

Robust LEO gateway-region planning for 6G NTN return paths under spatial service-priority uncertainty

Runzhe Liu

2025082583@yn.jtc.edu.cn

Yunnan Communications Vocational and Technical University

Dabin Wang

Yunnan Communications Vocational and Technical University

Yun He

Yunnan Communications Vocational and Technical University

Article

Keywords: 6G-oriented NTN, LEO gateway regions, return-path opportunity, minimax regret, spatial uncertainty, resilience

Posted Date: October 27th, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-10371009/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Robust LEO gateway-region planning for 6G NTN return paths under spatial service-priority uncertainty

Runzhe Liu^{1,*}, Dabin Wang¹, Yun He¹

¹ School of Transportation Information Engineering, Yunnan Communications Vocational and Technical University, Kunming, China

* Correspondence: 2025082583@ynjtc.edu.cn; ORCID: 0009-0009-7124-4508

Abstract

Ubiquitous connectivity requires ground-segment decisions that remain useful under uncertain spatial service priorities. We examine gateway-region planning for a sampled low-Earth-orbit return path, requiring user-uplink and feeder-downlink closure through the same satellite and epoch. Four unit-sum priority maps over 368 aggregated demand nodes were evaluated against 400 demand-independent candidate regions and three independently propagated reference architectures. Under a generic 0-dB C/N closure threshold and a fixed 1%-annual-exceedance propagation design state, the prospective Walker architecture attained a mean opportunity rate of 99.31% across maps with one gateway, a minimum-map rate of 99.16%, maximum normalized regret of 0.0737%, and a 4.5-min worst-node sampled gap. Two gateways raised the mean to 99.98%, shortened the gap to 0.5 min, and retained 77.95% worst-case $N-1$ absolute opportunity. Minimax regret tied with simpler objectives in the Walker and dense architectures; in the sparse architecture, in the main 24-h frozen-ephemeris window, it reduced one-gateway maximum regret from 2.4544% for an equal mixture to 1.0945%. Architecture constrained the opportunity envelope, whereas spatial priority most affected choices under sparse or margin-limited conditions. Within the analyzed 24-h reference window, cross-architecture comparisons assess how architecture constraints alter the value of demand-robust planning.

Keywords: 6G-oriented NTN; LEO gateway regions; return-path opportunity; minimax regret; spatial uncertainty; resilience

Introduction

The International Telecommunication Union identifies ubiquitous and resilient connectivity as objectives for IMT-2030 and beyond [1]. Low-Earth-orbit (LEO) non-terrestrial networks (NTNs) can contribute time-varying access opportunities over territories where terrestrial expansion is difficult, and 3GPP studies provide reference assumptions for propagation, mobility and NTN link evaluation [2,3]. The ground segment remains a separate architecture decision. A gateway region is useful only when a satellite path can connect a demand location to that region at the relevant epoch, and high average opportunity does not establish short interruption gaps or resilience to losing a gateway.

Gateway-planning studies have examined geographic coverage, traffic estimates, latency, reliability, gateway diversity and joint placement-routing decisions [4–10]. Feeder-link switching and constellation-failure analyses further show that orbital motion and gateway availability cannot be reduced to a static terrestrial distance [11–13]. These studies clarify the physical side of gateway planning. They commonly condition the objective, however, on one demand map or one traffic estimate. At national scale, population concentration, inverse-density inclusion, structural need and allocation-derived priority can support different planning questions. None of these maps is measured traffic, and their disagreement can change which geographically separated nodes receive value in the objective.

The choice of communication metric also matters. A spatially weighted mean can remain close to one while a low-priority node experiences a long contiguous gap. A second selected gateway can raise nominal opportunity without providing two simultaneously eligible gateway regions, and a nominally strong portfolio can lose much of its value after one region is removed. These quantities address distinct design questions: average opportunity describes the overall envelope; a minimum-map value protects against priority-map mismatch; a worst-node gap retains temporal order; simultaneous gateway opportunity measures available regional alternatives; and $N-1/N-2$ stress measures post-selection exposure. Combining them into one uncalibrated score would obscure which property changed.

Spatial-priority uncertainty differs from link uncertainty. Orbital geometry and propagation determine whether a same-satellite path exists. A priority map determines how the resulting opportunities are valued across nodes. Optimizing a single map can give a portfolio that is physically sound yet sensitive to an unverified planning objective. Averaging maps avoids choosing one map but can conceal a large relative loss for another. A scenario-wise comparison is more interpretable when each loss is normalized by the best value achievable for that map under the same architecture and budget.

Minimax regret provides this normalization [14,15]. For every priority map, it computes a scenario-specific ideal and then minimizes the largest relative shortfall. Robust facility-location models use related principles for uncertain demand [16], but an NTN formulation must preserve satellite identity and epoch: the user and gateway links are useful together only when the same satellite closes both hops at the same time. The formulation also has to separate architectures. Two public ephemeris sets and a prospective Walker construction differ in altitude, inclination, density and temporal geometry, but their satellites, gateway equipment and feeder resources cannot be pooled without an explicit interoperability model.

The comparison has two levels. Within an architecture, changing the priority map tests objective uncertainty while holding the physical opportunity tensor fixed. Across architectures, rebuilding the tensor tests how the same planning logic behaves under a different orbital envelope. Identical selected coordinates are neither required nor expected; the relevant comparison is whether constraint severity changes priority sensitivity across independently solved systems.

This study uses the geography of mainland China as a national-scale demand case with 368 aggregated demand nodes. The main pool contains 400 candidate gateway regions selected independently of the four priority maps. A prospective Walker-type constellation is the primary future-oriented reference architecture. Dense and sparse public-ephemeris sets provide architecture contrasts with different orbital opportunity patterns in the main reference window. Their source-group names are retained only in the Methods and Supplementary Information; the analysis does not reproduce operator beam patterns, scheduling, capacity, spectrum use or operational performance.

The analysis considers the return path only. The implemented path is user uplink $\rightarrow$ satellite $\rightarrow$ gateway feeder downlink. A valid sample represents a same-satellite two-hop return-path opportunity. Forward feeder uplink and user downlink parameters are absent, so bidirectional connectivity is outside the evidence. This definition permits physically meaningful tests of gateway-region screening, temporal continuity and failure exposure without attaching unsupported throughput or end-user availability claims.

