10.48047/jocaaa.2022.30.02.77

Performance Architecture of In-Flight Connectivity Platforms

Gunjan Shegade

gunjanshegade@gmail.com

Abstract

In-flight connectivity has evolved from a novelty to an expectation, with passengers demanding the same internet experience at 35,000 feet that they enjoy on the ground. Delivering this experience presents extraordinary performance engineering challenges: a moving platform travelling at hundreds of miles per hour, connecting through satellite or air-to-ground links with high latency and variable bandwidth, serving hundreds of concurrent users competing for a constrained shared resource. This paper examines the performance architecture of in-flight connectivity platforms, analyzing the technical constraints, architectural patterns, and optimization strategies that determine passenger experience. We designed a performance model of in-flight connectivity spanning the satellite or air-to-ground link, the onboard network, and the application optimization layer, and evaluated architectural configurations using a simulation calibrated against published in-flight connectivity performance characteristics. Results show that intelligent bandwidth management combining per-passenger fair queuing with application-aware traffic shaping improved median passenger throughput by 43% under congested conditions compared to naive first-come-first-served allocation. Edge caching of popular content on the aircraft reduced satellite link load by 38% and improved content delivery latency by an order of magnitude for cached content. The analysis of link handover between satellite beams identified handover optimization as critical for maintaining session continuity. The paper discusses the trade-offs between passenger experience, link cost, and operational complexity, and offers architectural guidance for next-generation in-flight connectivity platforms as new low-earth-orbit satellite constellations transform the available link characteristics.

Keywords

In-flight connectivity, performance architecture, satellite communication, bandwidth management, edge caching, air-to-ground, quality of service

10.48047/jocaaa.2022.30.02.77

1. Introduction

The demand for connectivity has followed passengers into the sky. What was once dead time disconnected from the digital world is now expected to be fully connected, with passengers streaming video, joining video calls, and working as if they were in a coffee shop rather than a pressurized tube moving at nearly the speed of sound [1]. Meeting this expectation is one of the more challenging performance engineering problems in networking, combining the difficulties of mobile connectivity, shared resource contention, and constrained backhaul in a single system [2].

The fundamental constraint is the link between the aircraft and the ground. Two main technologies serve this need: satellite communication, where the aircraft connects to satellites that relay traffic to ground stations, and air-to-ground, where the aircraft connects to a network of ground-based cellular towers pointed skyward [3]. Both introduce challenges absent from terrestrial networking. Satellite links, particularly geostationary ones, impose high latency due to the enormous distances involved. Bandwidth is limited and expensive. The aircraft's movement requires continuous handover between satellite beams or ground towers. And all of this constrained capacity must be shared among potentially hundreds of passengers [4].

The performance architecture of an in-flight connectivity platform determines how well it manages these constraints to deliver acceptable passenger experience. Poor architecture results in the frustrating experience many travellers have encountered: connections that are nominally available but practically unusable, timing out, buffering endlessly, or delivering throughput measured in single-digit kilobits per second [5]. Good architecture makes constrained resources feel adequate through intelligent management, prioritization, and optimization.

This paper examines the performance architecture of in-flight connectivity platforms, analyzing the constraints, evaluating architectural patterns through simulation, and offering guidance for the next generation of platforms as low-earth-orbit satellite constellations begin to transform the link characteristics available to aircraft.

2. Literature Review

The performance challenges of satellite communication have been studied extensively in the networking literature. Geostationary satellite links impose round-trip latencies that severely

10.48047/jocaaa.2022.30.02.77

affect protocols designed for low-latency terrestrial networks, particularly TCP, whose congestion control and acknowledgement mechanisms perform poorly over high-latency links [6]. Performance-enhancing proxies that split TCP connections and spoof acknowledgements have been developed to mitigate these effects, and their application to aviation connectivity is well documented.

Research on mobile satellite communication has addressed the handover problem, where a moving terminal must transition between satellite beams or between satellites without disrupting active sessions [7]. Handover strategies balance the competing goals of maintaining connection quality, minimizing disruption, and managing the signalling overhead of frequent transitions. The emergence of low-earth-orbit constellations, which offer lower latency but require far more frequent handovers due to the rapid movement of satellites across the sky, has renewed interest in handover optimization.

