p)}$. This is an exact *lower bound* on failure, because outages formed by several overlapping jammers can only add failures.

Computing Λ_1 exactly requires the route's geometry. For a long straight route, far from its endpoints, a jammer at perpendicular distance h with denial radius d covers a chord of length $2\sqrt{d^2-h^2}$, which exceeds ℓ^* when $d \geq \sqrt{h^2+\ell^{*2}/4}$. Counting such jammers along the route gives the additive approximation

$$\Lambda(p) = \int_p ds \int dh\, \lambda_J(s,h)\, \Pr\{d \geq \sqrt{h^2+\ell^{*2}/4}\}, \qquad \Pr\{\text{fail}\} \approx 1 - e^{-\Lambda(p)}, \tag{3}$$

with $\lambda_J(s,h)$ the jammer density at along-route position s and offset h, and $d = d(P)$ the denial radius of a jammer of EIRP P, random through the EIRP distribution f_P used in (4). It is additive along the route (with an orientation-dependent line density) and therefore usable in route search. On bent routes, a jammer can be counted on several segments, and chords are truncated at the endpoints, so (3) approximates Λ_1 but is not a bound; its accuracy must be measured.

What it gives. With 59 dB of protection (a sparse effective field), (3) predicted the failure of 150 routes on 2,000 fresh fields with mean absolute error 0.008 (maximum 0.026); for the two selected routes it over-predicted slightly (0.288 versus 0.277 and 0.257 versus 0.232, the latter about 2.6 binomial standard errors), consistent with double counting on bent routes. In dense fields, it under-predicted (0.89 versus 0.98 at 45 dB), because overlapping jammers also form long outages. Routes selected by (3) were identical to those selected by expected exposure at 59 dB and showed no statistically detectable difference at 45 dB (Table IV).

C. Chance constraints

Question. If routes are constrained to keep unaided distance below a threshold with a stated probability, does the constraint hold on new fields, and what does it take?

What to expect. A route's unaided distance is a sum over its cells, and neighboring cells are denied together by the same jammer, so the terms are positively correlated. Treating cells as independent understates the variance of the sum and

TABLE II

COASTING-LAW ESTIMATORS: PREDICTED MEDIAN / ACTUAL RMS POSITION ERROR (M) AFTER 60 S OF COASTING, AND COVERAGE OF THE PREDICTED 2σ ENVELOPE (TARGET 0.95); 100 RUNS PER SCENARIO. EACH ESTIMATOR'S IDENTIFIED MODEL RUNS IN ITS OWN FILTER, SO THE ACTUAL ERRORS DIFFER ACROSS COLUMNS.

Scenario	Datasheet	Inflation	ALS	Gain-first	Oracle
Stationary, vibration $\times 3$	1.5/4.4 (0.47)	4.1/3.9 (0.95)	4.0/3.9 (0.95)	3.9/4.0 (0.95)	3.9/3.9 (0.96)
Colored (Gauss–Markov) errors	1.5/8.4 (0.29)	7.7/6.9 (0.96)	7.5/7.0 (0.96)	7.4/7.1 (0.96)	7.2/6.9 (0.97)
Aiding noisier than assumed (calibration)	1.5/4.4 (0.47)	34.9/8.5 (1.00)	4.7/4.8 (0.96)	4.7/4.7 (0.97)	3.9/3.9 (0.96)

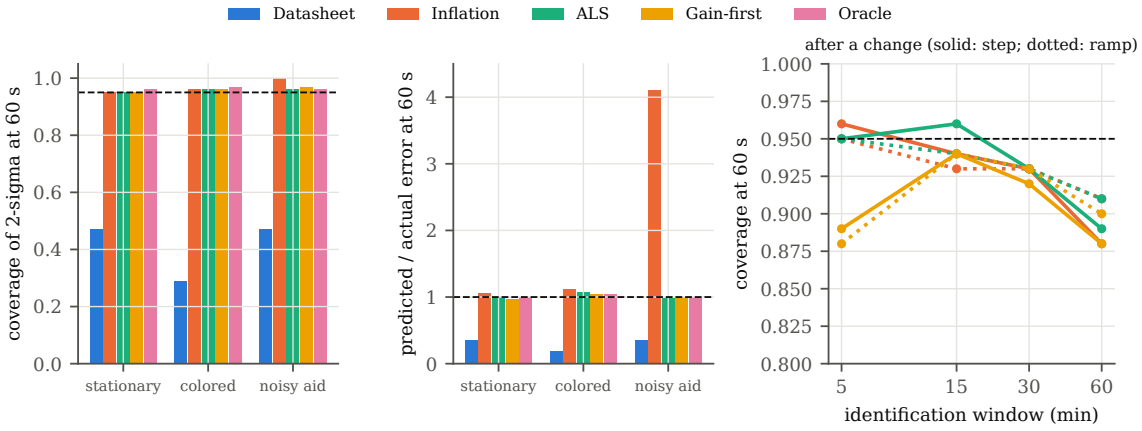


Fig. 1. Coasting-law estimators: coverage of the predicted 2σ envelope at 60 s (left) and ratio of predicted to actual error (middle) in three stationary scenarios; coverage versus identification window after a change of conditions (right).

TABLE III

IDENTIFICATION ON A SLIDING WINDOW AFTER A CHANGE OF CONDITIONS: COVERAGE OF THE PREDICTED 2σ ENVELOPE AT 60 S OF COASTING (TARGET 0.95) AND, IN PARENTHESES, MEDIAN PREDICTED OVER ACTUAL RMS ERROR; 100 RUNS PER SCENARIO.

Change	Method	Window (min): coverage (predicted/actual)			
		5	15	30	60
Step, 20 min before	Gain-first	0.89 (1.03)	0.94 (1.10)	0.92 (0.94)	0.88 (0.80)
	ALS	0.95 (1.14)	0.96 (1.14)	0.93 (0.98)	0.89 (0.82)
	Inflation	0.96 (1.11)	0.94 (1.09)	0.93 (0.93)	0.88 (0.84)
	Oracle		0.97 (1.14)		
Ramp over last 60 min	Gain-first	0.88 (1.02)	0.94 (1.04)	0.93 (0.96)	0.90 (0.86)
	ALS	0.95 (1.12)	0.94 (1.08)	0.93 (1.00)	0.91 (0.89)
	Inflation	0.95 (1.12)	0.93 (1.03)	0.93 (1.03)	0.91 (0.93)
	Oracle		0.97 (1.14)		

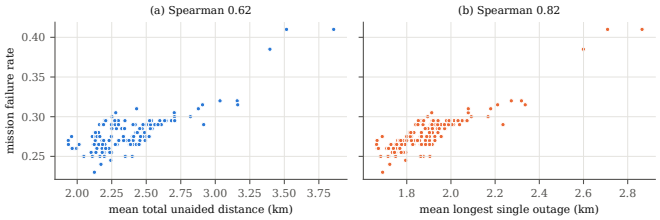


Fig. 2. Mission failure versus (a) mean total unaided distance and (b) mean longest outage, for 150 routes (Study A, no camera).

should make the constraint optimistic; assuming full correlation overstates it and should give safe but long routes. The exact covariance should lie in between.

A chance constraint keeps exposure below a threshold with a stated probability. The chance-constrained route minimizes the length subject to $\mathbb{E}[X] + z_\eta \sqrt{w^\top \Sigma w} \le \lambda$, which is the Gaussian form of $\Pr\{X \le \lambda\} \ge \eta$. Here, $X = \sum_n w_n D_n$, where D_n indicates that cell n is denied and w_n is the length of route in that cell, Σ is the covariance of the D_n, λ the exposure threshold, and z_η the standard normal quantile at confidence η. Neighboring cells are denied together by the same jammer, so the denials are correlated, and the covariance Σ matters. For the disc (Boolean) jammer model, it is exact: $\mathrm{Cov}(D_n, D_m) = (1 - p_n)(1 - p_m)(e^{\Lambda_{nm}} - 1)$, with p_n the denial probability of cell n and Λ_{nm} the expected number of jammers covering both cells. The variance $w^\top \Sigma w$ couples every pair of cells on the route, so it does not add up edge by edge as a label search requires; the route search therefore uses an additive approximation: each cell carries its own variance plus its summed covariance with the cells within about 11 km of it along a straight line (the largest over four grid directions), which is what a straight route through the cell would accumulate.

What it gives. Over 12 thresholds and confidence levels, the covariance-informed search met its target on held-out fields in 8 of 10 feasible cases; treating cells as independent met it in 2 of 12, and the fully correlated model in 6 of 6 feasible cases (2–14 km longer). The audit (Table V) shows that the additive approximation's standard deviation is within 6% of the exact one (variance within 12%), and the exact standard deviation 5% below the empirical one (variance about 10% low); the misses occurred at high confidence, where both the Gaussian tail and the 10% variance shortfall of the exact model make the constraint optimistic; we have not separated the two

TABLE IV
ROUTES SELECTED BY EXPECTED EXPOSURE (X) AND BY OUTAGE-RUN RISK (3), EVALUATED ON 2,000 FRESH FIELDS (PAIRED SIGN TEST).

Case, length budget	X-route: length, failure	Run-risk route: length, failure (predicted)	Paired p
59 dB, industrial IMU, $\leq$145 km	144.8 km, 0.277	144.8 km, 0.277 (0.288)	1.00
59 dB, industrial IMU, $\leq$162 km	148.9 km, 0.232	148.9 km, 0.232 (0.257)	1.00
45 dB, industrial IMU, $\leq$145 km	144.8 km, 0.989	144.8 km, 0.989 (0.929)	1.00
45 dB, industrial IMU, $\leq$162 km	151.4 km, 0.972	149.8 km, 0.979 (0.890)	0.14
45 dB, tactical IMU, $\leq$145 km	144.8 km, 0.982	144.8 km, 0.982 (0.894)	1.00
45 dB, tactical IMU, $\leq$162 km	151.4 km, 0.963	149.8 km, 0.969 (0.845)	0.19

TABLE V
AUDIT OF CHANCE-CONSTRAINED ROUTES.

Quantity	Value
Additive / exact s.d. (cov.-informed)	1.019 (1.00–1.06)
Exact / empirical s.d.	0.948
Feasible, Gaussian tail, exact moments	15/28
Feasible, Cantelli, exact moments	2/28
Cantelli-feasible, target violated	0

contributions. A distribution-free Cantelli bound was feasible for only 2 of 28 routes.

Insight. The mechanism is the shape of the coasting law. Inertial uncertainty grows faster than linearly with time (roughly as t^2 to t^3), so one outage of 4 km costs far more than four of 1 km, although both add 4 km of unaided distance; a fix in between resets the growth. Under a reset model, failure is therefore a statement about the longest outage; with fixes of different quality (Study B), aiding quality and the uncertainty accumulated before an outage matter too. Total unaided distance is a usable surrogate here, and with a camera an excellent one, but it can prefer one long gap to many short ones. Risk models and constraints should be stated in outage duration or navigation covariance, and the calibration of a chance constraint on X concerns X, not failure.

VI. OFFLINE ROUTE SELECTION

Question. Before flight, the jammer positions are unknown. How should that uncertainty be represented, and which selection criterion gives a route that performs well on fields not seen during planning?

A. Algorithms and uncertainty representation

Before flight, the planner must choose one route without knowing where the jammers are. Multiobjective search finds the routes that trade length against exposure, and the uncertainty can enter it in two ways: as a set of sampled jammer fields, on which each route's exposure is averaged, or as the analytic Poisson prior, which gives each cell's denial probability directly.

What to expect. The search algorithm should affect only speed, since exact methods return the same front. The representation should matter more: a finite sample of fields contains chance gaps, places where no sampled jammer happens to be, and a search that minimizes exposure on the sample is rewarded for finding them.

What it gives. LazyLTMOA* [15] returned the same Pareto sets as NAMOA* [16] three to four times faster (1.4 versus 4.0 s; 18.5 versus 66.6 s). A front built from 40 sampled jammer fields overfit: its least-exposure route had 0.06 km expected exposure on those fields and 2.7 km on new ones, because it threaded gaps that existed only in the samples. 200 sampled fields gave 2,127 near-tied routes in 860 s; the analytic Poisson prior gave 12 routes in 0.4 s (Fig. 3). These timings are for a small graph: 15,000 cells with eight-connected moves (about 1.2×10^5 directed edges) and three objectives.