The study makes three contributions. First, it connects multiple spatial service-priority maps to a time-varying same-satellite return-path tensor built from user-uplink and feeder-downlink budgets. Second, it formulates architecture-specific gateway-region planning as minimization of normalized regret relative to map-specific physical ideals and reports exact-enumeration or mixed-integer solver evidence. Third, it quantifies how priority uncertainty propagates to sampled opportunity, contiguous gaps, simultaneous gateway opportunities, budget response and $N-1/N-2$ stress, then tests whether those conclusions transfer across three separately modelled reference architectures (Fig. 1a–c).

Results

Spatial priority maps create directionally different objectives

The four planning-priority maps assigned contrasting values to the same 368 demand nodes (Fig. 2a–d). Density and inverse density were exact reverse rankings (Spearman $\rho = -1.000$). Density was also negatively associated with structural need ($\rho = -0.718$) and calibrated priority ($\rho = -0.808$), whereas structural need and calibrated priority were positively associated ($\rho = 0.782$; Fig. 2e). The maps encode directional disagreement, not small perturbations of one national ranking.

Macroregional shares show where the objectives diverge (Fig. 2f). Density assigned 29.75% of its national weight to the East, while inverse density and structural need placed their largest shares in the Northwest (21.94% and 19.89%). Calibrated priority placed 25.73% in the Southwest and 21.82% in the Northwest. These values describe planning emphasis only. They do not estimate traffic volumes, offered load or observed network demand.

Candidate generation did not use these weights. The main 400-region pool was selected from a 2,850-region geographic universe by provincial seeding and deterministic farthest-first ordering. Nested 200- and 800-region pools

were tested for both public-ephemeris architectures, and the complete 2,850-region universe was additionally enumerated for the sparse architecture. The roads/railways dataset was used only as a construction-accessibility descriptor in the Supplementary Information. This separation prevents a map used in the objective from also determining which candidate regions are allowed to compete.

The map disagreement also separates concentration from inclusion. A density optimum rewards repeated opportunity in already concentrated nodes, whereas inverse density increases the value of geographically dispersed low-density nodes. Structural need and calibrated priority partly follow the inverse-density direction but retain independent rankings, as shown by correlations below one. The four maps were therefore kept as separate scenarios; replacing them with one blended surface before optimization would remove the uncertainty the study is designed to measure.

Communication constraints replace the static-distance proxy

The model ladder isolated how increasingly specific communication assumptions changed the one-gateway result (Fig. 4a–c). M0, a 400-km static-distance negative control, produced a 12.41% mean opportunity across maps, a 1,440-min worst-node gap and 21.23% maximum normalized regret for every architecture because it contained no orbital information. M1 used same-epoch geometry only and saturated all three architectures at 100% mean opportunity with no sampled gap. Geometry alone was too permissive to discriminate candidate regions.

Adding the two link budgets in M2 changed both the opportunity envelope and location set. For the prospective Walker architecture, the one-gateway mean fell to 99.31%, the minimum across maps to 99.16%, and the worst-node gap became 4.5 min; maximum normalized regret was 0.0737%. The dense public-ephemeris architecture retained 100.00% after rounding, with an unrounded mean of 99.995% and a 1.0-min gap. The sparse architecture fell to 96.72%, with a 37.5-min gap and 0.0901% maximum regret. Thus, the same-satellite link tensor produced different performance and selected regions from the static control.

M3 conditioned feeder closure on the fixed 1%-annual-exceedance rain/cloud design state. The rounded M2 and M3 values were identical for the Walker and dense architectures. Its effect was concentrated in the sparse architecture, where link margins were most constrained.

For the sparse public-ephemeris reference in the main 24-h window, mean opportunity declined from 96.72% to 87.25%, the worst-node gap increased from 37.5 to 689 min, and maximum regret increased from 0.0901% to 1.0945%. These values are conditional on the frozen ephemeris and the selected 24-h window; temporal decision convergence was evaluated separately for the Walker reference only. The propagation state became decision-relevant where orbital and link margins were already constrained.

Complete optimal-set comparisons changed the interpretation of location stability (Fig. 4d). For $K = 1$, the Walker $M0 \rightarrow M1$ and $M1 \rightarrow M2$ optimal sets had zero Jaccard overlap, while $M2 \rightarrow M3$ had Jaccard 1. For $K = 2$, the $M1 \rightarrow M2$ Jaccard range was 0–1 because multiple tied optima existed; $M2 \rightarrow M3$ was exactly 1/3. A single deterministic tie-break would have hidden this range. Figure 4f displays selected screening anchors, while location claims use the full enumerated optimum sets.

Link margins and temporal checks support the sampled sequence

The C/N-margin distributions explain the architecture ordering (Fig. 3a,b). Under the 0-dB reference threshold, the sparse architecture had the largest mass below closure for both the user-uplink and clear-state feeder-downlink samples. The dense set had the largest link-eligible satellite counts, and the prospective Walker lay between the dense and sparse opportunity patterns (Fig. 3c). These are reference-link results; the threshold was not mapped to an MCS or BLER target.

The user and feeder samples also show why a single visibility flag is insufficient. A satellite can exceed the elevation mask for a demand node while failing the feeder budget at a candidate region, or the reverse. Equation (4) retains only the intersection for the same satellite and epoch. The dense architecture had enough simultaneous candidates for this intersection to remain almost complete; the sparse architecture did not. The model thereby attributes the opportunity loss to joint closure instead of assigning it to one hop by inspection.

The contiguous-gap CCDF retained temporal ordering and exposed differences that an average alone would miss (Fig. 3d). With one gateway in M3, the worst-node sampled gaps were 1.0 min for the dense architecture, 4.5 min for the Walker architecture and 689 min for the sparse architecture. The sparse tail persisted across hundreds of minutes.

The metric is computed from zero runs of the binary opportunity sequence and is not an annual weather-outage statistic.

Decision convergence was evaluated by rebuilding the Walker tensor and reoptimizing at 10, 30 and 60 s (Fig. 3e). For $K = 1$, mean opportunity across maps was 99.304%, 99.309% and 99.308%, respectively; the selected candidate index was 115 at all three resolutions. Maximum regret ranged from 0.0737% to 0.0814%. The worst-node gap changed from 3.5 to 4.5 and 9.0 min because a run is quantized by the sampling interval. For $K = 2$, mean opportunity remained between 99.979% and 99.980%, while the number of strict optima changed from seven at 10/30 s to nine at 60 s.

