F A R O S: Multistrom Clusters of Satellites

Abstract

The design and visualization of a satellite communication network for space objects identification will embrace cargo transfers and human missions in space. This study investigated the creation of a multistrom cluster of satellites keeping as core idea future potential dynamic rearrangement and the establishment of bidirectional communication between satellites and the earth. Model simulations of satellite networks in several orbits were developed in Python. A 3D Earth interactive representation program to increase and decrease number of orbits, change altitude, and configure number of satellites was created. It incorporated Dijkstra's algorithm to find the shortest path for communication between satellites. Based on the simulations, it was found that the backbone is effective in space because it reduces latency when the time taking to move from one hop to the other is more than 6.8 ms. What is more, a backbone in space is not subject to the seabed disturbance that has severed multiple submarine cables simultaneously, with repairs taking days[3]. These characteristics suggest that satellite systems would be less exposed to natural catastrophes, countries will be interconnected because of global access to the network of space, and space will be endorsed by a decentralized model of space governance.

Keywords

Multistrom, Inter Satellite Link, Data Handling, Data Migration, Shortest-Path Routing, Backbone in Space.

Introduction

Traffic in space constitutes a serious and growing problem. Space surveillance networks currently track approximately 40,000 objects in orbit, of which about 11,000 are active payloads, while debris larger than 1 cm is estimated to exceed 1.2 million[1]. At around 550 km, the altitude band preferred by communication constellations, the density of threatening debris is now of the same order as that of active satellites[1]. Less than a century has passed since the first successful satellite launch on October 4, 1957. Yet in only six years, between 2017 and 2022, more than 300 constellations representing over one million satellites were filed with the International Telecommunication Union — over 115 times the number then operational[2]. Even allowing for the fact that many of these filings are speculative and will never fly, near space is set to become considerably more congested in the coming decades. Space presents fewer natural obstacles in communication compared to Earth; cables on the seabed are exposed to the physical environment in a way orbital links are not, and their routes are constrained by geography and landing rights[3]. This research utilizes distance as an important factor in calculations for aiming communication speed. In communications there are multiple information exchanges between data processing systems. For example, for the dispersion and assembly of data, response time plays an important role. If there is intense movement in a region of space during data assembly, this would entail automated rearrangement of communication satellites (access points) and data satellites (data points). An access point is a communication satellite that relays traffic; a data point is a satellite that stores it. This cannot be applied on Earth because data centers cannot reduce the physical

distance between them. That is, backbone development in space consists of moving data centers.
The term multistrom is important to note at this point. Stroma is used to denote something that embraces something else. Stroma is a distinct context at a distinct altitude. The term multi-shell does not adequately present what the research intends to consider. This research is not talking about new phenomena – but about new networks; so, it is not talking about a shell. Take as an example a slice of bread with butter and marmalade spread on top. The bread is a different stroma from the butter or the marmalade stroma. This research is talking about first stroma, second stroma, and third stroma. Stromata differ in both context and purpose. This is the reason the research refers to orbits as low and medium for example: the difference indicates the altitude and the context – as well as some features, for example GPS satellites are in medium Earth orbit. In space networks, this research will examine different stromata with different characteristics. One stroma is satellites at a relatively low altitude, and another stroma is backbones at a higher altitude. This research is not solely referring to a stroma based on its altitude – because then it would talk about blended networks where backbones, communication satellites, and data are mixed into a single orbit. In terms of the previous example, the marmalade, butter, and bread are all blended into a mixer turning into paste. Obviously, that is possible, but the character of every ingredient is lost. That is why this paper preserves the term multistrom.

The architecture of this research is premised on keeping all network elements in orbit, so requests need not descend to the surface and return. In a decentralized network, response time interacts with processing and storage capacity, which this study does not model. Even spacing of backbone modules is investigated using the Fibonacci sphere concept, which distributes modules uniformly so that updates or a change in module count do not disrupt the network. This distribution is simulated in Python using NumPy[4] and Matplotlib[5].

Related Work

Routing in low Earth orbit constellations has been studied since the current generation of systems was proposed. Prior work established that satellites linked to one another can deliver lower latency than terrestrial infrastructure over long distances[6, 7], addressed how those links should be arranged as the constellation moves[8, 9, 10], and examined how internet routing itself would have to change to accommodate such a network[11]. Configurations without inter-satellite links have also been studied[12]. This study takes the low orbital stroma as given and asks a narrower question: whether adding a higher layer above it reduces delay, and under what conditions. Its contribution is a threshold rather than a comparison, together with the finding that architectures returning to the ground are geometrically unavailable on intercontinental routes rather than merely slower. Simulation frameworks modelling orbital motion in greater detail than the model used here are available[13], and a comparison against one is identified as further work. The proposal to place storage and processing in orbit rather than returning data to the surface has been advanced elsewhere[14], and the Data Point element of this architecture belongs to that line of work.

