maargin

WHITE PAPER · April 2026

# Designing a Frugal LEO Constellation *with Integrated Orbital Micro Data Centers*

Distributed by

ACCESS HUB

www.accesshub.world

**Aniruddh Rao Kabbinale**
Founder & CEO, Maargin

**Arjuna Sathiaseelan**
Advisor, Maargin

aniruddh@maargin.in | www.maargin.in

| 324 Satellites | $400–450M | 70–90% | 25 TFLOPS |
|---|---|---|---|
| LEO Constellation | Total CAPEX | Cost Reduction | On-Orbit Compute |

# Contents

## Executive Summary

This white paper presents the technical architecture and frugal innovation principles behind Maargin's low-cost Low Earth Orbit (LEO) satellite constellation, designed to deliver communications and orbital compute services primarily focused on India while offering global coverage.

Through strategic design choices — including higher-altitude LEO orbits at 1,000–1,200 km to minimise satellite count, low Size, Weight and Power (SWaP) optical communication modules, Software-Defined Radio (SDR) payloads, radiation-tolerant Commercial Off-The-Shelf (COTS) electronics with shielding, and integrated micro data centers forming a compute grid in space — the constellation delivers a fully integrated "frugal stack" that combines broadband connectivity and orbital computing in a single system.

Unlike mega-constellations such as Starlink, Project Kuiper, and OneWeb which optimise for global retail coverage at 300–600 km altitude, Maargin explicitly targets a higher-altitude, lower satellite count architecture coupled with distributed orbital micro data centers tailored to the needs and fiscal constraints of emerging markets.

The resulting 324-satellite, three-shell constellation is projected to require approximately **$400–450 million** in total capital expenditure for full deployment — representing a **70–90% reduction** relative to published cost ranges for comparable LEO broadband constellations. This budget enables secure national-scale connectivity for India and partner regions at a fraction of current mega-constellation costs, while providing on-orbit compute capacity suitable for AI/ML inference, low-latency edge services, and thermal-aware workload management.

**The Problem**

Mega-constellations cost $5–10B+, locking out emerging markets. Over 2 billion people still lack reliable internet.

**The Approach**

Frugal engineering at constellation scale — higher altitude, COTS hardware, SDR payloads, orbital AI compute.

**The Outcome**

National broadband + orbital edge compute for India and partners at $400–450M CAPEX, $55–70M annual OPEX.

## 1 Introduction

The global satellite communications industry has witnessed transformative growth over the past decade, driven primarily by large Low Earth Orbit (LEO) broadband constellations targeting global consumer and enterprise connectivity markets. Mega-constellations such as Starlink, Amazon's Project Kuiper, and OneWeb have pioneered vertically integrated architectures that combine mass-produced satellites, reusable launch vehicles, and sophisticated ground networks to deliver low-latency broadband at planetary scale. Starlink alone had deployed more than 9,300 satellites by early 2026, while Kuiper is ramping towards its planned 3,236 satellites and Eutelsat OneWeb operates roughly 650 LEO satellites [1, 2].

Despite their impressive technical achievements, these systems typically require capital expenditures in the range of $5–10 billion per constellation. High satellite counts, relatively short design lifetimes, and the need for continuous replenishment mean that annual operating expenditures can reach $1–2 billion for very large constellations. Such cost profiles are sustainable for well-capitalised global operators but create substantial barriers for emerging space companies and regional operators that wish to develop sovereign space-based connectivity capabilities [3, 4].

**India's space sector** offers a prominent example of how frugal principles can be applied at scale. The Indian Space Research Organisation (ISRO) has repeatedly demonstrated the ability to deliver complex missions at a fraction of the cost of comparable Western programmes. The establishment of the Indian National Space Promotion and Authorisation Centre (IN-SPACe) has opened ISRO infrastructure to private actors and enabled new forms of Public–Private Partnerships (PPP) [5, 6, 8].

A second macro-trend reshaping constellation design is the growing interest in **orbital computing and edge processing in space**. For broadband constellations, the convergence of connectivity and computing opens a dual-use design space: satellites can act simultaneously as communication nodes and as edge-compute resources within a distributed orbital fabric.