Start-epoch and window tests were kept separate from sampling resolution (Fig. 3f). Three 24-h starts gave Walker $K = 1$ means of 99.309%, 99.296% and 99.289%; a 72-h window gave 99.298%. The $K = 2$ mean remained 99.979–99.980%. Re-evaluating every selected portfolio on the common 10-s grid gave the same $K = 1$ metrics and zero maximum regret for all $K = 2$ selections within numerical tolerance. Temporal sampling affected gap quantization and tie-breaking more than the opportunity ordering in this reference case.

Robust objectives matter conditionally

Cross-map evaluation identifies the cost of optimizing the wrong priority map (Fig. 5a). In the Walker $K = 1$ case, applying the density optimum to inverse-density weights produced 0.1600% normalized regret; the largest structural-need shortfall was 0.0963%. The minimax solution limited the largest shortfall to 0.0737%. These are small effects relative to the architecture-induced opportunity differences, but they are computed against map-specific physical ideals and have a consistent interpretation.

Minimax regret was not uniquely better in every setting. Communication baselines are summarized in Fig. 4e, whereas the four non-equivalent planning objectives are compared in Fig. 5b and Table 2. For Walker $K = 1$, minimax regret, equal mixture, unweighted-node optimum, unnormalized max-min and greedy marginal availability all selected portfolios with the same maximum regret of 0.0737%. Dense-architecture values were below $10^{-8}\%$ for the same objectives. In the sparse $K = 1$ case, minimax regret and the unweighted-node optimum attained 1.0945%, whereas equal mixture and unnormalized max-min attained 2.4544%. The feeder-ranking, p-median and k-center controls had sparse-architecture maximum regrets of 18.85%, 15.40% and 3.23%, respectively. The robust objective was useful when the opportunity surface exposed map conflict, not as a universally dominant rule.

Random feasible portfolios provided a scale reference. Across 1,000 one-gateway draws, mean maximum regret was 27.21% for Walker, 20.65% for dense and 40.22% for sparse; the corresponding 2.5–97.5% intervals were 1.34–82.60%, 0.28–77.71% and 3.21–89.84%. Exact p-median and k-center solutions were evaluated on the same return-path tensor after solving their geographic objectives, so their weaker regret does not imply an error in those baseline optimizations.

Tie-aware enumeration further limits location claims. Walker M3 had one strict $K = 1$ minimax optimum and seven strict $K = 2$ optima; the 10^{-4} epsilon set contained 18 $K = 2$ portfolios. Dense M3 had 3,849 strict and 9,645 epsilon-optimal $K = 2$ portfolios, while sparse M3 had one strict and two epsilon-optimal portfolios (Fig. 5f). Across architectures, $K = 1$ optimum locations had zero Jaccard overlap. For $K = 2$, Walker and dense optimum sets had a Jaccard range of 0–1, whereas every comparison involving the sparse set had zero overlap (Fig. 5d). Figure 5e shows one deterministic single-map optimum per map to make the geographic divergence visible; all location-stability statements use the enumerated optimum sets. Architecture transfer concerns the method and evidence pattern, not identical selected coordinates.

Cross-architecture comparison within the main analysis window

The same priority maps produced different opportunity envelopes after each architecture was solved independently (Fig. 5c and Table 2). At $K = 1$, mean opportunity spanned 87.25–100.00% across architectures, while the minimum-map value spanned 85.45–99.99%. At $K = 2$, the mean range narrowed to 96.50–100.00%, but simultaneous ≥ 2 -gateway opportunity still ranged from 46.58% to 97.50%. The gap and redundancy metrics therefore retained architecture differences after average opportunity approached saturation.

Location transfer was weaker than metric transfer. No strict $K = 1$ optimum was shared by a pair of architectures. At $K = 2$, Walker and dense optimum sets could contain identical location pairs, but their minimum Jaccard was zero

because each architecture also admitted architecture-specific optima. Sparse optima had no location overlap with either other architecture. These results support reuse of the planning procedure and reporting logic, while rejecting a universal list of gateway regions.

The prospective Walker one-factor tests show that its main setting is one point within a broad architecture range (Fig. 6f). Halving the population to 324 satellites reduced $K = 1$ mean opportunity to 94.63%, while doubling it to 1,296 raised the mean to 99.85%. Raising altitude to 1,200 km reduced the mean to 94.10% and produced a 1,440-min worst-node gap for $K = 1$. Changing inclination to 45° raised the mean to 99.76% but yielded a 15-min gap, whereas 70° gave 98.07% and a 26.5-min gap. Mean opportunity and temporal continuity can thus order architecture variants differently.

Across these variants, changes in architecture could enlarge or suppress the decision effect of the same four maps. The spatial scenarios were fixed, so this interaction cannot be attributed to retraining or rescaling demand. It follows from how each architecture reshapes the candidate-by-node opportunity tensor before optimization.

Budget and failure stress expose different planning limits

Increasing the gateway budget compressed nominal opportunity differences (Fig. 6a,b). Walker mean opportunity rose from 99.31% at $K = 1$ to 99.98% at $K = 2$ and 100.00% at $K = 5$; its worst-node gap fell from 4.5 to 0.5 and then 0 min. Dense opportunity reached 100.00% at $K = 2$. The completed sparse-architecture budgets were $K = 1,2,10,21$: mean opportunity rose from 87.25% at $K = 1$ to 96.50% at $K = 2$ and 98.69% at $K = 10$, while its gap fell from 689 to 47 and 7.5 min. Additional gateways could not exceed the finite reference-link opportunity envelope in the sparse case.

The sparse curve plateaued at 98.69% from $K = 10$ through $K = 21$, and its worst-node gap remained 7.5 min. In contrast, Walker and dense reached 100.00% mean opportunity by $K = 10$. The sparse plateau reflects node–epoch pairs that remain infeasible under the reference-link assumptions. A larger budget cannot substitute for a different link threshold, propagation state or orbital architecture.

Simultaneous gateway opportunities distinguished nominal closure from geographic redundancy (Fig. 6c). At $K = 2$, the mean rate of epochs with at least two eligible selected gateways was 97.50% for dense, 78.97% for Walker and 46.58% for sparse. These values count gateway alternatives for the same node and epoch; they do not assert independent satellite, capacity or equipment resources.