Research Question

This study asks how realistic it is to develop near-space infrastructure connecting every hop - satellite - through a radio-based network largely independent of terrestrial infrastructure. Specifically, it examines end-to-end delay, redundancy, altitude, and the minimum number of orbital planes, backbones and satellites required, with the ultimate purpose of finding the shortest path for data migration using Dijkstra's algorithm from 300 km to 2,000 km.

Methodology

The study is a controlled simulation experiment. Three comparisons are conducted. The first isolates the effect of a medium Earth orbit backbone by holding endpoints, low Earth orbit stroma, routing algorithm and delay model constant and varying only the presence of the higher layer, with a control condition in which the additional satellites are placed in the lower stroma instead. The second varies the number of orbital planes while holding total satellite count constant. The third compares the meshed constellation against other in-space architectures — a bent-pipe relay, a ground-relay chain, and a single relay in medium Earth orbit. In every case the routing graph is rebuilt from scratch after any change in topology, and shortest-path computation is refused on a stale graph.

All data are generated by simulation; no human or animal subjects are involved. Ground stations are placed at Sydney (33.8688° S, 151.2093° E) and Tokyo (35.6762° N, 139.6503° E), separated by a central angle of 70.38° corresponding to a great-circle distance of 7,826 km. Satellite positions are generated from circular orbits at 550 km with 53° inclination unless otherwise stated, distributed across orbital planes in a Walker configuration with inter-plane phasing. In Walker notation i: T/P/F, the baseline constellation is 53°: 240/12/1. Where altitude is varied, the sweep runs from 300 km to 2,000 km.

Satellite position in an orbital plane of radius a, at in-plane angle θ, inclination i and right ascension of the ascending node Ω, is given by

$$x = a \cdot \cos\theta \cdot \cos\Omega - a \cdot \sin\theta \cdot \cos i \cdot \sin\Omega, \quad y = a \cdot \cos\theta \cdot \sin\Omega + a \cdot \sin\theta \cdot \cos i \cdot \cos\Omega, \quad z = a \cdot \sin\theta \cdot \sin i \quad \text{(Equation 1)}$$

with the right ascension of the ascending node for plane p of a constellation of N_p planes given by

$$\Omega_p = 2\pi p / N_p \quad \text{(Equation 2)}$$

For the time-stepped analysis, the in-plane angle advances at the mean motion of a circular orbit, obtained from the orbital period

$$T_{orb}{}^2 = (4\pi^2 / GM) \cdot a^3 \quad \text{(Equation 3)}$$

where G is the gravitational constant, 6.674×10^{-11} m^3 kg^{-1} s^{-2}, M is the mass of the Earth, and a is the orbital radius. At 550 km this gives a period of 95.5 minutes.

A link between two nodes is admitted only if it satisfies every one of four constraints. The straight segment between two satellites must not intersect the Earth; at 550 km this constraint begins to bind above approximately 5,400 km of separation. A ground-to-satellite link requires a minimum elevation angle of 10° above the local horizon[15]. Inter-satellite links are limited to 2,500 km. Finally, the link budget must satisfy a required carrier-to-noise ratio with margin. The link-budget constraint is evaluated during graph construction, so an edge exists only if it is both geometrically and radiometrically feasible.

Because the Earth is a sphere, an observer at altitude h can view only a finite portion of the surface before curvature blocks the line-of-sight[16]. The central angle φ between the observer and the tangent point, and the corresponding surface arc S, are

$$\varphi = \arccos(R / (R + h)), \quad S = R\varphi \quad \textit{(Equation 4)}$$

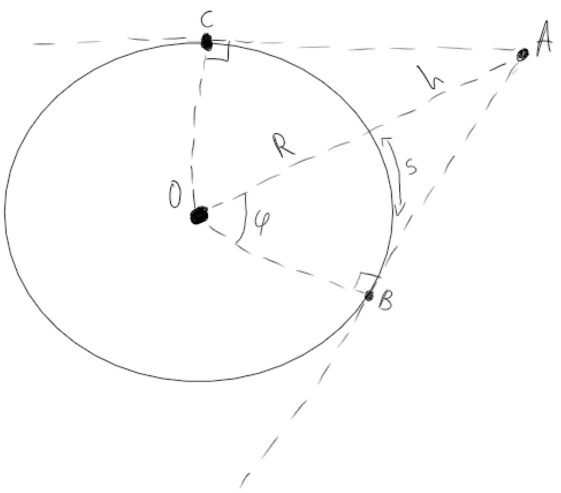

Figure 1 | Horizon geometry. S is the arc length along the Earth's surface to the horizon point, R is the radius of the Earth, and h is the observer's altitude above the surface. The quantity arccos(R/(R+h)) is the central angle φ between the observer and the tangent point, measured in radians.