**Frugal Constellation Design**

A frugal methodology for LEO at 1,000–1,200 km explicitly tailored to the coverage and budget constraints of emerging markets.

**Dual-Use Architecture**

Broadband communications combined with a distributed layer of orbital micro data centers — connectivity and compute in a single system.

**Quantified CAPEX/OPEX Model**

Cost savings of 70–90% benchmarked against representative LEO broadband constellations for comparable regional coverage.

**Practical Deployment Roadmap**

Grounded in current Indian PPP mechanisms and recent national space policy for a phased, risk-managed deployment.

# 2 Frugal Innovation in Satellite Design

## 2.1 Frugal Innovation Framework

Maargin's constellation and payload architecture is organised around five system-level frugal design levers rather than isolated component optimisations:

**1. Orbit & Coverage Co-Design**

Operating at 1,000–1,200 km reduces satellite count for target coverage while accepting modestly higher latency, lowering both manufacturing and launch requirements at constellation scale.

**2. SWaP Minimisation**

Low-mass optical inter-satellite links and compact phased-array RF front-ends preserve link performance while compressing per-satellite launch mass by 40–50%.

**3. COTS Avionics with Shielding**

Screened commercial off-the-shelf electronics with fault-tolerant system design avoid expensive radiation-hardened parts — reducing avionics cost by 80–90%.

**4. Software-Defined Flexibility**

SDR payloads enable protocol evolution and workload reconfiguration in orbit without hardware redesigns.

**5. Integrated Orbital Micro Data Centers (Orbital Compute)**

Deploying **25 TFLOPS** compute nodes on ≈108 satellites creates a distributed orbital AI/ML fabric — enabling on-orbit inference, workload management, and edge processing without a separate satellite programme. Combined, all five levers support an estimated **order-of-magnitude reduction** in capital expenditure relative to commercial LEO broadband baselines, while maintaining the coverage, throughput, and latency characteristics required for national and regional connectivity services.

## 2.2 Low-SWaP Optical Communication Modules

Optical inter-satellite links (OISLs) enable high-speed data transfer between satellites without ground-based relays. Traditional modules weigh 15–25 kg; Maargin's design achieves **7–11 kg** — a 40–50% weight reduction [12] — through three innovations: (i) an efficient in-house Pointing, Acquisition, and Tracking (P&AT) mechanism that reduces structural and actuation overhead; (ii) shared coarse-alignment elements with the RF subsystem, avoiding dedicated optical alignment hardware; and (iii) operation at 1,550 nm (C-band near-infrared), balancing eye safety, atmospheric transmission, and component availability. The system targets 10 Gbps–100 Gbps bidirectional links over distances up to 4,000 km.

## 2.3 Software-Defined Radio Payloads

Traditional satellite communication systems employ hardware-centric architectures with custom system-on-chip designs, requiring 12–24 month development cycles and approximately $1 million in non-recurring engineering (NRE) per design [16]. Maargin adopts SDR payloads built on generic SoCs and RF chains.

| Parameter | Traditional Hardware | SDR Approach |
|---|---|---|
| Development time | 12–24 months | Weeks |
| Unit cost | $10,000 | $20,000 |
| NRE cost | $1,000,000 | Minimal |
| Flexibility | Fixed at launch | Reconfigurable in orbit |

SDR payloads trade a slightly higher unit cost for dramatically shorter development cycles, negligible mission-specific NRE, and full in-orbit reconfigurability — including frequency-band changes and protocol updates via software upload. The RF payload connects SDRs to phased-array antennas operating primarily in Ku-band with planned Ka-band capability.

### 2.4 Radiation-Tolerant COTS Electronics

Traditional solutions employ radiation-hardened parts costing up to $200,000 per satellite [17]. Maargin's approach uses screened COTS components with targeted shielding, reducing costs to **$20,000–40,000 per satellite** — an 80–90% cost reduction [12]. The strategy uses automotive- and medical-grade COTS parts manufactured on larger process nodes ($> 28$ nm), which naturally tolerate 30–50 krad or higher total ionising dose (TID), supplemented by board-level shielding and system-level fault management (redundant circuitry, error-correcting codes, watchdog mechanisms). At 1,000–1,200 km altitude, with appropriate shielding, the effective TID at critical components is kept below 10 krad (Si), consistent with screened COTS tolerance [17, 18].

