

Supplementary Information

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

Authors: Runzhe Liu, Dabin Wang and Yun He

This supplement reports the demand-scenario audit, parameter provenance, public-ephemeris audit, physical validation, optimization evidence and sensitivity results supporting the Article. The implemented communication direction is user uplink to satellite to gateway feeder downlink. Every opportunity metric therefore concerns a sampled same-satellite two-hop return path.

S1. Scope, notation and metric conventions

The sets are demand nodes I , candidate gateway regions G , satellites S_c for architecture c , sampled epochs T , priority maps M and reference architectures C . Architectures are propagated, tensorized and optimized separately. No inter-satellite link, cross-architecture pooling, feeder uplink or user downlink is assumed. Candidate coordinates are regional screening anchors.

For a selected set X , $F_{m,c}(X)$ is the unit-sum map-weighted return-path opportunity. The single mean reported in the Article is $|M|^{-1} \sum_m F_{m,c}(X)$; the accompanying minimum is $\min_m F_{m,c}(X)$. A contiguous opportunity gap is a zero run in the original binary epoch sequence. Runs touching the half-open window boundaries are retained in truncated form and are not wrapped. The gap CCDF denominator is all zero runs across 368 nodes for one portfolio. Simultaneous gateway opportunity counts epochs with at least two eligible selected regions and does not imply independent satellites, equipment or capacity.

S2. Spatial service-priority scenarios and demand-side audit

The communication experiments read the analysis-ready file `data_intermediate/demand_nodes_priority_maps.csv` (SHA-256 882ef08556ed9ec76b4159a7e398189b3b552a08ec8fb7ce695c09346fe7c819). It contains 368 coordinate-referenced nodes and four non-negative maps, each normalized to sum to one. Density combines percentile-ranked settlement and economic concentration inputs. Inverse density reverses that ranking. Structural need combines remoteness-related features. Calibrated priority is an allocation-derived planning score and is not measured traffic.

The released Source Data archive contains this table, its analysis-ready city source table, the cleaned 2,850-region candidate master and all nested candidate pools. `metadata/source_metadata.csv` records the public source identity, retained version or package label, access route, study role and known source terms; `metadata/variable_dictionary.csv` defines the released columns; and `metadata/cleaning_lineage.csv` links each input, script and output stage. These are the exact cleaned inputs used by the communication pipeline.

The held-out 2024 allocation outcome contained 193 positives among 368 nodes (52.4%). This prevalence is the average-precision no-skill baseline. The paired provincial bootstrap used the same sampled provinces for both comparators and reran imputation, scaling, logistic hyperparameter selection, fitting and province-grouped out-of-fold Platt calibration. Deterministic upstream geospatial feature construction was held fixed.

Supplementary Table S1 | Demand-side model and refitting audit.

Metric	Value	Unit
2024 sample size	368	nodes
2024 positives	193	nodes
2024 prevalence / AP baseline	0.524457	fraction
Fixed logistic AP	0.903385	AP
Bootstrap requested / successful / failed	1000 / 1000 / 0	replicates
Paired AP difference mean	0.084618	AP
Monte Carlo SE of mean difference	0.001853	AP
Paired AP difference 2.5%-97.5%	[-0.046214, 0.191620]	AP
Residual Moran's I	0.370171	index

Metric	Value	Unit
Moran permutation P	0.0002000	5000 permutations

Supplementary Table S2 | Macroregional prediction audit.

Region	Score	N	Positive	Prevalence	AP	AP lift	Brier
Central	Calibrated priority	49	24	0.490	0.904	1.846	0.184
Central	Inverse density	49	24	0.490	0.779	1.591	0.230
East	Calibrated priority	77	21	0.273	0.415	1.523	0.204
East	Inverse density	77	21	0.273	0.276	1.013	0.223
North	Calibrated priority	36	19	0.528	0.858	1.626	0.238
North	Inverse density	36	19	0.528	0.843	1.597	0.190
Northeast	Calibrated priority	36	13	0.361	0.736	2.037	0.185
Northeast	Inverse density	36	13	0.361	0.810	2.243	0.207
Northwest	Calibrated priority	62	42	0.677	0.955	1.410	0.172
Northwest	Inverse density	62	42	0.677	0.905	1.336	0.199
South	Calibrated priority	54	21	0.389	0.836	2.150	0.171
South	Inverse density	54	21	0.389	0.746	1.919	0.180
Southwest	Calibrated priority	54	53	0.981	1.000	1.019	0.106
Southwest	Inverse density	54	53	0.981	1.000	1.019	0.277

The Southwest contained 53 positives among 54 nodes, so its AP is interpreted together with prevalence and AP lift rather than as a stand-alone performance result.

Supplementary Table S3 | Sixteen-city cellular-data audit.

