Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Orbital vs. Terrestrial AI Computing: Lifecycle Carbon Break-Even Analysis and Workload Placement

Abdelbassette Chenna

a.chenna@univ-batna2.dz

University of Batna 2

Djallel Eddine Boubiche

Higher National School HNS-RE2SD

Research Article

Keywords: Sustainable computing, Orbital data center, Lifecycle carbon, Multi-criteria decision, Workload taxonomy, Radiative forcing

Posted Date: November 27th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11815555/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Orbital vs. Terrestrial AI Computing: Lifecycle Carbon Break-Even Analysis and Workload Placement

Abdelbassette Chenna[a,*], Djallel Eddine Boubiche[b]

[a] *LASTIC Laboratory, University of Batna 2, Batna, Algeria*
[b] *LEREESI Laboratory, Higher National School HNS-RE2SD, Batna, Algeria*
[*] Corresponding author E-mail address: a.chenna@univ-batna2.dz

Abstract

Orbital AI computing has been proposed as an alternative to terrestrial infrastructure, yet published assessments of its carbon footprint range from near neutrality to emissions an order of magnitude above ground-based systems. This paper develops a unified lifecycle carbon break-even framework and traces that divergence to two accounting choices: the launch-emissions boundary and the climate metric applied to rocket emissions. Applying the climate-equivalence multiplier to the combustion-related share alone reproduces both published positions within one model, separated by a ratio of 11.5, and places the break-even grid intensity between 0.06 and 0.70 kg CO_2/kWh. A two-stage placement rule then screens workloads for operational feasibility, from data provenance and latency, before weighing carbon competitiveness under Monte Carlo uncertainty propagation. A geometric correction follows: dawn-dusk orbits near 650 km are not eclipse-free. The open OrbitOrEarth calculator reproduces every reported figure and lets readers substitute their own assumptions.

Keywords

Sustainable computing; Orbital data center; Lifecycle carbon; Multi-criteria decision; Workload taxonomy; Radiative forcing

1. Introduction

The electricity and water demands of artificial intelligence are growing faster than terrestrial data centers can sustainably meet. Data centers and data-transmission networks already drew on the order of 1 to 1.5 percent of global electricity in 2021 and released roughly 0.9 percent of energy-related greenhouse gases. Their local burden is acute where grids are dirty or water is scarce. They account for about 7 percent of national electricity in Singapore and 21 percent in Ireland, and in the United States they consumed on the order of 620 billion liters of water in 2019 [1]. The current surge in AI training and inference is accelerating this trajectory sharply. US data-center electricity is projected to rise from 176 TWh in 2023, about 4.4 percent of national use, to between 325 and 580 TWh by 2028 [1]. Global data-center demand is expected to roughly double, toward 945 TWh by 2030 [2]. Projections of this kind have been revised downward before, when efficiency gains turned out to offset demand growth [3], so the trajectory matters more than any single figure on it. The binding constraint is increasingly the carbon intensity of the grid a data center can reach. This raises a question that until recently sounded fanciful: should some AI compute leave the grid entirely?

Several operators now propose exactly that. In orbit, sunlight is nearly uninterrupted and deep space is a vast cold sink into which waste heat radiates without water. Within two years, Google's Project Suncatcher (with radiation-tested accelerators and a sun-synchronous design point) [4], SpaceX, Starcloud, the EU-backed ASCEND feasibility study [5], and China's Star-Compute and Three-Body constellation [6] have moved orbital data centers from speculation toward engineering. These concepts converge on sun-synchronous low-Earth orbit.

In its dawn-dusk variant the spacecraft sees near-continuous sunlight, so operational carbon approaches zero, the property on which every optimistic assessment rests.

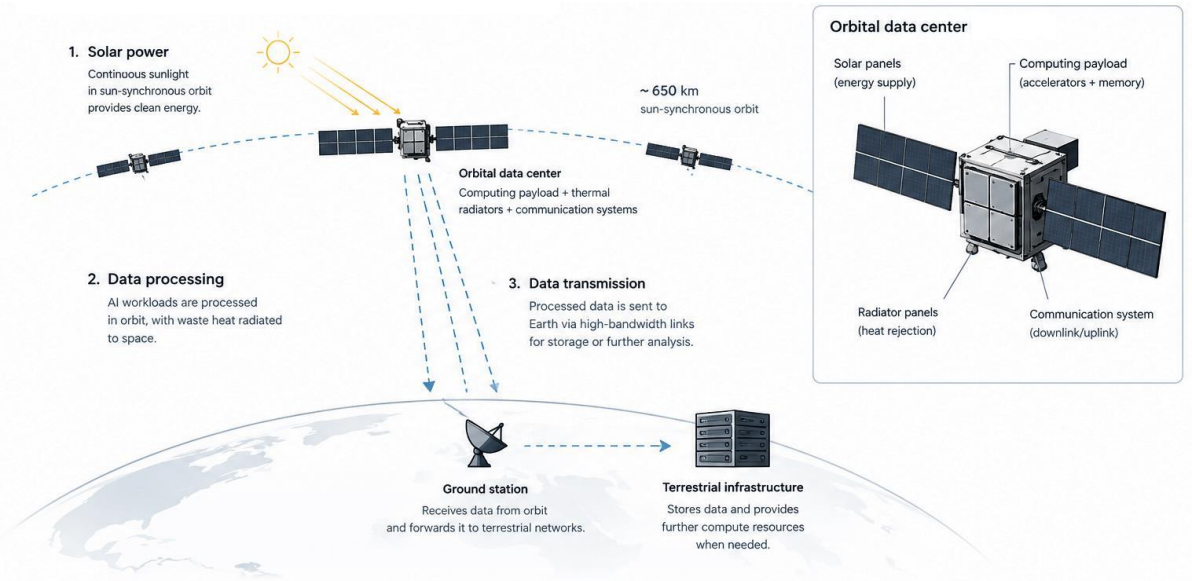


Fig. 1 *Orbital data-center deployment*

Yet the assessments that should guide such investments disagree by an order of magnitude. An optimist pole, in a Nature Electronics analysis from NTU, Zhejiang, and KAIST [7], holds that orbital data centers can be carbon-neutral within a few years of operation. That verdict uses a lifecycle carbon-usage-effectiveness metric in which operational carbon is zero. A pessimist pole, the "Dirty Bits in Low-Earth Orbit" study from Saarland University [8], finds that in-orbit systems are up to an order of magnitude worse than terrestrial equivalents, even under optimistic assumptions, once the embodied carbon of launch and re-entry is fully counted. Both rest on the same physics, both are peer-reviewed, and their conclusions are opposite. A decision-maker reading this literature cannot tell whether orbit is a climate solution or a climate liability.

We argue that this is not a genuine scientific contradiction. We reproduce both poles within a single per-kilowatt lifecycle model. Their disagreement reduces to a small set of technical parameters, chiefly the displaced grid intensity, how launch emissions are accounted, hardware lifetime, and radiator mass per kilowatt, and, decisively, to a choice of climate metric. Counting only the direct CO2 of propellant combustion places the break-even near 0.06 kgCO2/kWh, where orbit beats most operating grids. Counting the full CO2-equivalent raises it to roughly 0.70 kgCO2e/kWh, where orbit beats only the most carbon-intensive grids; the difference is the high-altitude radiative forcing that atmospheric-science work shows dominates rocket climate impact [9]. One specification choice matters here, and we make it explicit. The high-altitude forcing is produced by combustion products and not by the vehicle's supply chain, so the equivalence multiplier is applied to the combustion share of the launch factor alone. The accounting boundary and the climate metric therefore enter the break-even additively rather than multiplicatively (Section 4.1). Across the realistic range of system mass these choices trace three break-even contours (Fig. 2). The optimist and pessimist poles are the outer two, and the carbon-intensive grids on which AI actually runs fall in the metric-dependent band between them. The two poles are sampling opposite ends of one continuum, not contradicting each other.

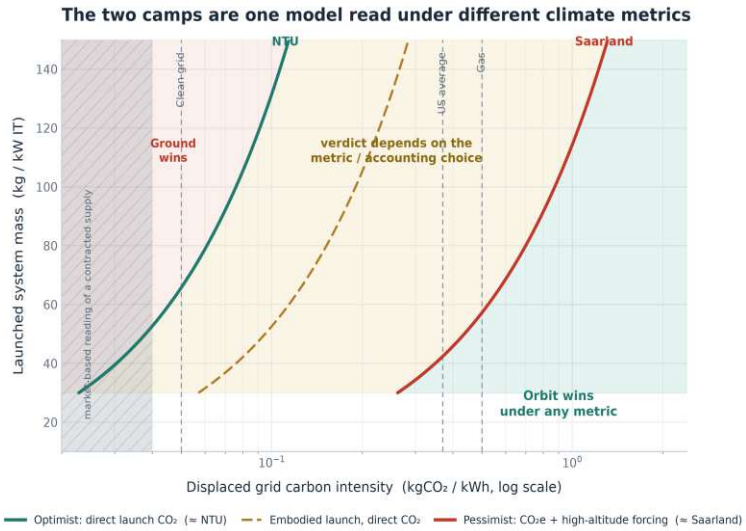


Fig. 2 *Break-even grid intensity versus launched mass under three climate metrics*

The framework also yields actionable guidance, and two cases stand out at opposite poles of the taxonomy. Processing Earth-observation data where it is captured, data born in orbit and under no strict latency budget, is the clearest setting for orbital placement. Its carbon margin rests on an avoided-downlink credit that we bound directionally but do not quantify. Interactive inference for ground-based users is the clearest setting against orbital placement, and it fails viability before its carbon is ever computed (full treatment in Section 5.4). What has been missing is a computing-side framework that adjudicates such cases under an explicit climate metric. Prior work performs lifecycle accounting, proposes orbital architectures, schedules and filters workloads across an orbital edge, or issues warnings, and one analysis places workloads between orbital and terrestrial clouds by carbon alone. None couples a workload's data provenance and latency tolerance to a stated accounting boundary and climate metric, and none has examined the disagreement itself as a metric choice.

This paper makes five contributions. C1, its spine, shows that the optimist-pessimist disagreement over orbital-compute carbon reduces to a small parameter set and, decisively, to the climate-metric choice (direct CO2 versus CO2-equivalent with high-altitude forcing), which reconciles two contradictory studies within one model. Applying the equivalence multiplier only to the combustion share makes the boundary and metric choices additive, and we report the resulting ratio in closed form. C2 is a per-kilowatt lifecycle break-even model with traceable, sourced parameters and Monte Carlo uncertainty propagation, which isolates the levers that determine the verdict. It performs no new primary measurement and relies throughout on figures published by others. Within that model we also correct a premise the orbital-compute literature treats as settled: the dawn-dusk reference orbit near 650 km is not eclipse-free, and Section 3.2 derives the eclipse season from orbital geometry and prices the duty-cycle loss it imposes. C3 is a two-axis workload taxonomy (data provenance crossed with latency class) together with a decision rule. The rule maps a workload, grid, launch vehicle, and metric to an orbit-or-ground recommendation, with a confidence label and a win probability derived from that propagation. We also report a negative result about our own rule: with the avoided-downlink credit left unquantified, the provenance axis changes no verdict in the evaluation of Section 5.5, so the operative gate there is latency alone. C4 evaluates carbon quantitatively together with water, latency, orbital debris, night-sky brightness, and stratospheric ozone. Those five enter directionally or ordinally rather than as quantified terms, and we report them as a dominance profile rather than a single weighted score. C5 is OrbitOrEarth, an open calculator that reproduces the break-even map, the sensitivity ranking, and the reconciliation of Section 5.2, letting either side replay its own assumptions.

The remainder of the paper is organized as follows. Section 2 positions the work against the optimist-pessimist spectrum and existing analyses. Section 3 defines the reference orbit and the per-kilowatt break-even model. Section 4 details the method and the sourced parameters. Section 5 reports the results: the break-even map, the metric-choice flip, the reconciliation of the two camps, the worked decisions, and a comparison with single-criterion placement rules. Section 6 discusses implications, including alternative orbital architectures, scale and cost, and Section 7 covers limitations and threats to validity. Section 8 presents OrbitOrEarth, the open deployment of the model as an interactive calculator, before Section 9 concludes.

2. Background and Related Work

Placing data centers in orbit has moved from a thought experiment to a set of funded engineering programs in roughly two years. This section traces that trajectory. It asks which programs are actually building toward orbital compute, why their published carbon assessments disagree so sharply, what physical constraint any such system must satisfy whatever its carbon accounting, and why that disagreement turns on a metric choice rather than a physical one. It closes by situating this paper's contribution against that body of work.

2.1. Real programs

Orbital compute is no longer speculative. Google's Project Suncatcher has committed to a dawn-dusk sun-synchronous orbit near 650 km and is radiation-testing its accelerators for a five-year mission [4]. SpaceX and Starcloud have announced convergence on the same sun-synchronous regime, drawn by its promise of continuous solar power. Neither has released a design description at the level of detail available for Suncatcher, so we treat their orbital regime as reported rather than sourced. We do not claim convergence across the whole field: only Suncatcher's orbit rests on programme documentation, and the Chinese constellation described below includes low-inclination shells. That convergence is not a coincidence. It reflects a shared engineering logic, explored further in Section 2.3, in which uninterrupted sunlight and radiative-only heat rejection jointly favor this orbital family over any alternative. China's Three-Body and Star-Compute constellation follows a related but distinct logic.

Rather than chasing continuous solar exposure, it is described as an Earth-observation computing system spread across dawn-dusk, sun-synchronous, and low-inclination orbits between 500 and 1000 km. We take that description from public program reporting rather than from an engineering source. The legal analysis of mega-constellations we cite alongside it [6] establishes the regulatory context in which such systems are being deployed, not their orbital parameters, and we do not rely on it for any technical quantity. We return to the born-in-orbit case in Section 5.4 in generic form, as an Earth-observation constellation rather than as a study of any one program, because the taxonomy turns on the origin of the data in orbit rather than on the details of a particular architecture. Taken together, these programs establish that the question this paper asks is no longer hypothetical: real capital is being committed to orbital data centers before the field has settled whether they help or hurt the climate. That unsettled state is the subject of the next section.

2.2. The carbon controversy as a spectrum, not a binary

The literature evaluating these programs has produced estimates spanning a full order of magnitude, and it is tempting to read this as two camps talking past each other. We argue instead that the optimist and pessimist positions are the extreme ends of a single continuum, one that Contribution C1 makes explicit by parameterizing the choices that separate them.

The optimist pole is set by Aili, Choi, Ong, and Wen (NTU, Zhejiang, and KAIST) [7], whose Nature Electronics Perspective proposes orbital edge and cloud data-center frameworks capable of carbon neutrality. Their case

rests on a lifecycle carbon-usage-effectiveness metric under which operational carbon is essentially zero, paired with a carbon-aware multicloud scheduler that places workloads across orbital and terrestrial sites by that metric alone. The pessimist pole is set by Ohs, Stock, Schmidt, Fraire, and Hermanns (Saarland University), whose "Dirty Bits in Low-Earth Orbit" study and accompanying ESpaS accounting tool reach the opposite conclusion. Even under optimistic assumptions, they find that in-orbit systems emit up to an order of magnitude more CO2-equivalent than terrestrial equivalents once embodied launch and re-entry are counted [8].

Between these poles sits a more heterogeneous middle ground. The ASCEND feasibility study, run by Thales Alenia Space under EU Horizon Europe funding, reaches a favorable verdict only conditionally: it presupposes a launch vehicle roughly ten times less emissive than those flying today [5]. A separate tether-based architecture contributes a detailed mass budget and a passive attitude-control scheme without taking a strong position on the carbon question either way [10]. We draw additionally on Marcy's assessment of the night-sky brightness produced by kilometer-scale orbital arrays near 500 km [11], and on the FAA's environmental assessment of the Falcon launch vehicle for the launch-emission factors used throughout our model [12]. None of this middle-ground work, however, resolves the optimist-pessimist gap; it simply occupies different points along it without asking why the endpoints diverge in the first place. Sections 2.3 and 2.4 take up that question.

A distinct and older body of computing work sits alongside this controversy and is easy to miss, because it does not price carbon at all. Orbital edge computing proposes nanosatellite constellations as a class of computer system in their own right, with inference performed on the spacecraft so that only results, not raw pixels, are downlinked [13], [14]. Subsequent work adds task filtering and scheduling that decide, per image and per pass, what is worth computing and what is worth transmitting [15]. Companion studies ask whether general compute in orbit is sensible at all [16], how downlink capacity can be pooled across a constellation [17], how serverless abstractions might be offered from low Earth orbit [18], and what latency and variability a real low-Earth-orbit service actually delivers [19]. This line matters here for two reasons. First, it establishes that placement decisions conditioned on data provenance and on downlink and latency budgets are not new; what is new in the present work is coupling them to an explicit climate metric. Second, it supplies the measured latency figures we adopt in Section 4.4 in preference to the geometric propagation floor. What it does not supply is any carbon criterion, which is why it cannot by itself adjudicate the disagreement Sections 2.2 and 2.4 describe.

2.3. The thermal wall

Whatever its carbon verdict, any orbital data center must first survive a physical constraint that has nothing to do with accounting. On the ground, a data center rejects waste heat mainly through convection: air or water carries heat away, and mechanical chillers or evaporative cooling towers do the rest. In vacuum, neither option exists. A spacecraft can lose heat only by radiating it away. That process is governed by the Stefan-Boltzmann law, in which radiated power scales with the fourth power of surface temperature. At a fixed rejection temperature the required radiator area is therefore linear in the heat load: doubling the heat a system must shed doubles the area required to shed it. What makes that linear coefficient large is the temperature at which the heat must be rejected. Here the fourth-power law cuts in the unhelpful direction, because electronics tolerate only a limited operating temperature before performance and reliability degrade. A black surface at 300 K radiates on the order of 0.46 kW per square meter. A kilowatt of continuous waste heat therefore requires on the order of two square meters of radiating surface, before margins, view factors, and solar and Earth-albedo backloading are counted. The practical consequence is that radiators, not solar arrays and not the compute payload itself, become the single largest mass item in an orbital data center's power system.

This is not a minor engineering detail. It is a structural fact that shapes every design surveyed in Section 2.1. Spacecraft thermal-control practice treats radiator area, and the structure that carries it, as a first-order driver of the power-system mass budget for high-power platforms [21]. Published small-spacecraft technology surveys

show the same pressure at the bus scale on which the current orbital-compute programs are being flown [22]. That share tends to grow, not shrink, as the power delivered per satellite increases, because heat generation scales with compute throughput while the deployable radiating area a given bus can carry does not scale as freely. The contrast with terrestrial practice is instructive. Reviews of thermal management for ground-based facilities catalog a wide range of convective, liquid, and evaporative strategies precisely because air and water remain available as heat sinks, an option orbital design cannot draw on at any power level. This is part of why the programs in the previous section converge on similar orbital regimes and similar power budgets rather than simply scaling up. Past a certain power level, the radiator mass required to keep pace with waste heat grows faster than the payload it is meant to support, so there is a practical ceiling on how much compute a single platform can carry before thermal management dominates the mass budget.