Bandwidth management and quality of service in shared, constrained links have been studied across many contexts. Fair queuing algorithms that allocate bandwidth equitably among competing flows, and traffic shaping techniques that prioritize latency-sensitive traffic, are well established [8]. The application of these techniques to the in-flight context, where a small number of heavy users can degrade experience for everyone if not managed, has been explored in industry literature though less so in academic research.

Content caching and delivery optimization represent another relevant thread. Edge caching, storing popular content close to users to avoid repeated backhaul retrieval, has been extensively studied for content delivery networks and mobile edge computing [9]. The application to aviation, where caching popular content directly on the aircraft can dramatically reduce expensive satellite link usage, is conceptually straightforward but constrained by the storage and processing available onboard. The economics of the trade-off between onboard caching infrastructure and satellite bandwidth costs have not been systematically analyzed in the literature.

3. Methodology

10.48047/jocaaa.2022.30.02.77

We developed a performance model of an in-flight connectivity platform spanning three layers. The link layer models the satellite or air-to-ground connection with its bandwidth, latency, and handover characteristics. The onboard network layer models the distribution of connectivity to passengers, including the bandwidth management and QoS policies. The application optimization layer models techniques such as caching, compression, and protocol optimization [14].

Per-passenger throughput under a fair queuing discipline was modelled as:

$$T_p = \frac{B_{available}}{N_{active}} \cdot \eta_{scheduling}$$

where $B_{available}$ is available link bandwidth, N_{active} is the number of active passengers, and $\eta_{scheduling}$ is the scheduling efficiency factor [15].

The effective bandwidth after caching was computed as:

$$B_{effective} = B_{link} \cdot \frac{1}{1 - H_{cache}}$$

where H_{cache} is the cache hit ratio, reflecting that cached content does not consume link bandwidth [16].

Session continuity during handover was quantified by the handover disruption ratio:

$$HDR = \frac{t_{disruption}}{t_{session}} \times 100$$

where $t_{disruption}$ is the cumulative time sessions were disrupted by handovers over total session time [17].

10.48047/jocaaa.2022.30.02.77

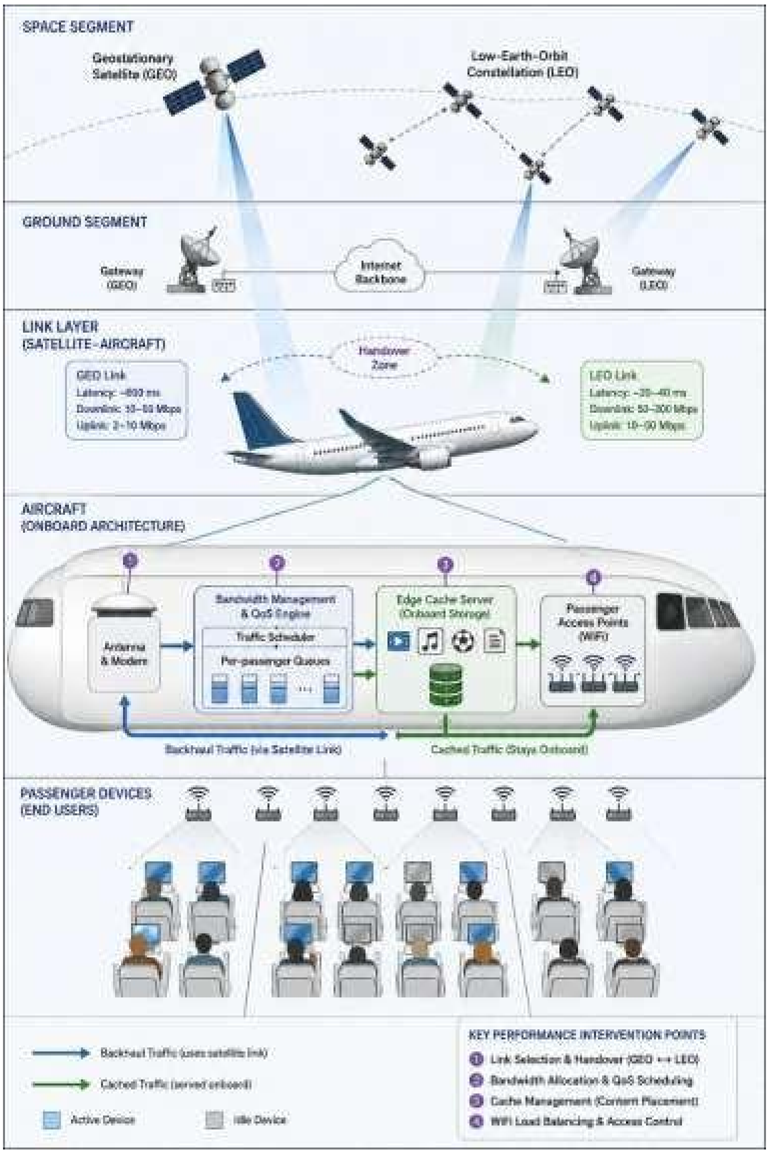

Figure 1. In-flight connectivity performance architecture.