Received power is computed from the Friis transmission equation

$$P_r = P_t \cdot G_t \cdot G_r \cdot (\lambda / (4\pi d))^2 \quad \textit{(Equation 5)}$$

where P_r is received power, P_t is transmitted power, G_t is transmitter antenna gain, G_r is receiver antenna gain, λ is the wavelength, and d is the distance between antennas. The thermal noise floor and the resulting signal-to-noise ratio are

$$P_N = k \cdot T_{sys} \cdot B, \quad SNR = P_r / P_N \quad \textit{(Equation 6)}$$

with k the Boltzmann constant, T_{sys} the system noise temperature and B the bandwidth. End-to-end delay is decomposed as

$$T_{e2e} = \Sigma_{links} (d_i/c + L/R_b) + \Sigma_{relays} (t_{proc} + t_{switch} + t_{queue}) \quad \textit{(Equation 7)}$$

Per-link terms are charged on every traversed link. Per-node terms are charged only at intermediate relay nodes; the destination does not forward the packet and is not charged. Table 1 lists every parameter with its value and unit.

Table 1 | Model parameters.

Parameter	Symbol	Value	Note
Speed of light	c	299,792.458 km/s	physical constant
Boltzmann constant	k	1.380649×10^{-23} J/K	physical constant
Earth radius	R	6,371 km	spherical approximation
Gravitational parameter	μ = GM	398,600.4418 km^3/s^2	Earth

Parameter	Symbol	Value	Note
Transmit power	P_t	5 W (6.99 dBW)	assumed design value
Antenna gain	G_t, G_r	3000 linear (34.77 dBi)	assumed design value
Carrier frequency	f	12 GHz (Ku band)	λ = 2.50 cm
System noise temperature	T_{sys}	290 K	standard reference temperature
Bandwidth	B	20 MHz	assumed
Required carrier-to-noise	C/N	10 dB + 3 dB margin	permits range to 10,551 km
Packet size	L	12,000 bits (1500 B)	Ethernet MTU
Link data rate	R_b	1 Gbps	assumed
Per-hop processing	t_{proc}	swept, 0 to 15 ms	parameter, not assumption
Per-node switching	t_{switch}	2 µs	assumed
Per-node queueing	t_{queue}	0	unloaded-network assumption
Ground elevation mask	—	10°	minimum above local horizon
Inter-satellite link range	—	2,500 km	LEO to LEO

The transmit power, antenna gain, data rate and switching delay are adopted design values rather than measured ones. Their influence on the reported results is negligible: transmission contributes 0.060 ms and switching 0.008 ms across the entire route, against approximately 29 ms of propagation, and the link budget permits ranges to 10,551 km, so the 2,500 km geometric constraint binds first. No reported path or delay is altered by the choice of these four values.

For each configuration the procedure is: generate satellite positions; place ground stations; construct the connectivity graph by testing every node pair against all four feasibility constraints; compute the shortest path with Dijkstra's algorithm[18] using Equation 7 as the edge weight; and decompose the resulting path into its delay components. For time-stepped analysis the entire procedure is repeated at each epoch, with the graph discarded and rebuilt rather than updated.

Uniform satellite spacing was investigated separately using an electrostatics-motivated technique. Satellites are treated as mutually repelling particles confined to a sphere, subject to a distance-dependent repulsive force

$$F = 1 / d^n, \text{ with } n = 2 \quad (\text{Equation 8})$$

Initial positions are taken from a Fibonacci sphere distribution, which gives a nearly uniform starting arrangement, and 50 satellites at 550 km altitude are then relaxed over 1,000 fixed iterations with a step size of 0.01 decaying linearly, with radial reprojection to the orbital radius at each step. No convergence criterion is applied; the routine terminates at a fixed iteration count. This force does not correspond to a specific physical phenomenon in orbital mechanics, and serves only as a tractable means of modelling spacing.

Delays are reported as medians with 5th and 95th percentiles across 200 constellation phases spanning a complete orbital period, rather than as single values from a fixed geometry. Redundancy is measured as the number of internally node-disjoint paths between the two ground stations, equivalently the minimum vertex cut, computed by unit-capacity maximum flow on the node-split graph. Availability is the fraction of sampled phases at which any path exists. The break-even processing delay is obtained by sweeping the per-hop processing term and locating the crossing point of the two architectures.

This study involves no human or animal participants and uses no personal data. All results derive from simulation. The code, the parameter set and the script corresponding to each figure are published in https://github.com/nasianikolova/Multistrom-Cluster-of-Satellites repository under an MIT license, and each figure caption states the script and commit hash used to generate it, so that every reported value can be independently regenerated.

Results

Propagation delay increases linearly with altitude, from approximately 1.00 ms at 300 km to 6.67 ms at 2,000 km for a vertical one-way path (Figure 2A). This quantity is the one-way propagation time from the surface to a satellite at that altitude and is not an end-to-end route delay; the two are distinguished throughout.

Hop count decreases with altitude as a step function rather than a smooth curve (Figure 2B). For a 24-satellite ring with a 2,500 km inter-satellite link range and endpoints 60° apart, the route requires six hops below 575 km, four hops between 575 km and 1,825 km, and two hops above. These transition altitudes are the design-relevant quantity[19]; an earlier continuous fit to these data obscured them and produced non-integer hop counts.

Signal-to-noise ratio declines with altitude as free-space path loss increases (Figure 2C), from 43.92 dB at 300 km to 27.45 dB at 2,000 km, remaining above the 13 dB required carrier-to-noise ratio with margin across the entire range.