Insight. Sampling turns route selection into a fitting problem, with the familiar failure of fitting: a rich route set fitted to few samples learns their accidents. More samples reduce the overfitting but multiply the near-tied routes and the computation. The analytic prior is the limit of infinitely many samples: it has no chance gaps, so a route that looks good under it is good on average on new fields, and it is two thousand times faster here. The cost of multiobjective search is driven by the number of non-dominated partial paths kept in OPEN, which grows with graph size, with the number of objectives, and with how many near-ties the objectives create; the 2,127 near-tied routes from 200 sampled fields are an instance of the last. On much larger road and mission-planning graphs, we have observed OPEN grow faster than sequential label-setting search can process it. At that scale the remedies are approximate search with ε-dominance [17], pruning by resource limits (as the endurance limit does here, Section VI-B), fewer or aggregated objectives, and parallel multiobjective shortest-path algorithms, including ordered parallel [18] and massively parallel [19] search and multi-attribute graph processing applied to maritime routing [20].

B. Endurance, lighting, and worst-case formulations

What to expect. The endurance limit removes the long detours around the dense band, so it should cost success, and more so as the limit tightens. A day camera gives map fixes only in daylight, so a route chosen for the day can rely on fixes that are absent at night.

Endurance limits trim the graph to cells with forward plus backward shortest distance within the limit [21]: at 90, 85, and 82 min, 49%, 30%, and 19% of cells remained, and the best success with a camera fell from 1.00 to 0.96 and 0.93. Within the limit, the best daytime route (161 km; success 1.00 by day, 0.67 at night) was not the most robust to lighting: a 149 km route achieved 0.95 by day and 0.78 at night. The worst-case

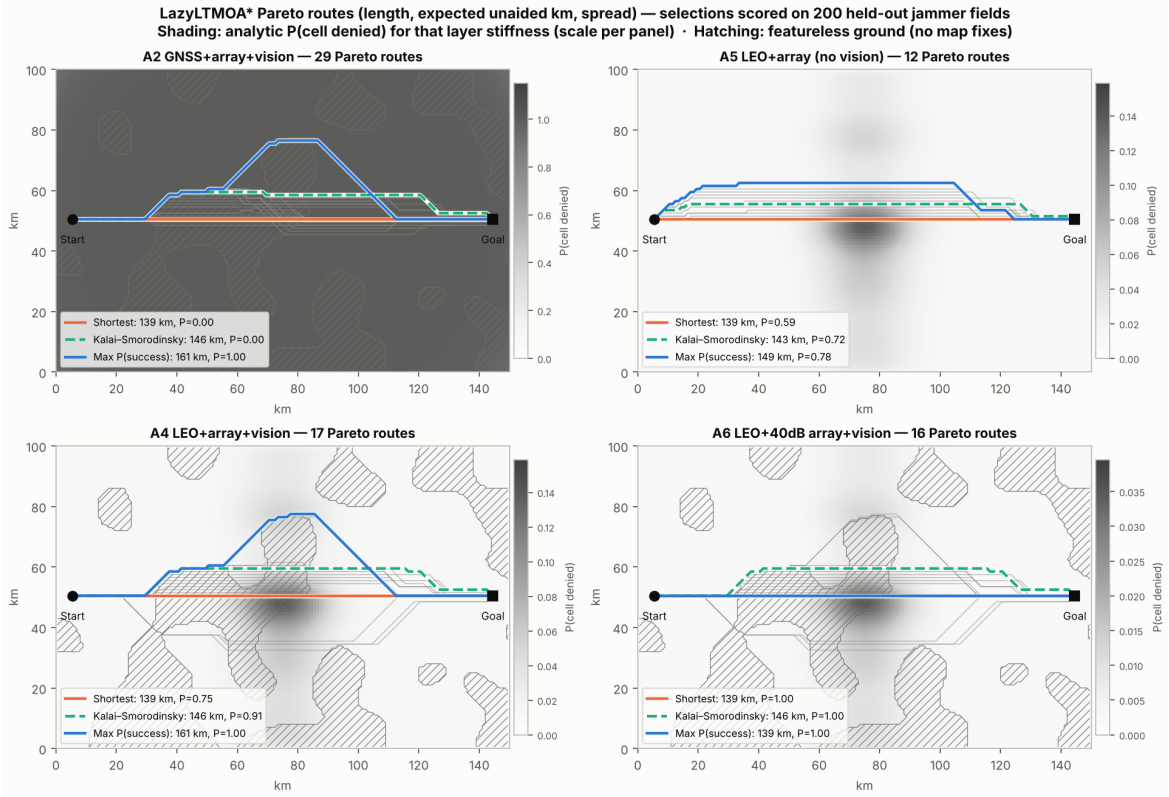


Fig. 3. Study A: Pareto routes (gray) for four user configurations, with the shortest, Kalai–Smorodinsky, and most reliable routes highlighted; shading is the analytic denial probability, hatching featureless ground. (Fronts computed without the endurance limit.)

(minmax) and minmax-regret routes were sought by Benders decomposition, which alternates between a master problem that proposes a route by label search and a subproblem that finds the sampled field on which that route does worst and adds it as a constraint (a cut) to the master; without a camera, it did not converge (gap 1.21–8.83 km) and its routes exceeded the endurance limit, so below we use the best feasible minmax route found among 1,340 candidates from LazyLTMOA*.

Insight. The best route by day is the best route *with* map fixes: it can cross long satellite-denied stretches because the camera bridges them over textured ground. At night, those fixes vanish, and the same stretches become outages. A route that is slightly worse by day but avoids such stretches is more robust to lighting; night operation is a degraded mode, and as in Section VIII, the route should be chosen for it as well.

C. Comparison on robustness metrics

Question. Which route-selection criterion gives the most robust route when judged by the metrics of robust routing used in [22]?

What to expect. A criterion built on the threat's statistical law uses information that holds for every future field; criteria built on sampled fields can exploit accidents of the sample. We therefore expect the analytic prior to generalize better, unless the law itself is wrong, which we test with mis-specified threats.

All routes satisfy the endurance limit. Because a nearly zero-exposure route exists in hindsight on almost every field (mean 0.015 km), worst regret approximately equals worst exposure and value of information approximately equals mean exposure here; the signal-to-noise ratio is a monotone function of the RMS. The Copeland score therefore uses only length, mean, worst, CVaR, RMS, and failure.

What it gives. Table VI. The analytic-prior route ranked first and had the lowest failure (0.22), but its 95% interval overlaps those of the next methods. Under mis-specified threats (Table VII: clustered jammers with the same mean density, 10 dB stronger jammers, twice the density, the dense band shifted 15 km), it remained best or tied; paired tests against the next-best route gave $p = 0.02$ (nominal), 0.67 (clustered), 0.84 (stronger), 0.02 (denser), and 0.04 (shifted). The minmax route's training worst case of 8.83 km became 13.5 km on new fields.

Insight. When the threat's statistical law is known, or approximately known, routes built on it generalize better than routes fitted to sampled scenarios, and this held under the mis-specifications tested. Minmax and regret over independently sampled fields optimize a rare-event statistic that new draws do not reproduce. This differs from weather routing, where the scenarios are the forecast ensemble and minmax hedges against forecast error [22]; against an adversary who chooses the worst case, minmax remains the right criterion (Section VII-F).

TABLE VI
ROBUSTNESS METRICS FOR ROUTE-SELECTION METHODS (CONFIGURATION WITHOUT CAMERA, 59 DB; 200 HELD-OUT FIELDS; UNAIDED DISTANCE IN KM; ALL ROUTES WITHIN THE 162 KM ENDURANCE LIMIT). COPELAND SCORE OVER LENGTH, MEAN, WORST, CVAR, RMS, AND FAILURE.

Method	Length (km)	Mean	Worst	$\text{CVaR}_{0.9}$	RMS	Failure [95% CI]	Copeland
Analytic prior: min mean	148.9	2.10	12.0	8.24	3.44	0.22 [0.17, 0.28]	8
200 sampled fields: min mean	162.0	2.19	10.3	8.74	3.63	0.27 [0.21, 0.34]	6
40 sampled fields: min mean	161.8	2.35	12.9	8.68	3.80	0.28 [0.22, 0.34]	3
Analytic prior: bargaining NS / CS	144.8	2.40	12.4	8.96	3.81	0.29 [0.23, 0.36]	2
Minmax exposure = minmax regret (best known)	156.4	2.30	13.5	9.14	3.85	0.26 [0.20, 0.32]	1
Chance constraint, covariance-informed	144.8	2.48	13.6	9.42	4.00	0.30 [0.24, 0.37]	-2
Analytic prior: bargaining KS	143.1	2.82	15.2	9.90	4.30	0.28 [0.22, 0.34]	-5
Training $\text{CVaR}_{0.9}$ selection	157.0	2.47	17.5	9.90	4.11	0.29 [0.23, 0.35]	-5
Shortest / independent-cell chance constr.	139.0	3.86	17.0	11.88	5.45	0.41 [0.34, 0.48]	-8

TABLE VII
MISSION FAILURE OF THE SELECTED ROUTES ON 1,000 FIELDS DRAWN FROM MISSPECIFIED THREAT GENERATORS.

Route (length)	nominal	clustered	power +10 dB	density $\times 2$	band shift
Analytic prior, min mean (149 km)	0.243	0.221	0.883	0.408	0.248
SAA-200: min mean exposure (162 km)	0.294	0.228	0.888	0.483	0.287
SAA-40: min mean exposure (162 km)	0.307	0.245	0.893	0.499	0.320
Minmax (best known) (156 km)	0.297	0.231	0.886	0.469	0.285
Training CVaR (157 km)	0.289	0.249	0.890	0.479	0.282
Bargaining NS (145 km)	0.285	0.265	0.921	0.459	0.283
Bargaining KS (143 km)	0.313	0.288	0.941	0.544	0.320
Chance constraint, covariance-informed (145 km)	0.309	0.282	0.931	0.493	0.288
Shortest (139 km)	0.427	0.383	0.974	0.696	0.462

VII. PROACTIVE ROUTING

Question. If the aircraft can detect jammers in flight and replan, how much does that raise mission success, and how elaborate must the replanning be?

A. Formulation

An offline route cannot react to jammers that turn out to lie on it. An aircraft that detects jammers, by direction finding with its array, can replan. This is a decision problem under partial observation: the aircraft never sees the whole field, only what it has sensed so far. The state at decision k is position, time since the last fix, distance flown, and belief b_k: the detected jammers and the intensity of undetected ones. A jammer is detected within 30 km, or only at the onset of denial (within 1.5 times its denial radius).

For the modeled Poisson field and ideal detection, the belief has an exact form; localization errors and false alarms, considered later, would require a different observation model. Because detection depends only on each jammer's own position and power, the same thinning argument (Section V-B) shows that the undetected jammers are again a Poisson field, with the prior intensity thinned by the event "not yet detected":

$$\lambda_k^{\mathrm{u}}(\mathbf{x}, P) = \lambda_J(\mathbf{x}) f_P(P)\, \mathbf{1}\{R(P) < d_k(\mathbf{x})\}, \qquad (4)$$

with $d_k(\mathbf{x})$ the distance from $\mathbf{x}$ to the nearest point where the aircraft has sensed and $R(P)$ the detection range of a jammer of power P. In words: a region the aircraft has looked at from close range is now known to be free of undetected jammers that are strong; weak jammers could still be there if the aircraft was not close enough to detect them. The posterior denial probability of a cell combines the detected jammers with this undetected intensity. Every 5 km, the policy chooses the next segment, subject to the endurance limit. The goal is to minimize the probability of failure; some policies optimize it directly by simulation, others a risk-weighted route cost that stands in for it (see Policies below).