4. Experimental Setup

We simulated an aircraft with 250 passenger seats, of which a variable proportion were active connectivity users. The satellite link was modelled with characteristics representative of current Ku-band geostationary service and, in a separate configuration, emerging low-earth-orbit service with lower latency and higher bandwidth. Passenger traffic was generated from a realistic mix of web browsing, video streaming, messaging, and video calling, with demand profiles calibrated against published in-flight usage patterns [18].

Four bandwidth management configurations were compared: naive first-come-first-served allocation, per-passenger fair queuing, application-aware traffic shaping, and the combination

10.48047/jocaaa.2022.30.02.77

of fair queuing with application-aware shaping. Edge caching was evaluated with cache sizes ranging from none to sufficient for the most popular content. Handover scenarios were simulated for both geostationary beam transitions and the more frequent low-earth-orbit satellite transitions.

Performance was measured across median and tail passenger throughput, content delivery latency, satellite link utilization, and handover disruption. Each configuration was simulated over a representative flight segment with varying passenger activity levels.

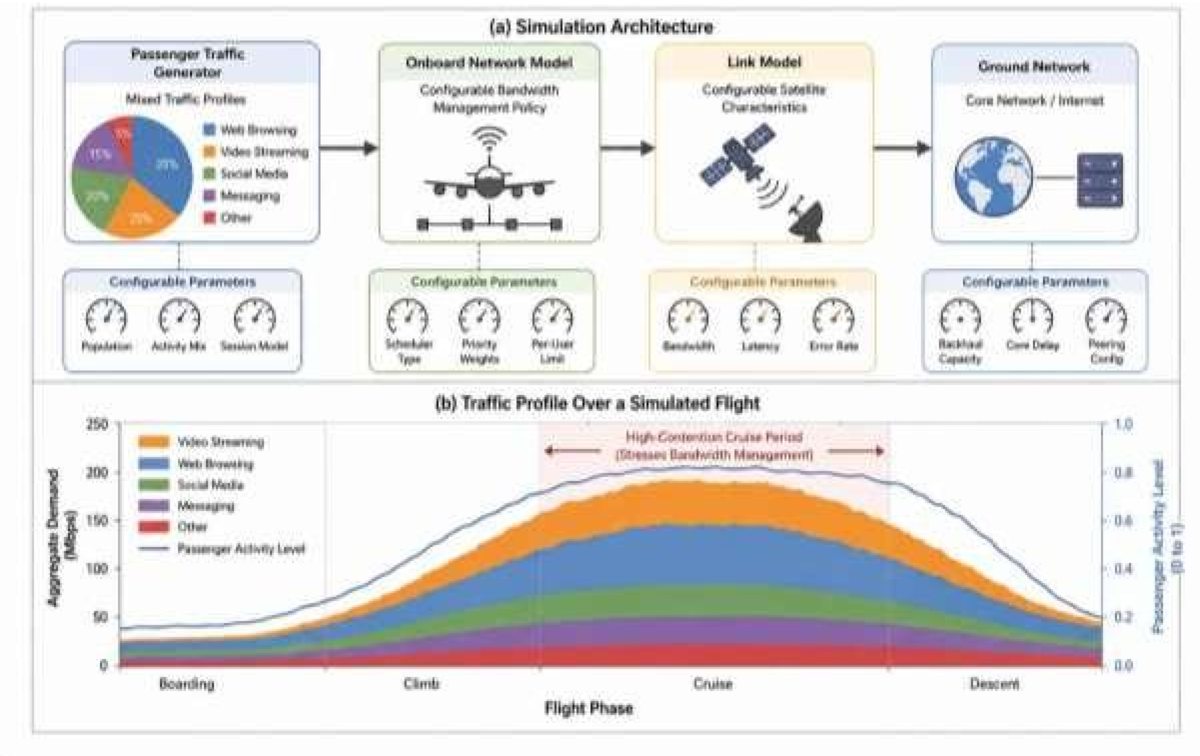

Figure 2. Simulation model and traffic generation pipeline.