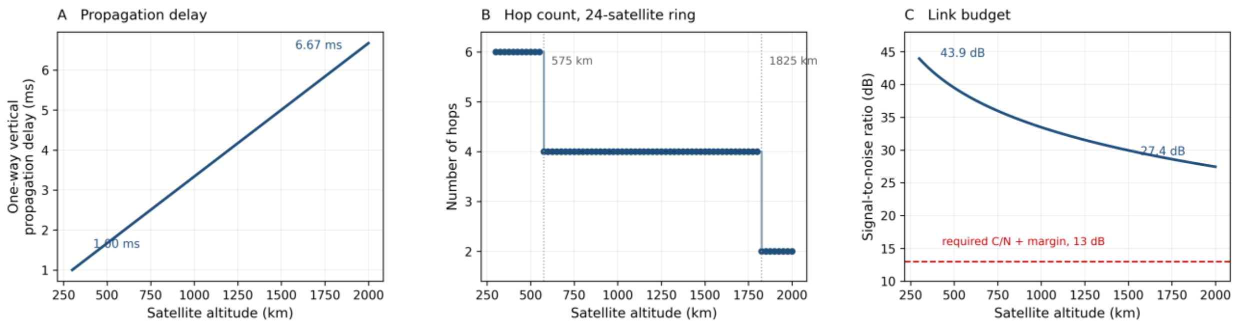

Figure 2 | Relationship between satellite altitude and network variables. A. One-way vertical propagation delay against altitude, 300 to 2,000 km. B. Simulated hop count against altitude for a 24-satellite ring, 2,500 km inter-satellite link range, endpoints 60° apart, 10° elevation mask. Markers are simulated values; hop count is integer-valued and changes at 575 km and 1,825 km. C. Signal-to-noise ratio against altitude in dB, with the 13 dB required carrier-to-noise ratio and margin marked. Generated by make_figure2.py, commit 405d8fe.

The three-dimensional topology of a multistrom constellation is shown in Figure 3, together with the Fibonacci-plus-repulsion distribution and a Walker configuration.

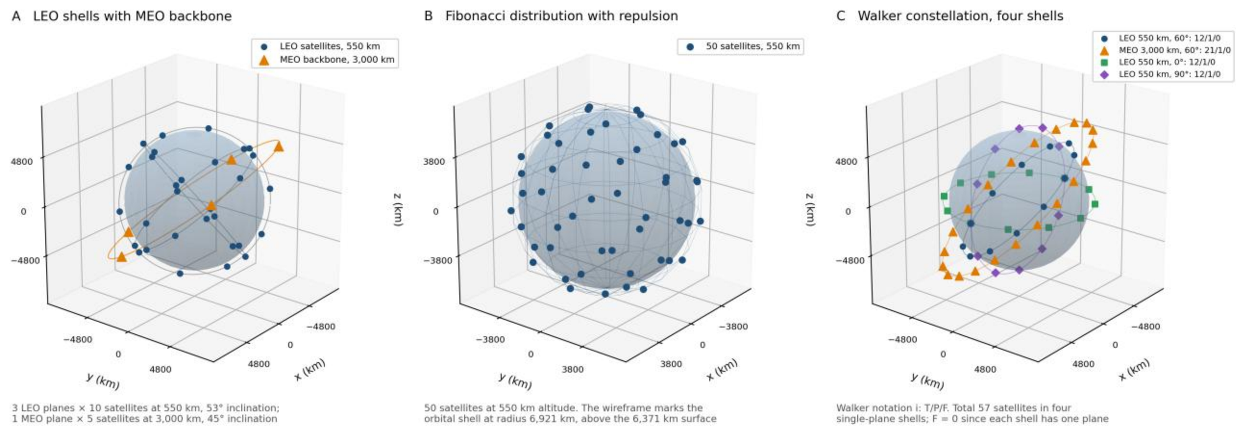

Figure 3 | Three-dimensional representations of Earth with LEO and MEO satellites. A. Three LEO orbital planes of ten satellites each at 550 km, together with one MEO plane of five satellites at 3,000 km. B. Fifty satellites at 550 km distributed by Fibonacci sphere initialization followed by repulsion relaxation. C. A Walker constellation of four single-plane shells: 12 satellites at 550 km and 60° inclination, 21 satellites at 3,000 km and 60° inclination, 12 satellites at 550 km and 0° inclination, and 12 satellites at 550 km and 90° inclination, in Walker notation 60°: 12/1/0, 60°: 21/1/0, 0°: 12/1/0 and 90°: 12/1/0. Generated by make_figure3.py, commit 405d8fe.

Under a controlled comparison in which endpoints, low Earth orbit stroma, routing algorithm and delay model are held constant and only the presence of the medium Earth orbit layer varies, the backbone does not reduce end-to-end delay (Table 2). At a per-hop processing delay of 1 ms the shortest-path algorithm does not select the backbone at all, returning the same five-hop route as the low Earth orbit configuration. A control condition placing the five additional satellites in the lower stroma rather than the higher one returns 32.78 ms, confirming that satellite count is not the operative variable.

