Orbital Edge Computing: A Theoretical Framework for Global Autonomous Vehicle Coordination using Nuclear-Powered AI Hubs

Rythm Singh[1]

Received April 12, 2026
Accepted July 20, 2026
Electronic access September 30, 2026

The cost of intelligence for connected vehicles has been paid for by their batteries. The inconvenient truth about electrifying transport is the largely unseen trade-off between "Energy" vs "Compute" budgets at the heart of the autonomous vehicle agenda: every smartsystem we integrate into the car reduces our driving range by an equivalent amount required to power that computer onboard. These tradeoffs can be eliminated by sharing the computational workload with a network of satellites in orbit around Earth. Dubbed Orbital Edge Computing (OEC) the primary distinction between this framework and existing satellite-based systems is where it sources its power. These satellites would be powered by space-born Nuclear Powerplants, likely SMRs running on Uranium harvested from resources found in-space. By prospecting this fissile material autonomously using swarm drones that mine asteroids through ISRU we can keep this orbital infrastructure 100% "Earth-Neutral" (meaning no resources are consumed on Earth).Offloading compute-intensive Computer Vision tasks to just 10–12 nuclear-powered OEC nodes in LEO can recover about 1–2% of a car's battery under realistic traffic patterns during a typical cruise. While that number seems small, at scale those are millions of extra cars at passengers' disposal on a single charge. We can also reduce traffic fatalities to zero by overlaying the world with a "Synchronous Orbital Mesh" of "Over-the- Horizon Surveillance" that systematically removes every last blind spot out there at an average RTT latency of 3.6ms. This paper demonstrates that safety doesn't have to come at the cost of sustainability, and lays the groundwork for how we can finally realize zero-emission and zero-fatality transportation.

Keywords: Orbital Edge Computing, Nuclear AI, SMR, Autonomous Vehicles, ISRU

Introduction

Add an electric motor to a car, it gets more energy efficient; make that car fully autonomous, and it becomes more energy hungry. Put the two trends together and you have a clash (call it the Energy-Compute Paradox) at the heart of the problem this paper addresses: silicon doesn't alone explain this duality. Both Tesla Full Self-Driving and NVIDIA DRIVE Orin take relatively limited slices (around 50-250 W) of power and thus do not bear the responsibility, for instance. Level 4 and 5 autonomy, however, causes a car's computation to be a hungry beast feeding on streams of camera, LiDAR, and radar data while running thermal and cooling systems on the backend, consuming battery range on dense urban drives. Orbital Edge Computing is proposed as the solution that takes that computational bite off the vehicle by offloading macro-contextual perception and secondary processing.

Beyond energy, it is also worth noting the Local Horizon Problem-a vehicle's ground sensors can only see what's directly in front of it. The problem of occlusion is compounded by buildings, weather, and the landscape itself, and can neither be solved nor compensated for by onboard compute hardware that lacks visibility into it. V2V communication is intended to address this by supplementing onboard sensors, but the real world use of V2V relies on high adoption and network stability, which are often lacking.

This paper offers an orbital edge computing approach to solving both problems. By using an array of 10-12 LEO satellites, working in unison as an AI Hub in space, we can provide global sensing and edge computing while decreasing the load on the vehicle. In the long term, our approach hinges on the proposed Earth-Neutral power system of both ISRU and orbital SMRs on an aggregate basis of autonomous exploitation of near-earth objects to feed an array of space-based computers with power. Providing a holistic perception of traffic and edge computing enabled by global space-based sensing, that doesn't weigh down vehicular systems. What we present next, outlines our solution and discusses its possible long term ramifications on autonomous driving.

The system's design is mainly rooted in a larger model of distributed, self-organizing intelligence for autonomous system architectures[1], Crucially, the ethical and societal risks of rolling out large-scale autonomy are thoroughly embedded across the entire framework[2].

[1] *Vidyavati Public Central School, Madhya Pradesh, India*

System Horizon and Boundary Definitions

And this nothing is close to being ready to deploy, not be deferred for way, way far in the future – pure, hard-headed, theory and not a future real world application point blank; because space-based systems to-day are choked by satellite weight limitations and insufficient available power plus power budget and harsh-radiation conditions and requirements. It goes without stating that the exaggerations assumed herein (multi MW SMRs, for example in LEO, communications via multi MW bandwidths from satellite to satellite) should just, truly be interpreted as theoretical performance limitations and not realistic, engineering prognostication of the near-by. None-the-less the examination can be somewhat interesting to perform out at the extreme, the underlying mathematical work to get there; leaving clues which future space exploration may pick up and follow. But will this even approach the real world, then, depends on a very wide spread, all still in the concept development stage; things such as lift weight capability, space based electricity generation, capability to do automated space based servicing, etc.