## 3 Constellation Design

### 3.1 Orbital Configuration

Maargin's space network employs a multi-orbit, multi-shell design combining Walker-Delta and Walker-Star constellation patterns [9]. The Walker methodology, widely adopted in GPS and Globalstar, provides symmetric orbital geometries that minimise station-keeping requirements and fuel consumption [10, 11]. The three-shell configuration is summarised below.

| Shell | Region | Altitude (km) | Inclination | Satellites |
|---|---|---|---|---|
| Shell 1 — Tropical | India / Equatorial | 1,000–1,200 | 20–40° | 144 |
| Shell 2 — Temperate | UK / Europe | 1,000–1,200 | 53° | 144 |
| Shell 3 — Polar | Polar regions | 1,000–1,200 | $> 80°$ | 36 |
| **Total** | | | | **324** |

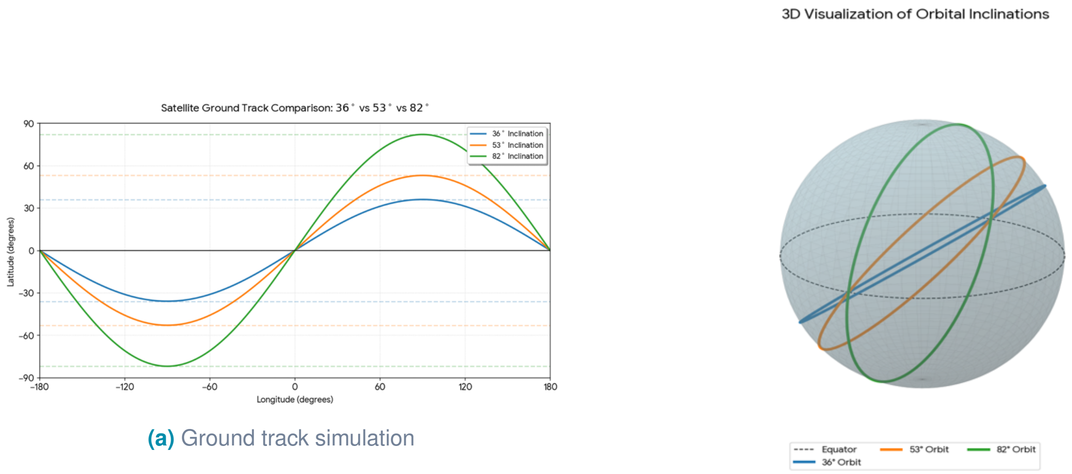

**(a)** Ground track simulation

**(b)** Orbital inclination geometry

**Figure 1:** Orbital attributes for inclinations of 36°, 53°, and 82°.

The constellation comprises 324 satellites: two inclined shells with 144 satellites each serving the Indian subcontinent and UK/European regions, and one polar shell with 36 satellites covering polar

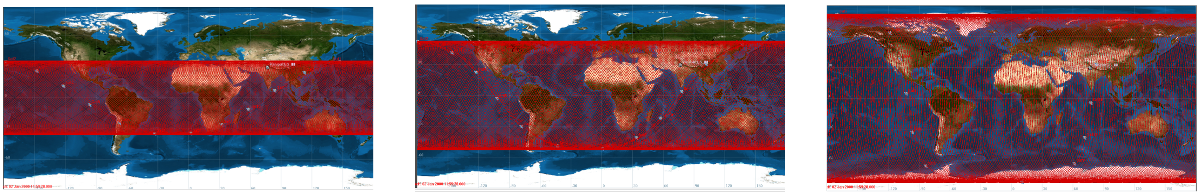

**(a)** 36° inclination **(b)** 53° inclination **(c)** 82° inclination

**Figure 2:** Ground tracks covered by the satellite constellation at various inclinations.

regions. At approximately 1,200 km altitude with a 25° ground station elevation angle, coverage extends to roughly ±15° latitude beyond the ground track edges.