For the purposes of this paper, we treat the thermal wall as an exogenous physical constraint rather than something we model from first principles. Reproducing detailed thermal design is outside our scope, and doing so would not change the conclusions we draw. Instead, we carry the radiator-mass constraint into the launched-system-mass parameter M used throughout our break-even model (Section 3.4). The reasoning is that any orbital data center serious enough to be worth evaluating will already have absorbed this constraint into its mass budget. We record M as an envelope rather than as a measured value: its range spans the published mass budgets of the programs surveyed above together with the design practice cited in the previous paragraph, and Table 2 labels it accordingly. This matters for how the rest of the paper should be read. The thermal wall bounds the physical side of the calculation, which is how heavy a system must be to deliver a kilowatt of compute in orbit. The next section shows that the metric controversy bounds the accounting side, which is how that mass is priced in carbon terms. The two constraints are independent of each other, and conflating them, treating a metric disagreement as if it were a physical one, is precisely the confusion this paper sets out to dissolve.

Two numbers bound what that envelope can be read to mean. At the 300 K reference the radiating surface is 2.18 square meters per kilowatt, so a floor of 30 kg/kW implies a radiator, its structure and its deployment mechanism together no heavier than about 13.8 kg per square meter if the radiator dominates the mass. We have no sourced areal density for a deployable radiator at this scale, so that floor is a declared envelope rather than a physical limit, and a lighter system cannot be excluded on thermal grounds alone. The rejection temperature and the backload move the area further than the areal density does. Raising the rejection temperature to 320 K cuts the required area to 1.68 square meters per kilowatt, while an emissivity of 0.9 against a 250 K effective sink raises it to 4.67, which at the same floor would imply about 6.4 kg per square meter. We hold the 300 K black-surface case throughout and record the sensitivity here, because the floor it supports enters every interval reported in Section 4.3.

One external anchor exists for the mass term, though not at our scale. The most detailed public assessment of megawatt-class orbital power is a system study of space-based solar power, which carries mass, launch and cost budgets for gigawatt-class platforms in geostationary orbit [23]. Its architecture is not ours, and its orbit is not ours, so we do not import its figures; we record it here as the nearest existing reference point for anyone who wants to test our envelope against a worked design.

A bottoms-up budget shows where that envelope sits, and we give one as an illustration rather than as a design. Taking the design case of the previous paragraph, 4.67 square meters per kilowatt at an areal density of 6 kg per square meter, the radiator and its structure are about 28 kg/kW. An array and its deployment at a raw-to-useful ratio of 1.15 add about 7, a battery sized for the eclipse season of Section 3.2 adds 7 to 13, power management and distribution about 8, avionics, attitude control and communications about 5, propulsion dry mass with disposal propellant about 12, and primary structure and mechanisms about 18. With a twenty percent system margin the total is near 105 kg/kW, which would place the three poles at 0.080, 0.201 and 0.924 rather than

0.061, 0.152 and 0.700. We keep 80 kg/kW as the central value because it is the figure the released artifact and the published programme descriptions support, and we report the rescaling because every threshold in this paper is strictly proportional to it: a reader who prefers the budget above should multiply the reported intensities by 1.3.

2.4. The carbon controversy is also a metric controversy

The thermal wall explains why orbital systems are heavy; it does not explain why estimates of their carbon impact disagree by an order of magnitude. That disagreement is rooted not in physics but in a largely invisible metric choice. Direct CO2 is not the dominant climate forcer of a rocket launch: black carbon and alumina particles injected directly into the stratosphere dominate the radiative forcing, while CO2 itself is a minor contributor [9]. Once this is recognized, the optimist-pessimist split resolves into a difference of accounting boundaries rather than a difference of underlying fact. The optimist pole counts only operational carbon, near zero in a dawn-dusk orbit, and amortizes launch emissions over the satellite's full operational life. The pessimist pole counts embodied launch and re-entry emissions under a CO2-equivalent metric that explicitly prices the high-altitude forcing the optimist leaves out. Contribution C1 formalizes exactly this distinction inside a single model: the two literatures are not contradicting each other about the physics of rocket launches; they are answering different questions about which of that physics to price. The empirical side of that disagreement has since been documented from several directions. An earlier model study of a proposed reusable launch system reached the same conclusion at a launch rate well below today's projections [24], a review of launch emissions reaches it from the atmospheric inventory side [25], and a life-cycle assessment of reusable launcher fleets finds the climate impact dominated by exhaust emissions, with earlier assessments underestimating it by orders of magnitude where high-altitude emissions were not characterised [26]. That last study is the empirical counterpart of the axis we formalise: it reports the launch factor itself as the dominant term, which is what makes the convention for pricing that term decisive rather than cosmetic. The state of the art on life-cycle assessment in the space sector is surveyed in [27].

2.5. Differentiation

Table 1 places the present work against the prior literature surveyed above. It is organized along five columns: whether the work performs lifecycle CO2 accounting, whether it proposes a physical architecture, whether it contains a decision rule for placing workloads, whether it evaluates multiple criteria rather than carbon alone, and whether it releases an open artifact. Reading across those columns clarifies both what exists already and where the actual gap sits.

The NTU/Nature Electronics work performs lifecycle accounting, through its carbon-usage-effectiveness metric, and proposes a partial architecture, and its carbon-aware multicloud scheduler places workloads between orbital and terrestrial sites. That placement rule uses carbon as its sole criterion, with no account of a workload's data provenance or latency sensitivity, and the work releases no open artifact. The Saarland/Dirty Bits study performs lifecycle accounting as well, and its ESpaS tool is openly released, which makes it valuable as a reproducible baseline. It proposes no architecture and contains no decision rule, since its purpose is measurement rather than placement, and it evaluates carbon exclusively. The orbital-edge-computing line surveyed in Section 2.2 is the converse again. It is rich in placement and scheduling logic, and that logic is already conditioned on data provenance and on downlink and latency budgets, but carbon does not appear in it as a criterion, so it cannot adjudicate the question this paper asks. The tether-based architecture is strong on physical design, with a detailed mass budget and attitude-control scheme. Its lifecycle accounting is only an envelope estimate rather than a full model, it offers no decision rule, and it is not released as an artifact. ASCEND sits closer to a complete system study, with both lifecycle accounting and an architecture, and it evaluates carbon alongside water as a partial

multi-criteria treatment. It still contains no decision rule mapping workloads to placement, and its results are not released in an open, reproducible form.

Table 1 Prior orbital-compute studies

Prior work	Lifecycle CO2	Architecture	Decision rule	Multi-criteria	Open artifact
NTU / Nature Electronics (optimist)	Yes (CUE)	Partial	Carbon-only placement	No	No
Saarland / Dirty Bits, ESpaS (pessimist)	Yes	No	No	No (carbon only)	Yes (ESpaS)
Tether architecture	Envelope only	Yes	No	No	No
ASCEND (Thales / EU)	Yes	Yes	No	Partial (carbon, water)	No
Orbital edge computing line (OEC)	No	Yes (nanosatellite OEC)	Yes (provenance- and downlink-aware scheduling)	No (carbon absent)	Yes (simulators, datasets)
This work	Yes (per-kW break-even; no new primary measurement)	No	Yes (workload-conditioned, metric-explicit)	Partial (carbon quantitative; water, latency, debris directional)	Yes (OrbitOrEarth)

None of these bodies of work, in other words, does what a decision-maker actually needs. That need is a rule that takes a workload's characteristics as input and returns a placement recommendation across more than one criterion, under a stated climate metric, backed by something a reader can independently rerun. The orbital-edge-computing line supplies the placement machinery but prices no carbon; the accounting studies price carbon but place no workloads. This work deliberately trades away two of the five columns to gain the other three. We perform no new primary measurement and propose no new architecture: our lifecycle accounting is a per-kilowatt break-even model built on figures published by others (Section 3) rather than a new measurement campaign. What we contribute is a decision rule that is explicitly workload-conditioned and metric-explicit, a multi-criteria evaluation that carries water, latency, and debris alongside carbon rather than carbon alone, and an open, reproducible calculator, OrbitOrEarth, that implements both. The gap this fills is not a gap in measurement or in engineering design, both of which the prior work covers reasonably well. Nor is it a gap in workload placement, which the orbital-edge-computing line covers on non-carbon criteria. It is a gap in translating a contested carbon measurement into a decision, which is where a builder or a policymaker actually needs help.

3. System Model and Assumptions

Every later result rests on the model set out here. This section gives the per-kilowatt lifecycle accounting, the reference orbit, and a precise climate-metric definition. It then tabulates each parameter with its sourced range, derives the break-even grid intensity, and defines the multi-criteria comparison that the decision rule will use.

3.1. Per-kilowatt lifecycle model

We compare one kilowatt of delivered compute capacity in orbit against the same kilowatt on the ground, over the useful lifetime of the hardware, and ask at what grid carbon intensity the two break even. The orbital side carries the carbon of placing and sustaining that kilowatt in space: the launch (propellant combustion and, optionally, the embodied carbon of manufacturing the vehicle), the satellite bus, the radiator, and the solar array.

Against this we credit the terrestrial data-center construction avoided by not building the equivalent capacity on the ground. That credit is an avoided emission rather than an emission of either side, and corporate reporting keeps such quantities outside the inventory itself; we net it inside the orbital numerator as a modeling convenience and report its leverage in closed form in Section 7. The terrestrial side carries the operational carbon of powering and cooling that kilowatt from the grid over the same period, scaled by the facility's power-usage effectiveness.

The embodied carbon of the compute accelerators themselves is assumed equal on both sides and cancels, an assumption we revisit in Section 7; for computing hardware that term is large enough to matter if the two sides ever diverge [28]. The functional unit is one kilowatt of installed capacity rather than one unit of delivered work. That is a declared choice, and it is not the unit an ICT lifecycle standard would default to. It matters because the two sides do not necessarily sustain the same utilization or the same refresh cadence. The decision variable is the break-even grid intensity I*: above it, orbit emits less; below it, the ground does. Two accounting conventions govern that comparison, and neither is neutral. The first is greenhouse-gas scope, which is defined relative to a reporting entity, so the answer changes with who is reporting. For a cloud operator buying launch as a service, grid electricity is Scope 2 and launch is Scope 3. Moving compute to orbit then converts a target-binding liability into a loosely bound one, for reasons of accounting boundary rather than of physics.

The scope assignment depends on the operator. For a vertically integrated operator that flies its own vehicles, launch is Scope 1 and that asymmetry disappears or reverses. Our model is scope-agnostic by construction: it counts both sides in absolute emissions and is silent on the category. A reader applying it inside a corporate target should be aware that the incentive the framework quantifies and the incentive an operator faces need not coincide. The second convention is the basis on which the displaced grid intensity is measured, and it carries more leverage than the metric axis this paper is otherwise about. Location-based accounting uses the physical grid average of the region, which is what our tabulated intensities represent. Market-based accounting uses the intensity of the contracted supply. For an operator holding renewable power-purchase agreements, that intensity is at or near zero under annual volumetric matching.

It is in the region of a lifecycle renewable intensity, a few hundredths of a kilogram per kilowatt-hour, if the contracted generation is itself assessed on a lifecycle basis. Either reading falls below our optimist break-even of 0.06, so under market-based accounting the ground wins under all three contours, unconditionally and without reference to the metric axis. The Scope 2 Guidance of the GHG Protocol requires both figures to be reported rather than one [29], and the market-based method is itself contested, on the ground that buying a certificate need not change physical generation [30]. The instruction is therefore not to fix the basis but to state it and run the comparison twice. We report the location-based case throughout because it tracks physical emissions. Two refinements we do not model would narrow the gap: a partially contracted operator reaches a residual-mix intensity rather than zero, and hourly rather than annual matching tightens what a contract can claim. Section 6 returns to this, and Table 2 carries the basis as a declared scenario axis.

3.2. Reference orbit

The reference regime is a dawn-dusk sun-synchronous low-Earth orbit. The nominal altitude is near 650 km, following Suncatcher, the most concrete program, and altitude is treated as a qualitative sensitivity axis across the 500 to 1600 km band. Continuous sunlight is what drives operational carbon toward zero. The strength of that premise needs stating more carefully than the orbital-compute literature usually does, because it does not hold unconditionally at our own reference altitude. Eclipse-free operation requires the beta angle to exceed a threshold set by the Earth's angular radius, which is 68.0 degrees at 500 km, 65.2 degrees at 650 km and 59.8 degrees at 1000 km. A sun-synchronous orbit is retrograde, with an inclination near 98 degrees at our reference altitude. For a dawn-dusk plane the beta angle follows the solar declination through the year, as the arcsine of

the sine of the sum of inclination and declination [31]. It reaches about 82 degrees at the equinoxes and falls to about 58.6 degrees at the June solstice, which is 6.6 degrees below the 65.2 degree threshold at 650 km. The reference orbit is therefore not eclipse-free, as Fig. 3 shows. It carries an eclipse season of roughly ninety days per year, reaching about twenty percent of an orbit in shadow at the worst, and averaging near four percent solar unavailability over the year. Eclipse-free dawn-dusk operation requires roughly 1,395 km if the criterion is the umbral shadow and about 1,443 km if the penumbra is counted, so the crossover is better reported as a band of 1,390 to 1,450 km. Either figure lies well above our reference, and Section 6 compares the alternatives.

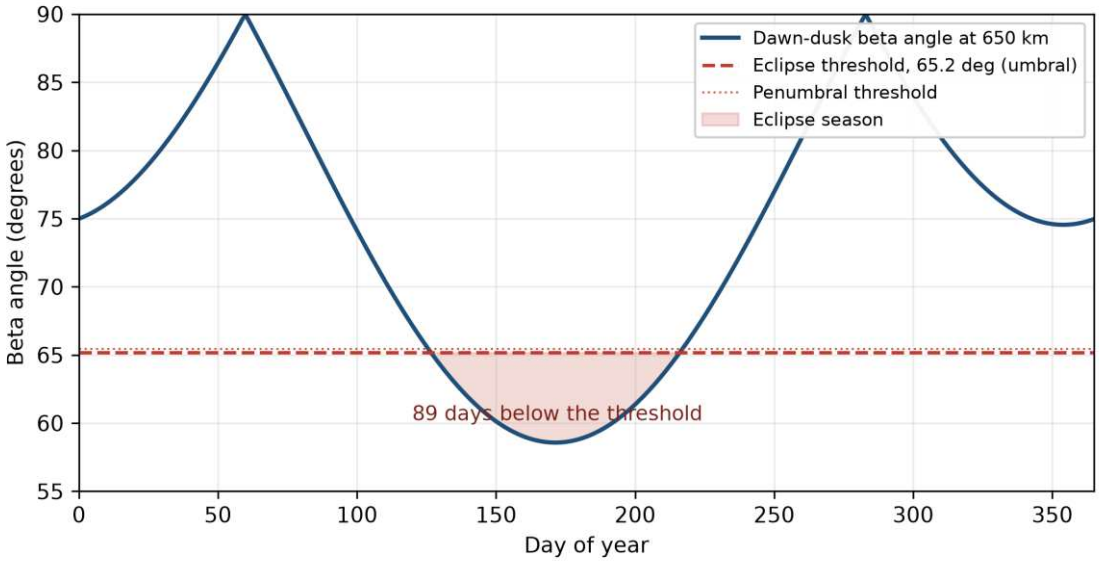


Fig. 3 *Dawn-dusk beta angle at 650 km through the year against the eclipse threshold*

The consequence for the carbon accounting is modest but real, and it runs against orbit. Treating the four percent as lost duty cycle raises the pessimist break-even from 0.700 to 0.729 kgCO2e/kWh. Meeting it with storage instead would add battery mass we do not carry in M. A worst-case eclipse of about twenty minutes needs roughly 0.33 kWh per kilowatt. At space-qualified pack energy densities, and at a depth of discharge chosen for five-year cycle life, that is of order 7 to 13 kg/kW, a tenth of the reference mass or more. Neither path is free, and the one we model is the duty-cycle path. We therefore state the premise as near-continuous rather than continuous illumination. An orbital-compute design that assumes literally zero operational carbon at 650 km is assuming something the geometry does not supply. The altitude band matters through four further channels: delta-V sets the launch factor and hence the numerator of Equation (1); the stratospheric portion of the ascent conditions the black-carbon forcing the multiplier prices; radiation dose sets the achievable hardware lifetime; and slant range sets the latency that gates stage one. This near-continuous illumination is the orbital analogue of a question already studied on the ground. Terrestrial siting studies have examined how far latitude and local renewable availability can push a facility toward comparably low operational carbon. They find that even favorable high-latitude sites cannot fully escape the seasonal and diurnal variability that a high dawn-dusk orbit escapes. We do not carry altitude as a numerical parameter in Table 2, and Section 7 records that as a limitation.

3.3. The metric definition

The launch factor f hides the paper's central choice, which we make explicit along two orthogonal axes. The first is the accounting boundary: f may count only the direct CO2 from propellant combustion, the optimist boundary, or it may add the embodied carbon of vehicle manufacturing and propellant production. The second is the climate

metric: the launch may be valued in direct CO2, or in CO2-equivalent that includes the high-altitude radiative forcing rockets uniquely produce. That forcing is dominated by black carbon injected into the stratosphere, with carbon dioxide itself a minor contributor [9], [33]. The climate and ozone response to that black carbon has since been modeled directly [34]. Converting it into a CO2-equivalent introduces a multiplier that we treat as an explicit scenario axis (Section 3.4) rather than a fixed constant. That multiplier is defined per unit of combustion CO2. It is built as a propellant-level black-carbon emission index times a black-carbon global warming potential, divided by the CO2 released per unit of propellant, so it applies to the combustion term in full rather than to the fraction of that term deposited above the tropopause. The derivation deserves printing rather than leaving the multiplier as a bare value. A black-carbon emission index of 0.030 to 0.050 kg per kg of kerosene, a black-carbon global warming potential of about 900 on a hundred-year horizon [35], and about 3.0 kg of CO2 per kg of kerosene give a multiplier between 9 and 15, which is why we take 10 as the central value. Two caveats follow from that construction. The potential is a tropospheric one, whereas stratospheric residence is longer and the implied potential larger, so the sweep ceiling of 50 in Table 5 is more likely to be low than high; the closed form gives 6.18 kgCO2e/kWh at a multiplier of 100. And the index is an empirical quantity specific to the propellant, so a methalox or hydrolox vehicle needs its own value rather than a rescaling of ours. Section 7 reports what the opposite convention would give. The two axes are not interchangeable, and keeping them separate requires one further decision that prior treatments leave implicit. The launch factor is not homogeneous. Part of it is combustion, released along the ascent trajectory and partly into the stratosphere; part of it is ordinary ground-level industrial CO2 from building the vehicle and producing its propellant. Only the first carries the high-altitude forcing the multiplier prices.