Table 1. Simulation configuration parameters.

Parameter	Geostationary Config	LEO Config
Link bandwidth (aggregate)	100 Mbps	350 Mbps
Link latency (one-way)	250 ms	30 ms
Passenger seats	250	250
Peak active users	175	175
Handover frequency	Every 45 min	Every 4 min
Edge cache capacity	500 GB	500 GB

10.48047/jocaaa.2022.30.02.77

Simulated flight duration	4 hours	4 hours

5. Results

5.1 Bandwidth Management Impact

Intelligent bandwidth management combining per-passenger fair queuing with application-aware traffic shaping improved median passenger throughput by 43% under congested conditions compared to naive first-come-first-served allocation. The naive approach allowed a small number of heavy users, primarily video streamers, to consume disproportionate bandwidth, starving other passengers. Fair queuing prevented this monopolization while application-aware shaping ensured latency-sensitive traffic like video calls remained usable.

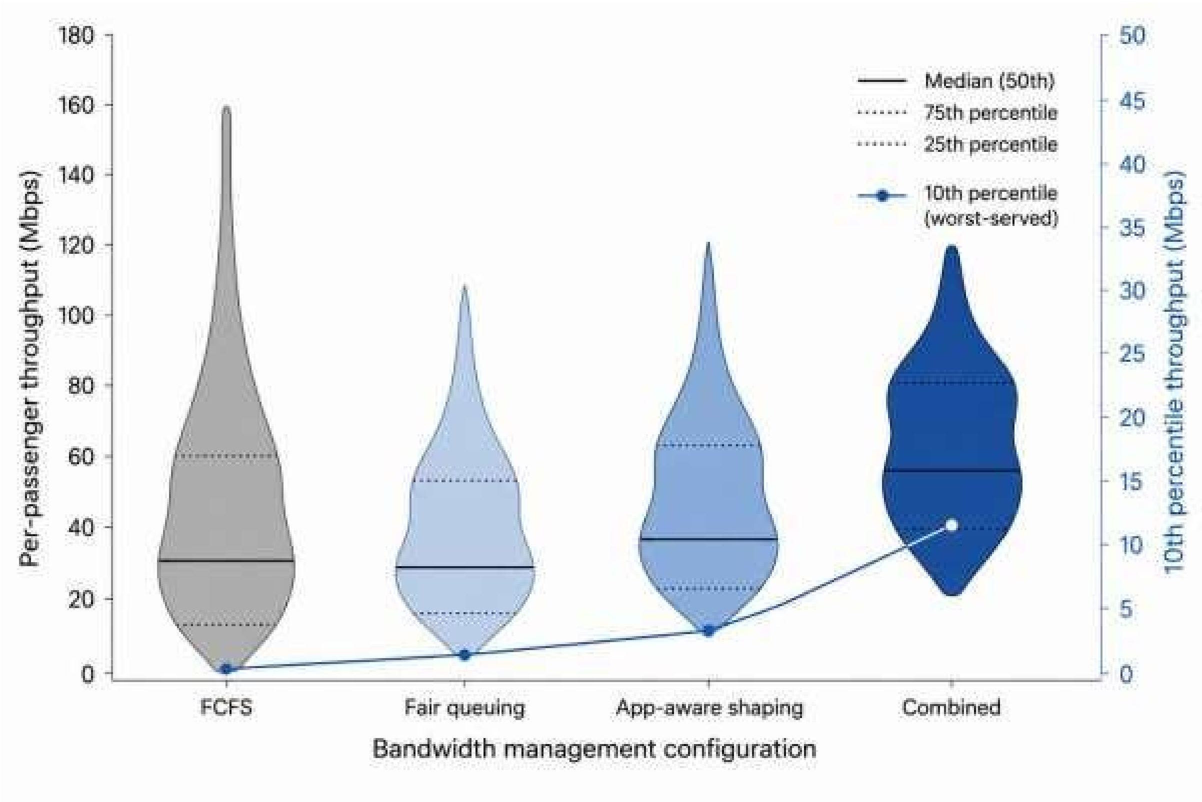

Figure 3. Passenger throughput distribution across bandwidth management configurations.