### 3.2 Altitude Selection

The 1,000–1,200 km orbit height represents a strategic sweet spot balancing multiple engineering constraints. Moving from 300 km to 1,200 km reduces the satellite count required per shell from 1,296 to 144 — an 89% reduction — with dramatic effects on both manufacturing and launch cost.

> *"Moving from 300 km to 1,200 km reduces satellite count per shell by 89% — from 1,296 to 144 — cutting shell manufacturing cost from $520M to $60M."*

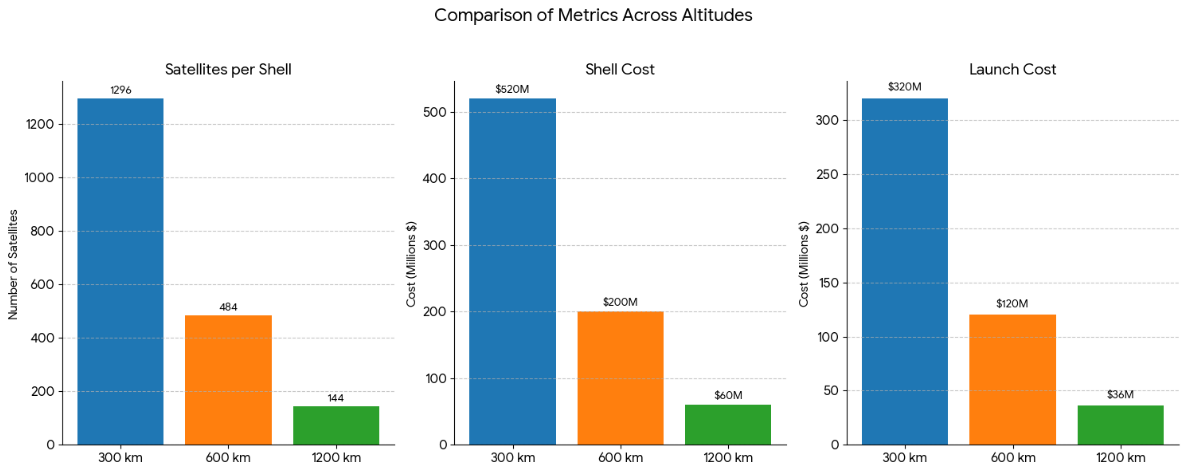

**Figure 3:** Constellation metrics (satellite count, shell cost, launch cost) across altitudes.

| Altitude | Satellites / Shell | Shell Cost | Launch Cost |
|---|---|---|---|
| 300 km | 1,296 | $520M | $320M |
| 600 km | 484 | $200M | $120M |
| **1,200 km** | **144** | **$60M** | **$36M** |

*Assumptions: satellite mass 250 kg, manufacturing cost $400,000/unit, launch cost $1,000/kg, 6-year satellite lifetime [12, 13, 14].*

Additional engineering benefits of the 1,000–1,200 km altitude include multi-year orbit stability without excessive station-keeping; manageable radiation exposure (total ionising dose reducible to below

10 krad (Si) at component level with appropriate shielding [17]); round-trip latency of 50–80 ms suitable for video conferencing, cloud services, and IoT; and reduced collision risk versus congested lower-altitude LEO bands.

### 3.3 Compact Satellite Design

Maargin targets a 200–250 kg satellite mass compared to 400+ kg for state-of-the-art broadband systems — roughly 50% mass reduction — with per-satellite manufacturing costs modelled at $400,000–500,000 versus approximately $1 million for 400 kg-class platforms.

**Indicative mass budget per satellite**

- **3–5 optical communication modules:** 35–50 kg total (7–11 kg each)
- **5–8 phased-array antennas:** 15–25 kg total
- **Compute payload** (on ≈one-third of satellites): 20–35 kg
- **Propulsion and bus:** remaining mass

**Average power:** 400 W (comms-only) | 600 W (with micro data center)

## 4 Micro Data Center Architecture

### 4.1 Why Orbital Computing?

Managing a 324-satellite constellation introduces operational challenges that traditional human-in-the-loop ground control cannot address efficiently. Autonomous anomaly detection across power, thermal, OISL, SDR, and attitude telemetry is required to pre-empt cascading failures. With 10 Gbps–100 Gbps optical mesh links carrying mixed traffic, the constellation also needs autonomous routing and downlink prioritisation.