Audit measure	N	Spearman rho	95% city-list interval	Population-controlled rho	95% city-list interval
5g logical records per 10k	16	0.068	[-0.480, 0.610]	0.005	[-0.617, 0.615]
5g logical records per 1000 km2	16	-0.115	[-0.570, 0.411]	-0.238	[-0.719, 0.370]
approx sites 10m per 10k	16	0.106	[-0.464, 0.599]	-0.021	[-0.585, 0.562]
approx sites 10m per 1000 km2	16	-0.185	[-0.648, 0.370]	-0.398	[-0.827, 0.178]
approx sites 100m per 10k	16	0.115	[-0.456, 0.638]	0.003	[-0.574, 0.609]
approx sites 100m per 1000 km2	16	-0.221	[-0.681, 0.348]	-0.422	[-0.803, 0.142]

Supplementary Table S3a | Raw counts and denominators for the 16-city audit.

City label	City ID	Population proxy (10,000 people)	Area (km^2)	5G observation records	Approx. sites, 10 m	Approx. sites, 100 m
北京	110000	2189.0	16406.4	210	60865	41611
上海	310000	2488.0	7985.7	1517	52684	35274
南京	320100	932.0	6587.9	692	13645	10347
苏州	320500	1275.0	8629.1	729	17837	14096
无锡	320200	746.0	4683.0	139	7305	6267
杭州	330100	1197.0	16876.7	2145	19394	13354
宁波	330200	942.0	10358.2	0	8060	7382
合肥	340100	937.0	11474.0	430	4248	3694
济南	370100	924.0	10222.0	32	5109	4600
青岛	370200	1011.0	11238.9	0	9803	8454
郑州	410100	1262.0	7515.9	17	6691	5536
武汉	420100	1233.0	8574.7	503	11817	9525

City label	City ID	Population proxy (10,000 people)	Area (km ²)	5G observation records	Approx. sites, 10 m	Approx. sites, 100 m
广州	440100	1874.0	7399.0	8477	43293	23987
深圳	440300	1763.0	2135.9	1917	31641	19411
东莞	441900	1048.0	2460.3	177	14891	12319
佛山	440600	952.0	3798.2	693	20329	13362

Supplementary Table S3b | Population- and area-normalized cellular audit values.

City label	Records per 10,000 people	Records per 1,000 km ²	10-m sites per 10,000 people	10-m sites per 1,000 km ²	100-m sites per 10,000 people	100-m sites per 1,000 km ²
北京	0.096	12.800	27.805	3709.823	19.009	2536.260
上海	0.610	189.966	21.175	6597.329	14.178	4417.170
南京	0.742	105.042	14.641	2071.234	11.102	1570.616
苏州	0.572	84.481	13.990	2067.068	11.056	1633.536
无锡	0.186	29.682	9.792	1559.895	8.401	1338.243
杭州	1.792	127.098	16.202	1149.157	11.156	791.268
宁波	0.000	0.000	8.556	778.129	7.837	712.673
合肥	0.459	37.476	4.534	370.228	3.942	321.945
济南	0.035	3.131	5.529	499.806	4.978	450.011
青岛	0.000	0.000	9.696	872.241	8.362	752.212
郑州	0.013	2.262	5.302	890.246	4.387	736.572
武汉	0.408	58.661	9.584	1378.117	7.725	1110.821
广州	4.523	1145.699	23.102	5851.213	12.800	3241.934
深圳	1.087	897.512	17.947	14813.865	11.010	9087.953
东莞	0.169	71.941	14.209	6052.405	11.755	5007.023
佛山	0.728	182.454	21.354	5352.250	14.036	3517.968

The 16 cities are the complete set in the available official key-city table, not a random national sample. These diagnostics were excluded from the communication results because they do not provide measured NTN traffic.

S3. Demand-independent candidate regions

The 2,850-region universe is stored in data_intermediate/candidate_regions_spatial_rank.csv (SHA-256 1974c03dc3aa7c2ce106d3edcce72e08f7473a1a3804fb8ad8b2e434c5e83b18). One seed nearest each provincial spherical centroid was retained before deterministic farthest-first ordering. This produced nested pools of 200, 400 and 800 regions without using priority weights, allocation outcomes or ephemerides. No administrative quota or 150-km spacing constraint was used in the main model. Road and railway features are interpreted only as construction-accessibility proxies.

S4. Reference architectures and public-ephemeris audit

The prospective Walker architecture uses a complete Walker-delta construction with study-selected inclination, satellite count, plane count and phase. It does not represent a named domestic deployment. The two public OMM groups are frozen orbital reference inventories. Active status was unavailable, so neither group is interpreted as an operational fleet. The source-group names are retained below solely for reproducibility.

Supplementary Table S4 | Public OMM snapshot and propagation audit.