We therefore split the launch factor into a combustion component and a non-combustion component, and we apply the multiplier to the combustion component alone. Weighting the whole embodied factor by a black-carbon multiplier would price a factory's stack emissions as though they had been injected at 20 km, which no reading of the atmospheric-science literature supports. One boundary of this split must be stated, because it cuts against us. The non-combustion component as we define it is the vehicle supply chain, and it is metric-neutral for the reason just given. It is not the only non-combustion term a full accounting would carry. Re-entry ablation deposits spacecraft metals in the stratosphere, which is both non-combustion and high-altitude, so it belongs to neither of our two components. Those metals are no longer hypothetical: they are measured in stratospheric aerosol particles, in proportions that match spacecraft alloys [36]. Our model has no re-entry term at all. Where we reproduce the pessimist pole in Section 5.2 we therefore reproduce a boundary narrower than the one that study actually adopted, and our 0.70 kgCO2e/kWh should be read as a lower bound on it. We treat this as a limitation rather than folding in an unsourced figure, and Section 7 records it alongside the other terms we leave directional.

One point of consistency must be stated, because it is easily mistaken for a sleight of hand. When the forcing metric is active we apply CO2-equivalent on both sides of the comparison, not CO2-equivalent to orbit and bare CO2 to the ground. The asymmetry in the result is physical, not metric. Rocket exhaust deposits black carbon directly into the stratosphere, where it forces strongly, whereas terrestrial generation has no equivalent high-altitude particulate term, so the ground's CO2-equivalent is essentially its CO2. A second qualification belongs here rather than in the limitations. The multiplier is not a physical constant but a metric construct. Like every CO2-equivalence it is conditional on a time horizon and on an equivalence convention: the choice among instantaneous radiative forcing, GWP-20, and GWP-100 changes it. The same problem has been litigated at length for short-lived climate forcers and for aviation's non-CO2 terms [37], [38], [39], [40]. We do not claim to have discovered that problem. We claim only that orbital-compute assessments have so far imported one answer to it without saying which. The metric choice is therefore a legitimate accounting decision: different stakeholders may reasonably count or discount a deeply uncertain stratospheric externality, and may reasonably differ on the

horizon over which they count it. Making it a visible axis rather than an embedded assumption is contribution C1.

3.4. Parameters and provenance

Table 2 lists the model parameters with central values, sourced ranges, and provenance; it mirrors the machine-readable parameter file released with the calculator. Three conventions in that table deserve stating. First, the launch factor is tabulated as two components, a combustion term and a non-combustion term, rather than as a pair of alternative totals. This makes the boundary choice a matter of including or excluding one row, it ensures by construction that the embodied factor can never fall below the direct factor, and it gives the equivalence multiplier a single unambiguous target. Second, the status column distinguishes parameters that are propagated from the one that is not. The forcing multiplier m is a scenario axis, held fixed within each reported distribution and swept between them, and it is deliberately excluded from the Monte Carlo draw for the reason given in Section 4.3.

Third, where the literature offered several published values for the same quantity, we applied a fixed priority order rather than averaging: a primary regulatory or standards document first, then a peer-reviewed measurement, then a peer-reviewed model estimate, then a program's own published figure, and finally a derivation from first principles. The direct launch factor follows that order down to the first tier, since it rests on a regulatory environmental assessment cross-checked against propellant stoichiometry. The non-combustion term rests on the third tier, since only a modeled envelope exists, and the mass term rests on the fourth and is labeled an envelope for that reason. No verified cradle-to-gate product carbon footprint has been published for any launch vehicle, which is itself a finding for the standardization agenda, and the space sector's own life-cycle guidance is where one would belong [41]. Where two sources of equal tier disagreed we took the wider interval rather than the midpoint. That biases the reported uncertainty upward, which is the conservative direction for a paper whose thesis is that the verdict is unsettled. The non-combustion launch term and the two second-order terms B and D carry the model's main remaining uncertainty (Section 7). Every parameter marked propagated is drawn over its tabulated range in the analysis of Section 4.3.

The combustion component deserves its own derivation, because it is a ratio of two quantities that have to be stated together: the CO2 a launch releases and the payload that launch delivers. A Falcon-class vehicle burns of order 112 t of kerosene in the configuration the regulatory environmental assessment documents [12], which at about 3 kg of CO2 per kg of kerosene is some 336 t of direct combustion CO2, while a full propellant load gives about 465 t. Delivered mass to a dawn-dusk sun-synchronous orbit near 650 km is about 15 t expendable and 9 to 10 t with first-stage recovery. The four pairings therefore span 21.5, 29.8, 33.6 and 51.7 kgCO2 per kg delivered. Our central value of 28 corresponds to the documented propellant load at close to expendable capacity, and the tabulated range of 17 to 63 covers the recovered configurations. A reader who takes a recovered flight as the reference case should use about 45 kgCO2/kg, which raises the direct pole from 0.061 to 0.098 and the CO2-equivalent pole from 0.700 to 1.069, and lowers the ratio of Corollary 1 from 11.5 to 10.9. We keep 28 because it follows the priority order stated above and because it is the figure the released artifact implements, and we print the derivation because this single ratio sets the absolute position of every contour in the paper.

Table 2 Model parameters, ranges and sources

Symbol	Value [low, high]	Unit	Quantity	Source	Status
M	80 [30, 150]	kg/kW	launched system mass per kW (radiator-dominated)	Suncatcher [4]; spacecraft thermal-control practice and	feasible envelope (§2.3); propagated

Symbol	Value [low, high]	Unit	Quantity	Source	Status
				small-spacecraft SoA (§2.3)	over [30, 150]
f_c (comb.)	28 [17, 63]	kgCO2/kg	combustion CO2 per kg launched, Falcon-class to dawn-dusk SSO	FAA EA [12] and RP-1 stoichiometry	sourced, propagated
f_n (non-comb.)	42 [28, 187]	kgCO2/kg	vehicle manufacturing and propellant production per kg launched	embodied total minus combustion; range spans the envelope reported by Saarland [8]	estimated, propagated
m (forcing)	10 {1, 10, 50}	dimensionless	CO2e multiplier, applied to the combustion component only (the metric axis)	[9], [33], [42]	scenario axis, not propagated (§4.3)
L	5 [2, 7]	yr	hardware useful lifetime	Suncatcher [4] (5-yr); refresh cadence	sourced, propagated
PUE	1.2 [1.1, 1.6]	dimensionless	terrestrial power-usage effectiveness	Uptime Institute 2024 [43]	sourced, propagated
U	0.7 [0.5, 0.9]	dimensionless	utilization (load factor)	Uptime 2024 [43]; AI clusters	sourced, propagated
H	8766	h/yr	hours per year	definition	exact
B	0	kgCO2/kW	orbital manufacturing embodied carbon	assumed to offset D; leverage reported in closed form (§4.3)	not propagated
D	0	kgCO2/kW	avoided terrestrial-construction carbon	assumed to offset B; leverage reported in closed form (§4.3)	not propagated
basis (grid)	location {location, market}	-	measurement basis for the displaced grid intensity	GHG Protocol Scope 2 Guidance [29]	scenario axis, not propagated (§3.1)

3.5. The break-even equation

We write M for the launched system mass per kilowatt, f for the launch carbon factor in kgCO2 per kg launched under the accounting and metric of Section 3.3, B for orbital manufacturing embodied carbon per kilowatt, and D for the avoided terrestrial-construction carbon per kilowatt. The orbital lifecycle carbon per kilowatt is then $M \cdot f + B - D$. Over a hardware lifetime of L years at utilization U, the same kilowatt on the ground delivers $H \cdot L \cdot U$ kilowatt-hours of useful compute, where H is hours per year. It emits $PUE \cdot H \cdot L \cdot U \cdot I_{grid}$ kilograms of CO2 at grid intensity I_{grid} and power-usage effectiveness PUE. Orbit emits less when $M \cdot f + B - D$ is below $PUE \cdot H \cdot L \cdot U \cdot I_{grid}$, that is, when the grid is dirtier than the break-even intensity

$$I^* = \frac{M \cdot f + B - D}{PUE \cdot H \cdot L \cdot U} \tag{1}$$

The verdict is the single comparison of I_{grid} against I^*. This expression is implemented unchanged in the OrbitOrEarth calculator (Section 5), so any reader can reproduce or perturb it.

Three terms are absent from this denominator and declared at unity, in the same style as B and D above. The first is the useful work delivered per watt in orbit relative to the ground, which we write eta: it carries any difference in accelerator performance per watt that mass, thermal and radiation constraints impose. The second is availability, written A, the fraction of installed capacity that remains serviceable over the mission, which no field replacement can restore in orbit. The third is an orbital power-usage effectiveness: the ground side is charged 1.2 for facility overhead, while power conversion, battery round trip, thermal-loop pumping, attitude control and communications all draw bus power that the orbital side is not charged for. All three enter multiplicatively, so the break-even scales as the reciprocal of their product. A performance-per-watt deficit of 30 percent raises the pessimist pole from 0.700 to 1.000 kgCO2e/kWh, an availability of 0.9 raises it to 0.777, an orbital overhead of 1.15 raises it to 0.805, and the first two together to 1.111. We hold all three at unity because we have no sourced value for any of them, and we state their leverage here so that a reader with a figure can rescale every threshold reported in this paper.

3.6. Multi-criteria metrics

We compare seven criteria as a dominance profile rather than a weighted score. Carbon is the per-kilowatt lifecycle break-even of Sections 3 and 4: orbit wins when the displaced grid intensity exceeds I* under the chosen metric, which is workload- and grid-dependent. The terrestrial side of that comparison deserves a word, because a generic data center is not a fair opponent. Best-practice facilities now combine renewable power-purchase agreements, direct-to-chip or immersion liquid cooling, and waste-heat reuse into district networks. Each of these moves the comparison against orbit rather than for it. Renewable procurement lowers the displaced grid intensity, which is the very quantity that must exceed I* for orbit to win. Immersion cooling pushes power-usage effectiveness toward the low end of our tabulated range and removes most evaporative water draw, which erodes orbit's clearest non-carbon advantage. Waste-heat reuse is a credit with no orbital analogue, since a radiator in vacuum has nobody to sell its heat to. Our parameter ranges partly absorb the first two, since PUE is swept from 1.1 to 1.6 and the grid intensity is the decision variable itself, but heat reuse is not represented at all. The reader should therefore treat our terrestrial baseline as an average facility on an average grid, and read the orbit-favorable verdicts as upper bounds that a best-practice terrestrial competitor would narrow.

Water is operational water per kilowatt. Orbit rejects heat radiatively and draws no operational water, whereas a terrestrial facility consumes water directly through evaporative cooling and indirectly through the water intensity of its grid. Orbit therefore holds a structural advantage. Its magnitude is sourceable from facility water-usage data [44], [45] but is not yet folded into the carbon break-even, and it is assessed on an operational boundary while carbon is assessed on a lifecycle one. Latency is the added ground-to-orbit-to-ground transit, roughly 20 to 50 ms round-trip for low-Earth orbit [46] plus handover jitter. It is a disadvantage for orbit on real-time interactive workloads and irrelevant for batch, so we report it separately for the two classes, whose verdicts differ. Debris and e-waste are ordinal: orbit adds collision and debris risk and an end-of-life burden through re-entry ablation, with no recovery or recycling, whereas terrestrial hardware is at least partly recyclable and ground-confined. Quantitative instruments for that axis exist, among them composite sustainability ratings for orbital missions [47], and the case for treating near-Earth space as an environment to be managed rather than a free resource is now argued explicitly [48], [49]. We keep the axis ordinal because we model no specific constellation. Two further criteria are ordinal and favor the ground for the same structural reason: large orbital arrays add reflected sunlight to the night sky, and launch and re-entry products perturb stratospheric ozone, an effect that no CO2-equivalent multiplier represents at any value. Finally, for born-in-orbit large-data workloads, processing in orbit avoids downlinking the raw data. That is a directional carbon and energy credit to orbit whose magnitude scales with data volume and downlink energy, and we treat it as a bounded directional adjustment, not yet folded into the base break-even.

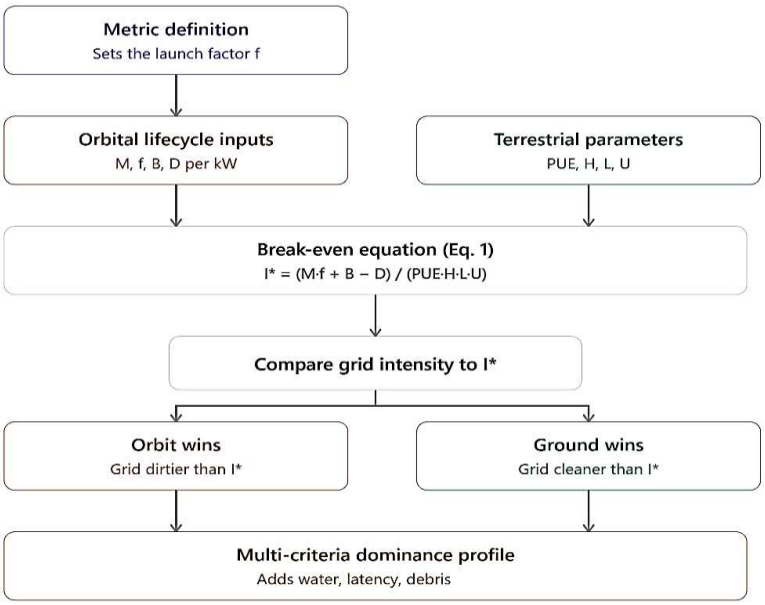


Fig. 4 *Per-kilowatt break-even model*

Fig. 4 summarizes this model end to end. The metric definition sets the launch factor entering the orbital-side inputs, these combine with the terrestrial parameters in the break-even equation, and the resulting comparison against grid intensity feeds forward into the multi-criteria dominance profile just described.

4. Method

Turning the model into a method requires four steps. We establish the analytical properties of the break-even, quantify how each parameter moves it through a sensitivity sweep and a Monte Carlo propagation, define the workload taxonomy that conditions the verdict, and assemble these into the two-stage decision rule applied in the results.

4.1. Analytical model and properties

The model admits a closed-form treatment, given here because the properties that follow are what make the framework's verdicts robust rather than tuned. Following established lifecycle accounting practice [50] [51] [52], the carbon of one delivered kilowatt splits into an orbital and a terrestrial side. With M the launched system mass per kilowatt, f the launch carbon factor (Section 3.3), B the orbital manufacturing embodied carbon, and D the avoided terrestrial-construction carbon, the two quantities are

$$C_{orb} = M \cdot f + B - D \tag{2}$$

$$C_{grd} = PUE \cdot H \cdot L \cdot U \cdot I_{grid} \tag{3}$$

the terrestrial side delivering $H \cdot L \cdot U$ kilowatt-hours over a hardware lifetime L at utilization U and hours-per-year H, with power-usage effectiveness PUE and grid intensity $I_{grid.}$ Equating C_{orb} and C_{grd} and solving for the grid intensity gives the break-even (1) of Section 3.5. The climate-metric choice of Section 3.3 enters through a single substitution,

$$f = m \cdot f_c + \beta \cdot f_n, \ \beta \in \{0,1\}, \ m \geq 1 \tag{4}$$

where the boundary index selects whether the non-combustion term is counted, and the multiplier m prices the high-altitude radiative forcing carried by the combustion term alone. The launch factor is the sum of a combustion component and a non-combustion component, and the substitution multiplies the first and leaves the second untouched, so the climate metric and the accounting boundary act on disjoint parts of the same quantity. This substitution is the formal locus of contribution C1, and its structure is what makes the two choices additive rather than multiplicative in Corollary 1 below.

Proposition 1 (Monotonicity of the break-even). Fix all parameters strictly positive with $M \cdot f + B - D > 0$. Then the break-even I^* is strictly increasing in the launched mass M and the launch factor f, and strictly decreasing in the power-usage effectiveness PUE, the operating hours H, the hardware lifetime L, and the utilization U:

$$\frac{\partial I^*}{\partial M} = \frac{f}{PUE \cdot H \cdot L \cdot U} > 0, \quad \frac{\partial I^*}{\partial f} = \frac{M}{PUE \cdot H \cdot L \cdot U} > 0 \tag{5}$$

$$\frac{\partial I^*}{\partial x} = -\frac{I^*}{x}, \quad x \in \{PUE, H, L, U\} \tag{6}$$

Proof. Write $I^* = N / (PUE \cdot H \cdot L \cdot U)$ with $N = M \cdot f + B - D > 0$. Since N is linear in M and in f, with $\partial N/\partial M = f$ and $\partial N/\partial f = M$, the first two derivatives are positive. For any denominator variable x in $\{PUE, H, L, U\}$, differentiation gives $\partial I^*/\partial x = -N / (x \cdot PUE \cdot H \cdot L \cdot U) = -I^*/x$, which is negative because N and every denominator are positive. The four denominator sensitivities are therefore all negative.

The reading is practical: the sign of every sensitivity is fixed in closed form, so the sweep and the Monte Carlo of Sections 4.2 and 4.3 quantify magnitudes only and never discover a direction. A heavier system or a dirtier launch always raises the break-even and makes orbit harder to justify; a longer lifetime or higher utilization always lowers it. The tornado ranking of Section 4.2 is thus a statement about magnitudes against signs that are known a priori.

Corollary 1 (The optimist-pessimist gap is additive in the two accounting choices). Take the reference case B = D = 0, so that I* is proportional to the launch factor. Write the launch factor as the sum of a combustion component and a non-combustion component. The optimist verdict counts the combustion component only, at multiplier one; the pessimist verdict counts both components and applies the multiplier to the combustion component alone. Their ratio is then the sum of the multiplier and the ratio of the non-combustion component to the combustion component,

$$\frac{I^*_{pess}}{I^*_{opt}} = m_p + \frac{f_n}{f_c} \tag{7}$$

independent of M, PUE, H, L, and U. The order-of-magnitude disagreement between the two studies is therefore not a modeling discrepancy. It follows from two explicit accounting choices, which enter additively rather than multiplicatively: the metric multiplier prices the combustion share, and the boundary adds a metric-neutral increment on top of it. Take central values of 28 kgCO2/kg for the combustion component and 42 kgCO2/kg for the non-combustion component, hence a component ratio of 1.5, with a forcing multiplier of 10. The ratio is then 11.5, which carries the break-even from 0.06 to 0.70 kgCO2/kWh as reported in Section 5.2. Contribution C1 is this corollary made explicit: fixing the boundary, the multiplier, and the quantity the multiplier applies to

dissolves the disagreement. The corollary also makes visible how much of the gap each choice owns, provided the two contributions are read as terms of a sum and not as ratios in their own right. At the central values the multiplier contributes 10 to that sum and the boundary contributes 1.5. A reader who rejects the forcing metric outright is therefore left not with a gap of 1.5 but with the ratio the boundary alone produces, namely 1 + 1.5 = 2.5. At the central values that residual is small enough to be settled by better launch-emission data rather than by argument. It is not robust across the tabulated ranges: the component ratio spans 0.44 to 11.0, so the metric-free gap spans roughly 1.4 to 12, and the 90 percent intervals for the direct and embodied boundaries overlap. The claim that the boundary axis is the tractable one therefore holds at the reference configuration and weakens away from it.