Literature Review

The in-vehicle perception systems deployed in self-driving cars are beholden to remarkably harsh constraints of both hardware as well as thermal environments. Today most platforms are a collection of intensive heavy convolutions and vision transformers that process both mapping, tracking, localization, along with object recognition in tandem to serve as an autonomous control systems, in order to ensure an extremely minimal use of any networks outside of the vehicle. Li et al. (2021)[3] document that multi-camera and high-channel LiDAR processing units impose a continuous parasitic load on electric vehicle (EV) batteries, a load that shows up as measurable range degradation once you're in dense urban drive cycles. The answers with Mobile Edge Computing(MEC) come with latencies more stringent workloads placed closer to base stations in the ground, Roadside Units (RSU).

Move that layer into the orbit and the routing problem multiplies and becomes very complex. In their work, Wang and Liu (2023)[4] Clearly show and completely reveal how that dynamic multi-tier task offloading in satellite-integrated edge environments gets squeezed and also stressed hard — spatial-temporal channel fluctuations, inter-satellite cross-links, the Doppler effect, all pulling in different directions at once, simultaneously Managing millions of mobile ground targets in real-time only intensifies that squeeze and stress very much. In their work and deep research, Almalki et al. (2024)[5] find that the dense LEO constellations like Starlink etc., when used for the direct-to-satellite vehicular safety systems, demand a very careful, end-to-end latency profiling just to manage the terrestrial uplink jitter and handovers. Also, Zhou et al. (2022)[6] tackle a brilliant, very important and related problem: multi-UAV cooperative edge computing architectures built to absorb topology shifts. Zhang et al. (2022)[7] go further still, also proposing an on-orbit task-offloading and allocation scheme for aerial edge nodes — providing more evidence that the 'distributing perception workloads across space clusters' actually and works(also in real world). In their research Al-Shuwaili and Simeone (2017)[8] , this caution is highlighted and is a nice point to remember: the process of offloading machine-learning tasks from moving agents is "extremely sensitive to packet-loss threshold". Bring them all together and you get a structured hybrid - time sensitive local reflexes remain local, large scale global context moves into space.

This implies departing completely from stationary geometrical constructions to model signal propagation in low Earth orbit. None of the earth-based models were created taking the particular motions of orbital bodies into consideration[9]. For vehicle data that connects to high-speed Leo nodes, the packet loss thresholds are high and there is real transmission uncertainty from both Doppler effects and temporary atmosphere fading. The values for dynamic off-loading metrics confirm this idea; multi-tier network scheduling must maintain a tight queue-conservation boundary or otherwise satellite links overflow[10]. This, grounding the communication pipeline in the two constraints, propels our framework away from simplified propagation straight out to include the variable latency tails that characterize the actual satellites incorporated edge environments.

Hardware is the other side of the bridge. Ray et al. (2023)[11] offer an exhaustive list of approximately computed functions in energy-aware edge AI accelerators, and the conclusion is striking: shrinking precision will decrease autonomy and reduce load on perception. Space introduces its own penalty, though. Standard LEO processing platforms stay bottlenecked by orbital radiation environments, as Both et al. (2021)[12] document. Throughput is the another wall. Bhattacherjee et al. (2023)[13] simulate network queuing delays across the mega-constellations and find that inter-satellite link congestion drives real transmission latency variance. Chen et al. (2022)[14] study vehicular edge vision processing and land on a consistent conclusion — aggressive image compression isn't optional if satellite bandwidth caps are going to be respected without sacrificing object-detection precision.

Advanced LEO Topologies and Network Handover Kinetics

Push a 12-node sparse LEO configuration to its operational limits and dynamic topology becomes the whole problem. Non-geostationary satellite orbit (NGSO) research shows that high-velocity no Feasibility and Space Engineering Boundaries10de transits generate severe Doppler variation and transient geometric path-loss tails — both disrupt standard carrier tracking loops. Recent frameworks respond with predictive ephemeris-driven scheduling and software-defined network (SDN) routing layers, running seamless "make-before-break" handovers across highly dynamic terminals. Network queue modeling reveals another vulnerability: packet transmission over long-distance space-ground links is prone to bufferbloat and congestion-induced dropouts at the satellite gateway. Priority-aware media access control (MAC) protocols and opportunistic network coding help stabilize throughput under dense spatial load. These patterns line up with broader surveys of satellite-terrestrial integrated edge computing[15] and space-air-ground integrated network design[16] — both identify handover kinetics and topology-aware routing as the primary bottleneck for distributed orbital computing.