Reference architecture	OMM group	Records	Unique IDs	Duplicate IDs	Epoch min / median / max UTC	Age min / median / max (h)	SGP4 non-zero	Download UTC	Snapshot SHA-256
dense public-ephemeris orbital reference set	starlink	10709	10709	0	2026-06-29T20:20:45.127392+00:00 / 2026-07-02T15:23:52.232928+00:00 / 2026-07-03T06:11:45.231360+00:00	5.804 / 20.602 / 87.654	3 records; 3019 states	2026-07-03T08:43:40+00:00	1ec38cc2c9b02fb3b3eb87e1a202addaa14435e564f446ef3316f93c52afb1b6

Reference architecture	OMM group	Records	Unique IDs	Duplicate IDs	Epoch min / median / max UTC	Age min / median / max (h)	SGP4 non-zero	Download UTC	Snapshot SHA-256
sparse public-ephemeris orbital reference set	oneweb	651	651	0	2026-07-02T00:01:53.905440+00:00 / 2026-07-02T15:43:20.038656+00:00 / 2026-07-02T20:43:09.352416+00:00	15.281 / 20.278 / 35.968	0 records; 0 states	2026-07-03T08:43:44+00:00	7f3b103bfd8bc3b676da5b687c6e29315c5ee4e9f4ef83e77abcb62fc4c6ec3c

The data owner is CelesTrak. The dense snapshot was acquired from <https://celestrak.org/NORAD/elements/gp.php?GROUP=starlink&FORMAT=csv>, and the sparse snapshot from <https://celestrak.org/NORAD/elements/gp.php?GROUP=oneweb&FORMAT=csv>, at the UTC times shown above. These live query endpoints identify the source but do not guarantee recovery of the historical snapshots. The exact CSV snapshots used here are therefore included in Source_Data.zip, together with their acquisition times and SHA-256 values, with CelesTrak attribution and usage-policy information retained in the package metadata.

Dense-set SGP4 code-1 states were quarantined before contact and link calculations. Sparse-set propagation had no non-zero status. Regional-culling checks found no omitted visible terminal-satellite contacts at the tested epochs.

S5. Parameter provenance and propagation semantics

Supplementary Table S5 gives the concise registry; the accompanying machine-readable file results/parameter_provenance_content_ready.csv preserves exact source, table/row, terminal role, conversion, classification, tested range and notes for every parameter. The 15.9-dB/K value originates from a 60-cm VSAT row and is transferred as a study-defined generic gateway receiver. It is not presented as an operational feeder-terminal specification. The feeder EIRP conversion is $4 \text{ dBW/MHz} + 10 \log_{10}(400 \text{ MHz}) = 30.0206 \text{ dBW}$.

Supplementary Table S5 | Directional reference-link, orbit and simulation parameter provenance.

Parameter	Value and unit	Terminal role / component	Exact source and table/row	Conversion	Sourced or study-defined	Tested range
time step	30 s	orbit and return-path tensor	study design; simulation grid	none	study-defined	10;30;60;600
time window	24 h	orbit and return-path tensor	study design; half-open UTC interval	none	study-defined	24;72
minimum elevation	10 °	user and feeder links	3GPP TR 38.821 V16.2.0; Table 4.2-2, row 'Min Elevation angle for both sat-gateway and user equipment'	none	sourced	5;10;15;20;30
user carrier	2 GHz	user uplink transmitter to satellite receiver	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC9 uplink, Frequency	none	sourced	fixed
user bandwidth	0.36 MHz	user uplink	3GPP TR 38.821 V16.2.0; Table 6.1.3.2-1, Channel bandwidth, handheld S-band uplink	360 kHz / 1000	sourced	fixed
user EIRP	-7 dBW	handheld user transmitter	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC9 uplink, TX EIRP	23 dBm - 30	derived from sourced value	fixed
user receiver G/T	1.1 dB/K	satellite user-link receiver	3GPP TR 38.821 V16.2.0; Table 6.1.1.1-1, LEO-600 S-band uplink G/T; Table 6.1.3.3-1, SC9 uplink	none	sourced	fixed
user gas loss	0.1 dB	user uplink	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC9 uplink, Atmospheric loss	none	sourced	0;0.1;0.3
user shadow margin	3 dB	user uplink	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC9 uplink, Shadow fading margin	none	sourced	0;3
user scintillation loss	2.2 dB	user uplink	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC9 uplink, Scintillation loss	none	sourced	0;2.2
feeder carrier	20 GHz	satellite transmitter to gateway-region reference receiver	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC6 downlink, Frequency	none	sourced	fixed
feeder bandwidth	400 MHz	feeder downlink	3GPP TR 38.821 V16.2.0; Table 6.1.3.2-1, System bandwidth; Table 6.1.3.3-1, SC6 downlink	none	sourced	fixed
feeder EIRP	30.0206 dBW	satellite feeder-link transmitter	3GPP TR 38.821 V16.2.0; Table 6.1.1.1-1, LEO-600 Ka-band downlink EIRP density	$4 \text{ dBW/MHz} + 10 \log_{10}(400 \text{ MHz}) = 30.0206 \text{ dBW}$	derived from sourced value	fixed

