DOCTORAL RESEARCH PROPOSAL · COMPUTER SCIENCE WITH AEROSPACE CO-SUPERVISION

Modular Orbital Data Centers: Replaceable-Node Architectures for Power, Heat, Networking, and Workload Placement

Two-page proposal · September 2026

1. Motivation and gap

Demand for large-scale digital compute is increasingly limited on the ground by electrical interconnection, water for cooling, land use, and permitting, not by the availability of accelerators. That pressure has moved orbital data centers (ODCs) from speculation to early flight. In November 2025 Starcloud-1 placed an NVIDIA H100 in low Earth orbit and reported both inference (Gemma) and a small on-orbit training run (NanoGPT on Shakespeare). In the same month Google published Project Suncatcher: solar-powered satellites carrying TPUs, free-space optical inter-satellite links (ISLs), and a dawn–dusk sun-synchronous orbit chosen for near-continuous illumination. Concept studies now speak of multi-megawatt tethered chains, integrated solar–compute–radiator panels, and even gigawatt-class arrays. These are *facilities*, not incremental improvements to a spacecraft computer.

The doctoral literature a computing or aerospace examiner would recognise does not yet treat that object. The mature thread is satellite edge computing (SEC): Magliarisi's orbit-aware, SGP4-informed scheduling for long-running services; Abreha's VNF mapping and on-board caching for 5G/6G non-terrestrial networks; Valente's resource allocation and in-orbit distributed learning for Earth observation. SEC assumes a largely fixed constellation of MEC nodes. It asks where to place a task, not how a plant grows by adding thermally isolated, optically networked, replaceable modules, nor how radiator area—not peak FLOPS—sets first-order capacity. Industry papers have started on the plant (formation packing, TPU radiation, tether mass budgets), but they are not yet an examined architecture with an open method, open assumptions, and a scaling law a third party can reproduce. This proposal sits in that gap.

2. Aim and research questions

Aim. To develop a reference architecture and scaling method for *modular* ODCs: how to partition compute, photovoltaic, and radiator functions across replaceable nodes so useful capacity can grow from a demonstrator-class spacecraft to the 1–100 MW regime without a new vehicle class each time.

1. **RQ1 (partitioning).** How should compute, solar collection, and heat rejection be allocated across modules so that adding a node increases facility capacity rather than merely adding bus mass and mutual thermal occlusion?
2. **RQ2 (interfaces).** What mechanical, electrical, thermal, optical, and software contract allows heterogeneous modules to join one facility, fail independently, and be spared or retired without taking the plant offline?
3. **RQ3 (growth path).** Under the same orbit, radiation, and launch-cost assumptions, which growth geometry—dense formation cluster, tethered node chain, or integrated panel array—minimises mass per useful watt at 1, 10, and 100 MW, and at what launch price ($/kg) does orbital placement undercut a terrestrial solar-plus-cooling plant for identified job classes?

Workload placement is a dependent question, not a fourth independent thesis: jobs are admitted only when they exploit data already in space, continuous SSO power, or terrestrial energy/water scarcity. Generic "move training to orbit" claims will be treated as hypotheses to falsify.

3. Approach

The unit of analysis is a *module* with a published interface, not a named company's satellite. Three published families supply the experimental settings: (i) dense km-scale clusters with optical ISLs (Suncatcher; 2026 cluster-geometry work mapping Clos-like fabrics onto formations); (ii) integrated solar–compute–radiator panels aimed at high watts per launched kilogram; (iii) tethered node chains with per-node radiative rejection on the order of kilowatts. Starcloud-1 is ground truth for COTS accelerator behaviour in LEO, not a MW-scale architecture.

Physics constraint. Vacuum admits no convection. Rejectable heat follows Stefan–Boltzmann, $P = \varepsilon \sigma A T^4$. For $\varepsilon = 0.9$, ideal one-sided flux is about 413 W/m² at 300 K and 1.3 kW/m² at 400 K: on the order of 2,400 m² of radiator per MW at silicon-friendly temperature, before Earth IR, albedo, atomic-oxygen coating damage, and cluster view-factor loss. Raising radiator temperature saves area and mass but collides with COTS GPU junction limits and with published radiation results (Google Trillium TPU tests: HBM as the weak link; NASA COTS GPU campaigns showing functional interrupts well before hard failure). The architecture must therefore co-design module size, radiator temperature, shielding mass, and software checkpointing. Heat, not interconnect marketing, is the primary figure of merit, expressed as m²/kW and kg/kW of sustained compute.