Uncertainty propagation. Parametric uncertainty is propagated by a Monte Carlo method [53], detailed in Section 4.3. Each propagated parameter is drawn from a triangular distribution over its tabulated range, and the 5th and 95th percentiles of the resulting distribution of I* give the reported interval and confidence label,

$$\theta_i \sim Tri(a_i, c_i, b_i), \quad CI_{90} = [Q_{0.05}(I^*), Q_{0.95}(I^*)] \tag{8}$$

with the accounting-and-metric choice of (4) treated as an explicit scenario rather than a distribution (Section 7).

Decision predicate. Writing viable(w) for the stage-one viability test of Section 4.5, where a workload is viable when its latency budget admits the orbital round-trip and its data-transport balance is favorable, the verdict is the two-stage map

$$D(w, g, v, m) = ground \;\; if \neg viable(w); \quad else \;\; orbit \Leftrightarrow g > I^*(v, m) \tag{9}$$

so that a workload failing viability is assigned to the ground before its carbon is ever computed, and only a viable workload reaches the break-even comparison under launch vehicle v and metric m.

4.2. Sensitivity sweep and tornado

We vary each parameter one at a time across its tabulated range while holding the others at their central values, recording the resulting break-even. Two conventions govern this analysis and both matter. The mass parameter is swept and drawn over its feasible range, 30 to 150 kg/kW, rather than over the wider published envelope. Section 2.3's radiative analysis excludes the light end, and propagating a region the paper's own physics rules out would bias every interval. The tornado diagram (Fig. 5) ranks the levers by response width, and the ranking is boundary-conditional. Under the direct-CO2 boundary hardware lifetime leads at 0.109 kgCO2/kWh, ahead of the combustion launch term at 0.100 and system mass at 0.091, with utilization and power-usage effectiveness far back at 0.038 and 0.021. Under the embodied boundary the non-combustion launch term dominates outright at 0.345, ahead of lifetime at 0.272 and mass at 0.228. Under the CO2-equivalent metric lifetime leads at 1.249, ahead of mass at 1.050 and the combustion term at 1.000. A one-at-a-time sweep weights each parameter by the width of its range rather than by the shape of its distribution. We therefore complement it with first-order and total-order Sobol indices over the same marginals (Table 3). The break-even is a product of independent factors, with the launch factor entering as a sum inside one of them, so those indices follow in closed form from the first two moments of the marginals; Table 3 reports the exact values rather than Monte Carlo estimates of them. The two methods do not agree, and the disagreement is informative.

Ranked by variance contribution, mass leads under the direct and CO2-equivalent boundaries, with total-order indices of 0.356 and 0.382. It falls behind the non-combustion launch term under the embodied boundary, 0.333 against 0.356, whereas the one-at-a-time sweep places it third, third and second in the same three columns. The reason is distributional rather than physical. What a variance decomposition weights is dispersion relative to the

mean, measured by the squared coefficient of variation, which over the feasible support is 0.081 for mass against 0.068 for the reciprocal of lifetime. The honest summary is therefore that no single lever dominates. Mass, lifetime and the relevant launch-factor component trade places according to the accounting boundary and the ranking method, so a builder should read the column that matches the standard they are held to rather than looking for one universal lever. What does not follow from Table 3 is any claim about statistical dependence between the inputs. The break-even is a product of independently drawn factors, and a product carries structural interaction between its factors even when those factors are independent, so the first-order indices do not sum to one. They sum to 0.917, 0.912 and 0.924 in the three columns, and that shortfall is fixed by the marginals alone. We report the indices for their ranking content only, and a reader should not read the residual as evidence of coupling. Coupling is addressed where it belongs, in the correlated propagation of Section 4.3. Utilization and power-usage effectiveness remain last under both methods and every boundary, which is the one unqualified conclusion available here: marginal facility-efficiency gains do not change the orbit-or-ground verdict.

Table 3 Sobol sensitivity indices of I*, computed in closed form

Accounting scenario	Total-order indices S_T (all propagated parameters)	Leading first-order S_1	Ranking note
Direct CO2	M 0.356 · f_c 0.329 · L 0.302 · U 0.068 · PUE 0.031 · f_n 0.000	M 0.304 (Σ S_1 = 0.917)	mass leads narrowly over the combustion term
Embodied	f_n 0.356 · M 0.333 · L 0.283 · U 0.064 · PUE 0.029 · f_c 0.027	f_n 0.302 (Σ S_1 = 0.912)	the non-combustion term leads; mass second
CO2e (m = 10)	M 0.382 · L 0.324 · f_c 0.235 · U 0.073 · PUE 0.033 · f_n 0.032	M 0.332 (Σ S_1 = 0.924)	mass leads; lifetime second

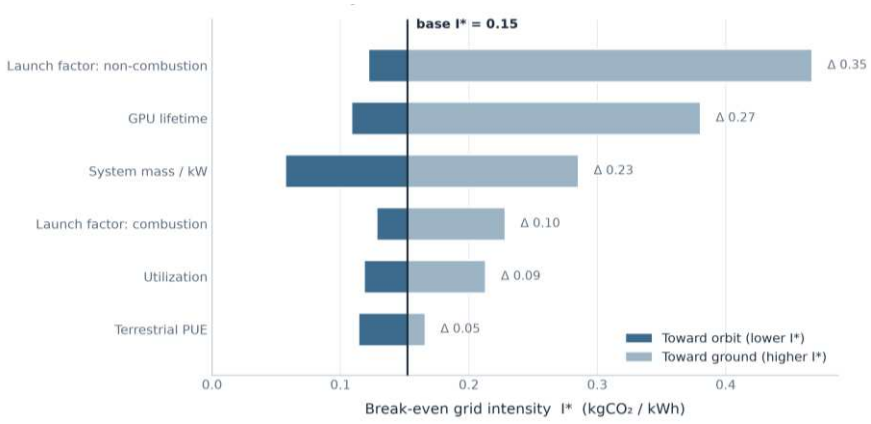


Fig. 5 *Tornado sensitivity of I**

4.3. Monte Carlo uncertainty propagation

To propagate the joint uncertainty rather than one lever at a time, we draw each propagated parameter from a triangular distribution over its tabulated range, with the mode at the central value, and evaluate I* for each draw. We report N = 200,000 draws under a fixed seed (20260822); a split-half comparison reproduces the reported 5th and 95th percentiles to within one percent. The triangular family is chosen because the tabulated inputs give a bounded range and a most-likely value and nothing further. We do not claim it is the true distribution of any parameter. The launch factor is tabulated as a combustion component plus a non-negative non-combustion component, so the sampler cannot generate an embodied factor below the direct factor. Mass is drawn over its

feasible support rather than the wider published envelope, so the sampler cannot generate a system lighter than Section 2.3 permits. That restriction matters: drawing mass from the bottom of the published envelope upward would lower every interval and shift the verdict in orbit's favor.

The resulting 90 percent intervals are 0.035 to 0.180 kgCO2/kWh under the direct boundary, 0.114 to 0.623 under the embodied boundary without forcing, and 0.455 to 2.175 under the CO2-equivalent metric. One consequence of the right skew should be stated rather than left for a reader to find. The point estimates we report throughout are computed at the central parameter values and sit below the medians of their own distributions, at the 29th, 14th and 22nd percentiles for the direct, embodied and CO2-equivalent scenarios respectively. They are modal rather than central estimates. A reader comparing a displaced grid against the interval rather than against the point estimate will reach a materially more pessimistic verdict for orbit, not a marginally different one. Algorithm 1 should be read with that in mind: it returns a verdict from the point estimate and a confidence label from the distribution, and those two quantities do not sit at the same place.

The parameters are not independent. Altitude drives launch CO2, forcing, hardware lifetime and latency jointly; heavier shielding and larger radiators plausibly buy lifetime; and the two launch components both scale with the vehicle and the mission profile. We therefore repeat the propagation under a Gaussian copula over the same marginals, with a declared rather than measured correlation structure, and report the result as a span between two structures rather than as a bound. Both structures are stated in full so that they can be reproduced: the two launch components are correlated at rank 0.80, mass and lifetime at rank plus 0.45 in the first structure and minus 0.45 in the second, and every other pair is left independent. Under the first structure, in which mass buys lifetime, the intervals narrow by 16 percent under the direct boundary, 8 percent under the embodied boundary and 11 percent under the CO2-equivalent metric. Under the second they widen by 16, 22 and 23 percent.

Three qualifications belong with that result. The declared correlations are rank correlations imposed through the copula, not Pearson correlations of the marginals. A Gaussian copula carries no tail dependence, which is a real limitation because the confidence labels are decided in the tails; a t-copula sensitivity would be the natural check, and we have not run one. Third, the two structures are chosen, not extremal, so the span indicates scale rather than bounding the effect. Two further points about coverage. The forcing multiplier is not drawn: it is a metric choice, not a measurable quantity, and we sweep it between scenarios instead. B and D are not drawn either, because we have no sourced range. We report their leverage in closed form, since the break-even shifts by the net imbalance divided by the denominator of Equation (1). A net imbalance of 1000 kgCO2/kW moves I* by 0.027 kgCO2/kWh, 45 percent of the optimist pole and 4 percent of the pessimist pole.

A fourth qualification belongs with the correlation result, because the choice of marginal family moves the intervals further than either structure does. On the same minimum, mode and maximum, uniform marginals raise the 95th percentile of the CO2-equivalent distribution from 2.17 to 3.38 kgCO2e/kWh, and PERT marginals lower it to 1.69. The family also moves the point estimates within their own distributions, from the 29th, 14th and 22nd percentiles under triangular marginals to the 39th, 26th and 34th under PERT. The modal-estimate caveat stated above is therefore conditional on the triangular family, and we report the alternatives rather than treating that family as neutral.

4.4. Workload taxonomy

The taxonomy has two axes, and both are chosen because they gate placement rather than because they are the only properties a workload has. The first is data provenance: data born in orbit, such as Earth-observation

imagery captured by the satellite, versus Earth-origin data, such as a user query or a training corpus stored on the ground. Provenance sets the data-transport cost, because processing where the data already lives avoids moving it. Born-in-orbit data processed in orbit avoids a large raw-data downlink, since Earth-observation satellites generate tens of terabytes per day each. Earth-origin data processed in orbit must first be uplinked, with no transport saving and a penalty if bulky.

The second axis is latency class, which we define by the budget within which a result must be returned to whoever consumes it. Real-time interactive work allows roughly 150 to 300 ms, as in conversational inference where a human waits on the response. Near-real-time work allows of order seconds, as in alerting or tip-and-cue tasking where a machine acts on the result but a short delay is tolerable. Batch work allows minutes to hours, as in model training or bulk Earth-observation processing where throughput matters and arrival time does not. Deferred work has no deadline at all, as in archival reprocessing. The boundaries are order-of-magnitude by design, and a workload is assigned to the tightest class its consumer actually requires rather than to the class its implementation currently achieves. Orbit adds a ground-to-orbit-to-ground transit whose floor is set by geometry and whose operational value is set by the constellation. The one-way free-space delay at 600 km is about 2 ms at zenith, so the propagation floor for a round trip to a spacecraft overhead is of order 4 ms, but measured low-Earth-orbit round trips sit in the range of roughly 20 to 50 ms [19].

The gap between the two figures is not slack. It is the sum of the mechanisms a geometric model omits, and we use the measured range throughout precisely because those mechanisms are what a latency budget has to absorb. They are four. Elevation angle matters, since a satellite near the horizon is several times farther away than one at zenith. Satellite visibility and handover matter, since a low-Earth-orbit spacecraft is in view of a given ground point for minutes at a time and the transition between spacecraft adds both delay and jitter. Ground-station and gateway availability matters, since a request may have to wait for a feeder link. Finally, inter-satellite routing matters, where a constellation relays traffic across multiple hops rather than serving it from the spacecraft overhead. Modeling these per constellation is a networking problem outside the scope of a carbon-accounting framework, and we do not attempt it. What the decision rule needs is a bound on the added delay, and the measured range supplies one. Two further remarks bound the taxonomy honestly. Latency is negligible for batch and deferred work but erodes the budget and adds instability for tight real-time loops with ground users, so the orbit-favorable corner is born-in-orbit and batch and the orbit-adverse corner is Earth-origin and real-time. And a workload has many properties these two axes do not capture, among them arithmetic intensity, memory footprint, communication-to-computation ratio, and tolerance of preemption. Those properties matter, but they determine whether a workload can run on a given orbital platform at all, which is a feasibility question about the hardware, rather than whether it should run in orbit or on the ground, which is the placement question this framework answers. A workload that fails on arithmetic intensity or memory fails identically wherever the accelerator sits.

A third limit on the taxonomy is that other gates bind before ours do, and we state them rather than implying that provenance and latency are the whole of placement. Data residency and sovereignty come first for regulated workloads: a spacecraft and its contents fall under the state of registry, so processing in orbit is a cross-border transfer with no settled theory of where the processing occurs, which disqualifies much public-sector and regulated enterprise work before carbon is considered. Cost is the second: at any launch price now in prospect, the capital per delivered kilowatt is far above its terrestrial equivalent, as Section 6 sets out. Availability and serviceability are the third, since nothing in orbit can be replaced in the field, so capacity decays with each permanent failure. Security practice is the fourth, since physical perimeter controls and media sanitization have no orbital equivalent short of deorbit. Time to capacity is the fifth, and it favors the ground today. Our framework addresses the two gates it can quantify from published data, and a workload that fails any of these five never reaches it.

4.5. The decision rule

The rule has two stages, and the interaction between its two inputs is what makes it more than a carbon threshold. Stage one establishes viability, independent of carbon. First comes latency feasibility. The test is not whether orbit's added round trip exceeds the whole budget, which for interactive work it does not. It is whether that round trip consumes enough of the budget to be disqualifying once the rest of the serving path is accounted for. We write it as a fraction: orbit is infeasible when the added round trip plus handover jitter exceeds a share tau of the workload's latency budget. Two things about that predicate must be stated plainly, because together they bound what it can decide. Tau is a stated convention, not a derived quantity; we use 0.2 on the reasoning that a serving path already spends most of its budget on inference, queueing and the terrestrial network. The jitter term, not the mean transit, is what usually binds, and we have no sourced value for it, so Table 2 does not carry one. At tau = 0.2 against a 150 ms budget the threshold is 30 ms, which the measured 20 to 50 ms transit range straddles. The predicate fires at the upper end and not at the lower, and where it fires depends entirely on a jitter term we cannot quote. We therefore do not claim that the predicate as specified rejects interactive workloads. What we claim is narrower and, we think, more defensible: the case against orbital serving for interactive work rests on the tail of the latency distribution rather than its mean, and a mean-based predicate of this form cannot express that.

The jitter term is the weakest input the rule has. Supplying a sourced jitter distribution, and re-expressing the gate as a tail quantile rather than a mean, is the first improvement we would make. Second comes the data-transport balance. Born-in-orbit data favors orbit by avoiding a downlink, while Earth-origin data whose uplink energy is not amortized by the compute it feeds favors the ground. A workload failing either test is assigned to the ground regardless of carbon. It receives no break-even interval and therefore no confidence label, only the reason for rejection.

Stage two evaluates the carbon break-even for viable workloads: compare the displaced grid intensity to I* under the chosen metric and accounting, with orbit winning when the grid exceeds I*. Where the location-based and market-based intensities differ, stage two is run under each and both verdicts are reported, following the dual-reporting requirement set out in Section 3.1. Confidence comes from the Monte Carlo distribution of I* described in Section 4.3, and we read two quantities from it: the confidence label defined next, and the probability that orbit wins, which is the share of draws in which the displaced grid exceeds the break-even. If the displaced grid lies above the 95th percentile, orbit wins for all but the most extreme five percent of draws, and the recommendation is orbit at high confidence. Below the 5th percentile the ground wins on the same reasoning, and inside the interval the verdict flips within the bulk of the distribution and is reported as uncertainty-dominated. The label describes robustness to parametric uncertainty at a fixed metric and nothing else: it is not a probability that orbit is greener, it does not span the metric scenarios, and it is undefined for workloads rejected at stage one. It should also be read against the modal-estimate caveat of Section 4.3. The verdict is taken from the point estimate and the label from the distribution, and the two part company in a band above each point estimate: 0.061 to 0.081 kgCO2/kWh under the direct boundary, 0.152 to 0.270 under the embodied boundary, and 0.700 to 1.008 under the CO2-equivalent metric. Inside those bands the point estimate returns orbit while most draws return the ground, which is why we report the win probability beside every verdict.

Algorithm 1 states this two-stage procedure formally, combining the break-even computation, the Monte Carlo confidence label, and the viability gate into the single routine implemented in the calculator.

Algorithm 1: Orbit-or-ground decision procedure

```
Input
  w          workload, with data provenance p and latency class ℓ
  g          displaced grid intensity; run once per measurement basis and report both (§3.1)
  v          launch vehicle
```

```
  (β, m)      accounting boundary and forcing multiplier: β = 0 direct, β = 1 embodied; m is a
                scenario, never drawn
  params      parameter ranges {M, f_c, f_n, B, D, PUE, H, L, U}, each with a central value and
                [low, high]
  τ           latency-share convention, default 0.2 (§4.5)
  defaults    downlink_credit(w) = 0 (§3.6) and handover_jitter = 0 (§4.5), the two terms this
                paper leaves unquantified
Output
  verdict     orbit or ground, or a non-viability reason returned by stage one
  I*, CI_90   break-even grid intensity and its 90% interval
  confidence  high, or low (uncertainty-dominated)
  p_orbit     probability that orbit wins

 1  // Stage one: viability, independent of carbon
 2  if orbit_round_trip(v) + handover_jitter > τ · latency_budget(ℓ) then
 3      return (ground, –, –, "non-viable: latency", –)
 4  if provenance(p) = Earth-origin and uplink_energy(w) not amortized by compute then
 5      return (ground, –, –, "non-viable: transport", –)

 6  // Stage two: carbon break-even, reached only when the workload is viable
 7  f ← m · f_c + β · f_n                        // Eq. (4): m prices the combustion share only
 8  I* ← (M · f + B – D) / (PUE · H · L · U)      // Eq. (1)
 9  if provenance(p) = born-in-orbit then
10      I* ← I* – downlink_credit(w)              // directional adjustment, §3.6

11  // Monte Carlo confidence, on the same quantity as the verdict
12  for i = 1 to N do                            // N = 200,000, seed 20260822
13      draw θ_i ~ Tri(low, central, high) for each propagated parameter; m held fixed
14      f_i ← m · f_c,i + β · f_n,i
15      I*_i ← (M_i · f_i + B – D) / (PUE_i · H_i · L_i · U_i)
16      if provenance(p) = born-in-orbit then I*_i ← I*_i – downlink_credit(w)
17  CI_90 ← [Q_0.05(I*_i), Q_0.95(I*_i)]          // Eq. (8)
18  p_orbit ← #{i : I*_i < g} / N                 // orbit wins where the grid exceeds I*

19  // Verdict, confidence label and win probability
20  if g > I* then verdict ← orbit else verdict ← ground   // Eq. (9)
21  if g > CI_90.high or g < CI_90.low then
22      confidence ← "high"
23  else
24      confidence ← "low (uncertainty-dominated)"
25  return (verdict, I*, CI_90, confidence, p_orbit)
```

5. Results

Applying the method to the parameters of Table 2 turns the analytical machinery of the model into a set of concrete answers about where orbital compute helps and where it hurts. This section moves in five steps. It maps the break-even across the full parameter space, then shows how the optimist and pessimist verdicts collapse into two accounting choices on a single model. It reports the multi-criteria dominance profile rather than compressing it into one score, works through representative deployment decisions, and finally compares the two-stage rule with the single-criterion placement rules it is meant to improve on.