Parameter	Value and unit	Terminal role / component	Exact source and table/row	Conversion	Sourced or study-defined	Tested range
feeder receiver G/T	15.9 dB/K	gateway-region reference receiver	3GPP TR 38.821 V16.2.0; Table 6.1.1.1-3, 60-cm VSAT G/T; Table 6.1.3.3-1, SC6 downlink	no conversion; terminal role transferred	study-defined transfer of sourced VSAT value	fixed
feeder gas loss	0.5 dB	feeder downlink	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC6 downlink, Atmospheric loss	none	sourced	0.1;0.5;1.0
feeder shadow margin	0 dB	feeder downlink	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC6 downlink, Shadow fading margin	none	sourced	0
feeder scintillation loss	0.3 dB	feeder downlink	3GPP TR 38.821 V16.2.0; Table 6.1.3.3-1, SC6 downlink, Scintillation loss	none	sourced	0;0.3;1.0
pointing loss	0 dB	both links	study design; generic reference-link assumption	none	study-defined	0;0.5;1.0
polarisation loss	0 dB	both links	study design; generic reference-link assumption	none	study-defined	0;0.5;1.0
implementation loss	0 dB	both links	study design; generic reference-link assumption	none	study-defined	0;1;2
generic C/N closure threshold	0 dB	both links	study design; reference closure threshold	none	study-defined	-3;0;3;6
rain attenuation	1 % annual exceedance	feeder downlink	ITU-R P.618-14 with P.837-8 rainfall data; fixed propagation design state	site and elevation calculation	sourced model with study-defined exceedance state	0.1;1;5
cloud attenuation	1 % annual exceedance	feeder downlink	ITU-R P.840-9; fixed propagation design state	site and elevation calculation	sourced model with study-defined exceedance state	0.1;1;5
Walker altitude	600 km	prospective Walker architecture	3GPP TR 38.821 V16.2.0; Table 4.2-2, LEO altitude	none	sourced range; study-selected main value	600;1200
Walker inclination	55 °	orbit	study design; prospective Walker architecture	none	study-defined	45;55;70
Walker satellites	648 satellites	orbit	study design; prospective Walker architecture	none	study-defined	324;648;1296
Walker planes	36 planes	orbit	study design; prospective Walker architecture	none	study-defined	18;36;72
Walker phasing	1 integer	orbit	Walker, J. Br. Interplanet. Soc. 37, 559-571 (1984); Walker-delta construction	none	study-defined application of sourced construction	0;1;17

The 0-dB value is a generic C/N closure threshold. No MCS, BLER or traffic class is inferred. M3 uses a fixed p%-annual-exceedance attenuation design state from ITU-R propagation models. It is not a time-correlated weather series, and the reported gap durations remain orbital-sequence diagnostics under that fixed state.

Supplementary Table S6 | ITU-R recommendation version audit.

Recommendation	Version	Approval	Status	Use in study	Checked UTC
P.618	P.618-14	08/2023	in force	implemented for long-term rain attenuation	2026-07-15T00:00:00Z
P.676	P.676-13	08/2022	in force	version checked; main tensor uses the fixed 3GPP gas-loss term, not a spatial P.676 calculation	2026-07-15T00:00:00Z
P.837	P.837-8	09/2025	in force	R0.01 digital map used by the P.618 calculation	2026-07-15T00:00:00Z
P.840	P.840-9	08/2023	in force	annual cloud-liquid maps and attenuation implemented	2026-07-15T00:00:00Z
P.1511	P.1511-3	08/2024	in force	topographic-height map implemented for gateway-region anchors	2026-07-15T00:00:00Z

S6. E0 physical and numerical validation

Supplementary Table S7 | Link-budget and propagation-input validation.

Check	Observed	Expected	Absolute error	Tolerance	Unit	Pass
fspl_user_reference	159.1	159.1	0	0.02	dB	True
fspl_feeder_reference	179.1	179.1	0	0.02	dB	True

Check	Observed	Expected	Absolute error	Tolerance	Unit	Pass
cn_user_reference	2.73614	2.8	0.0638578	0.15	dB	True
cn_feeder_reference	8.59917	8.5	0.0991672	0.15	dB	True
candidate_topography_finite	400	400	0	0	regions	True
candidate_r001_finite	400	400	0	0	regions	True
candidate_cloud_finite	400	400	0	0	regions	True

The independent orbit-geometry suite passed 21 of 21 checks. Maximum TEME-to-ECEF disagreement against the audit implementation was 0.0183 km, and the regional cull omitted 0 visible contacts.