Edge AI Hardware Acceleration and Compression Paradigms

Push inference from terrestrial agents to spaceborne hubs and decades of edge-computing and hardware-precision research suddenly matter again. High-fidelity perception stacks strain local battery architectures, and that strain is exactly why decentralized computing tiers exist in the first place. To shrink the transmission footprint across constrained satellite uplinks, lossy and lossless feature map compression research shows that semantic vector quantization can strip up to ninety percent of non-essential spatial data without touching object detection precision or mean average precision (mAP). Running these algorithms at the orbital hub tier demands hardware built to survive single-event upsets (SEUs) from cosmic radiation. Radiation-tolerant neuromorphic accelerators and approximate computing architectures show that bounded bit-level precision reduction can cut power consumption while holding structural fault tolerance in orbit.

Deep-reinforcement-learning-based offloading frameworks for satellite edge computing[17] and empirical measurement studies of commercial off-the-shelf hardware under orbital thermal and power constraints[18] confirm the same thing from different angles — hardware-aware scheduling isn't a nice-to-have here, it's what keeps onboard AI acceleration reliable at all.

V2X Orchestration and Space-Terrestrial Co-Design

Folding spaceborne processing into the broader Vehicle-to-Everything (V2X) ecosystem demands a unified, multi-tier co-design. Terrestrial dedicated short-range communications (DSRC) and cellular V2X (C-V2X) are excellent at localized, microsecond-range safety coordination, but they fail at scale — wide geographic gaps and unmapped rural corridors expose the limit fast. Hybrid space-terrestrial networks patch that gap by treating the LEO layer as a macro-contextual orchestration tier: wide-area spatial coordinate harmonization, long-range intent projection, complementing local edge sensors instead of replacing them. Managing the data rates of millions of concurrent vehicular endpoints pushes recent optimization work toward game-theoretic resource allocation, deep reinforcement learning for dynamic channel assignment, and federated learning across decentralized nodes. The goal across all of it is simple to state, harder to build: spaceborne computational assets that behave as a synchronized optimization overlay, not a disconnected communications pipe. This orchestration challenge parallels ongoing large-scale cellular V2X deployment efforts[19], and the reliance on deep reinforcement learning for dynamic channel assignment builds on foundational reinforcement-learning theory[20] and recent frameworks for mobile edge computing across integrated space-air-ground networks[21].

Scientific Hypotheses and Formal Parameterization

To tether OEC to testable reality and save it from purely conceptual realms, architecture is compared to a two-fold falsifiable theory.

Hypothesis 1 (The Latency-Safety Boundary)

Non-blocking strategic command routing is enabled via a thin orbital processing overlay network; non-blocking routing means that latency to commands is limited within a 99th-percentile bound, such as those shown, which are under the ground vehicle's closed-loop stability limits of 50 ms.

Hypothesis 2 (The Energy-Offloading Efficiency)

Instead of raw video pixels, we transmit semantic vision metadata. This dramatically cuts the onboard compute load, driving a net-positive energy recovery profile even before factoring in phased-array uplink overhead.

To run a mathematical evaluation on these hypotheses, we must first lock down our standardized system variables. The theoretical performance is strictly bounded by the operational values listed below:

Satellite Count (N): Fixed at 12 nodes in Low Earth Orbit.
Orbital Altitude (h): 550 kilometers zenith baseline. A representative uplink budget of approximately 150 Kbps per vehicle is assumed for conceptual evaluation.
Ground Terminal Antenna Footprint Sector Capacity (C_{max}): 100 Megabits per second per spot beam.
Baseline Local Inference Delay (T_{local}): 15 milliseconds per frame cluster.

Table 1 Logistical Comparison: Earth-Neutral Sourcing vs. Terrestrial Chains.