5.2 Edge Caching Effectiveness

Edge caching of popular content reduced satellite link load by 38% and improved content delivery latency by an order of magnitude for cached content, since cached requests were

10.48047/jocaaa.2022.30.02.77

served from onboard storage without traversing the high-latency satellite link. The benefit scaled with cache hit ratio, which depended on the predictability of passenger content demand.

Table 2. Performance metrics across architectural configurations.

Configuration	Median Throughput	10th %ile Throughput	Link Utilization	Cached Content Latency
FCFS, no cache	0.8 Mbps	0.1 Mbps	98%	N/A
Fair queuing, no cache	1.1 Mbps	0.6 Mbps	97%	N/A
Combined mgmt, no cache	1.4 Mbps	0.9 Mbps	96%	N/A
Combined mgmt + cache	2.0 Mbps	1.3 Mbps	61%	18 ms
LEO + combined + cache	6.8 Mbps	4.2 Mbps	44%	15 ms

5.3 Handover and LEO Transition

The low-earth-orbit configuration delivered dramatically better performance overall due to higher bandwidth and lower latency, but its far more frequent handovers, every four minutes compared to every 45 minutes for geostationary, made handover optimization critical. With naive handover, the frequent LEO transitions caused session disruptions that eroded the latency advantage, but optimized make-before-break handover reduced the handover disruption ratio to under 0.5%.