B. Policies

The policies range from simple to elaborate. *Static:* the offline route. *CE-MPC:* plans around the detected jammers only, ignoring undetected ones. *OLOF-MPC:* plans the best fixed route for the current posterior and applies its first segment. CE-MPC and OLOF-MPC minimize a route cost, $\sum_e L_e\,(1 + \beta\,\bar{c}_e)$ over route edges e of length L_e, with $\beta = 20$ and cell cost $c = \mathbf{1}\{\text{known jammer}\}$ for CE-MPC and $c = p + \sqrt{p(1-p)}$ for OLOF-MPC, where p is the posterior denial probability; the second term is the mean plus one standard deviation of a cell's denial indicator, a risk surrogate rather than the failure probability itself, and the unit term is the length penalty. This surrogate is not monotone in p: it peaks at $p = (1 + 1/\sqrt{2})/2 \approx 0.85$ and falls to 1 at $p = 1$, so a cell certain to be denied costs less than one denied with probability 0.9. We checked the consequence on 200 fields per case. The routes OLOF-MPC planned contained cells with $0.85 < p < 1$ in at most 0.6% of its decisions, and replacing c by its monotone envelope (held at its peak value for $p > 0.85$) left success unchanged on every field in all three cases (0.915, 0.270, and 0.390). We therefore recommend the monotone envelope for subsequent use; the tables report the original penalty, with which the comparisons were run, and the outcomes are the same. The risk term

itself matters: the risk-neutral penalty $c = p$ with the same weight gave nearly identical success with the LEO layer (0.920 against 0.915, one field in 200) but lower success in the two dense cases (0.175 and 0.255; it lost on 25 and 34 fields and won on 6 and 7, paired sign tests $p < 0.001$). *Rollout:* candidate first segments from several plans are scored over 24 posterior samples with common random numbers, with an open-loop base (follow the candidate) or a closed-loop base (CE-MPC simulated to the goal). *MCTS:* POMCP-style search over four plan-following macro-actions, three levels deep, 48 simulations per decision, branching on the number of new detections [23]. *AND-OR:* the four plans averaged over five sampled worlds, then CE and OLOF compared over three fresh worlds at the second level [24]. Rollout, MCTS, and AND-OR search, the look-ahead methods, estimate the failure probability itself, by propagating the navigation uncertainty along simulated trajectories in sampled worlds. The closed-loop ones (and all rollouts of Section VII-D) choose the first segment s that minimizes

$$J(s) = \frac{1}{M} \sum_{m=1}^{M} \left[F_m + 0.1 \min(r_m, 1.5) \right] + 10^{-4} L_s, \quad (5)$$

over M sampled worlds, where $F_m = 1$ if the simulated flight violated the requirement or exceeded the length budget and 0 otherwise, $r_m = \sigma_{\max,m}/16.7$ m is its peak navigation uncertainty relative to the requirement, and L_s is the length of the candidate segment in km. The peak-uncertainty term (at most 0.15 per sample) prefers larger margins, and the length term is a small penalty. Neither is a strict tie-breaker: with $M = 24$, one more sampled failure adds $1/24 \approx 0.042$ to J, while the peak term of two candidates can differ by up to 0.15, so the objective can accept an extra sampled failure in exchange for lower peak uncertainty. Failure minimization therefore does not have lexicographic priority; it is weighted against the margin. We kept the margin term because, with 24 samples, a failure count alone cannot distinguish a candidate whose flights peak just below the requirement from one with ample margin. The open-loop rollout scores each candidate plan by its sampled failure fraction plus 10^{-4} per km of the whole plan. OLOF-MPC is itself a form of rollout: replanning the best open-loop plan and applying its first segment is one-step look-ahead with the open-loop optimal plan as base policy [25], [26]. The rollout policies above differ from it in their base policies, their candidate sets, and in estimating the failure probability by simulation rather than minimizing a surrogate cost. A *clairvoyant* reference knows all jammers in advance; it searches only three cost weightings, so it understates what full knowledge allows.

What to expect. Replanning should help most where the offline route is most exposed and detection comes early enough to turn away. Look-ahead methods can anticipate future detections, which simple replanning cannot: they can, for example, prefer a path that will reveal whether a region is clear before committing to it. But they are only as good as the belief they sample from, and what they add is the value of information that future detections bring. If the belief is exact, detections come late, and the option set is small, that value, and hence the gain from deeper search, may be small.

TABLE VIII
MISSION SUCCESS OF PROACTIVE POLICIES ON THE SAME 80 HELD-OUT FIELDS (MEAN ROUTE LENGTH, KM), BY PROTECTION AND DETECTION MODE.

Policy	59 dB, onset	45 dB, onset	45 dB, 30 km
Static (offline route)	0.70 (144)	0.03 (152)	0.03 (152)
CE-MPC	0.78 (141)	0.26 (155)	0.39 (152)
OLOF-MPC	0.90 (149)	0.28 (157)	0.40 (155)
Rollout, open-loop base	0.89 (149)	0.23 (155)	0.42 (153)
Rollout, closed-loop base	0.91 (149)	0.31 (156)	0.41 (154)
MCTS (POMCP)	0.91 (149)	0.30 (155)	0.42 (154)
AND-OR tree	0.86 (149)	0.31 (156)	0.42 (155)
Clairvoyant reference	1.00	0.50	0.50

C. Results

What it gives. Table VIII and Fig. 4 (same 80 fields for all policies). With the LEO layer and onset detection, replanning raised success from 0.70 (Wilson 95% interval 0.59–0.79) to 0.90–0.91 (0.81–0.96) for the belief-aware policies; against CE-MPC (0.78), closed-loop rollout and MCTS each won on 17 fields and lost on 6 (sign test $p = 0.03$). In the dense and weakly protected case, the static route failed almost always (0.03); replanning reached 0.23–0.31 with onset detection and 0.39–0.42 with 30 km detection, against a clairvoyant 0.50; differences among closed-loop policies were not statistically significant ($p \geq 0.34$). Computation per flight was about 0.4 s for CE-MPC, 1.2 s for OLOF-MPC, and 90–130 s for rollout with a closed-loop base, MCTS, and AND-OR search. The belief matters: one that ignores power-dependent detection overstates the risk of unexplored regions, and in the dense onset case OLOF-MPC with such a belief succeeded on only 0.08 of 200 fields, against 0.27 with (4).

Insight. Why simple replanning sufficed: with the exact posterior, OLOF-MPC penalizes cells in increasing order of their denial probability over the range its routes actually meet (the non-monotone region above 0.85 affected less than 1% of its decisions and none of its outcomes), so it routes away from regions where undetected jammers are likely, and it does so again every 5 km as detections arrive. What look-ahead adds is anticipation of those detections, and here they arrive late (at the onset of denial, or within 30 km) and the choice at each step is among a few plans, so there is little information to anticipate and little freedom to exploit it. The belief, in contrast, changes every decision: one that wrongly treats unexplored regions as dangerous presumably drives the aircraft into detours that the endurance limit cannot afford.

D. When to replan

Question. The policies above sense and replan every 5 km. Does the timing of decisions matter, and should replanning follow a fixed interval or the detections themselves?

What to expect. A denied stretch costs the aircraft its length divided by its speed in unaided time, so failure is governed by

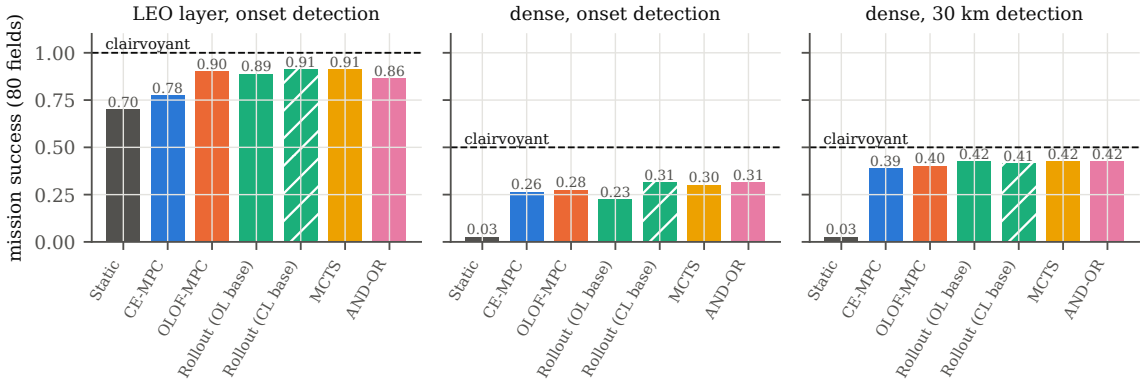


Fig. 4. Mission success by policy in three cases (80 fields); dashed: clairvoyant reference.

the coasting length $\ell^* = V\,t^*$ (Section III-D), and the decision interval Δs matters through the ratio $\Delta s/\ell^* = \Delta s/(V\,t^*)$. Speed does not cancel: at a fixed spatial interval, a faster aircraft has a longer coasting length, so the same interval is relatively shorter. With the industrial IMU of Study A, $\ell^* = 3.24$ km, or 108 s at 30 m/s, while a 5 km decision interval is 167 s. A jammer first detected just after a decision can therefore produce an outage longer than ℓ^* before the next one. With onset detection, a jammer is seen only within 1.5 times its denial radius, so the lead time is short, and either a shorter interval or replanning at the moment of detection should help. With 30 km detection the lead is long, and the interval should matter less.

We reran OLOF-MPC on the same 200 fields per case, replanning every 5 km (the policy above), every 2 km, and every 1 km, sensing at decision points, and in an event-triggered version that senses every 1 km cell and replans when a new jammer is detected, and otherwise every 5 or 2 km. The event-triggered version changes two things at once: it senses more often, and it acts on what it senses as soon as it senses it. To separate them, a third version senses every 1 km cell, exactly as the event-triggered one does, but replans only every 5 or 2 km.

What it gives. Table IX shows results for varying replanning rules. With the LEO layer and onset detection, replanning every 2 km raised success from 0.915 to 0.990 and every 1 km to 1.000. Event-triggered replanning succeeded on all 200 test fields with the same number of replans as the 5 km policy (29 per flight, against 140 for the 1 km interval) but five times as many sensing looks (140 against 29). Sensing alone did not help: sensing every 1 km while replanning every 5 km gave the same outcome on every field as sensing only at decision points (success 0.915, and the same peak error on all 200 fields), and event-triggered replanning won 17 of those fields and lost none (paired sign test $p < 10^{-4}$). The gain came from acting on a detection when it occurred, not from detecting more often. In the dense case with onset detection, the gains were smaller but significant: 0.270 to 0.325–0.340 (paired sign tests $p \le 0.008$). Here too sensing every 1 km without replanning on detection left success at 0.270 (the same peak error on 197 of 200 fields), and replanning on detection won 13 fields and lost none ($p < 0.001$). With 30 km detection, the rules differed little (0.390 to 0.395–0.420). Only sensing every 1 km with replanning every 2 km gained nominally (6 wins, no losses, $p = 0.031$), and that does not survive correction for the six rules compared with the 5 km policy (Bonferroni threshold 0.008); the others had $p \ge 0.12$. This is the one case where sensing between decisions can change a later plan, because a jammer seen 30 km ahead is still relevant at the next decision: with replanning every 5 km, sensing every 1 km changed the peak error on 51 fields and the outcome on 3, all gains ($p = 0.25$).

TABLE IX
MISSION SUCCESS OF OLOF-MPC VERSUS SENSING AND REPLANNING RULE (200 FIELDS PER CASE; MEAN REPLANS PER FLIGHT IN PARENTHESES). INTERVAL RULES SENSE ONLY AT DECISIONS (29, 70, OR 140 LOOKS PER FLIGHT); THE LAST FOUR ROWS SENSE EVERY 1 KM CELL (ABOUT 140 LOOKS PER FLIGHT).