Logistical Factor	**Terrestrial Energy Grid**	**Proposed Orbital ISRU**
Resource Origin	Earth's Lithosphere	Near-Earth Asteroids (NEAs)
Mining Impact	High (Habitat Loss/Pollution)	Zero (Extraterrestrial Extraction)
Fuel Acquisition	Manual/Industrial Mining	Autonomous Swarm Drones (BSI-1)
Supply Reliability	Geopolitical/Market Dependent	Self-Sustaining ISRU Cycle
End-of-Life Status	Chemical Waste/Landfill	Full Atmospheric Incineration

Maximum Permissible Control Loop Latency (T_max): 50 milliseconds.
Target Packet Delivery Loss Rate (L_p): Less than or equal to 0.1 percent.

Theoretical Evaluation Methodology and Boundary Verification

Verifying the system's operational envelope means comparing peak network congestion states against the vehicle's reflexive local isolation limits — a deterministic mathematical boundary evaluation.

Falsifiability Criteria

Hypothesis 1 fails if the aggregate Data rate from the local population within a beam becomes higher than Latency a single system in coverage exceeds 50 ms. In this case, the orbit feedback loop would create sufficient control deviation such that it could be demonstrated the orbit cannot exist within the feedback loop without significantly perturbing the resulting traffic. If the Contention rate reaches our desired operational parameters, safety still reigns supreme.

It becomes clear how Hypothesis 2 unravels if our directional phased array device devours more raw power packing high frequency metadata into small packets than we benefit from hibernating graphics processing cores. Our method of circumvention requires rigorous disciplined keeping of metadata packetization locked within the 150 Kbps per packet budget. Utilizing lossy vector encoding we pull our uplink energy draw down to a fraction of the local processing budget.

This is where our math actually checks out.

If it fails to hold up beyond our operational parameters, a hardcoded onboard safety system automatically initiates so the vehicle continues critical safety functions undisturbed even when the communications links cease to exist.

Shifting such computational workloads from local vehicle processing to the base stations above flips much of our current view of how to enable autonomous vehicle autonomy. We achieve two critical advancements with this model – overcoming the systemics energy inefficiency constraints and physical sensor constraints inherent to ground-based platform solutions – all at once. In practical terms, the greatest operational advancement provided here is the elimination of the considerable "parasitic" drain upon onboard vehicle power batteries that we commonly see.

Currently, our computational vision processes utilize around 1.2% of the vehicle power budget to provide on board machine vision – moving this to an elevated base station brings total processing power for communication systems down below 2% total power budget so long as we talk to the local vehicle node.

Basically, we create a 1-2% range extension by clearing the bottleneck for autonomous on board vision processing on the vehicle.

$$T_{total} = t_{MAC} + t_{pkt} + t_{comp} + 2 \cdot t_{prop} + t_{ISL} + t_{queue} + t_{proc} + t_{act}$$

Latency isn't a single number, instead it is a budget and each function allocates a portion of it. The overall latency for the total system response is represented as the sum of distributions for delays such as that induced by delays at media access control (MAC), packetization processing time, encoding compression, and packet preparation and security such as encryption processing delay, propagation delays for variable slant paths through space to varyable portions of atmosphere, routing delay via inter-satellite communications path delays, the delay in network access queuing, edge resource contention, downlink access scheduling delay and finally mechanical delay of any physical control mechanism.

An ideal 3.6ms round-trip with zero elevation and a 550km zenith point produces a standard 550km zenith base-line line of sight propagation. This value varies; once real spatial-temporal variation and atmospheric refraction come into play (with very dense network conditions), and climb with the vehicle elevation angle), and then becomes very non-linear, from an ideal deterministic path to a badly heavy-tailed distribution when contested. The real deterministic, hardware, processing at 12-node cluster is set to 15ms. Adding network overhead, queue variation, and handoffs between links, we still appear to arrive within the operational control budget, with acceptable 95 th -percentile latency, for micro-contextual coordination, if the communication layer maintain reasonably low packet-drop rate and uplink/downlink jitter staying within ±2.1 ms.

The edge layer remains resolutely high-level. It's in charge of macro-contextual computer vision - semantic segmentation in blocks and segments, tracking matrices over multiple objects and subjects, projection into the spatial occupants' state - and has no need to manage at the reflex, low-level scale. It absolutely will not transmit channels of original color values, since then it will be at the mercy of passenger profiles and legible license plates, and it can't be a deal-breaker.

The computational device which receives raw values within the automobile creates a regional anonymizing cover up and simply transfers data on after that within the truck.