Failures were introduced after selection, with surviving gateways allowed to reassociate (Fig. 6d). At $K = 2$, worst N–1 absolute opportunity was 96.35% for dense, 77.95% for Walker and 53.98% for sparse. Relative retentions were 96.35%, 77.97% and 56.00%. N–2 necessarily reduced every two-gateway portfolio to zero. At the common $K = 10$ budget, worst N–2 absolute opportunities were 100.00%, 99.99% and 96.36% for dense, Walker and sparse, respectively. The nominally near-complete Walker $K = 2$ solution remained exposed to losing one gateway.

Threshold and architecture tests show where these conclusions cease to hold (Fig. 6e,f). At $K = 2$, increasing the generic C/N threshold from 0 to +3 dB reduced mean opportunity to 86.34% for Walker, 94.51% for dense and 0.65% for sparse; +6 dB reduced them to 25.29%, 76.12% and 0.10%. Raising minimum elevation from 10° to 30° produced 58.70%, 67.71% and 87.90%. In the prospective Walker one-factor tests, the $K = 1$ mean ranged from 94.10% (1,200-km altitude) to 99.85% (1,296 satellites), and the $K = 2$ mean ranged from 97.98% to 100.00%. Architecture assumptions determine the envelope within which priority uncertainty can influence gateway regions.

Discussion

The results separate two sources of planning variation. Architecture and link constraints set the attainable return-path opportunity envelope, while spatial priorities determine which locations are valuable within that envelope. The difference is clearest in the sparse reference architecture: its $K = 1$ mean was 87.25%, its worst-node gap was 689 min and its maximum regret was 1.0945%. In the dense architecture, several objectives were numerically tied and thousands of $K = 2$ portfolios were optimal. A demand-robust objective adds little when almost every node-epoch pair is feasible; it becomes informative when communication constraints reveal trade-offs among nodes. This conditional result is more useful than a global method ranking. It indicates that robust demand treatment should be paired with a diagnostic of physical constraint severity, because an unconstrained opportunity matrix can make all

reasonable spatial objectives look equivalent. For 6G-oriented architecture studies, the practical sequence is to validate the physical closure tensor first, quantify whether priorities produce distinct ideals second, and select a robust objective only when the resulting regret is material.

The same-satellite construction is the communication mechanism behind that result. M1 geometry saturated the problem, whereas M2 link closure changed opportunity and location, and the fixed M3 propagation state further constrained the sparse architecture. The static M0 control selected different regions and could not represent temporal continuity. This progression supports a return-path planning claim without extending it to throughput, latency or bidirectional access. It also shows why spatial optimization and communication modelling must be evaluated together: a sophisticated robust objective cannot recover information absent from the physical tensor. Conversely, adding a link budget does not guarantee a decision change; Walker and dense M2–M3 values were unchanged at the reported precision. A model component is relevant only where it alters closure, rankings or uncertainty, so the ladder provides a direct check against attributing every result to propagation complexity.

Priority uncertainty affected decisions in two distinct ways. Cross-map regret measured performance loss, while complete optimal-set comparisons measured location ambiguity. Small regret did not imply a unique map-invariant region: Walker $K = 2$ had seven strict optima, and dense $K = 2$ had 3,849. Conversely, the sparse architecture had a unique strict solution but larger regret and lower opportunity. Reporting a single tie-broken map would overstate precision in the dense case, while reporting only objective values would hide the sparse case’s geographic rigidity. The epsilon sets add a practical layer: 18 Walker and 9,645 dense portfolios were within the stated tolerance, but only two sparse portfolios were. Where many near-equivalent portfolios exist, later site investigations can use land, fibre, power and permitting criteria without materially changing the communication objective. Where the set is narrow, those omitted criteria may force a measurable opportunity trade-off.

Continuity and post-selection failure stress add information beyond the mean. The Walker $K = 2$ portfolio attained 99.98% mean opportunity and a 0.5-min sampled gap, yet its worst N–1 absolute opportunity was 77.95%. Increasing to five regions raised worst N–2 opportunity to 99.62%. The simultaneous ≥ 2 -gateway rate of 78.97% anticipated this exposure more clearly than the nominal mean, although it is not a reliability probability. These results do not make failure exposure part of the optimization; they quantify how a nominal minimax portfolio behaves after removals. A future failure-aware model could place N–1/N–2 terms inside the objective once failure probabilities, gateway independence and regional hazards are calibrated. Such a model should retain both absolute post-failure opportunity and relative retention because a high retention ratio can still correspond to poor absolute performance.

Several boundaries limit external interpretation. The national geography is a demand case study, and selected coordinates are regional screening anchors rather than deployment-ready sites. The public ephemeris sets are reference architectures, not evaluations of operator networks or candidates for domestic adoption. The prospective Walker parameters are transparent study settings, not a representation of a complete Chinese constellation. The link model omits the forward path, beam scheduling, gateway concurrency, interference, spectrum coordination, capacity, terrestrial backhaul and time-correlated weather. The fixed annual-exceedance attenuation state supports design-state sensitivity but not annual network availability. In addition, the priority maps are planning scenarios, so their weights should not be interpreted as offered load or user counts. Future work should add calibrated antenna patterns, four-direction link budgets, gateway concurrency, time-resolved propagation, verified terrestrial infrastructure and offered-load conservation before operational planning conclusions are drawn. Those additions would allow opportunity-based screening to develop into a capacity- and reliability-aware ground-segment model without changing the present distinction between physical architecture uncertainty and spatial-priority uncertainty. Independent site feasibility, spectrum rights and equipment compatibility would still be required before any regional anchor could enter an engineering design.

Temporal decision stability was evaluated for the prospective Walker reference. The dense and sparse public-ephemeris results characterize the main frozen 24-h window and should not be interpreted as multi-epoch or operational performance estimates.

Within those boundaries, the present framework offers a reproducible intermediate layer between orbital link analysis and detailed gateway engineering. It can reject demand-sensitive choices that fail under an alternative priority map, expose portfolios whose near-complete mean hides long gaps or weak N–1 performance, and show when architecture

assumptions overwhelm the spatial objective. The framework supports regional screening before capacity- and reliability-aware gateway engineering.

Methods

Study units and spatial-priority scenarios

Let I denote 368 aggregated demand nodes and G a demand-independent main pool of 400 candidate gateway regions. The candidate universe contained 2,850 regions. Candidate ordering used provincial seeding followed by deterministic farthest-first geographic selection; the four objective maps were not used in this step. The mainland-China geography supplies national-scale variation in node density and remoteness but does not imply a domestic operator plan.