Replanning rule	59 dB, onset	45 dB, onset	45 dB, 30 km
Every 5 km (as above)	0.915 (29)	0.270 (29)	0.390 (29)
Every 2 km	0.990 (70)	0.325 (70)	0.410 (71)
Every 1 km	1.000 (140)	0.340 (141)	0.410 (141)
Sense 1 km, replan every 5 km	0.915 (29)	0.270 (29)	0.405 (29)
Sense 1 km, replan every 2 km	0.990 (70)	0.325 (70)	0.420 (71)
On detection, else every 5 km	1.000 (29)	0.335 (32)	0.405 (45)
On detection, else every 2 km	1.000 (70)	0.340 (72)	0.395 (79)

We then gave the three look-ahead methods of Table VIII the same rule: replan on detection, else every 5 km. Their internal simulations followed it too: the simulated aircraft senses every 1 km cell and replans when it detects a new jammer, both in the base-policy tail and in the simulated segments of MCTS and AND-OR search, so the future each method evaluates is the one its own policy would produce. A method whose simulated aircraft reacted only every 5 km would misjudge exactly the outages that event-triggered replanning avoids. We added a fourth method, rollout with OLOF-MPC itself as the base policy: the candidate first segments (OLOF-MPC's own among them) are scored over the same 24 posterior samples by following each candidate and then simulating event-triggered OLOF-MPC to the goal. The rollout theorem guarantees that this policy does no worse than OLOF-MPC in expected cost when the candidates' expected costs are

TABLE X
DENSE CASE (45 DB), FIRST 80 FIELDS: MISSION SUCCESS WITH REPLANNING ON DETECTION, ELSE EVERY 5 KM; W/L ARE FIELDS WON AND LOST AGAINST OLOF-MPC WITH THE SAME RULE, AND Δ IS THE PAIRED DIFFERENCE IN SUCCESS WITH ITS 95% INTERVAL [27].

	Onset			30 km		
Policy	Success	W/L	Δ [95% CI]	Success	W/L	Δ [95% CI]
OLOF-MPC, every 5 km	0.275	—	—	0.400	—	—
OLOF-MPC, on detection	0.363	—	—	0.413	—	—
Rollout (CE-MPC base)	0.375	3/2	$+0.013$ [-0.05, $+0.08$]	0.438	5/3	$+0.025$ [-0.05, $+0.10$]
MCTS	0.375	2/1	$+0.013$ [-0.04, $+0.07$]	0.400	1/2	-0.013 [-0.07, $+0.04$]
AND-OR search	0.400	4/1	$+0.038$ [-0.03, $+0.10$]	0.425	3/2	$+0.013$ [-0.05, $+0.08$]
Rollout (OLOF-MPC base)	0.375	2/1	$+0.013$ [-0.04, $+0.07$]	0.425	1/0	$+0.013$ [-0.03, $+0.05$]
Rollout (min of CE-, OLOF-MPC)	0.363	1/1	0.000 [-0.05, $+0.05$]	0.413	1/1	0.000 [-0.05, $+0.05$]
Rollout (min of four bases)	0.388	2/0	$+0.025$ [-0.02, $+0.07$]	0.425	2/1	$+0.013$ [-0.04, $+0.07$]

evaluated exactly, OLOF-MPC's own segment is among the candidates, and the simulated base policy is the one executed [26]. Our implementation meets the last two conditions but estimates expected cost from 24 samples and drops candidates early, so the guarantee does not carry over to it, and whether it improves on OLOF-MPC is an empirical question. The cost it minimizes, (5), is also not the failure probability alone, so an improvement in expected cost need not mean fewer failures. Its computation was reduced without changing any decision: the undetected-jammer exposure is linear in the undetected intensity, which only loses the newly cleared cells, so each replan convolves only those cells (agreement with the full computation to 10^{-15}); the routing graph is built once; and a simulated flight stops when no denied cell is reachable within the remaining length budget. One approximation was used: an alternative segment that did not beat OLOF-MPC's segment in any of the first 8 samples was dropped. Finally, we used several base policies at once, which, with exact evaluation, keeps the improvement guarantee with respect to each of them [26]: a candidate segment is scored by the smallest, over the base policies, of its mean cost (5) over the 24 samples. The minimum is taken over the means, not sample by sample; a per-sample minimum would let the policy choose its base with hindsight of the sampled world, which no flyable policy can do. The minimum of noisy means is biased low, increasingly so with more bases. We used two sets of bases: CE-MPC and OLOF-MPC, the two policies of Table VIII; and these two plus a risk-averse and a risk-tolerant planner (cell cost $c = 4p$ and $c = p/4$ with $\beta = 20$), each simulated with the same event-triggered rule (CE-MPC replans only on detection, as it ignores undetected jammers). A candidate was dropped early only if it failed to beat OLOF-MPC's segment under every base. Computation (1–5 min per flight with one base, 3–7 min with four) limited the check to the dense cases, the only ones where OLOF-MPC still failed often, on the first 80 fields of each.

What it gives. Table X provides the results of various policies. With onset detection, all six look-ahead policies gained from event-triggered replanning as OLOF-MPC did (0.363–0.400 against 0.275 for OLOF-MPC every 5 km on the same fields; paired sign tests $p \leq 0.021$), but none improved detectably on OLOF-MPC with the same rule (0.363): rollout on CE-MPC, MCTS, and AND-OR search won 2 to 4 fields and lost 1 or 2 ($p \geq 0.37$). Rollout on OLOF-MPC won 2 fields and lost 1; with the minimum over CE-MPC and OLOF-MPC, it won 1 and lost 1, and with the minimum over four base policies, it won 2 and lost none ($p = 0.50$). With 30 km detection, where the replanning rule itself made no detectable difference (0.400 and 0.413), the look-ahead policies reached 0.400–0.438, again with no detectable difference from OLOF-MPC (rollout on CE-MPC, MCTS, and AND-OR search: 1 to 5 wins, 0 to 3 losses, $p \geq 0.72$); rollout on OLOF-MPC won 1 field and lost none, and the two multiple-base versions won 1 and 2 and lost 1 each. Adding bases did not separate the rollouts from one another: against rollout on OLOF-MPC alone, the two-base version lost one field in each case and won none, and the four-base version won 2 and lost 1 (onset) and won 1 and lost 1 (30 km). The paired differences in success against OLOF-MPC with the same rule ranged from -0.013 to $+0.038$; their 95% intervals (Table X; adjusted Wald intervals for paired proportions [27]) all contained zero, with upper limits of $+0.05$ to $+0.10$, so with 80 fields gains of a few percentage points can be neither shown nor excluded.

Insight. The remaining failures of the 5 km policy with the LEO layer were not a failure of planning but of timing: the aircraft decided less often than its IMU could coast. What matters is reacting within the detection lead, and that is achieved most cheaply by replanning on detection rather than by replanning more often. This gives a design rule that ties the sensor and the decision rule to the IMU. What must stay below the coasting length is the distance flown without an adequate fix, from entering a denied region to leaving it. Detection lead shortens that distance and reaction delay lengthens it: approximately $L_{\text{unaided}} \approx \max(0,\, d_{\text{react}} - d_{\text{lead}}) + d_{\text{exit}} < \ell^*$, where d_{lead} is how far before the denial boundary the jammer could first be detected, d_{react} the distance flown from that point until the maneuver begins, and d_{exit} the unaided distance that remains after the maneuver begins (zero if the maneuver avoids denial altogether). The reaction distance has two parts: detection latency, up to one sensing interval (1 km for the event-triggered rule; for a periodic rule that senses only at decisions, up to the decision interval), and the delay from detection to a replan that acts on it (none for the event-triggered rule, up to the decision interval for a periodic rule that senses between decisions). When the lead is long (30 km detection), the reaction delay is absorbed at any reasonable

interval. In the dense case, the clairvoyant reference, which searches only three cost weightings, reached 0.50; it shows what was achievable with full knowledge in that restricted search, not an upper bound on what any reaction could achieve. Giving the look-ahead methods the same event-triggered rule, in execution and in their internal simulations, did not separate them from OLOF-MPC in the dense case (Table X), and rollout with OLOF-MPC as its own base, alone or with CE-MPC, risk-averse, and risk-tolerant planners as further bases, found no detectable improvement at the tested sample budgets (net differences of at most 2 fields in 80). This does not show that OLOF-MPC is optimal; it shows that these look-ahead implementations did not find a better policy. With the exact posterior, the gain lay in when to decide, not in how deeply to search.

TABLE XI
MISSION SUCCESS WITH IMPERFECT SENSING (200 FIELDS).

Sensing	Static	CE	OLOF	OLOF, P_d-aware
LEO layer (59 dB), onset detection				
perfect	0.705	0.790	0.915	0.915
$P_d = 0.8$	0.705	0.735	0.885	0.890
$P_d = 0.5$	0.705	0.695	0.855	0.845
2 km localisation error	0.705	0.640	0.770	0.770
false alarms	0.705	0.795	0.910	0.910
all three	0.705	0.625	0.805	0.805
Dense, weak protection (45 dB), 30 km detection				
perfect	0.025	0.445	0.390	0.390
$P_d = 0.8$	0.025	0.400	0.365	0.345
$P_d = 0.5$	0.025	0.310	0.300	0.305
2 km localisation error	0.025	0.065	0.110	0.110
false alarms	0.025	0.140	0.160	0.160
all three	0.025	0.025	0.060	0.070

E. Imperfect sensing

Question. How do missed detections, localization errors, and false alarms change these results?

Real direction finding misses jammers, mislocates them, and reports phantoms. We added missed detections (probability of detection P_d per look), 2 km localization error, and false alarms (one per 1,000 km^2 of sensed area per look), and an OLOF-MPC variant whose posterior accounts for P_d.

What to expect. The three errors act differently. A missed jammer leaves the belief no worse than before the look, since its region keeps the prior risk of undetected jammers. A false alarm causes only a detour. A mislocated jammer is the most harmful: the planner confidently avoids the wrong place and may route through the right one.

What it gives. Table XI (200 fields) shows the results. With the LEO layer, OLOF-MPC degraded gracefully with missed detections (0.915, 0.885, 0.855 at $P_d = 1$, 0.8, 0.5) and was unaffected by false alarms; 2 km localization error reduced it to 0.77, and CE-MPC fell below the static route (0.64 versus 0.705). In the dense case, localization error and false alarms (about three phantom jammers per look within 30 km) were severe for all policies. Modeling P_d in the posterior gave no gain.

Insight. The ordering follows from what each policy believes. CE-MPC trusts the detected positions exactly and ignores everything else, so a 2 km localization error sends it confidently into denial, below even the static route. OLOF-MPC keeps the risk of undetected jammers, which partly covers a jammer that is not where it was reported, and so degrades less. A likely reason that modeling P_d added nothing is that a jammer missed on one look is usually detected on a later, closer one, so the belief is corrected before it matters. For sensor design, localization accuracy is worth more than detection probability, as long as it is reasonably high.

F. Adaptive attacker

A random jammer field is the wrong model for an adversary who knows the route. An attacker places K jammers of 10 kW on a 6 km grid near the corridor to defeat the route of a design with 1,000 LEO satellites, a 16-element C-band array, and no camera; the defender may randomize over front routes (double oracle with a greedy attacker best response [28]).

Question. What if the jammers are placed by an adversary who knows the design, and perhaps the route?

What to expect. Against an informed adversary, the random-field results are optimistic because jammers are placed where they hurt. The array model should matter most: the attacker's question is how many jammers it takes to overwhelm the array, and the two models answer it differently.

What it gives. With the null-counting array model, an attacker who knew the route defeated it with 24 jammers, and randomizing routes did not help. With the element-level (MVDR) array model, random placement never defeated the best route for up to 96 jammers; a known fixed route fell at 32 jammers, but the defender could switch to a front route that the greedy attacker could not defeat even with 96 jammers.

Insight. Under null counting, an array of M elements is defeated by any M jammers, wherever they are, so the attacker needs only to outnumber the elements; route choice cannot help. Under the element-level model, protection depends on where the jammers lie relative to the satellite: jammers close in direction to each other share a null (Part I found that ground jammers seen from the air put most of their power into a few spatial patterns, with a participation ratio of 13–14 for 128–192 jammers on 64 elements), and jammers far from the satellite's direction cost the array little. The attacker must then place jammers where the route will be, which makes the route's secrecy and the defender's freedom to switch routes valuable.