Table 2 | End-to-end delay decomposition, Sydney to Tokyo, at 1 ms per-hop processing delay.

Configuration	Sats	Hops	Length (km)	Prop (ms)	Trans (ms)	Proc (ms)	Total (ms)
A. LEO only, 40 satellites	40	5	8,690.6	28.99	0.060	4.00	33.06
B. LEO 40 + MEO backbone 5	45	5	8,690.6	28.99	0.060	4.00	33.06
C. LEO only, 45 (control)	45	5	8,607.4	28.71	0.060	4.00	32.78
Backbone route, if forced	45	3	12,781.4	42.63	0.036	2.00	44.67

The backbone route removes two relay hops at a cost of 4,090 km of additional path length and 13.64 ms of additional propagation delay. Sweeping the per-hop processing delay locates the point at which this exchange becomes favorable: the two architectures break even at 6.80 ms per hop (Figure 4C). Below that value the low Earth orbit route is faster; above it the backbone is faster. The threshold is sensitive to inter-satellite link range, falling to 2.70 ms when that range is restricted to 2,000 km, and at 1,000 km the lower layer no longer forms a connected path at all, so the backbone ceases to be an optimization and becomes a precondition for connectivity. Transmission contributes 0.060 ms and switching 0.008 ms across the entire route, both negligible against approximately 29 ms of propagation.

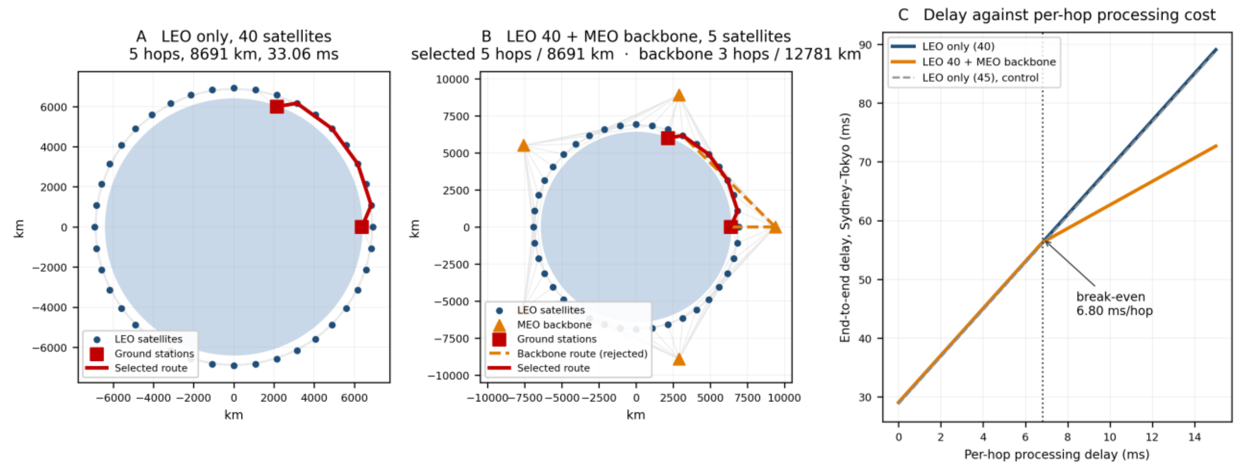

Figure 4 | Controlled comparison of end-to-end delay with and without a MEO backbone. All panels route between identical ground stations at the Sydney to Tokyo central angle of 70.38°, corresponding to 7,826 km great-circle distance, over an identical LEO stroma of 40 satellites at 550 km. A. LEO-only configuration: the selected route uses 5 hops over 8,691 km for a total modelled delay of 33.06 ms. B. The same stroma with a five-satellite MEO backbone at 3,000 km. The router still selects the LEO route, shown solid; the backbone route, shown dashed, reaches the destination in 3 hops but traverses 12,781 km. C. Total end-to-end delay against per-hop processing delay for both configurations, with a control in which the five additional satellites are placed in LEO. The curves cross at 6.80 ms per hop. Generated by make_figure4.py, commit 405d8fe.

Varying the number of orbital planes while holding total satellite count constant shows that network performance improves with plane count only up to a point (Table 3). At 240 satellites, path availability rises from 22 per cent with two planes to 90 per cent with ten, mean delay falls from 51.40 ms to 36.86 ms, and mean node-disjoint path count rises from 1.7 to 3.4. Beyond ten planes all three metrics plateau and redundancy declines slightly. The 480-satellite case reproduces the same pattern, indicating a structural effect rather than sampling noise.

Table 3 | Availability, delay and redundancy against orbital plane count, 240 satellites, averaged over 120 random constellation orientations.