Earth Observation (EO) satellites generate massive data volumes that create downlink bottlenecks when processed on the ground [20]. Processing data in orbit enables transmission of results rather than raw data, reduces latency, and enables time-sensitive applications to be handled near the collection point.

Maargin deploys micro data centers on approximately one-third of the constellation (≈**108 satellites**), creating a tightly connected compute layer built atop the optical ISL connectivity infrastructure.

| | | | | |
|---|---|---|---|---|
| 25 TFLOPS | 2 TB | 200 W | 20–35 kg | 108 |
| Compute per node | Persistent Storage | Avg. Power Draw | Module Mass | Compute Satellites |

## 4.2 Micro Data Center Specifications

**Per-unit micro data center specification**

- **Compute:** 25 TFLOPS
- **Storage:** 2 TB persistent
- **Mass:** 20–35 kg
- **Power:** 200 W average
- **Cooling:** 0.5 $\mathrm{m}^2$ radiator surface

**Primary workloads:** Autonomous anomaly detection, workload management and scheduling, communications protocol conversion and routing optimisation.

| Workload | Description |
|---|---|
| **AI/ML Inference** | Anomaly detection, traffic routing |
| **Comms Control** | Autonomous routing, downlink prioritisation |
| **Workload Management** | Intelligent scheduling and resource allocation across the orbital compute fabric |
| **Thermal Scheduling** | Intensive compute during eclipse for peak radiative cooling |

## 4.3 Thermal Management

Without atmospheric convection, satellites rely exclusively on thermal radiation for heat dissipation. For 200 W thermal loads from the micro data centers, Maargin employs:

- **Passive radiators:** 0.5 $\mathrm{m}^2$ high-emissivity panels per compute module.
- **Phase Change Materials (PCM):** absorb heat during high-power periods; release during eclipse.
- **Duty-cycle scheduling:** intensive compute scheduled during eclipse phases (external temperature $\approx -100$ to $-150°$C), when radiative cooling is most efficient.
- **Active thermal control:** heat pipes and pumped fluid loops transport heat from components to radiators [22].

Satellites spend roughly 30–35% of each orbit in Earth's shadow, providing a natural cooling window that Maargin's scheduler exploits. The 200 W-per-compute-module approach aligns with the distributed paradigm recommended in recent data-centre-in-space discussions [24] rather than attempting monolithic orbital facilities.

# 5 Ground Segment Infrastructure

Maargin's constellation relies on approximately 20–30 ground stations distributed across tropical, temperate, and polar regions. LEO constellations require architectures that differ substantially from GEO systems, with electronically steerable phased-array antennas capable of tracking fast-moving satellites, rapid handoffs every 8–12 minutes per pass, and high-speed fibre backhaul to internet Points of Presence [25, 26].

Maargin's higher operational altitude allows larger coverage footprints per satellite, reducing the required ground station count relative to lower-altitude constellations.

| System | Ground Stations | Antennas/Station | Antenna Size |
|---|---|---|---|
| Starlink | 150+ | 4–12 | 2.8–3.0 m |
| Amazon Kuiper | 300 (planned) | 4–10 | 5.4 m |
| OneWeb | 50+ | 10–15 | 2.4–3.5 m |
| **Maargin** | **28–30** | **8–10** | **2.5–3.0 m** |

### Ground station equipment

Each Maargin ground station includes Ku/Ka-band phased array antennas (uplink 27.5–30.0 GHz; downlink 17.8–19.3 GHz), high-speed fibre backhaul, gateway electronics for signal processing and routing, and interfaces for space situational awareness integration.

### Frequency Plan

- **Optical ISLs:** 1,550 nm for satellite-to-satellite links at 10 Gbps–100 Gbps
- **Ku-band:** 10.7–12.7 GHz downlink / 14.0–14.5 GHz uplink for user terminals
- **Ka-band:** 17.8–19.3 GHz downlink / 27.5–30.0 GHz uplink for gateways

# 6 Economic Analysis & Deployment Strategy

## 6.1 Capital Expenditure Breakdown

Total constellation deployment cost is estimated at approximately $400–480 million for 340 satellites (324 operational plus 16 on-orbit spares).