Method. (1) A system model coupling mass, solar input, radiative rejection, ISL topology, and a time-varying orbit (dawn–dusk SSO versus inclined LEO with eclipse). (2) A discrete-event simulation of module addition, failure, and sparing, with a thermal-aware scheduler that treats radiator headroom as a first-class resource. (3) One high-fidelity geometry per family (reduced-order thermal network; optical link budget using published 1.6 Tb/s-class ISL demonstrations where applicable). (4) An economic sensitivity on launch $/kg, downlink $/bit, and coating degradation over a five-year LEO life, recomputing commercial opex claims (including Starcloud's published 40 MW ten-year comparison) under transparent assumptions. No MW plant will be flown in this doctorate; validity is against physics, published mass budgets, radiation data, and Starcloud-1-class flight evidence. Hardware-in-the-loop is opportunistic (thermal-vac, optical bench, beam time), not a gate.

4. Expected contributions

C1. A reference modular ODC architecture: module types, interface contract, and failure domains that examiners in both computer systems and spacecraft engineering can inspect. C2. A scaling law stating when to add a module, densify a cluster, or raise radiator temperature, with m²/kW and kg/kW as primary metrics. C3. An open placement and economic model that states which job classes belong in orbit and at what launch price the comparison with terrestrial plants flips—treating vendor opex figures as claims, not results. The work is not another VNF-placement or EO-image scheduler with "orbital data center" in the title.

5. Work plan, supervision, and risks

Year	Output and pass/fail test
Y1	Critical survey that separates SEC from ODC; frozen system model and one geometry per family. *Pass:* W/kg and m²/kW for two architectures under the same orbit.
Y2	Core method (interface + scheduler, or fabric map coupled to thermal network). *Pass:* a peer-reviewed slice (systems, networks, or AIAA/IEEE aerospace).
Y3	Coupled evaluation with eclipse, ISL dynamics, failure, and launch-price sensitivity. *Pass:* results that change a design choice (module size, radiator temperature, packing).
Y4	Thesis: architecture, method, limits, and an examiner-readable argument across CS and aerospace.

Supervision. Primary in computer systems or distributed systems; co-supervisor in orbital mechanics or spacecraft thermal control. Without the second, the environment model cannot be examined. **Risks.** Flight data remain sparse: mitigate by bounding claims to physics and to published radiation/flight numbers. Commercial architectures will move: mitigate by keeping the interface and scaling law vendor-neutral. Scope creep into rad-hard microelectronics or 5G NTN backhaul: out of scope by RQ design.

6. Selected references

1. Google, "Exploring a space-based, scalable AI infrastructure system design" (Project Suncatcher), 2025.
2. IEEE Spectrum, "Why thermodynamics rules future orbital data centers," 2026.
3. Starcloud, Starcloud-1 mission (H100 in LEO; Gemma inference; NanoGPT training), launched 2 Nov. 2025.
4. D. Magliarisi, *Decentralized and Orbit-Aware Scheduling for Computational Tasks in Satellite Edge Computing*, PhD thesis, Sapienza University of Rome, 2025.
5. H. G. Abreha, *Optimized Service Provisioning for Next-Generation Applications with Satellite Edge Computing*, PhD thesis, University of Luxembourg, 2025.
6. F. Valente, *Design and evaluation of strategies to minimise operating costs and support in-orbit distributed learning within satellite networks for Earth observation*, PhD thesis, Sapienza University of Rome.
7. "Designing dense satellite clusters for distributed space-based datacenters," arXiv:2605.15335, 2026.
8. "Reduced-mass orbital AI inference via integrated solar, compute, and radiator panels," arXiv:2604.07760, 2026.
9. "Tether-based architecture for solar-powered orbital data centers," AIAA 2026-1237, 2026.
10. CNBC, "Nvidia-backed Starcloud trains first AI model in space," 10 Dec. 2025; CACM, "Datacenters go to space."

Institution, supervisors, and applicant details to be inserted. This document is a standalone two-page research proposal, not a company white paper.