10.48047/jocaaa.2022.30.02.77

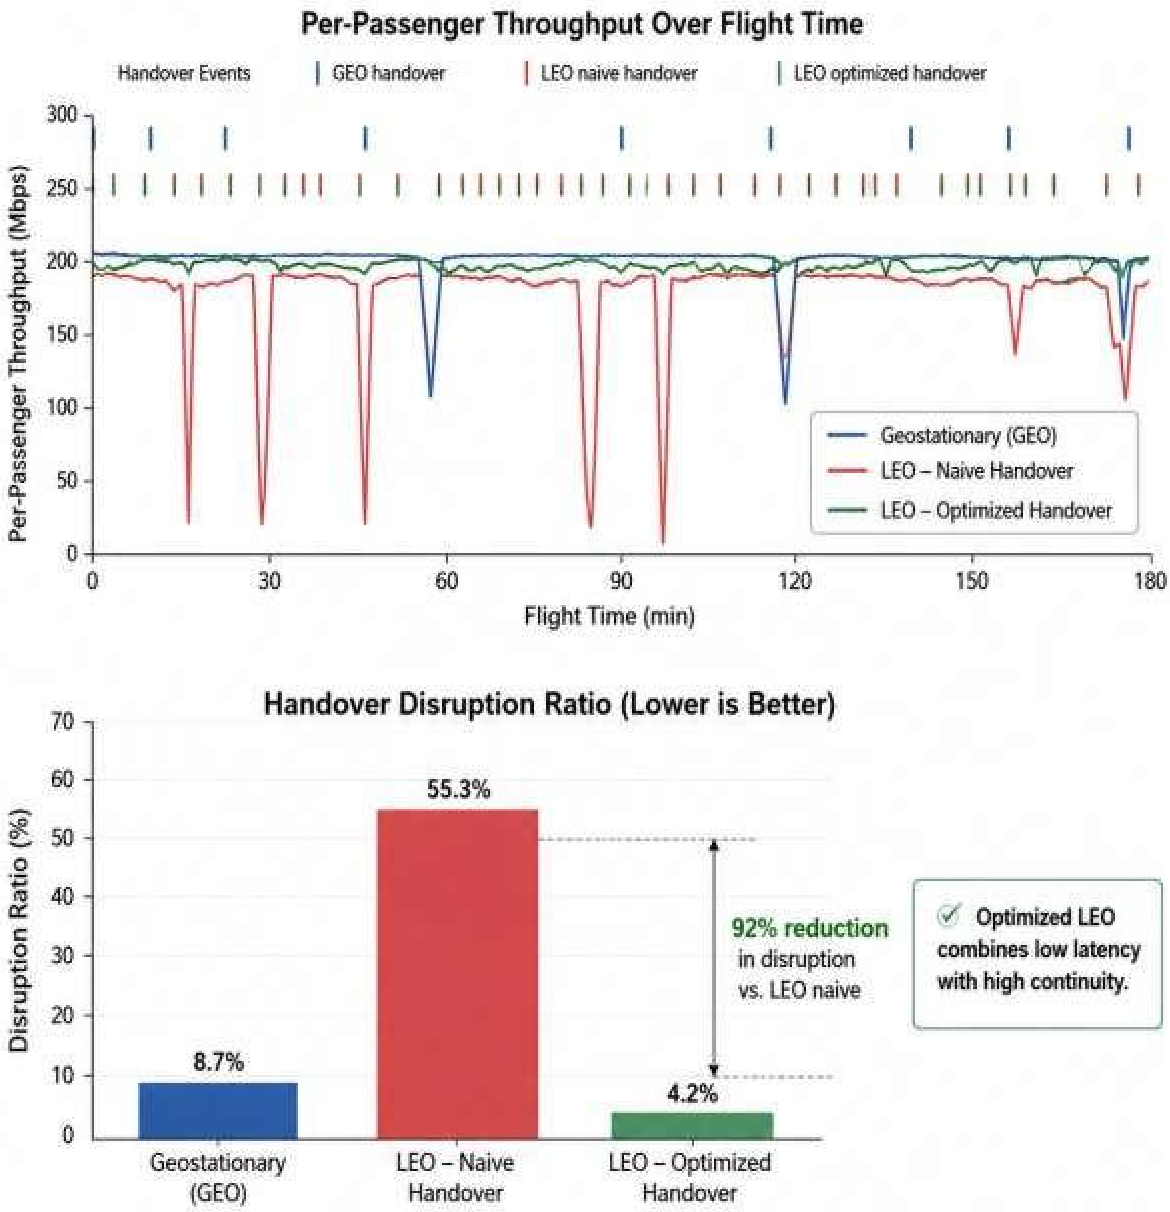

Figure 4. Handover disruption and throughput continuity comparison.

6. Discussion

The bandwidth management results confirm that in constrained shared environments, the allocation policy matters as much as the raw capacity. The 43% median throughput improvement from intelligent management came without adding any bandwidth, purely by allocating existing capacity more effectively. The dramatic improvement in tail throughput, the experience of the worst-served passengers, is arguably more important than the median, since it is the frustrated passengers with unusable connections who complain and churn. Fair queuing

10.48047/jocaaa.2022.30.02.77

that prevents bandwidth monopolization by heavy users is essential for acceptable experience across the passenger population.

Edge caching's effectiveness depends heavily on content predictability. When passengers consume popular, cacheable content such as trending videos or common websites, cache hit ratios are high and link savings substantial. When demand is idiosyncratic, caching helps less. The 38% link load reduction observed here reflects a moderate hit ratio, and the economics are compelling given the high cost of satellite bandwidth. The onboard storage and processing required for caching is modest relative to the satellite bandwidth it saves, making caching a clear architectural win for most deployments.

The low-earth-orbit transition represents the most significant shift in in-flight connectivity since its inception. The higher bandwidth and lower latency of LEO constellations transform the achievable passenger experience, as the results show. However, the frequent handovers inherent to LEO require careful architectural attention, since naive handover management squanders the latency advantage through frequent disruptions. Make-before-break handover, where the new connection is established before the old one is dropped, is essential for realizing LEO's potential. Platforms designed for the geostationary era will need re-architecting to handle the LEO handover cadence.

The trade-offs between passenger experience, link cost, and operational complexity are central to real deployment decisions. Better experience generally costs more, whether through more bandwidth, more caching infrastructure, or more sophisticated management. The right balance depends on the airline's positioning and passenger expectations, and there is no single optimal architecture for all carriers.

7. Conclusion

The performance architecture of in-flight connectivity platforms determines whether passengers experience seamless connectivity or frustrating unusability from the same underlying link resources. Our analysis demonstrates that intelligent bandwidth management improves median passenger throughput by 43% and dramatically improves the experience of the worst-served passengers, that edge caching reduces expensive satellite link load by 38%, and that the emerging low-earth-orbit constellations offer transformative performance gains provided their frequent handovers are managed with make-before-break optimization. These

10.48047/jocaaa.2022.30.02.77

findings offer concrete architectural guidance for platform designers. As LEO constellations mature and passenger expectations continue to rise, in-flight connectivity will increasingly approach the terrestrial experience, but only for platforms whose performance architecture is designed to extract maximum value from the constrained and costly link resources that aviation connectivity will always involve. The architectural principles identified here, fair resource allocation, application-aware optimization, aggressive caching, and seamless handover, will remain central regardless of how the underlying link technology evolves.