S7. Communication-model ladder

Supplementary Table S8 | One-gateway M0-M3 model ladder.

Architecture	Model	Mean across maps (%)	Minimum map (%)	Maximum regret (%)	Worst-node gap (min)
Prospective Walker	M0	12.406	10.190	21.2298	1440.0
Prospective Walker	M1	100.000	100.000	<10 ⁻⁸	0.0
Prospective Walker	M2	99.309	99.157	0.0737	4.5
Prospective Walker	M3	99.309	99.157	0.0737	4.5
Dense public-ephemeris reference	M0	12.406	10.190	21.2298	1440.0
Dense public-ephemeris reference	M1	100.000	100.000	<10 ⁻⁸	0.0
Dense public-ephemeris reference	M2	99.995	99.993	<10 ⁻⁸	1.0
Dense public-ephemeris reference	M3	99.995	99.993	<10 ⁻⁸	1.0
Sparse public-ephemeris reference	M0	12.406	10.190	21.2298	1440.0
Sparse public-ephemeris reference	M1	100.000	100.000	<10 ⁻⁸	0.0
Sparse public-ephemeris reference	M2	96.721	96.362	0.0901	37.5
Sparse public-ephemeris reference	M3	87.251	85.450	1.0945	689.0

M0 is the 400-km static-distance negative control. M1 uses same-epoch geometry. M2 requires user-uplink and feeder-downlink closure through the same satellite. M3 adds the fixed propagation design state to the feeder hop.

S8. Optimization baselines and certificates

Equal mixture maximizes the arithmetic mean of the four map utilities. The code label `maximum_average_availability` instead maximizes the equal-node mean and is reported as the unweighted-node optimum. Normalized max-min is algebraically identical to minimax regret and is omitted. Unnormalized max-min remains distinct because it maximizes the minimum raw map utility.

Supplementary Table S9 | One-gateway baseline comparison on M3.

Architecture	Method	Maximum regret (%)	Mean regret (%)	Optimization evidence
Prospective Walker	equal mixture	0.0737	0.0223	complete enumeration of all gateways
Prospective Walker	greedy marginal availability	0.0737	0.0223	deterministic greedy baseline; no optimality claim
Prospective Walker	Unweighted-node optimum	0.0737	0.0223	complete enumeration of all gateways
Prospective Walker	minimax regret	0.0737	0.0223	complete enumeration of all gateways
Prospective Walker	Unnormalized max-min	0.0737	0.0223	complete enumeration of all gateways
Prospective Walker	feeder link ranking	65.3474	54.9824	deterministic ranking by M3 feeder time availability and available-satellite tie break

Architecture	Method	Maximum regret (%)	Mean regret (%)	Optimization evidence
Prospective Walker	p median	8.9955	7.5361	ground-distance control; Optimization terminated successfully. (HiGHS Status 7: Optimal)
Prospective Walker	k center	0.3963	0.2965	ground-distance control; exact discrete radius; previous-radius status=2
Dense public-ephemeris reference	minimax regret	<10 ⁻⁸	<10 ⁻⁸	complete enumeration of all gateways
Dense public-ephemeris reference	equal mixture	<10 ⁻⁸	<10 ⁻⁸	complete enumeration of all gateways
Dense public-ephemeris reference	Unweighted-node optimum	<10 ⁻⁸	<10 ⁻⁸	complete enumeration of all gateways
Dense public-ephemeris reference	Unnormalized max-min	<10 ⁻⁸	<10 ⁻⁸	complete enumeration of all gateways
Dense public-ephemeris reference	greedy marginal availability	<10 ⁻⁸	<10 ⁻⁸	deterministic greedy baseline; no optimality claim
Dense public-ephemeris reference	feeder link ranking	15.3632	11.2072	deterministic ranking by M3 feeder time availability and available-satellite tie break
Dense public-ephemeris reference	p median	5.7826	4.7412	ground-distance control; Optimization terminated successfully. (HiGHS Status 7: Optimal)
Dense public-ephemeris reference	k center	<10 ⁻⁸	<10 ⁻⁸	ground-distance control; exact discrete radius; previous-radius status=2
Sparse public-ephemeris reference	equal mixture	2.4544	0.6615	complete enumeration of all gateways
Sparse public-ephemeris reference	greedy marginal availability	2.4544	0.6615	deterministic greedy baseline; no optimality claim
Sparse public-ephemeris reference	Unweighted-node optimum	1.0945	0.8965	complete enumeration of all gateways
Sparse public-ephemeris reference	minimax regret	1.0945	0.8965	complete enumeration of all gateways
Sparse public-ephemeris reference	Unnormalized max-min	2.4544	0.6615	complete enumeration of all gateways
Sparse public-ephemeris reference	feeder link ranking	18.8535	11.5021	deterministic ranking by M3 feeder time availability and available-satellite tie break
Sparse public-ephemeris reference	p median	15.4036	11.2435	ground-distance control; Optimization terminated successfully. (HiGHS Status 7: Optimal)
Sparse public-ephemeris reference	k center	3.2315	2.5833	ground-distance control; exact discrete radius; previous-radius status=2

Supplementary Table S10 | Random feasible one-gateway portfolios.