**Phase 1 — Shell 1**
**$200M**
Design, qualification,
first shell + ground segment

**Phase 2 — Shell 2**
**$100M**
Recurring manufacturing
& launch

**Phase 3 — Shell 3**
**$100M**
Full coverage
polar shell

| Cost Category | Amount (M$) |
|---|---|
| Satellite design, testing & qualification | 80–100 |
| Satellite manufacturing (340 units @ $400–500k) | 136–170 |
| Launch services (at $1,000–1,500/kg) | 80–100 |
| Licensing, spectrum, ITU filings, legal | 5–10 |
| Ground stations (30 stations + backhaul) | 60–100 |
| **Total CAPEX** | **400–480** |

This compares favourably with mega-constellation cost ranges: Starlink is estimated above $10 billion, OneWeb's Phase 1 required approximately $3.4 billion, and Amazon Kuiper's projected investment approaches $10 billion [4].

### 6.2 Comparative Cost Analysis

| Parameter | State-of-the-Art LEO | Maargin |
|---|---|---|
| Satellite mass | 400+ kg | 200–250 kg |
| Manufacturing cost / satellite | $1,000,000 | $300,000–400,000 |
| Launch cost / satellite | $400,000–600,000 | $250,000–375,000 |
| Shell deploy (manufacturing) | $520 M | $60 M |
| Shell deploy (launch) | $320 M | $36 M |
| **Cost reduction** | — | **88–90%** |

The cumulative effect of multiple frugal levers — higher-altitude shells with fewer satellites, low-SWaP optical ISLs, SDR-based RF payloads, radiation-tolerant COTS electronics, and integrated orbital micro data centers — delivers transformative cost reductions without proportional performance sacrifice.

### 6.3 Operational Expenditure

Annual operating costs post-deployment are estimated at $55–70 million per year. The 6-year satellite design life implies replacing approximately one-sixth of the fleet annually; at $400,000 manufacturing plus $250,000 in launch cost per satellite, this yields approximately $35 million per year in replenishment expenditure.

| Operating Cost Category | Annual Cost (M$) |
|---|---|
| Satellite replenishment (1/6 constellation, 6-year life) | 30 |
| Collision avoidance, command/control, insurance | 5–10 |
| Ground station operation and maintenance | 5–10 |
| Backhaul network connectivity and data centres | 10 |
| Miscellaneous operational costs | 5–10 |
| **Total Annual OPEX** | **55–70** |

### 6.4 Deployment Roadmap and Partnership Model

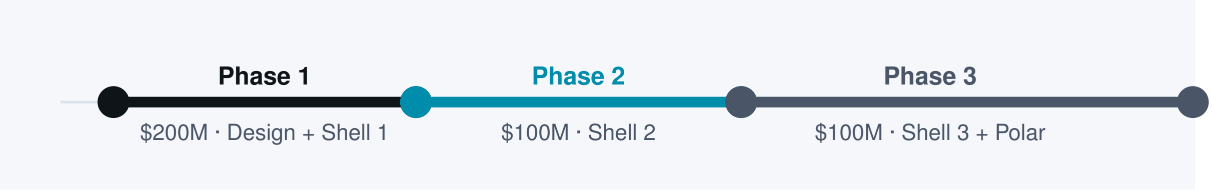

Phase 1 can be anchored by IN-SPACe and national space agencies using existing models for co-investment in LEO constellations and shared use of test and launch infrastructure. Phases 2 and 3 accommodate a wider set of private operators, regional partners, and service providers as risk is retired and demand is validated — a template adaptable to any emerging space nation.

## 7 Technical Performance

### 7.1 Coverage and Latency

At 1,200 km altitude with a 25° minimum elevation angle, individual satellites cover approximately 30° latitude swaths. The three-shell Walker configuration provides:

- **Low-inclination shell (20–40°):** Indian subcontinent and equatorial regions.
- **Medium-inclination shell (53°):** UK, Europe, and temperate zones.
- **Polar shell ($> 80^\circ$):** Polar regions and inter-shell connectivity.