The planning set M contained density, inverse density, structural need and calibrated priority. Each non-negative map was normalized so that $\sum_{i \in I} w_i^{(m)} = 1$. Density and inverse density intentionally represent opposite concentration choices. Structural need combines the retained remoteness-related inputs, and calibrated priority is derived from the allocation model described in Supplementary Section S2. Equal-node weights were used as an evaluation control. The complete classifier, calibration, bootstrap, Moran's I and data-audit material is reported in the Supplementary Information because those analyses construct priority scenarios rather than communication links.

Reference architectures and orbit propagation

The architecture set C contained a prospective Walker reference, a dense public-ephemeris reference and a sparse public-ephemeris reference. The Walker main setting used 600-km altitude, 55° inclination, 36 planes, 18 satellites per plane and phase factor 1. Altitude was anchored to the public LEO reference range in 3GPP TR 38.821; inclination, satellite count, plane count and phasing are study settings tested over the ranges in Table 1. The construction follows the Walker-delta convention [24].

The dense and sparse snapshots were downloaded from the CelesTrak OMM endpoints on 3 July 2026 [23]. They contained 10,709 and 651 unique NORAD identifiers, respectively. Active status could not be established consistently from the snapshots, so the source groups are treated as orbital reference sets. Duplicate handling, epoch distributions, ephemeris age, file SHA-256 values and SGP4 status codes are reported in Supplementary Table S4. Three dense-set objects produced SGP4 code-1 states and 3,019 affected object-epochs were quarantined; the sparse set produced no propagation error.

All architectures were propagated independently with SGP4 [17] over a half-open UTC interval beginning 3 July 2026 12:00:00. The main interval was 24 h at 30-s resolution. Independent 10-, 30- and 60-s runs rebuilt the tensor, scenario ideals and $K = 1, 2$ solutions. Additional tests used two later 24-h start epochs and a 72-h window. Satellites, links and gateway resources were never combined across architectures.

Directional reference-link budgets

The implemented direction is user uplink followed by satellite feeder downlink. For a link with carrier frequency f , slant range d and bandwidth B , free-space path loss was

$$L_{\text{FS}} = 92.45 + 20\log_{10}(f_{\text{GHz}}) + 20\log_{10}(d_{\text{km}}) \quad \text{dB.} \quad (1)$$

Carrier-to-noise density and carrier-to-noise ratio were calculated as

$$C/N_0 = \text{EIRP} + G/T - L_{\text{FS}} - L_{\text{other}} + 228.6 \quad \text{dB} - \text{Hz}, \quad (2)$$

$$C/N = C/N_0 - 10\log_{10}(B_{\text{Hz}}) \quad \text{dB.} \quad (3)$$

L_{other} includes the applicable gas, shadow, scintillation, pointing, polarization, implementation, rain and cloud terms. Closure required elevation at least 10° and $C/N \geq 0$ dB. The 0-dB value is a generic C/N closure threshold and is not mapped to an MCS or BLER. Thresholds of -3, 0, +3 and +6 dB and elevations of 5°, 10°, 15°, 20° and 30° were recomputed.

The user-uplink reference used 2 GHz, 0.36 MHz, -7 dBW EIRP, 1.1 dB/K satellite-receiver G/T , 0.1-dB gas loss, 3-dB shadow margin and 2.2-dB scintillation loss from the cited 3GPP rows [3]. The feeder downlink used 20 GHz

and 400 MHz. Its 30.0206-dBW EIRP follows $4 \text{ dBW/MHz} + 10\log_{10}(400 \text{ MHz})$. The 15.9-dB/K receiver value originates from a 60-cm VSAT row; transferring it to a generic gateway-region reference receiver is a study assumption, not a sourced operational gateway specification. The complete row-level provenance and conversions appear in Supplementary Table S5.

M3 used site-dependent rain and cloud attenuation at a fixed 1%-annual-exceedance design state following ITU-R P.618-14, P.837-8 and P.840-9 [18,20,21]. ITU-R P.1511-3 supplied the topographic input [22]. P.676-13 was version-audited [19], while the main gas-loss constants came from the 3GPP reference rows. No minute-by-minute weather sequence was generated. Temporal gaps in M3 arise from orbital ordering under a fixed propagation state.

Same-satellite return-path tensor

Let S_c be the satellites in architecture c and T the sampled epochs. User and feeder margins are defined as $M = C/N - \gamma$, where γ is the generic C/N closure threshold, and are denoted $M_{ist}^{\text{user},c}$ and $M_{sgt}^{\text{feed},c}$. In the implementation, margins were set to $-\infty$ whenever the corresponding elevation was below $e_{\min}$. The binary return-path tensor is

$$A_{igt}^c = \mathbb{1}\{\exists s \in S_c: M_{ist}^{\text{user},c} \geq 0 \wedge M_{sgt}^{\text{feed},c} \geq 0 \wedge e_{ist}^{\text{user},c} \geq e_{\min} \wedge e_{sgt}^{\text{feed},c} \geq e_{\min}\}, \quad \forall i \in I, g \in G, t \in T, c \in C. \quad (4)$$

The same satellite must satisfy both inequalities at the same epoch. Inter-satellite links are absent. Because feeder uplink and user downlink were not modelled, A_{igt}^c is a two-hop return-path opportunity and not bidirectional connectivity.

For a selected gateway set $X_c \subseteq G$, node opportunity is

$$a_{i,c}(X_c) = \frac{1}{|T|} \sum_{t \in T} \mathbb{1}\left\{ \sum_{g \in X_c} A_{igt}^c \geq 1 \right\}. \quad (5)$$

Map-weighted performance, the reported mean across the four maps and the minimum across maps are

$$F_{m,c}(X_c) = \sum_{i \in I} w_i^{(m)} a_{i,c}(X_c), \quad \bar{F}_c(X_c) = \frac{1}{|M|} \sum_{m \in M} F_{m,c}(X_c), \quad F_c^{\min}(X_c) = \min_{m \in M} F_{m,c}(X_c). \quad (6)$$

Accordingly, every single “mean opportunity” in the abstract, figures and Table 2 is $\bar{F}_c$; the minimum-map value is $F_c^{\min}$.