Architecture	Replicates	Mean maximum regret (%)	2.5%-97.5% (%)
Prospective Walker	1000	27.21	1.34-82.60
Dense public-ephemeris reference	1000	20.65	0.28-77.71
Sparse public-ephemeris reference	1000	40.22	3.21-89.84

S9. Complete optimal-set analysis for K=1 and K=2

Supplementary Table S11 | M3 minimax and common single-map optimum counts.

Architecture	K	Portfolios enumerated	Strict minimax optima	Epsilon minimax optima	Common strict single-map optima	Common epsilon single-map optima
Prospective Walker	1	400	1	1	0	0
Prospective Walker	2	79800	7	18	7	18
Dense public-ephemeris reference	1	400	1	1	1	1
Dense public-ephemeris reference	2	79800	3849	9645	3849	9645
Sparse public-ephemeris reference	1	400	1	1	0	0

Architecture	K	Portfolios enumerated	Strict minimax optima	Epsilon minimax optima	Common strict single-map optima	Common epsilon single-map optima
Sparse public-ephemeris reference	2	79800	1	2	0	0

Strict equality used an absolute objective tolerance of 10^{-12} . The minimax epsilon set used an absolute 10^{-4} objective allowance; maximization objectives used a relative 10^{-4} allowance. Displayed coordinates use lexicographic candidate-index tie-breaking, while all stability conclusions use complete sets.

S10. Decision convergence and temporal sensitivity

Supplementary Table S12 | Independently rebuilt and optimized Walker time tests.

Experiment	Start UTC	Window (h)	Step (s)	K	Selected indices	Mean (%)	Minimum (%)	Max regret (%)	Gap (min)	Strict optima	Epsilon optima
main_30s	2026-07-03T12:00:00Z	24	30	1	115	99.309	99.157	0.0737	4.5	1	1
main_30s	2026-07-03T12:00:00Z	24	30	2	122;165	99.979	99.970	$<10^{-8}$	0.5	7	18
decision_10s	2026-07-03T12:00:00Z	24	10	1	115	99.304	99.152	0.0814	3.5	1	1
decision_10s	2026-07-03T12:00:00Z	24	10	2	122;391	99.979	99.970	$<10^{-8}$	0.333	7	18
decision_60s	2026-07-03T12:00:00Z	24	60	1	115	99.308	99.157	0.0769	9	1	1
decision_60s	2026-07-03T12:00:00Z	24	60	2	122;165	99.980	99.972	$<10^{-8}$	1	9	19
epoch_plus_24h	2026-07-04T12:00:00Z	24	30	1	115	99.296	99.144	0.0900	4.5	1	1
epoch_plus_24h	2026-07-04T12:00:00Z	24	30	2	122;352	99.979	99.970	$<10^{-8}$	0.5	7	18
epoch_plus_48h	2026-07-05T12:00:00Z	24	30	1	115	99.288	99.135	0.1012	4.5	1	1
epoch_plus_48h	2026-07-05T12:00:00Z	24	30	2	122;352	99.979	99.970	$<10^{-8}$	0.5	7	18
window_72h	2026-07-03T12:00:00Z	72	30	1	115	99.298	99.146	0.0883	4.5	1	1
window_72h	2026-07-03T12:00:00Z	72	30	2	122;391	99.979	99.970	$<10^{-8}$	0.5	6	18

Supplementary Table S13 | Portfolio replay on the common 10-s reference grid.

Native decision experiment	K	Mean (%)	Minimum (%)	Max regret (%)	Tie-broken Jaccard	Optimal-set Jaccard	Pairwise Jaccard range
main_30s	1	99.304	99.152	0.0814	1.000	1.000	1.000-1.000
main_30s	2	99.979	99.970	$<10^{-8}$	0.333	1.000	0.333-1.000
decision_10s	1	99.304	99.152	0.0814	1.000	1.000	1.000-1.000
decision_10s	2	99.979	99.970	$<10^{-8}$	1.000	1.000	0.333-1.000
decision_60s	1	99.304	99.152	0.0814	1.000	1.000	1.000-1.000
decision_60s	2	99.979	99.970	$<10^{-8}$	0.333	0.778	0.333-1.000
epoch_plus_24h	1	99.304	99.152	0.0814	1.000	1.000	1.000-1.000
epoch_plus_24h	2	99.979	99.970	$<10^{-8}$	0.333	1.000	0.333-1.000
epoch_plus_48h	1	99.304	99.152	0.0814	1.000	1.000	1.000-1.000
epoch_plus_48h	2	99.979	99.970	$<10^{-8}$	0.333	0.750	0.333-1.000
window_72h	1	99.304	99.152	0.0814	1.000	1.000	1.000-1.000
window_72h	2	99.979	99.970	$<10^{-8}$	1.000	0.857	0.333-1.000

S11. Budget, failure and communication sensitivity

Supplementary Table S14 | M3 budget response.