Planes	Satellites per plane	Availability (%)	Mean delay (ms)	Mean disjoint paths
2	120	22	51.40	1.7
4	60	60	39.67	2.7
6	40	76	38.51	3.0
8	30	86	38.31	3.3
10	24	90	36.86	3.4
12	20	82	36.60	3.3
20	12	82	36.67	3.1

Because satellites move along their orbits and the Earth rotates beneath them, the selected route changes over time even though the constellation itself does not reconfigure. Over one complete orbital period of 95.5 minutes, sampled at 120 epochs with the graph rebuilt at each, a 240-satellite constellation in 12 planes maintained a path at 100 per cent of epochs (Figure 5). End-to-end delay varied between 34.49 ms and 42.24 ms with a mean of 38.44 ms and a standard deviation of 1.26 ms, and hop count varied between five and seven. The route changed 77 times, giving a mean route lifetime of 1.24 minutes before the selected path differs.

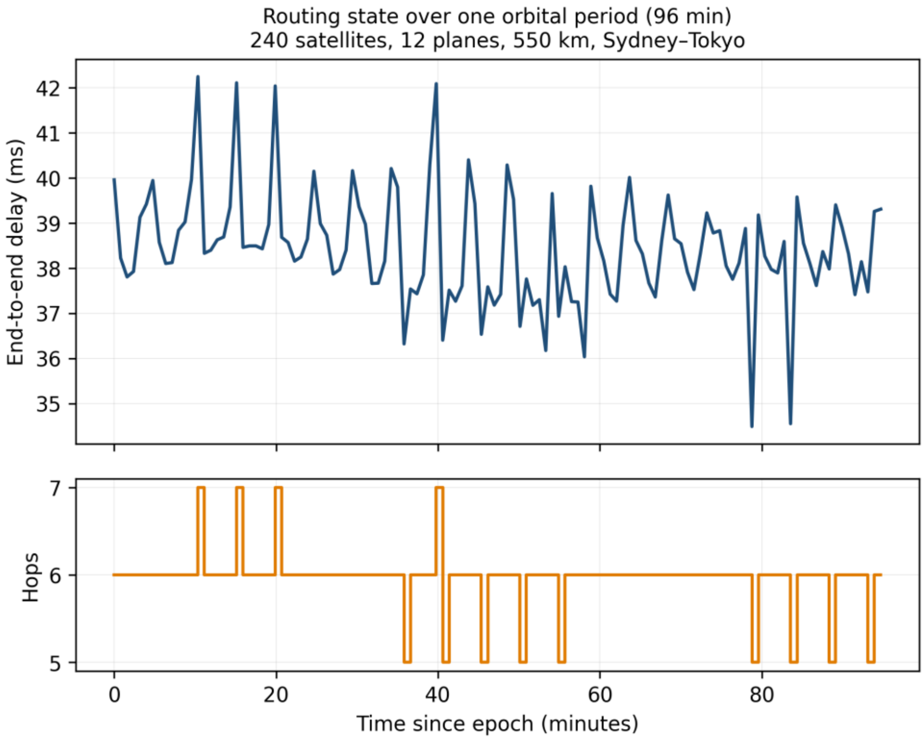

Figure 5 | Routing state over one orbital period. A constellation of 240 satellites in 12 planes at 550 km, routing between Sydney and Tokyo, sampled at 120 epochs across one orbital period of 95.5 minutes. The connectivity graph is rebuilt from scratch at each epoch. Upper panel, end-to-end delay; lower panel, hop count. Generated by make_figure5.py, commit 405d8fe.

Three alternative in-space architectures were tested. A bent-pipe architecture without inter-satellite links is not merely unavailable on this route but geometrically impossible: at 550 km with a 10° elevation mask a single satellite spans at most 3,329 km, against a 7,826 km endpoint separation. A ground relay does not resolve this either, since two such legs reach only 6,657 km, and a chain of two intermediate stations, while giving an idealized 34.35 ms, is never available against the actual constellation because each leg requires one satellite visible from both its endpoints at once.

A single relay at 3,000 km does span the route, since its visibility footprint reaches 8,444 km, at an upper-bound delay of 38.04 ms. This assumes a satellite positioned optimally and is therefore a bound rather than an achievable value. The meshed constellation achieved a median of 38.35 ms with a 5th to 95th percentile range of 35.12 ms to 40.05 ms, available at every sampled phase. Inter-satellite links are therefore a requirement rather than an optimization on intercontinental paths.

Discussion

This research set out to visualize a multistrom environment in which satellites and backbones communicate regardless of the orbit to which they belong. Data exist at every altitude, and the architecture treats the orbital segment as a single connected medium rather than a set of isolated stromata. Quantitative results are reported from 300 km to 2,000 km, with a backbone stroma at 3,000 km.

The central finding is that a higher orbital layer does not reduce end-to-end delay by itself. It exchanges physical path length for hop count, and whether that exchange is favorable depends on how expensive in terms of delay each hop is. The break-even value of 6.80 ms per hop is a design criterion.

Performance improves with orbital plane count only up to approximately ten planes, after which availability, delay and redundancy saturate. This connects the routing analysis back to the congestion problem: beyond that point, additional satellites add orbital density without measurably improving the network.