Insight. Closing the loop is worth as much as a large hardware step, and the belief must be right before the search is deep: with the exact posterior for the modeled Poisson field and ideal detection, the tested look-ahead implementations showed no statistically detectable improvement over OLOF-MPC, which needed one-hundredth of the computation, with either replanning rule; this does not establish equivalence, and holds for these implementations and scenarios (small option sets, compressed observations in MCTS). Late reaction caused most of the remaining failures with the LEO layer: replanning as soon as a jammer is detected removed them (Section VII-D); and localization accuracy mattered more

than detection probability. Against an informed attacker, route choice and secrecy matter; "more jammers than nulls" is not a reliable rule (Part I).

VIII. Architecture Design

Question. Which combinations of satellite layer, terrestrial radio navigation, antenna, inertial unit, and camera keep the aircraft within its navigation requirements, nominally and after losing any one source domain, and at what cost?

A. Elements and combination rules

Table XII lists the elements and Fig. 5 their tiers. Three rules govern how the elements combine. *Position information adds across sources*: independent fixes enter the filter as added information, so a coarse source still helps a precise one. The navigation model of Study B simplifies this for the two absolute radio sources: it uses an eLoran fix, and with a camera a map fix and visual odometry, only when no LEO fix is available (a selection rule, Section III-D), not both at once. Keeping eLoran active alongside LEO would also let the filter learn eLoran's slowly varying bias before a satellite outage, so the selection rule is, if anything, conservative for designs with both. *Jamming tolerance adds only across separate frequency bands* (Part I): a jammer must spread its power over every band the aircraft uses, so tolerances in separate bands add, whereas two receivers in one band are jammed together. *Failure domains follow dependencies, not hardware*: a cellular network timed by GNSS shares the GNSS domain, although its transmitters are on the ground.

Failure domains in practice. Real failure domains can be less obvious than a list of sources suggests. India's NavIC regional satellite system was designed with seven satellites, three more than positioning needs. After clock failures on several satellites (the clocks of one type, from one supplier), the loss of a replacement satellite to a launcher failure, and a second-generation satellite stranded by a propulsion failure, it had three active satellites by 2026 and could no longer provide standalone positioning [5], [29], [30]. Two lessons carry over to architecture design. Sources that share components or suppliers can fail together and then form one failure domain, not several. And a design sized at its minimum has no margin for loss; the N$-$1 criterion used below is the architectural form of that margin. Commercial constellations add a dependency that is not technical: the operator decides who may use the signals, in what form, and for how long. In May 2026, for example, Starlink withdrew local access to its terminals' location data with a month's notice and no stated reason [31], then restored it for subscribers to its Priority plans only [32]. That episode concerns a data interface on the operator's own terminals, not the radio signals that a signal-of-opportunity receiver uses, but it shows how completely access is the operator's to set. However many satellites it has, a navigation layer run by one operator is one failure domain, and an architecture that relies on signals of opportunity should treat their continued availability as an assumption to be backed up, not a given.

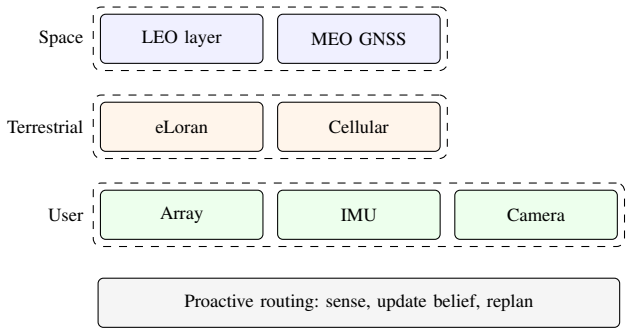


Fig. 5. Tiers and the decision layer.

B. Study B

The candidates combine a LEO layer (1,000, 5,000, or 10,000 satellites) or GNSS alone, with or without eLoran, a user antenna (patch, 7-element L1, 16- or 64-element C-band array), an IMU grade (industrial, tactical, navigation), and a camera (none, day, thermal): 252 designs. LEO denial uses the element-level MVDR array model of Part I (null depth about 30 dB, protection computed toward a representative satellite at 45° elevation), on the actual jammer fields. Navigation error is propagated as in Section III-D.

Routes and designs are chosen on separate data. Candidate routes come from fronts built on training fields. For each design, every candidate route is scored on 200 validation fields by four performance metrics: nominal success (failure by day or by night counts as failure), the worst success after losing one domain, length, and LEO 5 cm formal-solution availability. All routes non-dominated in these four are kept, so the front is computed over architecture design–route pairs, not over designs with one pre-selected route. The front and every qualification statement use the validation fields only; the 200 test fields are used only to report counts.

What to expect. Nominal success needs only enough protection to keep the satellite layer available most of the time, or a good enough IMU to bridge its outages. Surviving the loss of a domain is harder: what remains must carry the whole flight, and each remaining domain has its own outages. A design with two domains loses the requirement whenever one is lost and the other happens to be out, so we expect N$-$1 to need three. Route choice should matter most after a domain is lost, because the remaining sources depend on the terrain (map fixes need texture) and the IMU must bridge whatever gaps the route leaves between fixes.

C. The running example

Consider the design of Section I: 1,000 LEO satellites with a 7-element array, eLoran, a thermal camera, and a navigation-grade IMU. On the direct 139 km route, it met the requirement on all 200 test fields, but it failed on 21 after losing the satellite layer. Fig. 6 shows why on one of those fields. With LEO lost, eLoran happened to be in one of its regional outages (it is out on about 10% of missions), so only the thermal camera remained. Over textured ground, map fixes hold the cross-track uncertainty near 9 m. The direct route then crosses 49 km of

TABLE XII
ARCHITECTURAL ELEMENTS (*LISTED FOR COMPLETENESS; NOT EVALUATED IN STUDY B).

Element	Accuracy	Failure domain	Strengths	Weaknesses
GNSS (MEO, L-band)	cm (carrier) to m	Space, L-band	Global, free	Weak; jammed and spoofed easily
LEO navigation layer	cm (fixed), dm–m	Space, own band	25–33 dB stronger; float in $\sim$1 min	Large constellation; orbit/clock products
eLoran (100 kHz)	$\approx$10 m (95%)	Terrestrial LF	Wide-area jamming impractical	Coarse; fixed transmitters; outages
User antenna array	—	On board	Suppresses many horizon jammers	Aperture; tail geometries
IMU	drifts as t^2–t^3	Self-contained	Cannot be jammed or spoofed	Bounded coasting; cost/mass at navigation grade
Camera (day / thermal)	10 m map fixes	On board, weather	Independent physics	Featureless ground; weather; lighting (day camera)
Proactive routing	—	Decision layer	Avoids sensed jammers	Detection range; localization accuracy
LEO signals of opportunity (e.g., Starlink)*	m (research)	Space, own band; operator	Thousands of satellites in orbit now; Ku band, far from GNSS	No navigation message; public orbits, unknown clocks; access set by operator

featureless ground, about 27 minutes without a map fix; visual odometry slows the inertial drift but cannot stop it, and the uncertainty reaches 68 m before the next fix. In all 20 fields where the direct route failed and the 161 km route did not, eLoran was out. The longer route avoids featureless ground entirely and peaks at 10.5 m. A 159 km route through the same featureless region failed as often as the direct one: what matters is where the route goes, not its length.

Insight. After losing the satellite layer, and with eLoran out, the camera is the only source of absolute position, and it is absolute only over textured ground. Between map fixes, even a navigation-grade IMU aided by visual odometry drifts, and 27 minutes is longer than it can bridge within the navigation requirements. The longer route buys resilience not by adding hardware but by staying within reach of map fixes. The same mechanism accounts for the count: at least 20 of the 21 failures occurred on fields where eLoran was out, which happens on about 10% of missions.

D. Results

What it gives. *The front.* Table XIII shows six representative design–route pairs; the full front of 39 pairs is in Appendix A. Of those 39, 22 use a route other than the one a lexicographic rule (nominal success first, then domain-loss success) would have chosen, so choosing one route per design in advance would have missed most of the front.

What it gives. *Nominal success and centimeter capability.* The IMU grade mattered: with 1,000 LEO satellites and a 7-element array, test failures were 57, 6, and 0 of 200 for industrial, tactical, and navigation grades. The satellite layer is denied in about a tenth of the cells, and the IMU must bridge each denied stretch. In Study A's terms, an industrial IMU reaches the requirement after about 3 km of coasting and a tactical one after about 6 km (Section III-D); a navigation IMU drifts far more slowly, so only it bridges the longest denied stretches. The cheapest design without validation failures was GNSS with eLoran, a thermal camera, and a navigation IMU (\$22M per year for 1,000 aircraft; 1/200 failures on test; no centimeter capability). Adding centimeter capability took 1,000 LEO satellites with a 7-element array and a navigation IMU (\$366M; 0/200), or a 64-element array with a tactical IMU (\$443M; 0/200); raising LEO 5 cm formal-solution availability further took larger constellations (5,000 satellites: 0.6% unavailable, \$1,749M). Because the cold start is a fixed interval, the fraction can favor a longer route: in the running example the direct route is without a 5 cm LEO solution for 146 s and the 161 km route for 97 s.

What it gives. *Losing one domain.* No design with two independent domains met N$-$1 at 2% on validation: with LEO and eLoran, losing either leaves the other's own outages (16–21 failures of 200). Exactly two designs qualified, both with 1,000 LEO satellites, eLoran, a thermal camera, and a navigation IMU (7- or 16-element array; \$372M and \$457M), and only on the 161 km route. On test, the 7-element design had no nominal failures and 1, 0, and 0 failures out of 200 after losing the satellite layer, eLoran, and the camera. One failure in 200 meets a 2% target as an empirical rate, but its 95% Wilson upper bound is 2.8% (3.7% simultaneously over the three domains), so the data do not establish a rate below 2% at that confidence. Measured as horizontal RMS instead of cross-track, the counts of these designs were unchanged. The argument that a third domain makes loss of the flight require two simultaneous outages holds only when either surviving domain can carry the flight alone on that route; the running example shows a route on which the camera alone could not.

Insight. The N$-$1 result follows from simple arithmetic on the outage rates. With two domains, losing one leaves the other, and the other is itself unavailable on about 10% of missions, so about 20 failures in 200 are expected; 16–21 were observed. With three domains, losing one leaves two, and a failure needs both of them out or unable to carry the flight at once, which is rarer by about another factor of ten, provided that each remaining domain can carry the flight alone on the chosen route. That proviso is where the IMU grade and the route enter.

E. Sensitivity to the array model

Question. The array model protects one representative satellite at 45° elevation, while positioning needs several satellites

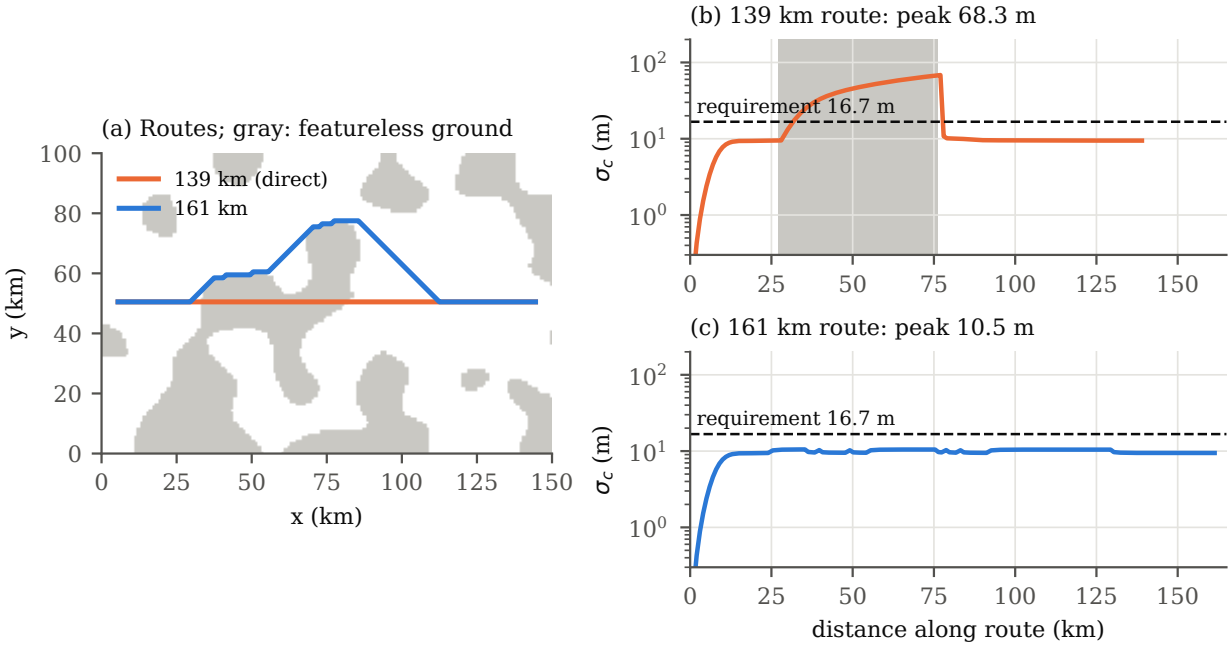


Fig. 6. Running example: the same architecture (1,000 LEO satellites, 7-element array, eLoran, thermal camera, navigation IMU) on two routes, on one test field where the satellite layer is lost and eLoran is in a regional outage. (a) The routes over the terrain; gray: featureless ground, where the camera gives no map fix. (b, c) Cross-track navigation uncertainty along each route; gray: featureless stretches; dashed: the requirement.