Minimax-regret gateway-region planning

For architecture c , priority map m and budget K , the scenario ideal is

$$F_{m,c}^* = \max_{X \subseteq G: |X|=K} F_{m,c}(X), \quad (7)$$

and normalized regret is

$$R_{m,c}(X) = \frac{F_{m,c}^* - F_{m,c}(X)}{F_{m,c}^*}. \quad (8)$$

The primary optimization is solved independently for each architecture:

$$\min_{X \subseteq G: |X|=K} \max_{m \in M} R_{m,c}(X). \quad (9)$$

No cross-architecture minimax problem is used in the main analysis.

The mixed-integer form uses gateway variables $x_g \in \{0,1\}$ and node-epoch closure variables $y_{it} \in \{0,1\}$. For a fixed architecture,

$$\sum_{g \in G} x_g = K, \quad y_{it} \leq \sum_{g \in G} A_{igt}^c x_g \quad \forall i, t, \quad y_{it} \geq A_{igt}^c x_g \quad \forall i, t, g. \quad (10)$$

With $r \geq 0$, the regret constraints are

$$r \geq 1 - \frac{\sum_i w_i^{(m)} |T|^{-1} \sum_t y_{it}}{F_{m,c}^*} \quad \forall m \in M, \quad \min r. \quad (11)$$

Identical node-epoch incidence rows were aggregated into weighted signatures before optimization. $K = 1$ and $K = 2$ were completely enumerated over 400 and 79,800 portfolios, respectively. Strict optima used an absolute objective tolerance of 10^{-12} . For minimax regret, the epsilon set used an absolute 10^{-4} objective allowance; for maximization objectives it used a relative 10^{-4} allowance. One displayed solution was selected by lexicographic candidate-index tie-breaking, while all location-stability claims use complete optimal sets. Set overlap was measured by $J(X, Y) = |X \cap Y| / |X \cup Y|$. Larger budgets were solved with HiGHS; main results with zero reported relative gap are labelled optimal, and any run lacking a certificate is labelled heuristic.

Baselines, continuity and failure metrics

Baselines used the same candidate pool, budget and tensor. They included random feasible portfolios; exact geographic p-median and k-center controls; feeder-link ranking; greedy marginal opportunity; the unweighted-node optimum $|I|^{-1} \sum_i a_{i,c}(X)$; a single-map optimum; the equal-mixture objective $|M|^{-1} \sum_m F_{m,c}(X)$; and unnormalized max-min $\max_X \min_m F_{m,c}(X)$. The implementation label maximum_average_availability denotes the unweighted-node objective and is reported only under that canonical name. The normalized max-min form was omitted because it is algebraically identical to minimax regret. Random results used 1,000 portfolios and seed 20260714.

For each node, let $b_{it} = \mathbb{1}\{\sum_{g \in X} A_{igt}^c \geq 1\}$. The longest sampled gap is

$$L_i = \Delta t \max\{\text{zero-run length in } (b_{it})_{t \in T}\}, \quad L_{\max} = \max_{i \in I} L_i. \quad (12)$$

If no zero run occurs for node i , $L_i = 0$. Runs touching the start or end of the half-open window were retained as observed, truncated runs; the sequence was not wrapped cyclically. Gap CCDFs use all zero runs across the 368 nodes for one portfolio as the denominator.

Simultaneous gateway opportunity was defined at node level by $d_{it} = \mathbb{1}\{\sum_{g \in X} A_{igt}^c \geq 2\}$ and aggregated as

$$D_{m,c}(X) = \sum_i w_i^{(m)} \frac{1}{|T|} \sum_t d_{it}, \quad \bar{D}_c(X) = \frac{1}{|M|} \sum_m D_{m,c}(X). \quad (13)$$

$\bar{D}_c$ counts simultaneous eligible gateway regions and does not assume independent satellite or capacity resources.

Failures were applied after selection. For $q \in \{1, 2\}$,

$$Q_c^{(q)}(X) = \min_{m \in MD \subseteq X, |D|=q} \min_{|D|=q} F_{m,c}(X \setminus D), \quad \rho_c^{(q)}(X) = \min_{m \in MD \subseteq X, |D|=q} \min_{|D|=q} \frac{F_{m,c}(X \setminus D)}{F_{m,c}(X)}. \quad (14)$$

Every N-1 removal and all N-2 pairs were enumerated for the reported budgets. Regional stress removed all selected regions in one macroregion. Surviving gateways reassociated at every epoch. The optimization itself remained nominal; these experiments are planning followed by post-selection failure stress testing.

Software and reproducibility

Computations ran on Windows 11 with Python 3.12.13, NumPy 2.3.5, pandas 3.0.1, SciPy 1.18.0, sgp4 2.27, Astropy 8.0.1, ITU-Rpy 0.4.0 and HiGHS 1.15.1. Each result row records its source file and SHA-256, generator, architecture, model, parameter scenario, random seed and solver status. Figures read only their corresponding Source Data CSV files and were exported as SVG, vector PDF and lossless 600-dpi PNG.

Data availability

Source Data underlying Figs. 1–6, the analysis-ready demand-node and candidate-region tables, nested candidate pools, exact frozen public OMM snapshots, compressed main-scenario optimization inputs, executed result tables, parameter provenance and file-level SHA-256 manifests are supplied with this Article in Source_Data.zip and the Supplementary Information. The archive also provides a variable dictionary, experiment catalogue, cleaning lineage and public-source metadata. Original providers retain rights in third-party material, and the corresponding source attribution and terms are recorded in the archive.

Code availability

The scripts used to prepare cleaned communication inputs, construct reference architectures, propagate orbits, calculate link margins, build return-path tensors, optimize gateway regions, analyze optimal sets and generate figures and documents are supplied in Supplementary_Software.zip. The included guide specifies the environment, execution order, random seeds, input-output contract and expected hashes. Figure and manuscript regeneration is self-contained from the released Source Data; full service-tensor reconstruction additionally uses the versioned public ITU-R component maps identified by official access page, expected path, byte size and SHA-256 in the package metadata.

Funding

This work was supported by the Scientific Research Fund Project of Yunnan Provincial Department of Education, China (grant no. 2024J1400). The funder had no role in study design, data analysis, interpretation, manuscript preparation or the decision to submit.

Author contributions

Runzhe Liu conceived the study, developed the methodology and software, conducted the analyses, prepared the figures and drafted the manuscript; Dabin Wang verified the data and results, reviewed the supplementary materials and revised the manuscript; and Yun He curated the data, supervised and validated the study, and revised the manuscript.

Competing interests

The authors declare no competing interests.