The impossibility results establish the architecture more strongly than any performance comparison. On an intercontinental route at 550 km, no architecture that returns to the ground is available, so a network of this reach must be connected in orbit. This confirms that all parts of a space network can exist in orbit and at different altitudes.

Delay is not a single value but a distribution. Over one orbital period the end-to-end delay varied between 34.49 ms and 42.24 ms, and the selected route changed on average every 1.24 minutes as satellites moved along their orbits and the Earth rotated beneath them, without the constellation reconfiguring. Any system built on such a network must cope with a path that changes roughly every minute and a delay that is never constant.

A topic not expanded on in this research is the modules that surround each backbone. These modules support the backbone by exchanging information and provide it with all the changes scientists wish to include in future works. The question is how they will orbit the backbone, return to the Earth and be replaced with a revised version of the backbone.

Conclusion

Putting satellites in a higher orbit does not make the network faster on its own. A backbone at 3,000 km cuts the number of satellites a signal passes through from five to three, but the signal has to travel 4,090 km further to get there. Whether that trade is worth making depends on how long each satellite takes to process a signal: below 6.80 ms per satellite the lower orbit performs better, above it the backbone does. Adding more orbital planes helps until about ten, after which it saturates.

The results replace a qualitative preference for higher orbital layers with a quantitative criterion for when such layers help. They also establish a bound on useful constellation density for intercontinental routing, which bears directly on orbital congestion. For network design, the finding that inter-satellite links are structurally required rather than merely advantageous at low altitude constrains the architecture of any comparable system.

All five objectives were met. End-to-end delay was quantified with and without the higher layer. The break-even processing delay was identified. Redundancy was measured as the number of node-disjoint paths between endpoints. The plane count at which performance saturates was determined. The constellation was compared against alternative in-space architectures. The initial expectation that a backbone would reduce delay was not supported, and that expectation has been replaced by the conditional result reported here.

Future work should introduce a maneuver and propulsion model with delta-v, transfer time and station-keeping cost, which would allow physical reconfiguration to be evaluated rather than assumed[17, 20]. Extending the analysis above 3,000 km, incorporating a traffic and queueing model, and comparing against an established open-source constellation simulator[21, 22, 23] would each strengthen the results. Measuring the per-hop processing delay of representative on-board hardware would convert the break-even threshold from a parametric result into a definite one.

All orbits are assumed circular and unperturbed; there is no eccentricity, no J2 perturbation and no atmospheric drag. No atmospheric, ionospheric or weather effects are modelled. There is no link capacity, queueing or traffic model, and queueing delay is set to zero under an explicitly unloaded-network assumption. This assumption is conservative against the hypothesis under test, since queueing scales with hop count and its inclusion would favor the backbone rather than the low Earth orbit route.

There is no propulsion, delta-v, maneuver, station-keeping or collision-avoidance model, and no power, mass or cost budget. The constellation does not physically reconfigure; what changes over time is the selected route as satellites move along fixed orbits. The orbital-plane analysis samples random constellation orientations rather than propagating orbits continuously, and this substitution is an approximation. The controlled comparison of Figure 4 is conducted in a single two-dimensional orbital plane rather than the full three-dimensional constellation described in the Methodology, which simplifies the geometry of that comparison. Quantitative results are confined to 300 to 2,000 km; the 3,000 km stroma appears only in the architectural comparison, and no altitude above that was simulated. Four link-budget parameters are taken as assumed design values and are not measured or independently sourced.

This research asked whether backbones can exist in space and what placing them there would achieve. The answer is that they can, and that the case for doing so rests less on speed than on resilience and reach. Orbital infrastructure is not exposed to the seabed disturbance that has severed multiple submarine cables simultaneously, with repairs measured in days[3], nor is it constrained by the geography and landing rights that determine terrestrial routing. The delay benefit is conditional, arising only where per-hop processing exceeds 6.80 ms. The reach benefit is not conditional, however; on intercontinental routes no architecture that returns to the ground is available, so a network of this span must be established in orbit or not established.