Architecture	K	Mean across maps (%)	Worst-node gap (min)	Simultaneous ≥ 2 -gateway opportunity (%)
Prospective Walker	1	99.309	4.5	0.000
Prospective Walker	2	99.979	0.5	78.973
Prospective Walker	5	100.000	0	99.943
Prospective Walker	10	100.000	0	100.000
Prospective Walker	21	100.000	0	100.000
Dense public-ephemeris reference	1	99.995	1	0.000
Dense public-ephemeris reference	2	100.000	0	97.496
Dense public-ephemeris reference	3	100.000	0	99.999
Dense public-ephemeris reference	5	100.000	0	100.000
Dense public-ephemeris reference	10	100.000	0	100.000
Dense public-ephemeris reference	21	100.000	0	100.000
Sparse public-ephemeris reference	1	87.251	689	0.000
Sparse public-ephemeris reference	2	96.502	47	46.577
Sparse public-ephemeris reference	10	98.692	7.5	98.664
Sparse public-ephemeris reference	21	98.692	7.5	98.692

The sparse reference has completed budgets K=1, 2, 10 and 21; no K=5 result is reported. Failure comparisons in the main figure therefore use the common K=2 and K=10 budgets. All N-1 removals and all N-2 pairs were enumerated.

Supplementary Table S15 | Post-selection failure stress with epoch-wise reassociation.

Architecture	K	Removal	Worst absolute opportunity (%)	Worst retention (%)
Dense public-ephemeris reference	2	N-1	96.351	96.351
Dense public-ephemeris reference	2	N-2	0.000	0.000
Dense public-ephemeris reference	10	N-1	100.000	100.000
Dense public-ephemeris reference	10	N-2	99.998	99.998
Prospective Walker	2	N-1	77.946	77.970
Prospective Walker	2	N-2	0.000	0.000
Prospective Walker	10	N-1	100.000	100.000
Prospective Walker	10	N-2	99.986	99.986
Sparse public-ephemeris reference	2	N-1	53.976	56.004
Sparse public-ephemeris reference	2	N-2	0.000	0.000
Sparse public-ephemeris reference	10	N-1	98.563	99.976
Sparse public-ephemeris reference	10	N-2	96.361	97.671

Supplementary Table S16 | Fixed-portfolio communication ablations and sensitivities at K=2.

Architecture	Setting	Mean across maps (%)	Worst-node gap (min)
Dense public-ephemeris reference	minimum_elevation_1 5deg	97.909	25.5
Dense public-ephemeris reference	minimum_elevation_2 0deg	90.073	1.44e+03

Architecture	Setting	Mean across maps (%)	Worst-node gap (min)
Dense public-ephemeris reference	minimum_elevation_3 0deg	67.709	1.44e+03
Dense public-ephemeris reference	minimum_elevation_5 deg	100.000	0
Dense public-ephemeris reference	no_feeder_link	100.000	0
Prospective Walker	minimum_elevation_1 5deg	98.892	3.5
Prospective Walker	minimum_elevation_2 0deg	94.070	1.44e+03
Prospective Walker	minimum_elevation_3 0deg	58.698	1.44e+03
Prospective Walker	minimum_elevation_5 deg	100.000	0
Prospective Walker	no_feeder_link	100.000	0
Sparse public-ephemeris reference	minimum_elevation_1 5deg	96.488	47.5
Sparse public-ephemeris reference	minimum_elevation_2 0deg	95.635	117
Sparse public-ephemeris reference	minimum_elevation_3 0deg	87.901	1.44e+03
Sparse public-ephemeris reference	minimum_elevation_5 deg	96.511	47
Sparse public-ephemeris reference	no_feeder_link	98.692	7.5

The shuffled-time diagnostic preserves marginal opportunity states but changes temporal order. It is used only to show that contiguous-gap metrics depend on ordering. The 600-s case is a coarse fixed-portfolio diagnostic, whereas the 10/30/60-s results in S10 rebuild and reoptimize the decision problem.

S12. Candidate-pool and Walker-architecture sensitivity

Supplementary Table S17 | Candidate-pool sensitivity.