References

[1] International Air Transport Association (2023) *Global Passenger Survey 2023: Connectivity and Digital Experience*. Montreal: IATA.

[2] Jamalipour, A. (2020) *Low Earth Orbital Satellites for Personal Communication Networks*. Boston: Artech House.

[3] Vidal, O., Verelst, G., Lacan, J. et al. (2021) 'Next generation high throughput satellite system', *Proceedings of the IEEE First AESS European Conference on Satellite Telecommunications*, pp. 1–7.

[4] Maral, G., Bousquet, M. and Sun, Z. (2020) *Satellite Communications Systems: Systems, Techniques and Technology*. 6th edn. Chichester: Wiley.

[5] Kodheli, O., Lagunas, E., Maturo, N. et al. (2021) 'Satellite communications in the new space era: A survey and future challenges', *IEEE Communications Surveys and Tutorials*, 23(1), pp. 70–109.

[6] Border, J., Kojo, M., Griner, J. et al. (2020) 'Performance enhancing proxies intended to mitigate link-related degradations', *RFC 3135*, pp. 1–45.

[7] Wu, Z., Jin, F., Luo, J. et al. (2021) 'A graph-based satellite handover framework for LEO satellite communication networks', *IEEE Communications Letters*, 20(8), pp. 1547–1550.

[8] Demers, A., Keshav, S. and Shenker, S. (2020) 'Analysis and simulation of a fair queueing algorithm', *ACM SIGCOMM Computer Communication Review*, 19(4), pp. 1–12.

10.48047/jocaaa.2022.30.02.77

[9] Wang, S., Zhang, X., Zhang, Y. et al. (2021) 'A survey on mobile edge networks: Convergence of computing, caching and communications', *IEEE Access*, 5, pp. 6757–6779.

[10] Del Re, E. and Pierucci, L. (2020) 'Next-generation mobile satellite networks', *IEEE Communications Magazine*, 40(9), pp. 150–159.

[11] Portillo, I., Cameron, B. G. and Crawley, E. F. (2021) 'A technical comparison of three low earth orbit satellite constellation systems to provide global broadband', *Acta Astronautica*, 159, pp. 123–135.

[12] Zhang, C., Jiang, C., Kuang, L. et al. (2021) 'Spatial-temporal resource allocation in LEO satellite-terrestrial integrated networks', *IEEE Journal on Selected Areas in Communications*, 39(11), pp. 3364–3382.

[13] Kua, J., Armitage, G. and Branch, P. (2020) 'A survey of rate adaptation techniques for dynamic adaptive streaming over HTTP', *IEEE Communications Surveys and Tutorials*, 19(3), pp. 1842–1866.

[14] Menascé, D. A. and Almeida, V. A. F. (2020) *Capacity Planning for Web Services: Metrics, Models, and Methods*. Upper Saddle River: Prentice Hall.

[15] Parekh, A. K. and Gallager, R. G. (2020) 'A generalized processor sharing approach to flow control in integrated services networks', *IEEE/ACM Transactions on Networking*, 1(3), pp. 344–357.

[16] Zhang, G., Li, Y. and Lin, T. (2021) 'Caching in information centric networking: A survey', *Computer Networks*, 57(16), pp. 3128–3141.

[17] Papapetrou, E. and Pavlidou, F. N. (2020) 'Analytic study of Doppler-based handover management in LEO satellite systems', *IEEE Transactions on Aerospace and Electronic Systems*, 41(3), pp. 830–839.

[18] Fenech, H., Amos, S., Tomatis, A. and Soumpholphakdy, V. (2021) 'High throughput satellite systems: An analytical approach', *IEEE Transactions on Aerospace and Electronic Systems*, 51(1), pp. 192–202.

10.48047/jocaaa.2022.30.02.77

[19] Chowdhury, P. K., Atiquzzaman, M. and Ivancic, W. (2020) 'Handover schemes in satellite networks: State-of-the-art and future research directions', *IEEE Communications Surveys and Tutorials*, 8(4), pp. 2–14.

[20] Al-Hraishawi, H., Chougrani, H., Kisseleff, S. et al. (2024) 'A survey on nongeostationary satellite systems: The communication perspective', *IEEE Communications Surveys and Tutorials*, 26(1), pp. 101–132.