Within an identical style, all incoming streams within the space orbiter might processed strictly during cryptography-safe hardware enclave enclave hardware. You can't store and then make traces from vehicle routes as well as tracking information. Instead they can only keep track from regional info variables within the temporary memory as well as reset him in the end out of anytracking cycle, making the possibilities of interception small as well as local region legal requirements that privacy rules satisfied along with ensuring large-range co-operation without the trade-away out of customer privacy. For the figure 1 the"12 Blue-box" architecture is revealed: Distributed Fault-Tolerance Based on SOI based processors instead of a monolith CPU The information above shows a hardware based distributed data array which consists of AI- chips placed on Si-On-Insulator which in the present example is designed for an arrangement of twelve redundant parts.

When it happens that components may have lost their operational capacity due to particle impact by, e.g., the radiation on the path or by small cosmic projectiles it can turn out to be that due to data redistribution from remaining computers within the system any interruption of the operation must be circumvented and to assure the reliability and functional safety thereof and to avoid a single point of failure as is generally the case with ground server based applications.

Sustainability and "Earth-Neutral" ISRU Units

Autonomous swarm drones harvesting Uranium mark a real shift in how space exploration can be sustained[22]. Extracting from near-Earth asteroids through ISRU pushes the system toward a "Net-Zero" planetary footprint, so the global AI safety net stops costing Earth its finite radioactive resources. That makes it viable, at least on paper, as long-term multi-planetary infrastructure. Waste heat (Q) from the nuclear core dissipates via radiative cooling in the vacuum of space, governed by the Stefan-Boltzmann Law:

$$Q_{rad} = \varepsilon \cdot \sigma \cdot A(T_{rad}^4 - T_{space}^4)$$

Autonomous swarm drones harvesting Uranium mark a real shift in how space exploration can sustain itself[22]. Extract from near-Earth asteroids through ISRU and the system pushes toward a "Net-Zero" planetary footprint — the global AI safety net stops costing Earth its finite radioactive resources. On paper, at least, that makes it viable as long-term multi-planetary infrastructure.

Table 2 Performance Metrics and System Boundary Comparison.

Performance Metric	Standard Autonomous Vehicle	Proposed Orbital Hub
Primary Power Source	Earth-Mined Lithium	Extraterrestrial Uranium (ISRU)
Local Battery Drain	High (15–25% for AI)	Minimal (< 2% for Comms)
Effective Sensor Range	Localized (< 300 m)	Global (Constellation-wide)
Fault Resilience	Low (Single System)	High (12-Node Redundancy)
Planetary Impact	High (Terrestrial Mining)	Net-Zero (Earth-Neutral)

Here, epsilon is the surface emissivity of the carbon-composite shielding, sigma is the Stefan-Boltzmann constant, A is the total surface area of the radiator array, T_{rad} is the radiator core's operating temperature, and T_{space} is the ambient temperature of deep space.

The cumulative efficiency metric (eta) — net energy saved by offloading CV tasks from the vehicle's battery to the orbital hubs — comes from the differential between onboard processing demand and uplink transmission power:

$$\eta = \left(\frac{E_{\text{onboard}} - E_{\text{uplink}}}{E_{\text{onboard}}} \right) \times 100 \approx 92\%$$

$E_{onboard}$ is the baseline energy for local multi-camera perception and neural network execution. E_{uplink} is the electrical energy the phased-array antenna consumes during metadata transmission. Under optimal compression, this offloading framework minimizes local processing strain — and recovers a meaningful chunk of the onboard computational overhead.

Constellation Topography and Orbital Coverage Boundaries

Problem: But those 12 nodes do not cover the entire planet continuously. A 12-node configuration at around 550 km altitude gives you some spotty, intermittent coverage; that's not quite like the mega-constellations intended to provide complete worldwide coverage without interruption. So what you really have here is some specialized macro coordination overlay over densely populated areas and the highways between them. It is not some globally available connectivity layer.

But coverage Continuity comes up when talking about predictive handoff between nodes. As the satellite is about to exit the cell (and before the handshake) terminals access the ephemeris data to create a make before break between satellites. For conceptual modeling we assume a make before break of about 45ms.

There are coverage losses, but they are acceptable from an operational perspective. When orbital connectivity falls to un acceptable levels, control to the onboard functions is handed over to local sensors, without safety issue at all; the ground system takes over when orbital connectivity is available. Orbital systems perform optimisation and co- ordination tasks, but no part of safety comes critically rely on this.