5.1. The break-even map

Fig. 2 plots three break-even contours in the plane of displaced grid carbon intensity, on a logarithmic horizontal axis, against launched system mass on the vertical axis. The three contours are direct launch CO2, the optimist boundary; embodied launch; and CO2-equivalent with high-altitude forcing applied to the combustion share, the pessimist boundary. They are not three competing predictions. They are the same equation evaluated at three points along the single metric axis defined earlier. Because I* is the grid intensity a workload must exceed for orbit to win, the orbit-favorable region for a given contour lies to the right of it and the ground-favorable region

to its left. Orbit therefore wins under every metric only to the right of the pessimist contour, and the ground wins under every metric only to the left of the optimist contour. Reading the figure vertically turns it into a design tool. Since I* is proportional to M by Equation (1), every contour is a ray through the origin in linear coordinates. A tenfold increase in launched mass per kilowatt therefore displaces the whole contour one decade to the right: the grid a system must displace to justify itself scales exactly with what it weighs. At the direct boundary the 80 kg/kW reference needs a grid of 0.061 kgCO2/kWh, and the heavy end of the tabulated range, 150 kg/kW, needs 0.114, which already excludes the cleanest national grids.

The metric choice moves those thresholds by an order of magnitude. Under the CO2-equivalent metric the same two masses require 0.700 and 1.312. We do not illustrate the light end of the tabulated range, because the radiative analysis of Section 2.3 excludes it. At a 300 K rejection temperature a kilowatt of waste heat needs of order two square meters of radiating surface before margins, and no plausible areal density brings the radiator alone, let alone the bus and array, down to the light end of the published envelope. Table 2 therefore tabulates and propagates M over its feasible range, 30 to 150 kg/kW, and contour positions below that floor are not evaluated.

Mass reduction and metric choice act on the same quantity, though not to the same extent. Stepping from the CO2-equivalent metric back to the embodied boundary moves the break-even by a factor of 4.6, whereas the feasible mass range spans a factor of 5 in total and only a factor of 2.7 below the reference configuration. A builder who cannot influence which metric a regulator adopts can move the verdict through mass, but not by as much. Between the outer contours sits a band whose width is a factor of 11.5 at the central configuration with the second-order terms set to zero, and the carbon-intensive grids on which most AI workloads run today fall inside that band. A verdict for those grids is therefore not settled by physics alone; it is settled by which contour a reader is willing to stand behind. This is Contribution C1 in a single figure. A fourth axis is marked on the figure rather than swept by it. The hatched band at the left edge is a market-based reading of the displaced grid, which sits at or near zero intensity and therefore left of all three contours. That is why Section 3.1 treats the measurement basis as carrying more leverage than the metric this figure sweeps.

5.2. Reconciliation of the two camps

We hold the physical system fixed at the reference configuration and feed the model each study's own accounting boundary and climate metric. The same per-kilowatt model then reproduces the direction and the approximate magnitude of each published verdict (Table 4). Three limits on that claim should be stated at once. The non-combustion launch term is taken from the Saarland study rather than derived independently, so the exercise shows that one model can span both positions, not that either was predicted from first principles. The combustion-share application of the multiplier is our convention, not one we can attribute to either study; Table 4 records it in the pessimist row because that is the reading we defend, but neither study states it. And our model carries no re-entry term while the pessimist boundary counts one, so what we reproduce is a corrected pessimist position rather than the published one. The reconciliation is in that sense normative: it shows what the disagreement reduces to once a defensible set of conventions is imposed, not what either author was doing.

Table 4 Reconciliation of optimist and pessimist verdicts

Camp	Launch accounting	Climate metric	Model I*	Verdict at real grids	Published claim
NTU (optimist)	direct CO2 (combustion)	CO2 only	0.06 kgCO2/kWh	orbit beats essentially all grids	carbon-neutral within years

Camp	Launch accounting	Climate metric	Model I*	Verdict at real grids	Published claim
Saarland (pessimist)	embodied launch (combustion plus production)	CO2e; forcing applied to the combustion share (our convention, not stated by that study)	0.70 kgCO2e/kWh	orbit loses to all but the most carbon-intensive grids; 14x worse at 0.050, 1.9x at 0.375 kgCO2/kWh	up to an order of magnitude worse
This work	explicit, swept	explicit, swept	0.06 to 0.70	depends on metric and grid	n/a

Under the optimist's choices, direct launch CO2 with no high-altitude forcing, the break-even falls to 0.06 kgCO2/kWh. That is below essentially every operating grid, so orbital compute sits at or under the ground baseline, which is the optimist's carbon-neutral conclusion. Under the pessimist's choices, embodied launch under a CO2-equivalent metric that prices the high-altitude forcing carried by the combustion share, the break-even rises to 0.70 kgCO2e/kWh at the reference multiplier of 10. At that threshold orbit is worse than the ground across almost the whole operating range. It is about 14 times worse against a very clean grid at 0.050 kgCO2/kWh and about 1.9 times worse against the US average at 0.375 kgCO2/kWh, the two grid intensities we use throughout for these comparisons. That is the pessimist's conclusion, and the factor of roughly 14 at clean grids is what reproduces its published order-of-magnitude characterization. Two qualifications belong here rather than in the limitations, because both cut against the reading just given. First, the reference multiplier is not the strictest case the parameter table admits. At the upper end of the swept range, m = 50, the break-even rises to 3.13 kgCO2e/kWh and the ratio to 51.5, above every operating grid, and orbit loses everywhere.

Second, our model carries no re-entry term, while the study we are reproducing counts re-entry within its boundary, so 0.70 understates that boundary by an amount we do not quantify (Section 3.3). With both qualifications in view, one consequence of the corrected metric application should still not be passed over. At the reference multiplier the threshold is 0.70 rather than a figure above every real grid, so a coal-dominated system above roughly 0.70 kgCO2e/kWh favors orbit. That flip is conditional on m, not a property of the CO2-equivalent metric as such. The pessimist verdict is best read as a statement about the grids most data centers actually reach at a stated multiplier, rather than as a universal one. The two studies differ in a third respect we do not model, namely amortization: the optimist spreads launch over the satellite's full operational life and the pessimist over the shorter hardware refresh cycle. That difference would widen the gap beyond the factor of 11.5 our two axes account for. We flag it as a third accounting choice a future standard would also have to pin down, rather than folding an unsourced figure for it into our results. Within the two axes we do model, the disagreement is not empirical: it is a choice of accounting boundary and climate metric, and one model spans both verdicts.

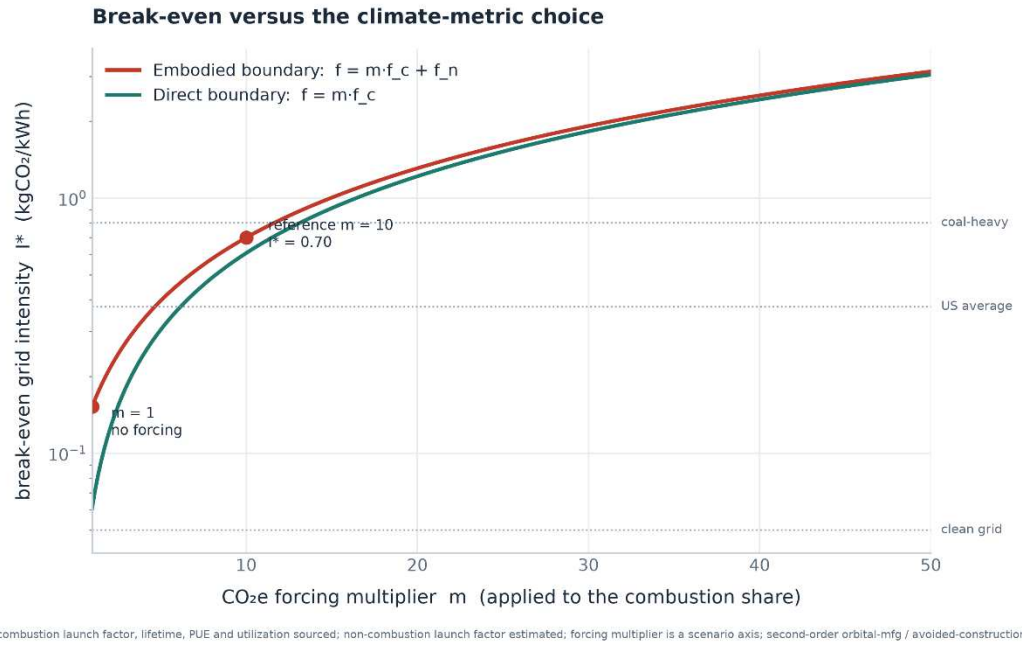


Fig. 6 *I* versus forcing multiplier m*

Two companion figures make this reconciliation concrete rather than purely algebraic. Fig. 6 sweeps the forcing multiplier continuously rather than jumping between the two published endpoints. It traces the direct and embodied boundaries as separate curves and marks the optimist and pessimist verdicts as points that fall on those curves, at the multiplier each study implicitly assumed. Seen this way, the gap between the two literatures is not a discontinuity to be bridged. It is a single continuous quantity that the two studies happened to sample at opposite ends. Table 5 gives the same sweep numerically at intermediate multipliers, which is where the decision-relevant structure lives, because a verdict does not move smoothly with m so much as flip at particular thresholds. At m = 1, the embodied boundary with no forcing, the break-even is 0.152 and orbit beats every grid above roughly 0.15 kgCO2/kWh, which includes the US average and most of Asia. By m = 5 the break-even has passed 0.375 and the US average has flipped to the ground. At the reference m = 10 only coal-dominated systems above 0.70 still favor orbit, and they do so in 32 percent of the propagated draws rather than in a majority. By m = 20 the break-even of 1.31 exceeds every operating grid and orbit loses everywhere.

The two published poles are therefore not merely far apart. They sit on opposite sides of two distinct threshold crossings, and a reader can locate their own position on this table by deciding only which multiplier they are willing to defend. Fig. 7 complements this with the Monte Carlo distribution behind the reported 90 percent intervals. Those intervals are 0.035 to 0.180 kgCO2/kWh under the direct boundary, 0.114 to 0.623 under the embodied boundary without forcing, and 0.455 to 2.175 under the CO2-equivalent metric. Read as win probabilities at the three reference grids of Table 5, 0.05, 0.375 and 0.80 kgCO2/kWh, orbit wins in 17, 100 and 100 percent of draws under the direct boundary. It wins in 0, 74 and 99 percent under the embodied boundary, and in 0, 2 and 32 percent under the CO2-equivalent metric. The distributions are right-skewed, so the interval bounds alone understate how much probability mass sits toward the pessimist regime. A grid intensity landing in the bulk of a distribution should be read as genuinely uncertainty-dominated, not as a coin flip between two equally likely point estimates. We deliberately do not adopt either pole as the paper's working conclusion. Committing to one would re-enact the choice this paper exists to expose, and Figs. 5 and 6 show precisely how sensitive any single point estimate is to a metric question the atmospheric-science literature has not settled.

Table 5 Break-even and verdict versus forcing multiplier

Forcing multiplier m	Break-even I* (kgCO2e/kWh)	Ratio to the optimist pole	Verdict at grids 0.05 / 0.375 / 0.80 kgCO2/kWh
1 (no forcing)	0.152	2.5	ground / orbit / orbit
2	0.213	3.5	ground / orbit / orbit
5	0.395	6.5	ground / ground / orbit
10 (reference)	0.700	11.5	ground / ground / orbit
20	1.308	21.5	ground / ground / ground
50 (sweep maximum)	3.133	51.5	ground / ground / ground

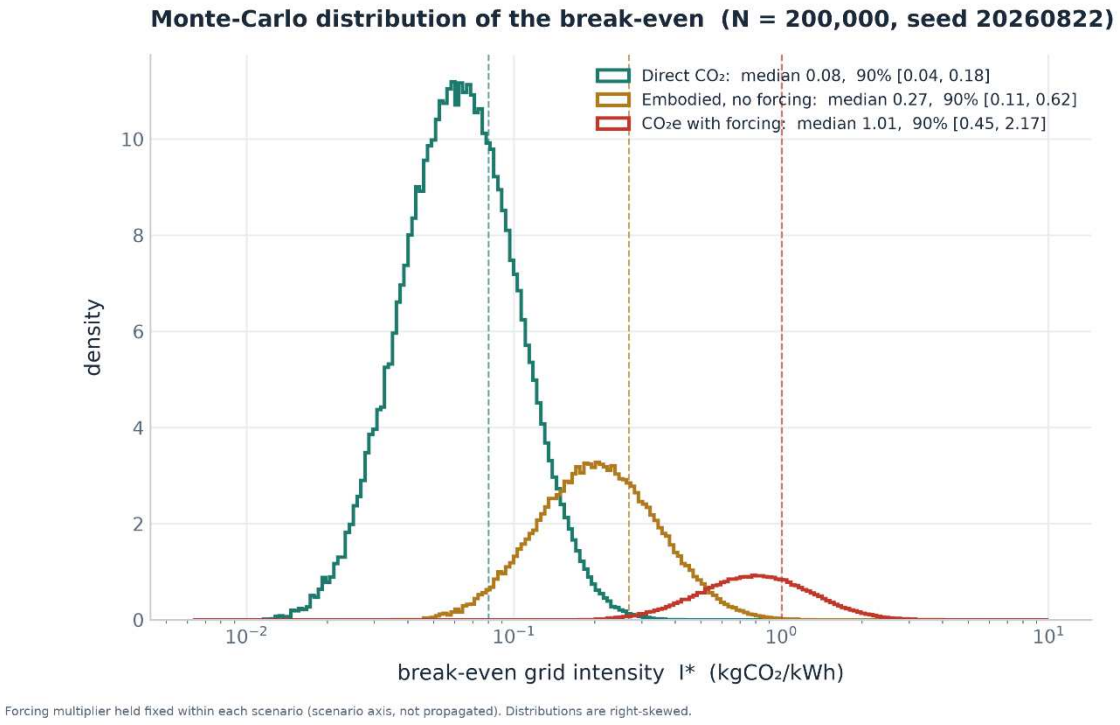


***Fig.** 7 Monte Carlo distribution of I**

5.3. Multi-criteria dominance, not a single score

Carbon, however, is only one axis a deployment decision has to clear, and it is the axis most contested. No criterion in the framework dominates universally, and collapsing all seven into one weighted score would obscure exactly the tension the paper is built to expose. Orbit holds a structural advantage on water and, for born-in-orbit workloads, on the avoided-downlink credit. The ground holds a structural advantage on debris and e-waste, and on latency for interactive workloads. Carbon alone is workload- and metric-dependent, as Sections 5.1 and 5.2 just established. One asymmetry in this table should be declared rather than left for a reader to find. Carbon is assessed on a full lifecycle boundary, including the embodied carbon of placing the system in orbit. Water is assessed on an operational boundary only, counting cooling and grid water but not the water embodied in manufacturing the launch vehicle, the spacecraft, or the terrestrial facility. Orbit's clearest win therefore rests

on the narrower of the two boundaries, and closing that gap would require a water-side lifecycle inventory we do not have. Because these criteria point in opposing directions rather than reinforcing one another, we report a dominance profile and let the workload, together with the decision-maker's own weighting of these concerns, determine the outcome (Table 6).

Table 6 Multi-criteria dominance profile

Criterion	Orbit	Ground	Favors
Carbon	continuous solar; born-in-orbit credit	grid-dependent	workload and metric-dependent
Water	approximately zero operational water	evaporative plus grid water	orbit (operational boundary only; see §5.3)
Latency (interactive)	adds 20 to 50 ms plus jitter	low	ground
Latency (batch)	irrelevant	irrelevant	tie
Debris and e-waste	adds collision and debris risk, peaking in the 800-1000 km band; re-entry ablation, no recycling	recyclable, ground-confined	ground
Night-sky brightness	adds reflected sunlight from large arrays	none	ground
Stratospheric ozone	launch and re-entry products; not representable as CO2e	none	ground

The profile is not a fixed scoreboard. It shifts with the workload, and that shift is the entire point of building it as a profile rather than a single number. For Earth observation, which is batch and born-in-orbit, the carbon verdict is itself metric-conditional: orbit wins at most grids under the direct and embodied boundaries, and loses under the CO2-equivalent metric unless the avoided-downlink credit closes the gap, as Section 5.4 sets out. Orbit wins water outright on the operational boundary; the ground wins debris, night-sky brightness, and stratospheric ozone; and latency is a tie. The recommendation for this workload is therefore robust to how the remaining criteria are weighted under the first two accounting boundaries, and conditional under the third. For interactive inference, by contrast, which is real-time and Earth-origin, the ground wins latency so decisively that the workload is rejected at stage one, before carbon or water are ever weighed. The same seven criteria, read against two different workloads, yield opposite verdicts from the same table. Declining to aggregate is a choice, and it should be defended rather than assumed. A weighted score would require commensurating a lifecycle carbon quantity, an operational water quantity, a latency budget in milliseconds, and three ordinal environmental risks. Any weight vector that did so would encode a normative position at least as contestable as the climate metric this paper spends its length making explicit.

Collapsing the seven criteria into one score would reintroduce, one layer up, precisely the hidden choice C1 exposes. What can be said without a weight vector is which verdicts survive any weighting at all, and that is the question a dominance profile is built to answer. For interactive inference the answer is unconditional: the workload fails stage one, so no weighting of the remaining six criteria can recommend orbit. For Earth observation under the direct and embodied boundaries neither alternative is dominated in the strict sense: orbit wins carbon and water and ties latency, while the ground wins the three ordinal risks. What survives is a weaker and more precise statement. Scoring each criterion 1 for a win, 0.5 for a tie and 0 for a loss, orbit is preferred by any weight vector that gives carbon and water jointly more weight than debris, night-sky brightness and ozone jointly. Under the CO2-equivalent metric neither alternative dominates: orbit wins water, the ground wins carbon

and the three ordinal risks, and the verdict then turns on how a decision-maker prices an unquantified water saving against a quantified carbon penalty. That last case is the honest boundary of the method, and naming it is more useful than manufacturing a score that would conceal it. A single weighted score averaged across all workloads would necessarily mislead one of them, which is why the decision has to be workload-conditioned rather than criterion-conditioned. That conditioning is what the two worked decisions below make concrete.