Expected round-trip latency is **50–80 ms** (geometric path-length estimate for 1,200 km orbits), consistent with reported OneWeb figures at similar altitude. This is sufficient for video conferencing, cloud services, and IoT connectivity, though higher than Starlink's 20–40 ms at ≈550 km.

### 7.2 Data Throughput

| Optical ISL Backbone | Ka-Band Gateway | Concurrent Users |
|---|---|---|
| **10–100 Gbps** | **≈20 Gbps** | **200–400** |
| per link between satellites | typical throughput per satellite | per satellite at standard rates |

### 7.3 Constellation Reliability

System-level reliability mechanisms include:

- **16 on-orbit spare satellites** (≈5% overhead) for immediate replacement of failed units.
- **Optical ISL mesh multi-path routing** around failed nodes — the mesh topology provides multiple alternative data paths.
- **20–30 globally distributed gateways** eliminating single points of failure in the ground segment.
- **Annual fleet refresh** — replacement of ≈one-sixth of the constellation per year for continuous technology refresh.

The Walker constellation pattern inherently provides redundancy through symmetric satellite distribution, enabling seamless handoffs during individual satellite outages.

## 8 Risks and Mitigation

| Risk Area | Description | Mitigation |
|---|---|---|
| **Optical pointing** | Maintaining laser alignment over 4,000 km at 7 km/s | Closed-loop P&AT design; RF backup links during acquisition |
| **Thermal management** | Dissipating 200 W in vacuum with limited radiator area | PCM, duty-cycle scheduling, proven heat-pipe technology |
| **COTS radiation** | Single-event upsets and long-term TID degradation | Redundant architectures, ECC, watchdog timers, planned reboots |
| **SDR complexity** | Software verification and update risks | Ground testing, staged rollouts, known-good configuration rollback |
| **Space debris** | Growing object density at LEO <1,200 km | Conjunction data messages, automated avoidance, EOL de-orbit protocol |
| **Spectrum licensing** | ITU filings and national regulatory approvals | Early regulator engagement; use of established Ku/Ka bands |
| **Technology obsolescence** | Rapid evolution of space and comms markets | 6-year life with annual refresh; software-upgradeable payloads |

### 8.1 End-of-Life and Debris Mitigation

Operating at 1,000–1,200 km implies substantially longer natural orbital lifetimes than lower LEO shells, placing greater emphasis on robust post-mission disposal (PMD). Maargin targets a PMD reliability of at least 0.9 per satellite. Each satellite carries a propulsion system sized for EOL manoeuvres that lower the orbit to an atmospheric re-entry corridor within 25 years of mission end, consistent with prevailing international guidelines.

The EOL sequence includes: controlled passivation (propellant depletion and energy source safing); a sequence of de-orbit burns to reduce perigee into a sub-25-year decay regime; and final telemetry updates to space situational awareness providers. On-board autonomy can trigger pre-planned disposal sequences in response to persistent anomalies, and propulsion subsystem redundancy provides margins for component-level failures.

## 9 Conclusion and Future Directions

> *"Maargin proves that frugal engineering can deliver national-scale broadband + orbital AI compute for under $500M — opening the space economy to the nations that need it most."*

Maargin's LEO constellation with integrated micro data centers demonstrates that frugal innovation can deliver large cost reductions — 70–90% — relative to state-of-the-art mega-constellations while maintaining competitive technical performance. The architecture's total deployment cost of approximately $400–450 million for 324 satellites makes it accessible to emerging space nations and private entities, advancing the democratisation of space-based communications and computing.

The design operationalises frugality along five tightly coupled dimensions: orbit and coverage optimisation that trades frontier latency for a higher-altitude, lower-satellite-count configuration; SWaP-minimised payloads; a COTS-centric avionics stack with targeted shielding; constellation-level software flexibility through SDR; and integrated orbital micro data centers as the fifth lever — creating a dual-use compute and connectivity fabric that generates additional value from on-orbit inference, distributed workload management, and low-latency edge services.