References

1. International Telecommunication Union. *Framework and overall objectives of the future development of IMT for 2030 and beyond*. Recommendation ITU-R M.2160-0 (2023). <https://www.itu.int/rec/R-REC-M.2160-0-202311-I/en>
2. 3rd Generation Partnership Project. *Study on New Radio (NR) to support non-terrestrial networks*. 3GPP TR 38.811, version 15.4.0 (2020). <https://portal.3gpp.org/desktopmodules/Specifications/SpecificationDetails.aspx?specificationId=3234>
3. 3rd Generation Partnership Project. *Solutions for NR to support Non-Terrestrial Networks (NTN)*. 3GPP TR 38.821, version 16.2.0 (2023). https://www.3gpp.org/ftp/Specs/archive/38_series/38.821/38821-g20.zip
4. Monzon Baeza, V., Ortiz, F., Lagunas, E., Abdu, T. S. & Chatzinotas, S. Gateway station geographical planning for emerging non-geostationary satellites constellations. *IEEE Network* **38**, 158–165 (2024). <https://doi.org/10.1109/MNET.2023.3321531>
5. Guo, J. *et al.* Gateway placement optimization in LEO satellite networks based on traffic estimation. *IEEE Trans. Veh. Technol.* **70**, 3860–3876 (2021). <https://doi.org/10.1109/TVT.2021.3065994>
6. Cao, Y., Guo, H., Liu, J. & Kato, N. Optimal satellite gateway placement in space-ground integrated networks. *IEEE Network* **32**, 32–37 (2018). <https://doi.org/10.1109/MNET.2018.1800016>
7. Cao, Y., Shi, Y., Liu, J. & Kato, N. Optimal satellite gateway placement in space-ground integrated network for latency minimization with reliability guarantee. *IEEE Wirel. Commun. Lett.* **7**, 174–177 (2018). <https://doi.org/10.1109/LWC.2017.2763149>
8. Iacovelli, G., Lacoste, C., Lagunas, E. & Chatzinotas, S. Ground segment design for Q/V band NGSO satellite systems via 10-norm minimization. *IEEE Wirel. Commun. Lett.* **14**, 756–760 (2025). <https://doi.org/10.1109/LWC.2024.3522559>
9. Cano, J., Israel, J. & Feral, L. AIRIS2: a smart gateway diversity algorithm for very high-throughput satellite systems. *IEEE Wirel. Commun. Lett.* **14**, 2189–2193 (2025). <https://doi.org/10.1109/LWC.2025.3565705>

10. Abe, Y. *et al.* Optimizing satellite network infrastructure: a joint approach to gateway placement and routing. In *2024 IEEE 99th Vehicular Technology Conference* 1–6 (IEEE, 2024). <https://doi.org/10.1109/VTC2024-Spring62846.2024.10683455>
11. Pizzi, S., Rinaldi, F., Iera, A., Molinaro, A. & Araniti, G. Tackling satellite mobility in LEO-based non-terrestrial networks: principles and enhancements of feeder link switch. *IEEE Veh. Technol. Mag.* **19**, 83–91 (2024). <https://doi.org/10.1109/MVT.2024.3449796>
12. Guo, B. *et al.* Resilience of mega-satellite constellations: how node failures impact inter-satellite networking over time? *IEEE Trans. Commun.* **73**, 15705–15721 (2025). <https://doi.org/10.1109/TCOMM.2025.3610221>
13. Zhang, L., Wang, L., Guo, J., Guo, J. & Xiao, Y. Gateway placement optimization for cascading survivability enhancement in low Earth orbit satellite networks. *Reliab. Eng. Syst. Saf.* **272**, 112613 (2026). <https://doi.org/10.1016/j.ress.2026.112613>
14. Kouvelis, P. & Yu, G. *Robust Discrete Optimization and Its Applications* (Springer, 1997). <https://doi.org/10.1007/978-1-4757-2620-6>
15. Assavapokee, T., Realff, M. J., Ammons, J. C. & Hong, I.-H. Scenario relaxation algorithm for finite scenario-based min-max regret and min-max relative regret robust optimization. *Comput. Oper. Res.* **35**, 2093–2102 (2008). <https://doi.org/10.1016/j.cor.2006.10.013>
16. Baron, O., Milner, J. & Naseraldin, H. Facility location: a robust optimization approach. *Prod. Oper. Manag.* **20**, 772–785 (2011). <https://doi.org/10.1111/j.1937-5956.2010.01194.x>
17. Vallado, D. A., Crawford, P., Hujsak, R. & Kelso, T. S. Revisiting Spacetrack Report No. 3. In *AIAA/AAS Astrodynamics Specialist Conference and Exhibit* (AIAA, 2006). <https://doi.org/10.2514/6.2006-6753>
18. International Telecommunication Union. *Propagation data and prediction methods required for the design of Earth-space telecommunication systems*. Recommendation ITU-R P.618-14 (2023). <https://www.itu.int/rec/R-REC-P.618>
19. International Telecommunication Union. *Attenuation by atmospheric gases and related effects*. Recommendation ITU-R P.676-13 (2022). <https://www.itu.int/rec/R-REC-P.676>
20. International Telecommunication Union. *Characteristics of precipitation for propagation modelling*. Recommendation ITU-R P.837-8 (2025). <https://www.itu.int/rec/R-REC-P.837>
21. International Telecommunication Union. *Attenuation due to clouds and fog*. Recommendation ITU-R P.840-9 (2023). <https://www.itu.int/rec/R-REC-P.840>
22. International Telecommunication Union. *Topography for Earth-space propagation modelling*. Recommendation ITU-R P.1511-3 (2024). <https://www.itu.int/rec/R-REC-P.1511>
23. CelesTrak. CelesTrak GP data: query and data formats. <https://celestrak.org/NORAD/documentation/gp-data-formats.php> (accessed 3 July 2026).
24. Walker, J. G. Satellite constellations. *Journal of the British Interplanetary Society* **37**, 559–571 (1984).

Figure legends

Figure 1 | Directional return-path architecture and independent robust-planning workflow. **a**, A sampled opportunity requires user-uplink and feeder-downlink closure through the same satellite at the same epoch; no ISL is modelled. **b**, The prospective Walker, dense public-ephemeris and sparse public-ephemeris architectures are propagated separately and produce independent tensors. **c**, Priority maps enter after tensor construction to define map-specific objectives and ideals, followed by one minimax-regret solve per architecture. Architectures do not share communication or gateway resources.