5.4. Worked decisions

Two cases illustrate the rule, one at each pole of the workload taxonomy, and together they show why carbon alone is an insufficient basis for either recommendation. They are illustrations rather than validations: both are constructed from the taxonomy's own extreme corners and neither is checked against a deployed system, so they demonstrate what the rule does, not that it is right.

i) Earth observation, where orbit wins:

A constellation captures imagery in orbit and the processing carries no hard latency deadline. At stage one, latency is irrelevant, and processing in orbit avoids downlinking raw data, a transport and energy saving that stage-one viability credits immediately. The workload clears the viability gate and is provenance-favored. At stage two the verdict depends on a quantity this paper does not measure, and the result of sizing that quantity is not the one we expected. Under the direct-CO2 boundary the break-even is 0.06 kgCO2/kWh, below the 0.10 to 0.40 kgCO2/kWh range we take for the grids an Earth-observation ground segment draws on, so orbit wins on the base model. Under the embodied boundary the break-even is 0.15 kgCO2/kWh, below most of that range at the point estimate. The embodied 90 percent interval is 0.114 to 0.623, however, so a displaced grid anywhere in our range falls inside it, and the verdict is uncertainty-dominated by the paper's own rule rather than high-confidence. Under the CO2-equivalent metric the break-even is 0.70 kgCO2e/kWh, above the whole range. Orbit would then win only if the avoided-downlink credit exceeded 0.70 minus the displaced grid, that is 0.30 to 0.60 kgCO2/kWh of effective break-even reduction. Expressed in carbon that is 11,000 to 22,000 kgCO2 per kilowatt, and the question is what could supply a credit of that size.

Not all of the avoided downlink can supply it. The dominant energy term in a raw-data downlink is the spacecraft transmitter, and under our own assumption of near-zero operational carbon in a dawn-dusk orbit that energy is carbon-free, so it contributes nothing to a carbon credit. What the credit can draw on is the carbon-bearing part of the avoided chain: ground-station electricity, backhaul, and any terrestrial processing the orbital system displaces. That part scales with station-hours and with the grid the ground segment draws on, not with the number of bits. Converting the required credit into total link energy therefore overstates the pool that could supply it, and we withdraw that comparison. The scale that remains is unfavorable to orbit. Section 7 carries the ground segment of our own architecture at 500 kgCO2 per kilowatt, which moves the break-even by 0.014 kgCO2/kWh; the credit this case needs is twenty to forty times that. A raw-data architecture carries a larger ground segment than a compute-in-orbit one, so the comparison is indicative rather than conclusive, but it is the only evidence available and it points one way. We therefore still report no verdict for this case under the CO2-equivalent metric, while recording that it leans against orbit rather than sitting symmetrically between the two outcomes. What would settle it is a carbon budget for the ground segment of a raw-data downlink path, station electricity and backhaul at their own grid intensities, rather than a joules-per-bit figure for the link as a whole. The Earth-observation advantage is real, but it is an advantage on the direct and embodied boundaries, on water, and on avoided data movement as an operational matter. It is not an advantage that survives the strictest climate metric.

ii) Interactive inference for ground users, where the ground wins:

A ground user issues a latency-sensitive query, with a budget around 150 to 300 milliseconds, and the input data is Earth-origin and small. Serving that query from orbit adds a ground-to-orbit-to-ground round trip of roughly 20 to 50 milliseconds plus handover jitter, for no provenance benefit, since the data already sat on the ground before the query was made. We do not claim the stage-one predicate settles this case. As Section 4.5 sets out, at tau = 0.2 the 30 ms threshold falls inside the measured transit range, so whether the gate fires depends on a jitter term we cannot source, and at a 300 ms budget it does not fire at all. The argument against orbital serving here is therefore not the mean transit but the tail. Handover and visibility windows make the upper quantiles of the latency distribution, not its average, the quantity an interactive service must budget against, and a constellation cannot remove that the way it can amortize a downlink. The recommendation is the ground, and it carries no statistical confidence label, because the rule never reaches the break-even computation that would produce one. What we can say is that the verdict is insensitive to every carbon assumption in this paper, which is a different and in some ways stronger claim than a percentile. This case remains the one a carbon-only analysis gets wrong: orbital inference could be near carbon-neutral on paper and still be the wrong place to run it, because carbon was never the binding constraint.

Read together, these two cases are the paper's central practical claim in miniature: the same seven-criterion profile and the same break-even equation, applied to two workloads that differ only in provenance and latency class, produce opposite recommendations. Neither recommendation would survive a single-criterion analysis that started and ended with carbon.

One case deserves separate treatment, because it shows where the transport gate has no force. The programs that motivate this work in Section 2.1 are conceived as general orbital compute serving Earth-origin workloads: training corpora and user traffic that originate on the ground and must be uplinked, which is the description the transport test is meant to catch. Applied quantitatively, it does not catch them. The transport gate asks whether the uplink energy is amortized by the compute it feeds, and the compute is very large. A kilowatt of orbital capacity delivers about 30,700 kWh over a five-year life at our reference utilization. Uplinking a training corpus of ten to a hundred terabytes costs, at link energetics between 0.1 and 10 microjoules per bit, something between a hundredth of a percent and a few percent of that. For the uplink to fail the amortization test, the corpus would have to exceed roughly one and a half petabytes at the most pessimistic link energetics we can justify, and far more at plausible ones. Earth-origin bulk training data is close to the best-amortized case a transport gate can encounter, not the worst. These programs therefore pass stage one on our own criterion, and the framework has nothing adverse to say about them until their displaced grid intensity and accounting basis are specified. The methodological point generalizes beyond this case. A viability gate stated qualitatively can appear to decide something it does not, and the transport gate as we have specified it does not bind for any Earth-origin workload whose data volume is small relative to the compute it feeds.

5.5. Comparison with single-criterion placement rules

The decision rule is only worth its complexity if it returns verdicts that simpler rules do not, so we compare it directly with the two single-criterion rules it is meant to improve on. The first is a carbon-only rule of the kind used by the carbon-aware orbital scheduler of the optimist study, which places a workload in orbit whenever the displaced grid exceeds the break-even. The second is a data-gravity rule of the kind implicit in the orbital-edge-computing line, which places born-in-orbit, latency-tolerant work in orbit and everything else on the ground, without reference to carbon. We apply all three to the same 36 cases: the four cells of the workload taxonomy crossed with the three accounting scenarios and the three reference grids of Table 5. Table 7 counts the orbit verdicts each rule returns. The avoided-downlink credit is not applied, for the reason given in Section 5.4.

The added round trip is 50 ms against a 150 ms budget, at the convention tau = 0.2 of Section 4.5. Under those settings the two-stage rule disagrees with the carbon-only rule in 10 of the 36 cases, and every one of them is a

real-time workload on a grid dirty enough that carbon alone would send it to orbit. It disagrees with the data-gravity rule in 9 cases. Four are cases in which that rule sends born-in-orbit batch work to orbit on grids too clean to justify it under the chosen metric. Five are cases in which it keeps Earth-origin batch work on the ground although the carbon comparison favors orbit and the transport gate does not bind. The comparison therefore makes precise what the rule adds and where that addition rests. Over the carbon-only rule it adds exactly the latency gate, and nothing else; at the lower end of the measured transit range, 20 ms, the gate does not fire and the two rules agree in all 36 cases. Over the data-gravity rule it adds the carbon comparison. The rule is thus the conjunction of the two baselines rather than a new criterion, and its advantage over carbon-only placement is exactly as strong as the latency evidence behind it, which Section 4.5 identifies as its weakest input.

Two features of Table 7 are results about our own rule rather than about the baselines. The born-in-orbit rows and the Earth-origin rows are identical, because the credit that would separate them is not quantified, so the provenance axis changes no verdict here and the operative gate is latency alone. And the counts are taken from the point estimate. Read at the median of the propagated distribution, each carbon-sensitive cell returns 4 of 9 rather than 5 of 9. The disagreement with the carbon-only rule then falls from 10 cases to 8, and the disagreement with the data-gravity rule stays at 9 while redistributing between its two sources.

Table 7 Orbit verdicts of three placement rules

Workload cell	Carbon-only rule (NTU-type)	Data-gravity rule (OEC-type)	Two-stage rule (this paper)
born-in-orbit, batch	5 / 9	9 / 9	5 / 9
born-in-orbit, real-time	5 / 9	0 / 9	0 / 9
Earth-origin, batch	5 / 9	0 / 9	5 / 9
Earth-origin, real-time	5 / 9	0 / 9	0 / 9

6. Discussion

The framework's first implication is that the orbit-versus-ground carbon debate has been miscast. It has been argued as an empirical question, is orbital compute greener, when our reconciliation shows it is a normative one at the current level of evidence. The answer is fixed by a choice of accounting boundary and climate metric, and the same model returns the optimist's and the pessimist's verdicts depending on that choice. The productive questions therefore shift. In place of who is right, they become which boundary and metric a given stakeholder should adopt, and for which workloads orbit is even viable before carbon is weighed at all. A dispute that looked factual is better understood as one about accounting and metric. Progress comes from making that choice explicit, not principally from collecting more launch-emission data, though tighter data does narrow the ranges.

For the two literatures we reconcile, the implication is conciliatory rather than corrective: neither camp is making an arithmetic error. The optimist who reports near carbon-neutral orbital compute has counted direct launch CO2 and amortized it over a long mission; the pessimist who reports an order-of-magnitude penalty has counted embodied launch and re-entry under a CO2-equivalent metric that prices the stratospheric forcing rockets produce. Both verdicts follow from the same physics under different, defensible accounting. Recognizing this dissolves an apparent contradiction in the literature and relocates the disagreement to where it actually lives.

For those building orbital compute, the sensitivity ranking is actionable, but it is actionable differently depending on which accounting a builder is held to. Hardware lifetime, system mass and the combustion launch term lead under the direct-CO2 and CO2-equivalent columns, with their order depending on whether levers are ranked by response width or by variance contribution. The non-combustion launch term dominates under the embodied

boundary, and utilization and power-usage effectiveness move the break-even least everywhere. Effort that extends radiation-tolerant hardware life, lightens radiators, or decarbonizes the manufacture of the launch vehicle therefore changes the verdict, while marginal facility-efficiency gains do not. That the ranking itself shifts with the accounting choice is a further practical consequence of C1: a builder optimizing against one standard may be optimizing the wrong lever for another. The reference orbit matters too. The operational-carbon advantage is specific to dawn-dusk sun-synchronous orbits, and altitude trades launch carbon, forcing, radiation, and latency against one another, so siting is multi-dimensional even within orbit.

For policy and standardization the implication is sharper. Because the verdict depends on the metric and the boundary, any carbon-accounting standard for orbital compute must fix those first, or comparisons across studies and providers will remain incommensurable. A standard that does not state its conventions leaves room for either verdict to be reported in good faith. One distinction sharpens that point. A compliance inventory counts the gases a reporting framework recognizes, and black carbon is not among them, so a multiplier of one is mandated there rather than chosen. A climate-impact assessment prices the forcing, where a multiplier above one is required. The same deployment can therefore be inventory-conformant at near neutrality and materially worse in forcing terms, with no bad faith on either side, which is the position aviation's non-CO2 terms have occupied for two decades. Concretely, five axes need pinning down, and this paper has had to choose on every one of them. The accounting boundary: combustion only, or combustion plus vehicle manufacturing and propellant production. The equivalence convention and its horizon: instantaneous radiative forcing, GWP-20, GWP-100, or a warming-equivalent formulation. The component the multiplier applies to: the whole launch factor, or the combustion share alone, a choice that moves our own pessimist pole by a factor of 2.2 and is the axis prior work leaves most thoroughly implicit. The amortization basis: satellite operational life or hardware refresh cycle.

And the greenhouse-gas scope and measurement basis under which the displaced grid is counted, since launch sits in Scope 3 while grid electricity sits in Scope 2. A market-based reading of a contracted renewable supply can return the opposite verdict from a location-based reading of the same physical grid (Section 3.1). One of these axes is less open than it looks. The amortization basis is settled differently by the two frameworks a reader might apply. Corporate reporting recognizes capital goods in the year of acquisition rather than depreciating them over an asset's life, while a product life-cycle assessment amortizes them over the functional unit. A standard must therefore say which framework governs, rather than leaving the period free. The natural venues are the greenhouse-gas accounting bodies that already maintain scope definitions for ICT services and the space-agency standardization tracks that maintain spacecraft environmental requirements, including the sector life-cycle guidance that already exists [41] and the debris-mitigation requirements that already bind disposal [54]. A minimal and immediately actionable step would be a disclosure line requiring any published orbital-compute carbon figure to state its position on all five axes. That alone would make the studies we reconcile here directly comparable, without settling which position is correct.

The method also generalizes beyond the orbital case, though we have not demonstrated it outside orbit and offer it as a conjecture rather than a result. The two-stage logic establishes viability from a workload's data provenance and latency tolerance, and only then evaluates carbon under an explicit metric. It has an obvious terrestrial counterpart: a training run whose corpus already sits in one region faces the same provenance gate, and a carbon-aware scheduler that ignores it can recommend a cleaner-grid region whose egress cost exceeds the carbon it saves. Whether the ordering pays off there is an empirical question we leave open. What does transfer without qualification is the narrower methodological point: where a placement verdict depends on a contested climate metric, the metric belongs in the decision as an explicit axis rather than inside a constant.

The reference orbit is one architecture among several, and a reader may reasonably ask whether it is the favorable one. The model can answer that for the illumination term directly, because eclipse exposure follows from orbital

geometry alone. We compute the sun-synchronous inclination for each altitude, derive eclipse statistics over one year, and treat mean solar unavailability as lost duty cycle. Table 8 compares the dawn-dusk plane at the altitudes the surveyed programs use with a noon-midnight plane at the reference altitude. Two results stand out. The orbital plane matters far more than the altitude. A noon-midnight sun-synchronous orbit at 650 km spends 36% of every orbit in shadow, which as lost duty cycle raises the pessimist break-even to 1.097 kgCO2e/kWh, whereas moving the dawn-dusk plane between 500 and 1,000 km changes it only between 0.741 and 0.711. And eclipse-free operation all year requires roughly 1,395 km, a crossover that moves by a few kilometers with the Earth-radius convention, and which is where the tether concept sits. That altitude is not free. Atmospheric drag is negligible there, so a failed platform has no natural decay path and stays a hazard for centuries, and the band sits inside a secondary maximum of the debris population [49]. An eclipse-free orbit buys a few percent of break-even at the price of a disposal problem that a sustainability framing cannot set aside [48].

Table 8 Illumination of alternative orbital architectures. Illumination only: the disposal propellant priced below reverses the altitude comparison.

Orbit variant	Altitude (km)	Eclipse days per year	Worst-case orbit fraction in shadow	Annual mean solar unavailability	Pessimist break-even (kgCO2e/kWh)
Dawn-dusk SSO	500	104	23.9%	5.5%	0.741
Dawn-dusk SSO (reference)	650	89	20.2%	4.0%	0.729
Dawn-dusk SSO	1000	57	12.4%	1.5%	0.711
Dawn-dusk SSO	1400	0	0.0%	0.0%	0.700
Dawn-dusk SSO (tether concept)	1600	0	0.0%	0.0%	0.700
Noon-midnight SSO	650	365	36.2%	36.2%	1.097

Illumination is not the only term that moves with altitude, and the others move the other way. Launch energy per kilogram rises with altitude and raises the launch factor. Radiation dose rises toward the inner belt and shortens hardware lifetime, one of the leading levers of Section 4.2. Debris flux peaks in the 800 to 1,000 km band [55], so the reference altitude sits below the peak while the eclipse-free altitudes sit above the primary peak and inside a secondary maximum. We do not have sourced altitude-resolved functions for these terms, so we cannot say which altitude minimizes the break-even. What we can say is that the dawn-dusk plane dominates the alternative plane on illumination by a wide margin, which is why the surveyed programs converge on it. Within the dawn-dusk family, the illumination gain from climbing is small, a few percent of the break-even, against penalties in launch energy, radiation and debris that the literature has not yet priced.

The altitude comparison also carries an omission that it cannot absorb, and pricing it reverses the comparison. Neither our model nor Table 8 carries a disposal term. A controlled deorbit from 650 km needs about 170 m/s, and one from 1,395 km about 350 m/s, which at a specific impulse of 220 s is 7.6 and 14.9 percent of wet mass respectively. Charged to the launched mass, the pessimist pole becomes 0.79 kgCO2e/kWh at 650 km with its eclipse season and 0.82 at 1,395 km without one. Climbing to the eclipse-free altitude therefore loses what Table

8 appears to show it gaining, before any of the other altitude-dependent terms are counted. The comparison is conservative in the orbit-favourable direction on both sides, since it charges no propellant for station keeping and none for collision avoidance.

The break-even is defined per kilowatt, which makes it scale-invariant by construction: doubling the installed capacity doubles both sides of Equation (1) and leaves I* unchanged. That invariance holds for a single platform at the margin and fails for a fleet, through at least three channels the model does not represent. The first is cross-section. At the idealised two square meters of radiating surface per kilowatt of Section 2.3, one megawatt of orbital compute carries of order 2,000 m2 of radiator, and a hundred megawatts of order 0.2 km2, before arrays; at the design case of Section 2.3 both figures roughly double. Collision exposure grows in proportion to area, and collective debris risk grows faster than the number of objects once fragmentation begins to feed itself [56], [57], so a fleet imposes a cost that no per-kilowatt figure captures. A sector-wide projection of this growth reaches the same conclusion from the inventory side, finding the climate impact of space activity rising faster than its mass [58], [59]. The second is launch cadence. The high-altitude forcing the multiplier prices is short-lived, so at a steady launch rate it approaches a plateau while a rising rate raises it in proportion. A pulse-based equivalence of the kind our multiplier embodies represents neither, and warming-equivalent metrics were developed for exactly this situation [38], [39]. The third is shared infrastructure: at 80 kg/kW a hundred megawatts means some 8,000 tonnes in orbit, served by ground stations and spectrum whose capacity scales with data volume rather than with compute. The paper's results should therefore be read as statements about the marginal platform, and a fleet-level assessment would need a model in which these three terms depend on the size of the fleet.