TABLE XIII
REPRESENTATIVE DESIGN–ROUTE PAIRS FOR A FLEET OF 1,000 AIRCRAFT (FULL FRONT IN APPENDIX A). TEST FAILURES ARE COUNTS OUT OF 200 TEST FIELDS, NOMINALLY AND AFTER LOSING THE SATELLITE LAYER, ELORAN, OR THE CAMERA; "—" MEANS THE DESIGN DOES NOT HAVE THAT SOURCE. 5 CM UNAV.: SHARE OF THE FLIGHT, AND SECONDS, WITHOUT A FORMAL 5 CM LEO SOLUTION, INCLUDING THE COLD START.

Role	Space	Terr.	Antenna	IMU	Camera	Cost	Route	Test failures (of 200)				5 cm unav.	
						($M/yr)	(km)	Nominal	Lose sat.	Lose eLoran	Lose cam.	(%)	(s)
Low-cost corridor navigation	GNSS only	eLoran	patch	navigation	thermal	22	161	1	1	18	21	100.0	5379
Adds centimeter capability	LEO 1k	—	L1 7-el	navigation	—	366	159	0	200	—	—	1.7	89
Tactical IMU with large array	LEO 1k	—	C-band 64-el	tactical	—	443	151	0	200	—	—	1.2	60
Resilient to losing any one domain	LEO 1k	eLoran	L1 7-el	navigation	thermal	372	161	0	1	0	0	1.8	97
Same hardware, direct route	LEO 1k	eLoran	L1 7-el	navigation	thermal	372	139	0	21	0	0	3.1	146
Further centimeter availability	LEO 5k	—	L1 7-el	tactical	day	1,749	158	0	200	—	0	0.6	29

in usable directions (Part I). Do the conclusions change if protection must hold toward satellites in other directions?

We recomputed LEO denial on the test fields with MVDR protection required toward two alternatives, *a low-elevation satellite* (20° above the horizon) and *the fourth-best of eight satellites* spread uniformly over the sky above 15°, and re-evaluated the selected routes of the key designs.

What to expect. Ground jammers appear just below the horizon. A satellite high in the sky is far from them in direction, and the array can null the jammers while keeping its beam on the satellite; a satellite low in the sky is close to them in direction, and keeping unit response toward it while nulling them takes large weights that amplify noise, so the array's gain toward that satellite collapses (Part I). A requirement that the best-placed satellites can meet should therefore change the results little, whereas one that depends on a particular low-elevation satellite should change them substantially.

What it gives. Requiring protection toward the fourth-best of eight satellites changed almost nothing: the share of denied cells stayed at 9–10%, the key designs had 0 or 1 nominal failures of 200 (0 with the representative satellite), and centimeter unavailability rose by at most 1.1 percentage points. Requiring protection toward a single satellite only 20° above the horizon is a harsher test, because the ground jammers lie just below the horizon and therefore close to it in direction (Part I): with 7- and 16-element arrays the denied share roughly tripled (29–32%), designs that rely on LEO alone failed on 53–63 of 200 fields, and the three-domain designs kept their

zero nominal failures only because eLoran and the camera carried the flight, while losing centimeter capability about half the time (45–51% unavailable). The 64-element array was barely affected (9 failures out of 200). This is a directional-availability sensitivity: it tests which satellite directions must survive the interference, not whether the surviving satellites give adequate geometry. The convergence curve from Part I assumes up to 48 tracked satellites, chosen by highest elevation rather than by an optimized selection, and does not adapt to the surviving subset, so the chain from surviving links to their geometry and to convergence is not modeled. Consequently, f_4 is a surrogate: it applies the nominal convergence curve of Part I, computed with up to 48 satellites tracked (the highest in elevation), whenever the satellite layer is available. It is not a convergence result for the satellites that survive the interference. Within that scope, the conclusions are insensitive to requiring four usable satellites anywhere in the sky but not to requiring protection toward satellites low in the sky.

F. Cost model and sensitivity

Question. How does the cheapest qualifying design change with the fleet size and with the share of infrastructure cost the operator bears?

Costs are annualized. The annual cost of a design is

$$C = q_S\, C_S + q_T\, C_T + n\, C_U, \tag{6}$$

where C_S is the space segment (LEO satellites at \$2.1M over 6 years, with a 25% payload premium for C-band; the Walker patterns labeled 1,000, 5,000, and 10,000 have 988, 4,976, and 10,024 satellites, and these exact counts are used), C_T the terrestrial segment (six eLoran transmitters at \$15M over 25 years), C_U the user equipment per aircraft over 5 years (patch antenna \$200, 7-element L1 array \$20k, 16-element C-band array \$12k, 64-element C-band array \$40k; industrial IMU \$1.5k, tactical \$15k, navigation \$80k; day camera \$2k, thermal camera \$10k; eLoran receiver \$3k; LEO receiver \$500), and n the fleet size. The shares q_S and q_T are the fractions of the space and terrestrial segments charged to this operator; they are below one when the infrastructure is shared with other users. GNSS is treated as free. The front above uses $n = 1{,}000$ and $q_S = q_T = 1$.

What to expect. Infrastructure costs are fixed, while user-equipment costs grow with the fleet. A small fleet that bears the full infrastructure cost should favor designs that put the capability in the aircraft; a large fleet, or a shared constellation, should favor designs that put it in infrastructure and carry cheaper equipment.

What it gives. Table XIV shows which designs qualified on the validation fields and at what cost, with their test counts. For no validation failures, GNSS with eLoran, a thermal camera, and a navigation IMU was the cheapest unless the space segment was shared or the fleet was large; then 1,000 LEO satellites with eLoran, a 16-element array, and a tactical IMU were cheapest, and at 100,000 aircraft with a 1% share, 5,000 satellites with eLoran, a 7-element array, and an industrial IMU were cheapest. Designs qualified on validation can still fail on test (here 1–3 of 200). For N−1 at 2%, only the two three-domain designs qualified at any fleet size and share.

Insight. The crossover is a ratio of fixed to per-aircraft cost. A 1,000-satellite layer costs about \$346M per year (988 satellites at \$2.1M over 6 years), more than the equipment of a 1,000-aircraft fleet, so at full cost the satellite-free design wins. A navigation IMU costs \$16k per aircraft per year, so at 100,000 aircraft, it alone costs \$1.6B per year; there a shared constellation that lets each aircraft carry an industrial IMU becomes cheaper. N−1 qualification, in contrast, is set by the outage arithmetic above, not by cost, so it selects the same designs at every fleet size.

Insight. Meeting the nominal validation criterion was comparatively inexpensive in the evaluated designs; resilience to the loss of any one source was not. Among the architectures evaluated, and under the assumed 10% per-mission outage rates of eLoran and the camera, the only qualifying family combined three independent domains (space, terrestrial low-frequency, on-board vision that works at night), an IMU good enough to bridge between them, and a route chosen for the degraded cases; other availability assumptions or candidate sets could admit two-domain designs. The running example shows the mechanism: when two domains are gone, the third must carry the flight, and whether it can depends on the route. Navigation-grade IMUs are heavy and costly, so for small aircraft the N−1 target may be out of reach; for corridor navigation alone, vision and inertial sensing without LEO meet the nominal target cheaply, and LEO's distinct contributions are centimeter accuracy and a domain independent of weather and terrain texture.

IX. Discussion

What the aircraft must get right. Its coasting law (identified in flight, with a window matched to how fast conditions change), its belief about the threat (the exact posterior for the modeled field and sensor), and its route under degraded conditions. Each mattered more than algorithmic sophistication: ALS and gain-first were equally accurate when ALS was given the true model structure, the tested look-ahead methods showed no detectable gain over OLOF-MPC, and expected-value routing on the analytic prior beat minmax over samples.

Hardware, estimation, and route together. The running example is the paper's argument in one picture. The same hardware is resilient on one route and not on another; a planner that evaluates hardware on a fixed route, or routes with a laboratory coasting law, would get both wrong.

Gain-first identification in this setting. Its value here was not superior accuracy in stationary conditions, where every adaptive method that estimated measurement noise jointly matched the oracle, but that it needs no IMU error model (only a polynomial filter order, which a multiple-model test can select from the data), keeps its covariances valid by construction, is robust to misjudged aiding noise, and has a streaming form with a window that can be matched to the rate of change [11]. A matched comparison with streaming versions of ALS is future work [11].

X. Limitations

- The theater, jammer field, terrain texture, outage probabilities, and costs are synthetic or illustrative; real threat

TABLE XIV
CHEAPEST QUALIFYING DESIGN VERSUS FLEET SIZE AND THE OPERATOR'S SHARES OF THE SPACE (q_S) AND TERRESTRIAL (q_T) SEGMENTS. A DESIGN QUALIFIES ON THE VALIDATION FIELDS WITH NO NOMINAL FAILURES AND, FOR THE SECOND REQUIREMENT, AT MOST 2% (4 OF 200) FAILURES AFTER LOSING ANY ONE SOURCE. TEST FAILURES ARE COUNTS OUT OF 200 TEST FIELDS, NOMINALLY AND FOR THE WORST SINGLE LOSS.