Figure 2 | Spatial service-priority scenarios over 368 aggregated demand nodes. **a–d**, Within-map percentiles for density, inverse density, structural need and calibrated priority; the same color scale applies to all maps. **e**, Spearman rank correlation. **f**, Macroregional share of each unit-map. The weights are planning priorities, not measured traffic.

Figure 3 | Reference-link margins, temporal continuity and decision convergence. **a**, User-uplink C/N margin. **b**, Clear-state feeder-downlink C/N margin. **c**, Mean numbers of user-visible and gateway-visible satellites over the 24-h main window. **d**, Zero-run CCDF for $K = 1$ M3 portfolios under the fixed propagation design state. Panels **e–f** apply only to the prospective Walker reference: **e** shows independently rebuilt and optimized 10-, 30- and 60-s tensors and portfolios, whereas **f** reports Walker start-epoch and 72-h-window checks. Opportunity is reported as the mean and minimum across four priority maps.

Figure 4 | Communication-model ladder and baselines. **a–c**, Mean opportunity, worst-node sampled gap and maximum normalized regret for one-gateway minimax portfolios under M0–M3. M0 is the shaded static-distance negative control. **d**, Min–max location Jaccard across complete optimal sets for successive Walker models. **e**, One-gateway baseline regret under Walker M3; the random interval is the 2.5–97.5% range across 1,000 portfolios. **f**, Lexicographically selected screening anchors for the three architectures.

Figure 5 | Priority mismatch, robust objectives and cross-architecture comparison. **a**, Regret when Walker $K = 1$ single-map optima are evaluated under every map. **b**, Maximum regret for four non-equivalent objectives and three architectures; Raw max-min denotes the unnormalized max-min objective, and Node mean denotes the unweighted-node objective. Exact zero values remain at zero. **c**, Mean opportunity for $K = 1, 2$. **d**, Jaccard ranges between strict scenario-specific optimum sets for the same priority map across each pair of architectures. Open symbols denote $K = 1$, filled symbols denote $K = 2$, and horizontal segments span the minimum to maximum Jaccard across the four priority maps. **e**, Selected single-map anchors. **f**, Strict and 10^{-4} epsilon-optimal portfolio counts for $K = 2$.



Figure 6 | Budget, failure and sensitivity results. **a–c**, Mean opportunity, worst-node sampled gap and simultaneous ≥ 2 -gateway opportunity versus budget; the sparse architecture has no $K = 5$ point because that run was not completed. **d**, Worst post-failure opportunity and retention after complete N–1/N–2 removals with reassociation at the common $K = 2$ and $K = 10$ budgets. Filled symbols show absolute opportunity and open symbols show retention; color and shape identify architecture. **e**, Separate $-3/0/+3/+6$ -dB threshold and 5° – 30° elevation tests with a common opportunity scale. **f**, One-factor Walker architecture sensitivity; open circles denote $K = 1$ and filled squares denote $K = 2$. Color, marker and line type identify architectures redundantly in panels **a–e**.

Table 1 | Principal directional-link, propagation, orbit and simulation parameters.

Parameter	Main value	Tested range	Unit	Component	Source
Sampling interval	30	10, 30, 60; 600 coarse check	s	Simulation	Study setting
Main time window	24	24, 72	h	Simulation	Study setting
Minimum elevation	10	5, 10, 15, 20, 30	°	Both hops	3GPP TR 38.821 Table 4.2-2 [3]
User carrier / bandwidth	2 / 0.36	Fixed	GHz / MHz	User uplink	3GPP TR 38.821 Tables 6.1.3.2-1, 6.1.3.3-1 [3]
User EIRP / receiver G/T	-7 / 1.1	Fixed	dBW / dB K ⁻¹	User uplink	3GPP TR 38.821 Tables 6.1.1.1-1, 6.1.3.3-1 [3]
User gas / shadow / scintillation	0.1 / 3 / 2.2	See Table S1	dB	User uplink	3GPP TR 38.821 Table 6.1.3.3-1 [3]
Feeder carrier / bandwidth	20 / 400	Fixed	GHz / MHz	Feeder downlink	3GPP TR 38.821 Tables 6.1.3.2-1, 6.1.3.3-1 [3]
Feeder EIRP	30.0206	Fixed	dBW	Feeder downlink	Derived from 4 dBW MHz ⁻¹ [3]
Feeder receiver G/T	15.9	Fixed	dB K ⁻¹	Generic reference receiver	Study transfer of 60-cm VSAT value [3]
Generic C/N threshold	0	-3, 0, +3, +6	dB	Both hops	Study setting
Propagation design state	1	0.1, 1, 5	% annual exceedance	Feeder downlink	ITU-R P.618-14, P.837-8, P.840-9 [18,20,21]
Walker altitude / inclination	600 / 55	600–1,200 / 45–70	km / °	Orbit	3GPP range plus study setting [3]
Walker satellites / planes	648 / 36	324–1,296 / 18–72	count	Architecture	Study setting
Main candidate pool / universe	400 / 2,850	200, 400, 800, 2,850	regions	Gateway screening	Demand-independent construction

Table 2 | Core M3 minimax-regret results under the 0-dB threshold and fixed propagation design state.


Architecture	K	Mean across maps (%)	Minimum map (%)	Maximum regret (%)	Worst-node gap (min)	Mean simultaneous ≥ 2 -gateway opportunity across maps (%)	Worst N-1 absolute (%)
Prospective Walker	1	99.31	99.16	0.0737	4.5	0.00	0.00
Prospective Walker	2	99.98	99.97	$<10^{-8}$	0.5	78.97	77.95
Dense public ephemeris	1	100.00	99.99	$<10^{-8}$	1.0	0.00	0.00
Dense public ephemeris	2	100.00	100.00	0	0.0	97.50	96.35
Sparse public ephemeris	1	87.25	85.45	1.0945	689.0	0.00	0.00
Sparse public ephemeris	2	96.50	96.38	0.1308	47.0	46.58	53.98

a Same-satellite two-hop return path**b****Architecture-specific tensor construction**

Each architecture is propagated and evaluated independently

c

Priority uncertainty enters the architecture-specific planning model



Outputs: opportunity, continuity, and post-selection resilience
 Architectures do not pool satellite link, capacity, or gateway resources

Figure 1




Figure 2




Figure 3




Figure 4




Figure 5




Figure 6

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- [SupplementaryInformation.pdf](#)
- [SupplementarySoftware.zip](#)
- [SourceData.zip](#)