Cost is outside the paper's scope, but the framework extends to it without new structure, and the extension is informative. Replacing the launch carbon factor with a launch price per kilogram in Equation (1) turns the carbon break-even into a cost break-even. Orbit is then cheaper to operate than the ground only if terrestrial electricity costs more than the launched mass times the launch price, divided by the same delivered-energy denominator. At the reference mass of 80 kg/kW, every 1,000 USD per kilogram of launch price requires terrestrial electricity to be dearer by 2.17 USD/kWh for orbit to break even on operating cost alone, before any difference in hardware, bus or ground-segment capital is counted. An independent system study of gigawatt-class space-based solar power reaches the same structure, with launch and manufacturing dominant and a levelized cost of 0.61 to 1.59 USD/kWh against 0.02 to 0.05 for terrestrial renewables [23]. We make no claim about what launch prices will be. The point is that the cost comparison is governed by the same product of mass and a per-kilogram factor as the carbon comparison, so the two can be combined through a carbon price in a single break-even, and mass reduction improves both at once.

None of this makes orbital compute good or bad; it makes the question well-posed. The framework's value is to convert a contradictory literature into a parameterized, metric-explicit decision. It also makes precise the quantities whose resolution would most narrow the verdict: the forcing multiplier, the born-in-orbit downlink credit, and the per-kilowatt water term. Those are the natural next steps (Section 7), and a reader can substitute their own estimates for them inside the open calculator.

7. Limitations and Threats to Validity

The central threat to validity is also the paper's thesis. The forcing multiplier, the ratio of CO2-equivalent to direct CO2 for the combustion share of a launch, is deeply uncertain and spans an order of magnitude, which is why we sweep it from 1 to 50 rather than drawing it. It depends on a climate-metric choice that the atmospheric-science literature has not settled, and it also bundles an empirical quantity with that choice, since the black-carbon emission index of a given propellant is measurable in principle. Section 9.3 asks for the empirical half to

be measured and the conventional half to be declared. Black carbon dominates rocket radiative forcing while CO2 is minor [9], [33]. Converting that forcing into a single CO2-equivalent number requires choosing a time horizon, instantaneous radiative forcing, GWP-20, or GWP-100, and an emission index, and these remain contested. A second, narrower judgement sits inside the first. We apply the multiplier to the combustion component of the launch factor and not to the manufacturing and propellant-production component, on the reasoning that only combustion products reach the stratosphere. This is the defensible reading, but it is a reading, and it is consequential. Applying the multiplier to the entire embodied factor instead would raise the pessimist break-even from 0.70 to 1.52 kgCO2e/kWh and turn the additive relation of Corollary 1 into a multiplicative one. A reader who prefers that convention can reproduce it in the calculator, and we report the alternative here so the choice is visible rather than buried. We therefore do not, and cannot, report a single carbon verdict for orbital compute: the verdict moves with the metric, and that movement is precisely contribution C1. This is a genuine limitation, since a decision-maker wanting one number will not find it here, but it is equally the finding.

The opposite convention bounds the other side of the same judgement. A reader who keys the multiplier to the stratospheric share of the combustion term alone, pricing the remainder as ordinary CO2, obtains a lower pole. At a stratospheric share of 0.35 the pessimist break-even falls from 0.70 to 0.34 kgCO2e/kWh and the ratio of Corollary 1 from 11.5 to 5.65; across a share of 0.25 to 0.40 that pole spans 0.29 to 0.37. This is a different convention rather than a correction of ours, because a multiplier keyed to the stratospheric share alone calls for a larger value on that share, black-carbon residence being longer there. Taken together the three readings put the pessimist pole between 0.29 and 1.52 kgCO2e/kWh at a fixed multiplier of 10, which is the honest width of the convention axis once the multiplier itself is held still.

Several second-order terms are set to zero rather than independently estimated, and we now report what that costs. The embodied carbon of orbital manufacturing and the avoided carbon of not building a terrestrial facility, B and D in the break-even, are assumed to roughly offset. Rather than invent a distribution for them, we give their leverage exactly: the break-even shifts by the net imbalance divided by the denominator of Equation (1), so a net imbalance of 1000 kgCO2/kW moves I* by 0.027 kgCO2/kWh. Against the optimist pole of 0.06 that is a 45 percent shift, and against the pessimist pole of 0.70 it is 4 percent. The optimist verdict is therefore materially conditional on B being approximately equal to D and should be read as such, while the pessimist verdict is robust to plausible imbalances. A reader with better estimates of either term can substitute them directly. Water enters the multi-criteria profile only directionally, since orbit draws essentially no operational water while a terrestrial facility consumes both cooling and grid water. The per-kilowatt magnitude is now sourceable from facility water-usage data [1], but it is not yet folded into a quantitative water break-even, and the water and carbon criteria are assessed on different system boundaries, as Section 5.3 states. The born-in-orbit downlink credit, which underwrites the carbon half of the Earth-observation worked case under the strictest metric, is likewise a directional adjustment. Processing data where it is generated avoids a large raw-data downlink, but we do not quantify that saving, which is why Section 5.4 reports that case as conditional on the credit rather than attaching a confidence label to it. Of the outstanding terms this is the one whose resolution would most change a reported verdict.

The headline numbers are specific to one launch vehicle and one orbital regime. The launch factors are anchored to a Falcon-class vehicle delivering payload to a dawn-dusk sun-synchronous orbit near 650 km. A cleaner launch system, a future reusable methalox vehicle or the roughly ten-times-less-emissive launcher that the ASCEND study [5] treats as a precondition, shifts the launch factor and therefore the break-even substantially. The model is parameterized for this, and the calculator lets a reader re-run any vehicle, but the specific intensities reported here are not vehicle-agnostic. The reference orbit carries its own caveats. Operational carbon is taken as approximately zero because a dawn-dusk sun-synchronous orbit receives near-continuous sunlight, which is a property of that specific regime and not of sun-synchronous orbits in general. Altitude couples to launch CO2,

high-altitude forcing, radiation dose and thus hardware lifetime, and latency, so moving the reference altitude moves several terms at once.

Several further assumptions deserve flagging. We assume the embodied carbon of the compute hardware is equal on the orbital and terrestrial sides and cancels. If orbital hardware requires radiation-hardening or a different refresh cadence this no longer holds, though radiation testing of commercial accelerators for multi-year orbital missions suggests the gap may be smaller than once assumed. The decision rule's viability thresholds are order-of-magnitude. The latency classes and the low-Earth-orbit round-trip figures we use are representative, and precise viability for a given workload depends on its actual latency budget and on the constellation's visibility and handover profile, which we do not model. Two couplings that Section 3.2 declares are also left unmodeled, and both cut against orbit at the top of the altitude sweep. Radiation dose rises steeply toward the inner proton belt, so the hardware lifetime L that our sensitivity analysis identifies as a leading lever is not independent of the altitude we sweep to 1600 km. A design standard for radiation-induced degradation [60] would be the natural basis for making L altitude-dependent, which we have not done. Debris flux likewise peaks in the 800 to 1000 km band that the sweep crosses [56], [57], [55], so the ordinal debris penalty in Table 6 understates the altitude dependence of the risk. Finally, our treatment of stratospheric externalities is asymmetric. We price launch black carbon quantitatively through the multiplier, but re-entry ablation products and launch-driven ozone depletion appear only as ordinal entries [36], [34], even though they arise from the same atmospheric physics and, unlike a radiative-forcing term, are not representable by a CO2-equivalent multiplier at any value.

Three terms that a full accounting would carry are absent from the orbital side of Equation (1). Because they all enter through the same channel, their leverage can be stated exactly rather than estimated: any additional orbital-side carbon shifts the break-even by its per-kilowatt magnitude divided by the denominator of Equation (1), that is by 2.72 x 10^-5 kgCO2/kWh per kgCO2/kW. First, the model carries no hardware redundancy. Fault-tolerant orbital operation requires cross-strapped avionics, spare compute lanes, and margin against single-event effects, none of which appear in the launched mass per kilowatt. Because I* is exactly linear in M, a redundancy factor of 1.15 raises the reference break-even from 0.70 to 0.81 kgCO2e/kWh, a factor of 1.3 raises it to 0.91, and a factor of 1.5 raises it to 1.05. Our mass envelope is therefore optimistic for any design that carries real redundancy, and a reader with a redundancy budget can rescale the reported thresholds directly.

Second, the equal-embodied-hardware assumption stated above carries a leverage worth quantifying. If radiation-hardened parts, additional screening, or a shorter qualification-driven refresh cycle make orbital hardware more carbon-intensive to produce, the difference enters as B. An asymmetry of 1000 kgCO2/kW moves I* by 0.027, which is 4 percent of the pessimist pole but 45 percent of the optimist pole, so this assumption is load-bearing for the optimist verdict and nearly irrelevant to the pessimist one. Third, we exclude the energy and carbon of the communication infrastructure that an orbital data center depends on, namely ground stations, feeder links, and any relay spacecraft. On the same leverage, 500 kgCO2/kW of ground-segment burden moves I* by 0.014, about 2 percent of the pessimist pole. The omission is therefore small against the metric axis but not negligible against the optimist pole, and it works in the same direction as the downlink credit of Section 5.4 works against it, since both concern the same ground segment.

Two assumptions about time deserve separate treatment. The first is utilization. We hold U constant over the mission, whereas real AI workloads vary substantially across the day and across a hardware generation. For a lifetime break-even this matters less than it appears. The delivered compute energy is H*L*mean(U), so I* depends on the mean of the utilization profile and not on its shape. A flat profile, a bimodal profile alternating between the endpoints of the tabulated range, and a linear ramp across it all return I* = 0.700 identically once normalized to the same mean. Load-dependent power-usage effectiveness does not reopen the question either, at least not under the simplest model of it: if fixed facility overhead is spread over a varying IT load, the energy-

weighted power-usage effectiveness is one plus the overhead divided by the mean utilization, so the profile shape cancels there as well. A second-order effect would require a PUE model whose load dependence is not of that form, and we do not have a sourced one. The constant-utilization assumption is therefore exact in its first-order effect and unquantified beyond it. The second assumption is hardware lifetime under a varying radiation environment. We treat L as a single distribution, whereas the dose that drives degradation varies with altitude, with inclination through the South Atlantic Anomaly, and with the solar cycle. A five-year lifetime at solar minimum near 650 km is not the same five years as at solar maximum near 1200 km. Since L is one of the two leading levers in the variance decomposition of Section 4.2, making it environment-dependent is among the more valuable extensions available, and the radiation-effects design methodology we cite is the natural basis for it.

One parameter does three jobs at once, and future work should split it. The hardware lifetime L sets the amortization period for a radiator-dominated structure that would outlive it, the comparison period against the ground, and the cadence at which accelerators are replaced. On the ground a refresh is an economic decision taken against continuous repair; in orbit it is a wear-out wall, and a replacement cycle would require launching the compute mass again, with the launch factor charged a second time. Our single-L treatment therefore flatters orbit in two ways at once, and since lifetime carries a total-order sensitivity index near 0.3 in Table 3 the simplification is load-bearing rather than cosmetic. A platform life distinct from a refresh interval is the modelling change we would make first.

Finally, one limitation concerns what this work has not been tested against. The framework has not been validated against operational data from a flown orbital compute platform. Inference has flown on small Earth-observation satellites, where an on-board neural network selects what is worth downlinking [61], but nothing has flown at the power scale the model addresses, and no operator has published the per-kilowatt mass, lifetime, and downlink figures the model consumes. Every parameter is drawn from published design studies, regulatory filings, and terrestrial analogues, so the paper should be read as a framework whose inputs are declared and replaceable rather than as a calibrated predictor. The first mission to publish a mass and energy budget would allow a direct confrontation, and we regard that as the most valuable single piece of future work.

The Monte Carlo analysis propagates parametric uncertainty, through independent triangular draws over the sourced ranges, but not structural uncertainty. The independence assumption is itself a limitation rather than a modeling convenience, for the reason just given: altitude drives several parameters jointly. The reported intervals are therefore not bounds in either direction: against the structure in which mass buys lifetime they are 8 to 16 percent too wide, and against its mirror image 16 to 23 percent too narrow, as Section 4.3 reports. The choice of accounting boundary and climate metric, the C1 axis, is not a distribution to be averaged over but a scenario to be made explicit, and we treat it as such, so the confidence intervals quantify uncertainty within a metric, not across metrics. We counter these limitations in four ways rather than hiding them. We report the verdict under every metric rather than a convenient one. We evaluate multiple criteria as a profile rather than collapsing them into a single contestable score. We state in closed form the leverage of the terms we do not propagate. And we release the model and parameters as an open calculator, so that any reader can replace our assumptions with their own.

8. OrbitOrEarth: an open deployment of the model

The model developed in the preceding sections is instantiated as OrbitOrEarth (Fig. 8), an open, self-contained calculator. It deploys the break-even and multi-criteria framework as an interactive decision tool rather than leaving it as a closed-form result available only in static tables and figures. This deployment serves a methodological purpose beyond convenience. The central claim of the paper is that the optimist-pessimist

disagreement reduces to an explicit choice of accounting boundary and climate metric. That claim is tested more rigorously by a tool that lets a reader impose the opposing choice and watch the model reproduce the opposing verdict than by prose description alone.

The calculator exposes every parameter of the model (Table 2) as a direct control. Continuous inputs cover the launched system mass, the hardware lifetime, the forcing multiplier, the power-usage effectiveness, the utilization, the displaced grid intensity on a location basis, and the contracted intensity on a market basis. Two discrete toggles correspond to the metric definition of Section 3.3: the accounting boundary (combustion-only versus embodied) and the climate metric (direct CO2 versus CO2-equivalent under high-altitude radiative forcing). Three presets instantiate the configurations examined in Section 5.2: an optimist preset reproducing the NTU accounting choices [7], a pessimist preset reproducing the Saarland accounting choices [8], and the reference configuration of Table 2. The computational sequence from a user action to a displayed recommendation runs in six steps, and it is the same sequence as Algorithm 1, so that the interactive and printed forms cannot diverge. First, the control values are read and merged over the parameter file, so any control the user has not touched retains its tabulated central value.

Second, the two toggles are resolved into the boundary indicator and the multiplier of Equation (4), and the launch factor is assembled from its combustion and non-combustion components. Third, the point-estimate break-even is evaluated in closed form from Equation (1). Fourth, the Monte Carlo loop is run over the propagated parameters at the current scenario, with the multiplier held fixed, and the 5th and 95th percentiles are extracted. Fifth, the displaced grid intensity is compared against the point estimate to produce the verdict, against the interval to produce the confidence label, and against the draws themselves to produce the win probability, following the mapping of Section 4.5. The comparison is run under both measurement bases and both verdicts are displayed, as Section 3.1 requires. Sixth, the contour map is redrawn by evaluating the closed form across the mass axis at the three metric settings, and the current operating point is placed on it. Steps three through six re-run on every control change, so the displayed verdict, interval, label and map are always mutually consistent; no result is cached across a parameter change. Selection of any preset propagates through this same pipeline. The transition between presets therefore renders the reconciliation reported in Table 4 as a live, user-controlled operation rather than a fixed outcome presented after the fact.

The calculator's principal output is the break-even contour map introduced in Section 1 and analyzed in Section 5.1. The three contours are generated by evaluating the closed-form expression of Section 3.5 at the current control settings, so any operating point can be located on the map and watched as it shifts with the accounting and metric assumptions. The calculator also reproduces the one-at-a-time sensitivity ranking of Section 4.2 and the two-stage decision procedure of Algorithm 1. The conformance suite that accompanies the release covers the closed form, the decomposition of Equation (4), Corollary 1 and its invariance conditions, the monotonicity signs of Proposition 1, the Monte Carlo intervals at the declared sample size and seed, the forcing-multiplier sweep, and the parameter file against Table 2. It also covers the results that earlier releases computed outside the code path: the closed-form indices of Table 3 and the sum of the first-order indices, the correlated propagation of Section 4.3 under both declared copula structures, the marginal-family sensitivity, the win probabilities of Sections 4.5 and 5.2, the alternative multiplier convention of Section 7, and the leverage of the two terms Section 3.5 declares at unity. Every quantity this paper reports of its own artifact is therefore checked by the suite rather than asserted.

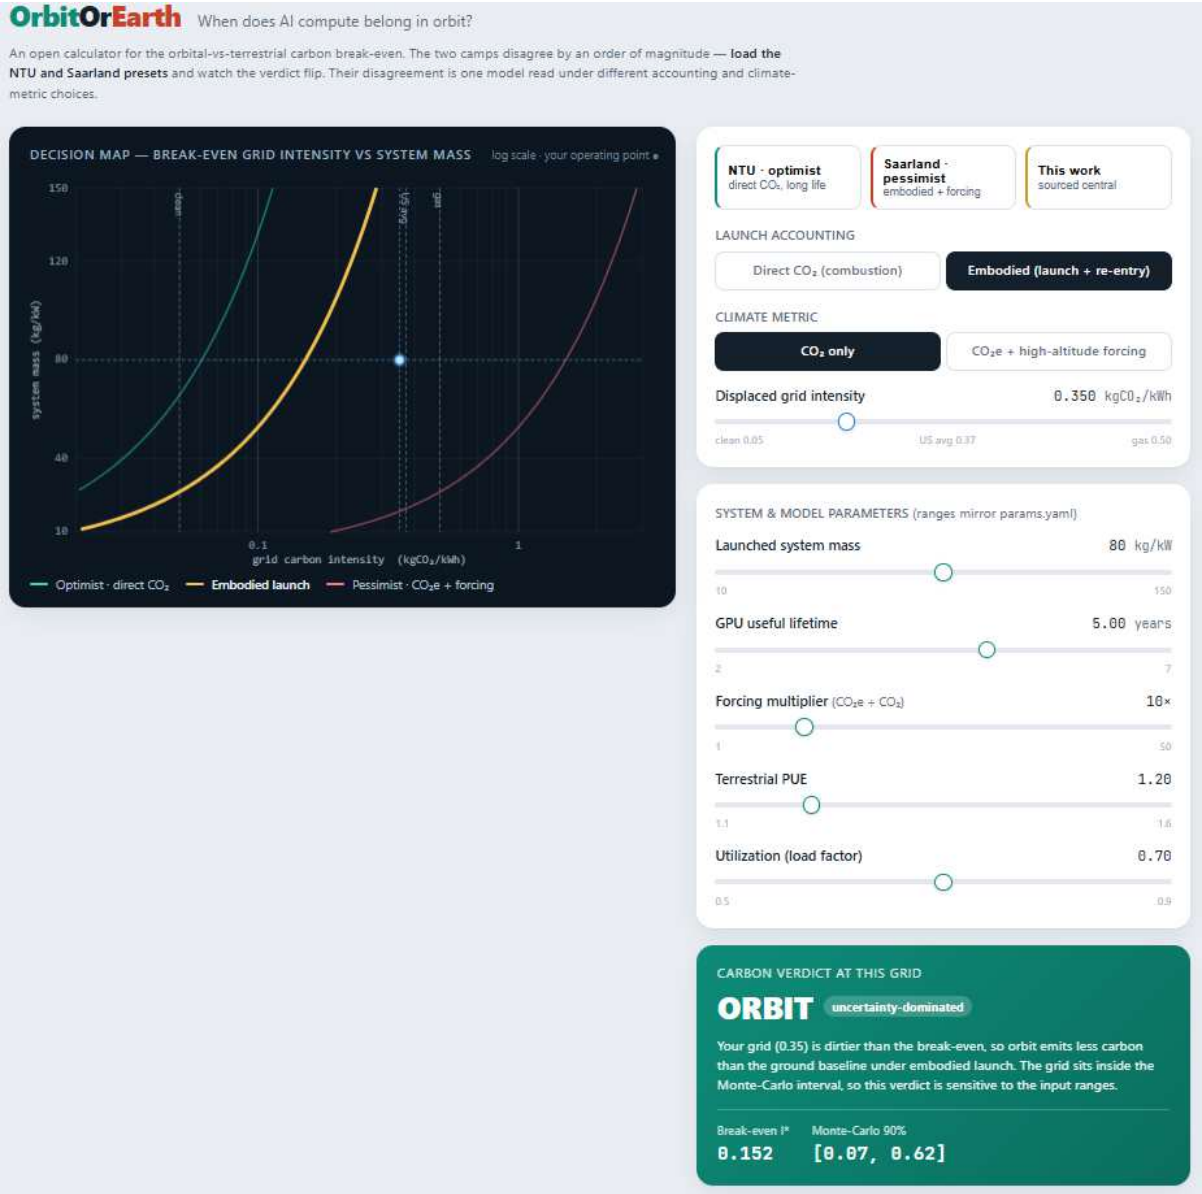


Fig. 8 *OrbitOrEarth open calculator*

To support independent verification, the model, its sourced parameter file, and the calculator implementation are released under an open license. Either camp, or any subsequent reviewer, can substitute alternative parameter assumptions and reproduce or contest any figure reported in this paper. The release carries a conformance suite that pins the implementation to this paper rather than to itself. Its 77 tests check the closed form of Equation (1) against an independent reimplementation, the decomposition of Equation (4), the additive ratio of Corollary 1 and its invariance in the denominator parameters, the monotonicity signs of Proposition 1, the Monte Carlo intervals at the declared sample size and seed, the forcing-multiplier sweep of Table 5, and the parameter file against Table 2. They also check the win probabilities and the band in which the point estimate and the median disagree, the sensitivity indices in closed form, both copula structures, the marginal families, and the thermal sizing behind the mass envelope. The suite covers the break-even results computed in the code path; the disposal, mass-budget and orbital-overhead figures introduced in this revision are stated in closed form and can be checked by hand from the leverage identities given with them. The suite passes in full at the released version, and it is