Requirement (validation)	Fleet	q_S	q_T	Cheapest qualifying design	Cost	Test failures (of 200)	
					($M/yr)	Nominal	Worst loss
Nominal	100	1	1	GNSS + eLoran; patch antenna; navigation IMU; thermal camera	5.5	1	21
Nominal	100	0.1	1	GNSS + eLoran; patch antenna; navigation IMU; thermal camera	5.5	1	21
Nominal	100	0.01	1	LEO 1k; C-band 64-el array; tactical IMU	5.4	0	200
Nominal	100	0.1	0.1	GNSS + eLoran; patch antenna; navigation IMU; thermal camera	2.2	1	21
Nominal	100	0.01	0.01	GNSS + eLoran; patch antenna; navigation IMU; thermal camera	1.9	1	21
Nominal	1,000	1	1	GNSS + eLoran; patch antenna; navigation IMU; thermal camera	22.2	1	21
Nominal	1,000	0.1	1	GNSS + eLoran; patch antenna; navigation IMU; thermal camera	22.2	1	21
Nominal	1,000	0.01	1	LEO 1k + eLoran; C-band 16-el array; tactical IMU	14.0	2	23
Nominal	1,000	0.1	0.1	GNSS + eLoran; patch antenna; navigation IMU; thermal camera	19.0	1	21
Nominal	1,000	0.01	0.01	LEO 1k + eLoran; C-band 16-el array; tactical IMU	10.5	2	23
Nominal	10,000	1	1	GNSS + eLoran; patch antenna; navigation IMU; thermal camera	190.0	1	21
Nominal	10,000	0.1	1	LEO 1k + eLoran; C-band 16-el array; tactical IMU	107.8	2	23
Nominal	10,000	0.01	1	LEO 1k + eLoran; C-band 16-el array; tactical IMU	68.9	2	23
Nominal	10,000	0.1	0.1	LEO 1k + eLoran; C-band 16-el array; tactical IMU	104.6	2	23
Nominal	10,000	0.01	0.01	LEO 1k + eLoran; C-band 16-el array; tactical IMU	65.4	2	23
Nominal	100,000	1	1	LEO 1k + eLoran; C-band 16-el array; tactical IMU	1,045.8	2	23
Nominal	100,000	0.1	1	LEO 1k + eLoran; C-band 16-el array; tactical IMU	656.8	2	23
Nominal	100,000	0.01	1	LEO 5k + eLoran; L1 7-el array; industrial IMU	521.0	3	38
Nominal	100,000	0.1	0.1	LEO 1k + eLoran; C-band 16-el array; tactical IMU	653.6	2	23
Nominal	100,000	0.01	0.01	LEO 5k + eLoran; L1 7-el array; industrial IMU	517.5	3	38
Nominal and any one loss	100	1	1	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	351.7	0	1
Nominal and any one loss	100	0.1	1	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	40.5	0	1
Nominal and any one loss	100	0.01	1	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	9.3	0	1
Nominal and any one loss	100	0.1	0.1	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	37.2	0	1
Nominal and any one loss	100	0.01	0.01	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	5.8	0	1
Nominal and any one loss	1,000	1	1	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	372.1	0	1
Nominal and any one loss	1,000	0.1	1	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	60.9	0	1
Nominal and any one loss	1,000	0.01	1	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	29.0	0	1
Nominal and any one loss	1,000	0.1	0.1	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	57.6	0	1
Nominal and any one loss	1,000	0.01	0.01	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	25.5	0	1
Nominal and any one loss	10,000	1	1	LEO 1k + eLoran; L1 7-el array; navigation IMU; thermal camera	576.4	0	1
Nominal and any one loss	10,000	0.1	1	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	257.8	0	1
Nominal and any one loss	10,000	0.01	1	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	218.9	0	1
Nominal and any one loss	10,000	0.1	0.1	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	254.6	0	1
Nominal and any one loss	10,000	0.01	0.01	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	215.4	0	1
Nominal and any one loss	100,000	1	1	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	2,545.8	0	1
Nominal and any one loss	100,000	0.1	1	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	2,156.8	0	1
Nominal and any one loss	100,000	0.01	1	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	2,117.9	0	1
Nominal and any one loss	100,000	0.1	0.1	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	2,153.6	0	1
Nominal and any one loss	100,000	0.01	0.01	LEO 1k + eLoran; C-band 16-el array; navigation IMU; thermal camera	2,114.4	0	1

data, terrain masking, and airborne jammers are not modeled.

- Failures are navigation-uncertainty violations, not corridor departures; integrity risk and protection levels are not computed, and spoofing is treated qualitatively.
- Re-convergence uses warm-restart curves from receivers that tracked for 900 s before the outage (applied after at least 120 s of tracking); after outages longer than 180 s it uses the slower of the 180 s warm-restart and cold-start curves, an extrapolation rule. The coasting study is single-axis; the planning studies use the true in-flight coasting law rather than an identified one. The coasting law is for straight, unaccelerated flight; a turn adds specific force that couples attitude error into position and can shorten the coasting time, which is not modeled.
- The outage-risk model counts single-jammer outages and underpredicts in dense fields; the reset model in Study A treats all fixes alike.
- The proactive policies are compared with 5 km replanning and small option sets. OLOF-MPC was also run with 2 and 1 km intervals and with event-triggered replanning on all cases; the look-ahead methods were run with event-triggered replanning only on the first 80 fields of the two dense cases, and not with 2 or 1 km intervals (Section VII-D). MCTS compresses observations, and the AND-OR tree reuses the root's undetected intensity at depth two; the attacker is greedy.
- Study B uses one array null depth, a representative-satellite array model (with the sensitivity above), assumed correlated-error magnitudes for eLoran and map fixes, and per-mission (not partial) outages; zero failures in 200 trials bound rates only below about 1.9%.

XI. Conclusion

Keeping an unmanned aircraft's navigation within a corridor under uncertain interference is a joint property of its sources, its inertial error growth, its route, and its in-flight decisions. The studies are complementary rather than one integrated demonstration: the identification study shows how a coasting law can be calibrated from aided flight records, demonstrated here on simulated records (validation on recorded aircraft data remains to be done), and the planning and architecture studies establish performance conditional on the supplied (true in-flight) coasting law. Within that scope, the elements that mattered were a coasting law taken from flight data rather than a datasheet, identified with a window matched to the rate of change, route risk measured by outage duration (exactly so under the reset model of Study A), routes selected on the threat's statistical law, replanning with the exact posterior for the modeled field and sensor, triggered by detection so that the aircraft reacts within its coasting length, and, for resilience to losing any one source, three independent domains with a capable IMU and a route chosen for the degraded case. Future work includes identification on board in streaming form, route search on outage risk with multi-jammer outages, partial outages and weather models, real threat data, integrity under spoofing, and route search on large graphs, where the growth of OPEN calls for approximate or parallel multiobjective search.

Appendix A
Full architecture front

Table XV lists all non-dominated design–route pairs (cost, validation mission failure, validation N-1 failure, validation centimeter unavailability) with validation mission failure at most 30/200, for a fleet of 1,000 and full infrastructure cost.

Appendix B
Navigation model of Study B

The cross-track error state is $\mathbf{x} = [\delta p,\ \delta v,\ b_a,\ \theta,\ b_g,\ b_L,\ b_E,\ b_M]^\top$: position and velocity error, accelerometer bias, tilt (which couples into acceleration through gravity g), gyroscope bias, and the correlated errors of the LEO, eLoran, and map fixes. In continuous time $\dot{\mathbf{x}} = F_c\mathbf{x} + \mathbf{w}$ with

$$\begin{aligned}&\delta\dot p = \delta v, \quad \delta\dot v = b_a + g\,\theta + w_v, \quad \dot b_a = w_a,\\ &\dot\theta = b_g + w_\theta, \quad \dot b_g = 0,\\ &\dot b_j = -b_j/\tau_j + w_j \quad (j \in \{L, E, M\}),\end{aligned}$$

and white-noise intensities $q_v = (k\,\mathrm{VRW})^2$, $q_\theta = (k\,\mathrm{ARW})^2$, $q_a = (0.2\,b_{a,0}/\sqrt{3600\,\mathrm{s}})^2$, and $q_j = 2s_j^2/\tau_j$, so that each b_j is stationary with standard deviation s_j. The vibration factor is $k = 3$; VRW, ARW, the residual accelerometer bias $b_{a,0}$, and the gyroscope bias of each IMU grade are listed in the code (industrial: 0.05 m/s/$\sqrt{\mathrm{h}}$, 0.15°/$\sqrt{\mathrm{h}}$, 0.1 mg, 1°/h; tactical: 0.03, 0.05, 0.05, 0.1; navigation: 0.01, 0.002, 0.025, 0.01). The correlated errors use $(s_L, \tau_L) = (s_b, 600\,\mathrm{s})$ from (2), $(s_E, \tau_E) = (4.6\,\mathrm{m}, 667\,\mathrm{s})$, and $(s_M, \tau_M) = (8.7\,\mathrm{m}, 67\,\mathrm{s})$, the eLoran and map correlation lengths of 20 and 2 km divided by the 30 m/s airspeed. Over a cell of duration Δt the model is discretized exactly by Van Loan's method [33], in his form with $-F_c$ in the upper-left block: with

$$\mathcal{M} = \begin{bmatrix} -F_c & Q_c \\ 0 & F_c^\top \end{bmatrix}, \qquad e^{\mathcal{M}\Delta t} = \begin{bmatrix} \cdot & \Phi_{12} \\ 0 & \Phi_{22} \end{bmatrix},$$

the transition matrix is $F = \Phi_{22}^\top$ and the process-noise covariance $Q = F\,\Phi_{12}$.

The scalar measurement models are

$$\begin{aligned}\text{LEO: } &z = \delta p + b_L + v, && \mathrm{Var}(v) = R_L(t),\\ \text{eLoran: } &z = \delta p + b_E + v, && \mathrm{Var}(v) = (2\,\mathrm{m})^2,\\ \text{map fix: } &z = \delta p + b_M + v, && \mathrm{Var}(v) = (5\,\mathrm{m})^2,\\ \text{visual odometry: } &z = \delta v + v, && \mathrm{Var}(v) = (0.3\,\mathrm{m/s})^2,\end{aligned}$$

with $R_L(t) = \max\{s^2(t) - s_b^2,\ (1\,\mathrm{mm})^2\}$ as in (2), applied in each cell in the order of Section III-D. Before take-off the filter is initialized by 10 minutes of 5 cm position fixes on the ground, with the source biases at their prior variances. The reported uncertainty is $\sigma_c = \sqrt{P_{11}}$, its peak over the flight compared with 16.7 m.

Appendix C
Experiment specification

Theater and fields. The theater is a 150×100 km grid of 1 km cells, flown at 30 m/s (33.3 s per cell). Jammer fields are drawn with independent seeds: training fields $i = 0, \ldots, 599$ with seed $20{,}000 + i$, validation fields $i = 600, \ldots, 799$ with seed $20{,}000 + i$, and test fields $i = 0, \ldots, 199$ with seed $5{,}000 + i$. The proactive-routing study (Section VII) draws its held-out fields with the same seeds $5{,}000 + i$ (the first 80 in Table VIII, all 200 elsewhere), and selects its static route on 150 fields with seeds $20{,}000 + i$. The per-mission eLoran outage (probability 0.1) uses seed $72{,}000 + i$ on validation and $71{,}000 + i$ on test field i; the camera outage (probability 0.1) uses seeds 73,000 (validation) and 74,000 (test). The same draws are used for every design (common random numbers).

Constellations. The LEO layers are the three-shell Walker patterns of Part I, Appendix B: 988, 4,976, and 10,024 satellites for the layers labeled 1,000, 5,000, and 10,000. LEO denial for each array uses the MVDR model of Part I with null depth of about 30 dB (channel decorrelation 10^{-3}).

Inertial units and initialization. Noise intensities and their units are given in Appendix B. The initial covariance is

$$\begin{aligned}\mathrm{diag}\big((0.05\,\mathrm{m})^2,\ (0.01\,\mathrm{m/s})^2,\ (b_{a,0}g)^2,\ 10^{-8}\,\mathrm{rad}^2,\\ b_{g,0}^2,\ s_b^2,\ s_E^2,\ s_M^2\big),\end{aligned}$$

with $b_{a,0}$ converted from mg and $b_{g,0}$ from deg/h to rad/s, followed by 600 s of 5 cm position updates at 1 Hz on the ground.

Recovery curves. All architecture results (Tables XIII, XIV, and XV, the array and correlation sensitivities, and Fig. 6) use the same rule: Part I's cold-start curve (two frequencies, 1 Hz, industrial IMU, 10 cm ionosphere and 2 cm orbit-clock residuals, 50 start times) at take-off and after outages preceded by less than 120 s of continuous tracking; otherwise Part I's warm-restart curve for the shortest simulated outage (15, 60, or 180 s) at least as long as the actual one; for outages longer

TABLE XV
NON-DOMINATED DESIGN–ROUTE PAIRS. FAILURES ARE COUNTS OUT OF 200 FIELDS. VALIDATION FAILURES DECIDE QUALIFICATION ("WORST LOSS": THE LARGEST COUNT AFTER LOSING ANY ONE SOURCE); TEST FAILURES ARE REPORTED NOMINALLY AND AFTER LOSING THE SATELLITE LAYER, ELORAN, OR THE CAMERA, WITH "—" WHERE THE DESIGN DOES NOT HAVE THAT SOURCE. 5 CM UNAV.: SHARE OF THE FLIGHT WITHOUT A FORMAL 5 CM LEO SOLUTION.