References

1. European Space Agency Space Debris Office. ESA's annual space environment report. ESA/ESOC, Darmstadt, Germany, 2025. https://www.sdo.esoc.esa.int/environment_report/Space_Environment_Report_latest.pdf
2. A. Falle, E. Wright, A. Boley, M. Byers. One million (paper) satellites. Science. Vol. 382, pg. 150–152, 2023, https://doi.org/10.1126/science.adi4639.
3. E. W. W. Chan, X. Luo, W. W. T. Fok, W. Li, R. K. C. Chang. Non-cooperative diagnosis of submarine cable faults. Passive and Active Measurement (PAM 2011), Lecture Notes in Computer Science. Vol. 6579, pg. 224–234, 2011, https://doi.org/10.1007/978-3-642-19260-9_23.
4. C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. Fernández del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, T. E. Oliphant. Array programming with NumPy. Nature. Vol. 585, pg. 357–362, 2020, https://doi.org/10.1038/s41586-020-2649-2.
5. J. D. Hunter. Matplotlib: a 2D graphics environment. Computing in Science & Engineering. Vol. 9, pg. 90–95, 2007, https://doi.org/10.1109/MCSE.2007.55.
6. M. Handley. Delay is not an option: low latency routing in space. Proceedings of the 17th ACM Workshop on Hot Topics in Networks (HotNets 2018). pg. 85–91, 2018, https://doi.org/10.1145/3286062.3286075.
7. M. Handley. Using ground relays for low-latency wide-area routing in megaconstellations. Proceedings of the 18th ACM Workshop on Hot Topics in Networks (HotNets 2019). pg. 125–132, 2019, https://doi.org/10.1145/3365609.3365859.
8. T. Klenze, G. Giuliari, C. Pappas, A. Perrig, D. Basin. Networking in heaven as on Earth. Proceedings of the 17th ACM Workshop on Hot Topics in Networks (HotNets 2018). pg. 22–28, 2018, https://doi.org/10.1145/3286062.3286066.
9. D. Fischer, D. Basin, K. Eckstein, T. Engel. Predictable mobile routing for spacecraft networks. IEEE Transactions on Mobile Computing. Vol. 12, pg. 1174–1187, 2013, https://doi.org/10.1109/TMC.2012.93.
10. D. Bhattacherjee, A. Singla. Network topology design at 27,000 km/hour. Proceedings of the 15th International Conference on Emerging Networking Experiments and Technologies (CoNEXT 2019). pg. 341–354, 2019, https://doi.org/10.1145/3359989.3365407.
11. G. Giuliari, T. Klenze, M. Legner, D. Basin, A. Perrig, A. Singla. Internet backbones in space. ACM SIGCOMM Computer Communication Review. Vol. 50, pg. 25–37, 2020, https://doi.org/10.1145/3390251.3390256.
12. Y. Hauri, D. Bhattacherjee, M. Grossmann, A. Singla. "Internet from space" without inter-satellite links. Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets 2020). pg. 205–211, 2020, https://doi.org/10.1145/3422604.3425938.
13. S. Kassing, D. Bhattacherjee, A. B. Águas, J. E. Saethre, A. Singla. Exploring the "Internet from space" with Hypatia. Proceedings of the ACM Internet Measurement Conference (IMC 2020). pg. 214–229, 2020, https://doi.org/10.1145/3419394.3423635.
14. B. Denby, B. Lucia. Orbital edge computing: nanosatellite constellations as a new class of computer system. Proceedings of the 25th International Conference on Architectural Support for Programming

Languages and Operating Systems (ASPLOS 2020). pg. 939–954, 2020, https://doi.org/10.1145/3373376.3378473.

15. J. M. Gongora-Torres, C. Vargas-Rosales, A. Aragón-Zavala, R. Villalpando-Hernandez. Link budget analysis for LEO satellites based on the statistics of the elevation angle. IEEE Access. Vol. 10, pg. 14518–14528, 2022, https://doi.org/10.1109/ACCESS.2022.3147829.
16. J. A. Christian. A tutorial on horizon-based optical navigation and attitude determination with space imaging systems. IEEE Access. Vol. 9, pg. 19819–19853, 2021, https://doi.org/10.1109/ACCESS.2021.3051914.
17. H. D. Curtis. Chapter 6 – orbital maneuvers. Orbital mechanics for engineering students. 4th ed., Butterworth-Heinemann, pg. 287–350, 2020.
18. E. W. Dijkstra. A note on two problems in connexion with graphs. Numerische Mathematik. Vol. 1, pg. 269–271, 1959, https://doi.org/10.1007/BF01386390.
19. C. Rao, E. Modiano. Minimum-hop constellation design for low earth orbit satellite networks. IEEE INFOCOM 2025 – IEEE Conference on Computer Communications. pg. 1–10, 2025, https://doi.org/10.48550/arXiv.2506.11995.
20. G. Di Pasquale, M. Sanjurjo-Rivo, D. Pérez Grande. Optimization of constellation deployment using on-board propulsion and Earth nodal regression. Advances in Space Research. Vol. 70, pg. 3281–3300, 2022, https://doi.org/10.1016/j.asr.2022.09.019.
21. Z. Lai, H. Li, J. Li. StarPerf: characterizing network performance for emerging mega-constellations. IEEE 28th International Conference on Network Protocols (ICNP 2020). pg. 1–11, 2020, https://doi.org/10.1109/ICNP49622.2020.9259357.
22. Z. Lai, W. Liu, Q. Wu, H. Li, J. Xu, J. Wu. SpaceRTC: unleashing the low-latency potential of mega-constellations for real-time communications. IEEE INFOCOM 2022 – IEEE Conference on Computer Communications. pg. 1339–1348, 2022, https://doi.org/10.1109/INFOCOM48880.2022.9796887.
23. Z. Lai, H. Li, Y. Deng, Q. Wu, J. Liu, Y. Li, J. Li, L. Liu, W. Liu, J. Wu. StarryNet: empowering researchers to evaluate futuristic integrated space and terrestrial networks. 20th USENIX Symposium on Networked Systems Design and Implementation (NSDI 2023). pg. 1309–1324, 2023