the mechanism by which a reader can establish that the code and the paper have not drifted apart. The artifact is permanently archived on Zenodo under the concept DOI 10.5281/zenodo.20946603 (https://doi.org/10.5281/zenodo.20946603), which always resolves to the current release. The version described here is v1.1.2. It restricts the launched-mass parameter to its feasible support, as v1.1.1 first did, and adds the quantities this paper reports that earlier releases did not compute: the win probability, dual-basis reporting, the closed-form indices, the correlated propagation, the marginal families, and the two denominator terms declared at unity. The archived release comprises the model implementation, the parameter file underlying Table 2, the scripts that generate the reported figures, the generated figures themselves, the conformance suite, and the calculator.

One subsection of the release is about how not to use it. The calculator makes it easy to produce a single favorable number, and the paper's own framing invites that: a user may choose a metric, a measurement basis, a counterfactual grid and an unquantified credit, and report the result as the carbon figure for an orbital deployment. That is the opposite of the intent. The defaults ship at the location-based grid, at the reference multiplier, with the downlink credit at zero and the second-order terms at zero, and the three contours are displayed together rather than one at a time. A verdict quoted from this tool should travel with the four choices that produced it, as the disclosure line of Section 6 requires of any published figure. A verdict that holds under all three contours is robust in the sense this paper can defend; any other is metric-determined, and saying so is part of reporting it.

9. Conclusion

This paper develops a lifecycle carbon break-even framework to assess when orbital AI computing may be environmentally competitive with terrestrial infrastructure. The analysis shows that the substantial disagreement between existing estimates largely reflects differences in launch-emissions boundaries and climate metrics. Under the modeled assumptions, the break-even grid intensity ranges from 0.06 to 0.70 kg CO_2/kWh, with a ratio of 11.5 between the optimistic and pessimistic cases. The results also show that launched system mass, hardware lifetime, and launch-emissions factors influence the break-even threshold, while parameter correlations modify the estimates without eliminating the gap between climate scenarios. Accounting for eclipses in the reference dawn-dusk orbit near 650 km further increases the pessimistic threshold from 0.700 to 0.729 kg CO_2e/kWh. These findings demonstrate that the environmental advantage of orbital computing depends strongly on the accounting boundary, climate metric, and terrestrial electricity source.

The framework combines carbon assessment with a workload-placement rule that evaluates operational feasibility before environmental performance. Under the modeled conditions, terrestrial infrastructure remains preferable for most workloads, particularly latency-sensitive inference serving ground users. Orbital computing may be more attractive for workloads originating in space, such as Earth-observation processing, although the potential benefit of avoided data transmission requires further quantification. The open OrbitOrEarth calculator enables users to explore the model assumptions and their effects on the resulting decisions.

These results should be interpreted as scenario-based estimates rather than calibrated predictions, because operational data from orbital computing platforms at the relevant scale are not yet available. The model also omits or simplifies several lifecycle factors, including re-entry emissions, hardware redundancy, ground-segment energy, and differences between orbital manufacturing and terrestrial infrastructure construction. Future work should prioritize measurements of launch-related climate effects, Earth-observation downlink emissions, orbital platform mass and lifetime, and workload latency under realistic operating conditions. Incorporating these

data into an altitude-resolved lifecycle model would strengthen the assessment and help establish when orbital computing can provide a credible environmental advantage over terrestrial alternatives.

References

1. Shehabi, A., Smith, S.J., Hubbard, A., Newkirk, A., Lei, N., Siddik, M.A.B., Holecek, B., Koomey, J.G., Masanet, E.R., Sartor, D.A.: 2024 United States data center energy usage report (2024)
2. International Energy Agency: Energy and AI (2025). https://www.iea.org/reports/energy-and-ai. Accessed 20 November 2026
3. Masanet, E., Shehabi, A., Lei, N., Smith, S., Koomey, J.: Recalibrating global data center energy-use estimates. Science 367(6481), 984-986 (2020)
4. Google Research: Towards a future space-based, highly scalable AI infrastructure system design (Project Suncatcher) (2025). https://research.google/. Accessed 20 November 2026
5. Thales Alenia Space: ASCEND: Advanced Space Cloud for European Net zero emission and Data sovereignty, feasibility study results (2024). https://ascend-horizon.eu/. Accessed 20 November 2026
6. Li, A.S.: Zigzagging Outer Space: The Legal Complexities of Mega-Constellations. Nebraska Law Review 104(2), 149 (2025)
7. Aili, A., Choi, J., Ong, Y.S., Wen, Y.: The development of carbon-neutral data centres in space. Nature Electronics, 1-11 (2025)
8. Ohs, R., Stock, G.F., Schmidt, A., Fraire, J.A., Hermanns, H.: Dirty bits in low-Earth orbit: The carbon footprint of launching computers. ACM SIGENERGY Energy Informatics Review 5(2), 26-33 (2025)
9. Ross, M.N., Sheaffer, P.M.: Radiative forcing caused by rocket engine emissions. Earth's Future 2(4), 177-196 (2014)
10. Bargatin, I., Jin, D., Alansari, Z., Raney, J.R.: Tether-based architecture for solar-powered orbital AI data centers. AIAA SCITECH 2026 Forum, 1237 (2026)
11. Marcy, G.W.: The impact of computing data centres orbiting Earth (arXiv preprint arXiv:2603.28829, 2026)
12. Federal Aviation Administration: Environmental assessment for SpaceX Falcon launch vehicle at Kennedy Space Center and Cape Canaveral Air Force Station (2020). https://www.faa.gov/. Accessed 20 November 2026
13. Denby, B., Lucia, B.: Orbital edge computing: Machine inference in space. IEEE Computer Architecture Letters 18(1), 59-62 (2019)
14. Denby, B., Lucia, B.: Orbital edge computing: Nanosatellite constellations as a new class of computer system. In: Proceedings of the 25th International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '20), pp. 939-954 (2020)
15. Denby, B., Chintalapudi, K., Chandra, R., Lucia, B., Noghabi, S.: Kodan: Addressing the computational bottleneck in space. In: Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3 (ASPLOS '23), pp. 392-403 (2023)
16. Bhattacherjee, D., Kassing, S., Licciardello, M., Singla, A.: In-orbit computing: An outlandish thought experiment? In: Proceedings of the 19th ACM Workshop on Hot Topics in Networks (HotNets '20), pp. 197-204 (2020)
17. Vasisht, D., Shenoy, J., Chandra, R.: L2D2: Low latency distributed downlink for LEO satellites. In: Proceedings of the ACM SIGCOMM 2021 Conference, pp. 151-164 (2021)

18. Pfandzelter, T., Bermbach, D.: Komet: A serverless platform for low-Earth orbit edge services. In: Proceedings of the 2024 ACM Symposium on Cloud Computing (SoCC '24) (2024)
19. Michel, F., Trevisan, M., Giordano, D., Bonaventure, O.: A first look at Starlink performance. In: Proceedings of the 22nd ACM Internet Measurement Conference (IMC '22), pp. 130-136 (2022)
20. Nadjahi, C., Louahlia, H., Lemasson, S.: A review of thermal management and innovative cooling strategies for data center. Sustainable Computing: Informatics and Systems 19, 14-28 (2018)
21. Gilmore, D.G. (ed.): Spacecraft Thermal Control Handbook, Volume I: Fundamental Technologies, 2nd ed. The Aerospace Press / AIAA (2002)
22. NASA Small Spacecraft Systems Virtual Institute: State-of-the-Art Small Spacecraft Technology, NASA/TP-20260003140. NASA Ames Research Center (2026)
23. NASA Office of Technology, Policy, Strategy: Space-based solar power. National Aeronautics and Space Administration (2024)
24. Larson, E.J.L., et al.: Global atmospheric response to emissions from a proposed reusable space launch system. Earth's Future 5, 37-48 (2017)
25. Dallas, J.A., Raval, S., Alvarez Gaitan, J.P., Saydam, S., Dempster, A.G.: The environmental impact of emissions from space launches: a comprehensive review. Journal of Cleaner Production 255, art. no. 120209 (2020)
26. Dominguez Calabuig, G.J., Wilson, A., Bi, S., Vasile, M., Sippel, M., Tajmar, M.: Environmental life cycle assessment of reusable launch vehicle fleets: large climate impact driven by rocket exhaust emissions. Acta Astronautica 221, 1-11 (2024)
27. Maury, T., Loubet, P., Serrano, S.M., Gallice, A., Sonnemann, G.: Application of environmental life cycle assessment (LCA) within the space sector: a state of the art. Acta Astronautica 170, 122-135 (2020)
28. Gupta, U., et al.: Chasing carbon: the elusive environmental footprint of computing. In: IEEE International Symposium on High-Performance Computer Architecture (HPCA), pp. 854-867 (2021)
29. World Resources Institute, World Business Council for Sustainable Development: GHG Protocol Scope 2 Guidance: An Amendment to the GHG Protocol Corporate Standard. World Resources Institute (2015)
30. Brander, M., Gillenwater, M., Ascui, F.: Creative accounting: a critical perspective on the market-based method for reporting purchased electricity (scope 2) emissions. Energy Policy 112, 29-33 (2018)
31. Vallado, D.A.: Fundamentals of Astrodynamics and Applications, 4th ed. Microcosm Press, Hawthorne, CA (2013)
32. Sheme, E., Holmbacka, S., Lafond, S., Lučanin, D., Frashëri, N.: Feasibility of using renewable energy to supply data centers in 60 north latitude. Sustainable computing: Informatics and systems, 96-106 (2018)
33. Ryan, R.G., Marais, E.A., Balhatchet, C.J., Eastham, S.D.: Impact of rocket launch and space debris air pollutant emissions on stratospheric ozone and global climate. Earth's Future 10(6), e2021EF002612 (2022)
34. Maloney, E.M., Portmann, R.W., Ross, M.N., Rosenlof, K.H.: The climate and ozone impacts of black carbon emissions from global rocket launches. Journal of Geophysical Research: Atmospheres 127, art. no. e2021JD036373 (2022)
35. Bond, T.C., et al.: Bounding the role of black carbon in the climate system: a scientific assessment. Journal of Geophysical Research: Atmospheres 118(11), 5380-5552 (2013)
36. Murphy, D.M., et al.: Metals from spacecraft reentry in stratospheric aerosol particles. Proceedings of the National Academy of Sciences 120(43), art. no. e2313374120 (2023)
37. Forster, P., Storelvmo, T., Armour, K., Collins, W., Dufresne, J.-L., Frame, D., Lunt, D.J., Mauritsen, T., Palmer, M.D., Watanabe, M., Wild, M., Zhang, H.: The Earth's energy budget, climate feedbacks, and

climate sensitivity. In: Climate Change 2021: The Physical Science Basis. Contribution of Working Group I to the Sixth Assessment Report of the IPCC, pp. 923-1054. Cambridge University Press (2021)

38. Allen, M.R., Shine, K.P., Fuglestvedt, J.S., Millar, R.J., Cain, M., Frame, D.J., Macey, A.H.: A solution to the misrepresentations of CO2-equivalent emissions of short-lived climate pollutants under ambitious mitigation. npj Climate and Atmospheric Science 1, art. no. 16 (2018)
39. Cain, M., Lynch, J., Allen, M.R., Fuglestvedt, J.S., Frame, D.J., Macey, A.H.: Improved calculation of warming-equivalent emissions for short-lived climate pollutants. npj Climate and Atmospheric Science 2, art. no. 29 (2019)
40. Lee, D.S., et al.: The contribution of global aviation to anthropogenic climate forcing for 2000 to 2018. Atmospheric Environment 244, art. no. 117834 (2021)
41. European Space Agency, Clean Space: Space system life cycle assessment (LCA) guidelines. European Space Agency (2016)
42. Ross, M., Mills, M., Toohey, D.: Potential climate impact of black carbon emitted by rockets. Geophysical Research Letters 37(24) (2010)
43. Uptime Institute: Global data center survey 2024 (2024). https://uptimeinstitute.com/. Accessed 20 November 2026
44. Mytton, D.: Data centre water consumption. npj Clean Water 4, art. no. 11 (2021)
45. Siddik, M.A.B., Shehabi, A., Marston, L.: The environmental footprint of data centers in the United States. Environmental Research Letters 16(6), art. no. 064017 (2021)
46. Lin, X., Hofstrom, B., Wang, Y.-P.E., Masini, G., Maattanen, H.-L., Ryden, H., et al.: 5G New Radio evolution meets satellite communications: Opportunities, challenges, and solutions. In: 5G and Beyond: Fundamentals and Standards, 517-531 (2021)
47. Rathnasabapathy, M., et al.: Space sustainability rating: designing a composite indicator to incentivise satellite operators to pursue long-term sustainability of the space environment. In: 71st International Astronautical Congress (IAC), paper IAC-20 (2020)
48. Lawrence, A., et al.: The case for space environmentalism. Nature Astronomy 6, 428-435 (2022)
49. Boley, A.C., Byers, M.: Satellite mega-constellations create risks in low Earth orbit, the atmosphere and on Earth. Scientific Reports 11, art. no. 10642 (2021)
50. Gupta, U., Elgamal, M., Hills, G., Wei, G.-Y., Lee, H.-H.S., Brooks, D., Wu, C.-J.: ACT: Designing sustainable computer systems with an architectural carbon modeling tool. In Proceedings of the 49th Annual International Symposium on Computer Architecture, 784-799 (2022)
51. International Organization for Standardization: Environmental management - Life cycle assessment - Principles and framework (2006)
52. Kline Jr, D., Parshook, N., Ge, X., Brunvand, E., Melhem, R., Chrysanthis, P.K., Jones, A.K.: GreenChip: A tool for evaluating holistic sustainability of modern computing systems. Sustainable Computing: Informatics and Systems 22, 322-332 (2019)
53. BIPM/JCGM 101: Evaluation of measurement data - Supplement 1 to the 'Guide to the expression of uncertainty in measurement' - Propagation of distributions using a Monte Carlo method. Joint Committee for Guides in Metrology (BIPM, 2008)
54. International Organization for Standardization: ISO 24113: Space systems - Space debris mitigation requirements. ISO (2023)
55. ESA Space Debris Office: ESA's Annual Space Environment Report, Issue 9, Rev. 1. European Space Agency (2025)

56. Kessler, D.J., Cour-Palais, B.G.: Collision frequency of artificial satellites: The creation of a debris belt. Journal of Geophysical Research: Space Physics 83(A6), 2637-2646 (1978)
57. Liou, J.-C., Johnson, N.L.: Risks in space from orbiting debris. Science 311(5759), 340-341 (2006)
58. Miraux, L., Wilson, A.R., Dominguez Calabuig, G.J.: Environmental sustainability of future proposed space activities. Acta Astronautica 200, 329-346 (2022)
59. Miraux, L.: Environmental limits to the space sector's growth. Science of the Total Environment 806, art. no. 150862 (2022)
60. ECSS Secretariat: Space engineering - Methods for the calculation of radiation received and its effects, and a policy for design margins, ECSS-E-ST-10-12C. ESA-ESTEC Requirements and Standards Division (2008)
61. Giuffrida, G., et al.: The Phi-Sat-1 mission: the first on-board deep neural network demonstrator for satellite Earth observation. IEEE Transactions on Geoscience and Remote Sensing 60, 1-14 (2022)

Statements and Declarations

Funding

The authors declare that no funds, grants, or other support were received during the preparation of this manuscript.

Competing Interests

The authors have no relevant financial or non-financial interests to disclose. The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Author Contributions

Abdelbassette Chenna: conceptualization, methodology, software, formal analysis, investigation, data curation, writing of the original draft, review and editing, visualization. Djallel Eddine Boubiche: conceptualization, validation, supervision, review and editing, project administration. Both authors read and approved the final manuscript.

Data Availability

OrbitOrEarth, including the break-even model, its sourced parameter file, the figure-generation scripts, the generated figures, the conformance suite, and the interactive calculator, is openly available and permanently archived on Zenodo under the concept DOI 10.5281/zenodo.20946603 (https://doi.org/10.5281/zenodo.20946603). The version corresponding to the results reported here is v1.1.2. No other data were used.

Declaration of generative AI and AI-assisted technologies in the manuscript preparation process

During the preparation of this manuscript, the authors used generative AI and AI-assisted technologies exclusively for English-language editing, including the correction of grammatical errors and the improvement of clarity and readability. All scientific contributions, including the research design, methodology, data analysis, interpretation of results, and conclusions, are the authors' own. The authors take full responsibility for the content of the manuscript.