Maargin's frugal LEO constellation delivers **70–90% CAPEX savings** versus state-of-the-art mega-constellations while maintaining competitive performance. The $400–450M total deployment cost makes sovereign space-based connectivity and orbital AI/ML compute accessible to emerging space nations.

The five-lever frugal architecture — orbit optimisation, SWaP-minimised payloads, COTS avionics with shielding, SDR-based software flexibility, and integrated orbital micro data centers — combined with a phased PPP deployment model provides a replicable template for any nation seeking sovereign LEO infrastructure.

## 9.1 Future Work Priorities

- **Prototype demonstration** of low-SWaP optical P&AT mechanisms and SDR payload integration in relevant orbital environments.
- **Micro data center workload scheduling** algorithms that exploit orbital thermal cycling to maximise compute efficiency under power and thermal constraints.
- **Inter-constellation cooperation** exploring federated computing meshes across multiple operators' satellite networks.
- **Next-generation gateway links** investigating E-band (60–90 GHz) for terabit-per-second-class backhaul.
- **AI/ML accelerators** for orbital edge inference and joint optimisation of constellation geometry, micro data center placement, and workload scheduling.

# References

## References

[1] C. Baugh, “SpaceX Starlink and Amazon Kuiper look set to dominate the satellite connectivity market,” Analysys Mason, July 2024.

[2] “Satellite Internet Technology,” *IJSRET*, vol. 11, no. 5, pp. 175–183, Oct. 2025.

[3] P. F. Bronson and B. G. Gladstone, “Schedule and Cost Estimating Analysis for LEO Satellite Constellations,” IDA Document NS D-33436, Apr. 2023.

[4] Comparison-of-Constellations, Maargin internal reference model, 2026.

[5] R. Radhakrishnan, “ISRO: The Frugal Frontier,” *Asian Scientist Magazine*, May 2016.

[6] “Experts Seek Fiscal Support and Critical Infrastructure Status for India’s Space Sector,” ETV Bharat, Jan. 2026.

[7] “Frugal Innovation: How Emerging Space Markets are Redefining the Space Race,” *Business Standard*, Mar. 2025.

[8] “IN-SPACe inks pact to build India’s first startups-led satellite constellation,” *Moneycontrol*, Jan. 2026.

[9] Maargin internal technical note on Walker Delta and Walker Star constellation configuration, 2026.

[10] “Satellite constellation — Walker constellation,” *Wikipedia*.

[11] “Communication Constellation Design — Minimum Number of Satellites,” arXiv:2410.03354, 2024.

[12] Frugal-Constellation modelling document, Maargin internal frugal constellation cost model, 2026.

[13] Maargin launch-cost model and price bands for current small-satellite launchers, 2026.

[14] “Emerging Technologies and Military Applications: Micro Satellite Constellations,” Concordia University, ca. 2020.

[15] “Optical Inter-Satellite Links,” MPB Communications, Dec. 2025.

[16] “Advantages of Software-defined Radio Payload,” YTTEK, Nov. 2024.

[17] “The road to high-performance and robust satellite compute,” Bloc Ventures, Oct. 2024.

[18] “Reduced SWaP, COTS in space, and budget constraints in the rad-hard world,” *Military Embedded Systems*, Jun. 2015.

[19] Maargin internal compilation of radiation-tolerant COTS vendor data sheets, 2025.

[20] R. Kumar et al., “Space Microdatacenters,” MICRO-56, 2023.

[21] “On-Orbit Data Centers: Mapping the Leaders in Space-Based AI Computing,” Cutter Consortium, Sep. 2025.

[22] “Super Cool Small Satellite Thermal Control System,” Advanced Cooling Technologies, Jun. 2025.

[23] “Top Strategies to Cool Down a Space Data Center,” Feb. 2026.

[24] “Elon Musk Has Grand Plans for Data Centers in Space. Experts Are Skeptical,” *PCMag*, Feb. 2026.

[25] “Ground Segment Consideration for LEO and MEO Satellites,” *Satellite Magazine*, Dec. 2016.

[26] I. del Portillo et al., “Optimal Location of Optical Ground Stations to Serve LEO Spacecraft,” IEEE, Mar. 2018.