Space	Terr.	Antenna	IMU	Camera	Cost	Route	Validation failures		Test failures				5 cm
					($M/yr)	(km)	Nominal	Worst loss	Nominal	Lose sat.	Lose eLoran	Lose cam.	unav.
GNSS only	eLoran	patch	industrial	—	5	139	16	200	21	21	200	—	100.0%
GNSS only	eLoran	patch	navigation	thermal	22	161	0	17	1	1	18	21	100.0%
LEO 1k	—	L1 7-el	industrial	thermal	352	145	22	200	28	200	—	67	2.0%
LEO 1k	—	L1 7-el	industrial	thermal	352	146	20	200	31	200	—	72	2.1%
LEO 1k	—	L1 7-el	tactical	—	353	157	5	200	6	200	—	—	1.7%
LEO 1k	—	L1 7-el	tactical	—	353	158	6	200	8	200	—	—	1.6%
LEO 1k	—	L1 7-el	tactical	—	353	159	10	200	5	200	—	—	1.7%
LEO 1k	eLoran	L1 7-el	industrial	—	354	139	8	117	13	21	112	—	3.1%
LEO 1k	eLoran	L1 7-el	industrial	day	355	145	4	200	11	200	67	11	2.0%
LEO 1k	—	L1 7-el	tactical	thermal	355	146	3	200	7	200	—	12	2.1%
LEO 1k	—	L1 7-el	tactical	thermal	355	154	1	200	1	200	—	15	2.1%
LEO 1k	—	L1 7-el	tactical	thermal	355	161	4	200	2	200	—	21	1.8%
LEO 1k	eLoran	L1 7-el	industrial	thermal	356	139	1	44	4	21	50	13	3.1%
LEO 1k	eLoran	L1 7-el	industrial	thermal	356	145	3	33	7	38	28	11	2.0%
LEO 1k	eLoran	L1 7-el	industrial	thermal	356	159	20	33	30	38	36	34	1.8%
LEO 1k	eLoran	L1 7-el	tactical	—	357	159	0	16	1	21	5	—	1.7%
LEO 1k	eLoran	L1 7-el	navigation	thermal	372	161	0	0	0	1	0	0	1.8%
LEO 1k	—	C-band 64-el	industrial	—	441	151	17	200	10	200	—	—	1.2%
LEO 1k	—	C-band 64-el	industrial	thermal	443	139	3	200	2	200	—	29	1.5%
LEO 1k	—	C-band 64-el	industrial	thermal	443	146	5	200	6	200	—	21	1.3%
LEO 1k	—	C-band 64-el	industrial	thermal	443	149	6	200	6	200	—	24	1.3%
LEO 1k	—	C-band 64-el	tactical	—	443	151	0	200	0	200	—	—	1.2%
LEO 1k	eLoran	C-band 64-el	industrial	—	445	139	0	25	0	21	29	—	1.5%
LEO 1k	eLoran	C-band 64-el	industrial	thermal	447	139	0	16	0	21	2	0	1.5%
LEO 1k	eLoran	C-band 64-el	industrial	thermal	447	147	3	33	3	38	4	4	1.3%
LEO 1k	eLoran	C-band 64-el	industrial	thermal	447	149	5	33	4	38	6	4	1.3%
LEO 1k	eLoran	C-band 64-el	tactical	—	448	151	0	16	0	21	0	—	1.2%
LEO 5k	—	L1 7-el	industrial	—	1,746	158	11	200	11	200	—	—	0.5%
LEO 5k	—	L1 7-el	industrial	day	1,746	158	7	200	13	200	—	13	0.6%
LEO 5k	—	L1 7-el	industrial	thermal	1,748	149	2	200	4	200	—	12	0.6%
LEO 5k	—	L1 7-el	industrial	thermal	1,748	158	3	200	8	200	—	13	0.6%
LEO 5k	—	L1 7-el	tactical	—	1,749	158	0	200	0	200	—	—	0.5%
LEO 5k	—	L1 7-el	tactical	day	1,749	158	0	200	0	200	—	0	0.6%
LEO 5k	eLoran	L1 7-el	industrial	—	1,750	139	0	27	3	21	38	—	0.8%
LEO 5k	eLoran	L1 7-el	industrial	thermal	1,752	139	0	16	0	21	4	3	0.8%
LEO 5k	eLoran	L1 7-el	industrial	thermal	1,752	142	0	33	1	38	6	2	0.7%
LEO 5k	eLoran	L1 7-el	industrial	thermal	1,752	158	1	33	6	38	8	7	0.6%
LEO 5k	eLoran	L1 7-el	tactical	—	1,753	158	0	16	0	21	0	—	0.5%
LEO 5k	eLoran	L1 7-el	tactical	day	1,753	158	0	16	0	21	0	0	0.6%

than 180 s the slower of the 180 s and cold-start curves; and the variance floor of (2). The 10,000-satellite layer uses the 5,000-satellite curves.

Code. The accompanying archive contains every script and a file `REPRODUCE.md` listing the commands, in order, that regenerate each table and figure.

REFERENCES

[1] E. D. Kaplan and C. J. Hegarty, *Understanding GPS/GNSS: Principles and Applications*, 3rd ed. Norwood, MA: Artech House, 2017.

[2] T. E. Humphreys, B. M. Ledvina, M. L. Psiaki, B. W. O'Hanlon, and P. M. Kintner, "Assessing the spoofing threat: Development of a portable GPS civilian spoofer," in *Proc. ION GNSS 2008*, 2008, pp. 2314–2325.

[3] T. G. R. Reid, A. M. Neish, T. Walter, and P. K. Enge, "Broadband LEO constellations for navigation," *NAVIGATION*, vol. 65, no. 2, pp. 205–220, 2018.

[4] F. S. Prol *et al.*, "Position, navigation, and timing (PNT) through low Earth orbit (LEO) satellites: A survey on current status, challenges, and opportunities," *IEEE Access*, vol. 10, pp. 83 971–84 002, 2022.

[5] K. R. Pattipati, "What low-Earth-orbit constellations buy for jam-resilient positioning: Geometry, spectrum, carrier phase, and antenna arrays (Part I)," 2026, working draft.

[6] International Loran Association, "Enhanced Loran (eLoran) definition document," International Loran Association, Tech. Rep., 2007.

[7] P. D. Groves, *Principles of GNSS, Inertial, and Multisensor Integrated Navigation Systems*, 2nd ed. Boston, MA: Artech House, 2013.

[8] R. T. Rockafellar and S. Uryasev, "Optimization of conditional value-at-risk," *Journal of Risk*, vol. 2, no. 3, pp. 21–41, 2000.

[9] E. B. Wilson, "Probable inference, the law of succession, and statistical inference," *J. Amer. Statist. Assoc.*, vol. 22, no. 158, pp. 209–212, 1927.

[10] L. D. Brown, T. T. Cai, and A. DasGupta, "Interval estimation for a binomial proportion," *Statist. Sci.*, vol. 16, no. 2, pp. 101–133, 2001.

[11] M. Macesker, K. R. Pattipati, D. Sidoti, and L. Zhang, "Streaming innovation whitening for adaptive Kalman filtering and covariance identification: Part I—Linear systems and exact identifiability," 2026, manuscript.

[12] B. J. Odelson, M. R. Rajamani, and J. B. Rawlings, "A new autocovariance least-squares method for estimating noise covariances," *Automatica*, vol. 42, no. 2, pp. 303–308, 2006.

[13] K. R. Pattipati, D. Sidoti, L. Zhang, and M. Macesker, "Gain-first design of α, α-β, and α-β-γ tracking filters: Ensemble and finite-record innovation whitening," 2026, manuscript.

[14] J. F. C. Kingman, *Poisson Processes*. Oxford University Press, 1993.

[15] C. Hernández, W. Yeoh, J. A. Baier, A. Felner, O. Salzman, H. Zhang, S.-H. Chan, and S. Koenig, "Multi-objective search via lazy and efficient dominance checks," in *Proc. Int. Joint Conf. Artificial Intelligence (IJCAI)*, 2023, pp. 7223–7230.

[16] L. Mandow and J. L. P. de la Cruz, "Multiobjective A* search with consistent heuristics," *Journal of the ACM*, vol. 57, no. 5, pp. 27:1–27:25, 2010.

[17] H. Zhang, O. Salzman, T. K. S. Kumar, A. Felner, C. H. Ulloa, and S. Koenig, "A*pex: Efficient approximate multi-objective search on graphs," in *Proc. Int. Conf. Automated Planning and Scheduling (ICAPS)*, vol. 32, 2022, pp. 394–403.
[18] L. Gold, A. Bienkowski, D. Sidoti, K. R. Pattipati, and O. Khan, "OPMOS: Ordered parallel algorithm for multi-objective shortest-paths," in *Proc. 39th ACM Int. Conf. Supercomputing (ICS '25)*, 2025, pp. 1146–1161.
[19] L. Gold, D. Sidoti, K. R. Pattipati, and O. Khan, "MPMOS: Massively parallel multi-objective shortest paths," in *Proc. 40th ACM Int. Conf. Supercomputing (ICS '26)*, 2026, pp. 395–408.
[20] ——, "Regular-MAGMO: Multi-attribute graphs for multi-objective graph processing with application to maritime routing," in *Proc. IEEE Int. Symp. Workload Characterization (IISWC)*, Boulder, CO, Sep. 2026.
[21] A. Bienkowski, L. Zhang, D. Sidoti, and K. R. Pattipati, "Intelligent search and multiobjective path planning with application to ship routing," *IEEE Journal of Oceanic Engineering*, vol. 51, no. 1, pp. 370–388, 2026.
[22] A. Bienkowski, D. Sidoti, L. Zhang, and K. R. Pattipati, "Robust multi-objective route planning for maritime navigation under weather uncertainty," 2026, IEEE Access, under revision.
[23] D. Silver and J. Veness, "Monte-Carlo planning in large POMDPs," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2010, pp. 2164–2172.
[24] M. Kearns, Y. Mansour, and A. Y. Ng, "A sparse sampling algorithm for near-optimal planning in large Markov decision processes," *Machine Learning*, vol. 49, no. 2–3, pp. 193–208, 2002.
[25] D. P. Bertsekas, *Dynamic Programming and Optimal Control, Vol. I*, 4th ed. Belmont, MA: Athena Scientific, 2017.
[26] ——, *Rollout, Policy Iteration, and Distributed Reinforcement Learning*. Belmont, MA: Athena Scientific, 2020.
[27] D. G. Bonett and R. M. Price, "Adjusted Wald confidence interval for a difference of binomial proportions based on paired data," *Journal of Educational and Behavioral Statistics*, vol. 37, no. 4, pp. 479–488, 2012.
[28] H. B. McMahan, G. J. Gordon, and A. Blum, "Planning in the presence of cost functions controlled by an adversary," in *Proc. Int. Conf. Machine Learning (ICML)*, 2003, pp. 536–543.
[29] ETV Bharat, "India's navigation satellite system NavIC goes offline amid critical satellite loss," https://www.etvbharat.com/en/technology/indias-navigation-satellite-system-navic-goes-offline-amid-critical-satellite-loss-enn26031403113/, Mar. 2026.
[30] Business Today, "Only three satellites left active: Centre acknowledges NavIC standalone navigation gap," https://businesstoday.in/india/story/only-three-satellites-left-active-centre-acknowledges-navic-standalone-navigation-gap-546285-2026-07-30, Jul. 2026.
[31] ISPreview, "Starlink shuts GPS location feature on satellite broadband service," https://www.ispreview.co.uk/index.php/2026/04/starlink-shuts-gps-location-feature-on-satellite-broadband-service.html, Apr. 2026.
[32] RV Mobile Internet, "Starlink reverses course, restores local GPS location API access for priority plans," https://www.rvmobileinternet.com/starlink-reverses-course-restores-local-gps-location-api-access-for-priority-plans/, May 2026.
[33] C. F. Van Loan, "Computing integrals involving the matrix exponential," *IEEE Trans. Autom. Control*, vol. 23, no. 3, pp. 395–404, 1978.