Architecture	Pool size	K	Selected candidate IDs	Maximum regret (%)	Solver evidence
Dense public-ephemeris reference	200	1	621224	0.0011	status 0; gap 0
Dense public-ephemeris reference	200	2	659011;610728	<10 ⁻⁸	status 0; gap 0
Dense public-ephemeris reference	400	1	620502	<10 ⁻⁸	status 0; gap 0
Dense public-ephemeris reference	400	2	620502;513227	<10 ⁻⁸	status 0; gap 0
Dense public-ephemeris reference	800	1	620502	0.0003	status 0; gap 0
Dense public-ephemeris reference	800	2	620502;513223	<10 ⁻⁸	status 0; gap 0
Sparse public-ephemeris reference	200	1	623024	0.4369	status 0; gap 0
Sparse public-ephemeris reference	200	2	371082;513330	0.1308	status 0; gap 0
Sparse public-ephemeris reference	400	1	513233	1.0945	status 0; gap 0
Sparse public-ephemeris reference	400	2	371082;513330	0.1308	status 0; gap 0
Sparse public-ephemeris reference	800	1	513224	0.8842	status 0; gap 0
Sparse public-ephemeris reference	800	2	632624;370687	0.0946	status 0; gap 0
Sparse public-ephemeris reference	2850	1	513222	0.1938	status 0; gap 0
Sparse public-ephemeris reference	2850	2	632624;370212	0.3067	status 0; gap 0

The dense reference was completed for 200, 400 and 800 regions. The sparse reference additionally completed the full 2,850-region universe. Candidate-pool expansion changes both the admissible portfolio set and the scenario-specific ideal values, so regret is recomputed within each pool.

Supplementary Table S18 | Prospective Walker one-factor sensitivity.

Scenario	Changed parameter	Architecture setting	K	Mean across maps (%)	Worst-node gap (min)	Maximum regret (%)	Selected-set Jaccard to main
china_ref	main	600 km; 55 deg; 648 satellites; 36 planes	1	99.309	4.5	0.0737	1.000
china_ref	main	600 km; 55 deg; 648 satellites; 36 planes	2	99.979	0.5	<10 ⁻⁸	1.000
china_ref_alt 1200	altitude	1200 km; 55 deg; 648 satellites; 36 planes	1	94.097	1.44e+03	0.1469	0.000
china_ref_alt 1200	altitude	1200 km; 55 deg; 648 satellites; 36 planes	2	97.975	5.5	0.2323	0.000
china_ref_inc 45	inclination	600 km; 45 deg; 648 satellites; 36 planes	1	99.756	15	0.0149	0.000
china_ref_inc 45	inclination	600 km; 45 deg; 648 satellites; 36 planes	2	100.000	0	<10 ⁻⁸	0.333
china_ref_inc 70	inclination	600 km; 70 deg; 648 satellites; 36 planes	1	98.068	26.5	0.1168	0.000
china_ref_inc 70	inclination	600 km; 70 deg; 648 satellites; 36 planes	2	99.852	1.5	<10 ⁻⁸	0.000
china_ref_n3 24	satellite_count	600 km; 55 deg; 324 satellites; 18 planes	1	94.633	15	0.3360	0.000
china_ref_n3 24	satellite_count	600 km; 55 deg; 324 satellites; 18 planes	2	98.897	5	0.2313	0.000
china_ref_n1 296	satellite_count	600 km; 55 deg; 1296 satellites; 72 planes	1	99.854	3	<10 ⁻⁸	0.000
china_ref_n1 296	satellite_count	600 km; 55 deg; 1296 satellites; 72 planes	2	100.000	0	<10 ⁻⁸	0.000

S13. Reproducibility and Source Data index

The master numeric registry is supplied as results_summary.csv in Source_Data.zip (the working-project copy is results_summary_content_ready.csv). Each row contains a result identifier, architecture, model, scenario, budget, method, metric, unit, originating file, source SHA-256, generator, seed and solver status. Main figures read only the corresponding files in source_data/.

Supplementary Table S19 | Main-figure Source Data files.

Figure	Source Data file	Rows	SHA-256
Figure_1_Source_Data.csv		11	ecd56dea988485def0d94d28f277a742b96e4134c3e563efd076314e5741565d
Figure_2_Source_Data.csv		1516	89b3d09a64d87fbf89536172425f2a976c5c3643792732aae920beb4c0a4b0c9
Figure_3_Source_Data.csv		161519	e811fa2bb11cf30b08ff7200337aa91bde5a647e2c8c0f74937d1d50657c557d
Figure_4_Source_Data.csv		84	b0eb689ece212859ac0edd204866edc6affaadff84789709efe8b0cc36b33c53
Figure_5_Source_Data.csv		182	c41be481fe0713933a29fec4ef38de8a6750e03b23add74528b43de27c936d5b
Figure_6_Source_Data.csv		147	004bce4bb4d18983a0623db558383faa519a0874669b5699a5418a85ed2945a1

Computations ran on Windows 11 (64-bit) using an Intel Core Ultra 9 285H processor with 16 cores and 31.5 GB RAM. The software environment and executable sequence are listed in reproducibility_README.md.