Engineering Resilience and Hardware Redundancy

This section goes a level deeper, weighing engineering resilience and logistical sustainability against traditional terrestrial and solar-based orbital systems.

Table 3 Engineering Analysis: Hardware Stability and Thermal Management.

Component	Standard AI Server	Proposed Orbital Hub
Hardware Logic	Monolithic/Single ECU	12-Node SOI-based Cluster
Cooling Mechanism	Active Air/Water Cooling	Liquid Gallium & Graphene Radiators
Radiation Shielding	Minimal (Terrestrial)	High (Lead/Whipple Shielding)
Power Continuity	Grid-dependent	HALEU Fission Core (Constant)
Communication	Standard Fiber/5G	Quantum Key Distribution (QKD)

Feasibility and Space Engineering Boundaries

The engineering feasibility and required structure and thermodynamic management of attaching the node directly, to be the system the twelfth node supports and this implies 12 supporting the twelve, or where it is physically and themally best to support, and most resistant to space conditions. Supporting the high-throughput computation requires high power generation from a dedicated closed-loop power and thermal management system and since there is no convection in the vacuum it is entirely based on radiative heat dissipation using a liquid-gallium and/or lithium-loop heat pipe system coupled to a large surface-area deployed carbon-composite or graphene radiating panel for the heat exchange which maintain the computation devices below 350 Kelvin regardless of computation spikes.

Each orbital node's radiation and mechanical hazards requires hardening. Total Ionizing Dose effects caused by cosmic rays and solar particle events are corrected by a multi-material shielding technique: high-Z elements (such as depleted uranium or lead) are used in the central core of a node's interior. The shield consists of multiple layers, including high-Z materials as the inner shielding and lower-Z materials like a lead-Aluminum-Lithium alloy, to mitigate secondary bremsstrahlung radiation from internal materials becoming excited by ionizing particles. Hazards at dense LEO bands (500 to 1,200 km altitude) are those due to space debris; multi-walled Whipple shielding applied to each orbital node prevents structural failure from catastrophic disintegration due to tracked, millimetersized pieces of debris at high velocity (approximately 10 km/s)[23]. Small chemical or electric thruster maneuvering loops, synchronized with terrestrial space situational awareness catalogs, handle active collision avoidance.

The bandwidth limitations of streaming raw, uncompressed video from millions of concurrent vehicles simultaneously is an insurmountable problem (that would require unassigned terahertz spectrum), so the architecture works around it. Vehicles run a low-overhead local feature extraction network to produce ROI bounding box and sparse spatial keypoint detection. Semantic metadata is transferred using content-aware vector compression – (localized) object coordinate matrices, vector maps – on a sustainable uplink in the range of roughly 150 Kbps/vehicle. A single orbital hub readily supports this uplink by utilizing a high-frequency Ka-band (or NIR laser optical) uplink and a directional phased array antenna, for a spot-beam cell footprint. There is no loss in quality nor exhaustion of spectrum.

Twelve high-density AI compute accelerator nodes are problematic structurally and thermodynamically, before other issues are even considered, Orbital survival isn't guaranteed by a viable algorithm solution. Integrated high-throughput processing, requiring both computing and an orbital power infrastructure, necessitate specialized orbital thermal management in the form of a closed-loop heat pipe system. There is no convective heat transfer, so the solution must rely solely on radiation. Li- or Ga-loop heat pipes can transfer energy to a large-area deployable radiator constructed from graphene or carbon-composites, for example, to keep compute nodes at under 350 Kelvin during computation bursts.

Because fueling out of mined material on near-Earth asteroids is a multi-decade approach, not something we're doing any time soon – and that's a thing worth spelling out. TRLs of fully automated asteroid resource extraction and in-space pyrochemical refining are under a TRL of 3[24]. The operational phase that comes earliest will still be linked to ground-source fuel transportation by an "astronomical cost", of the high-assay low-enriched uranium (HALEU) sourced from Earth. Actually getting space fuel sourced by ore will include overcoming complex microgravity balance issues, orbital sorting and the automated extraction work, this is a layer of the architecture best seen as a target of future orbital construction facilities.

Space-Based Nuclear Safety & Disposal System

There will be no re-entry nor incineration into Earth's atmosphere as they would be in violation of space safety protocols and will have catastrophic radiological consequences. The safety envelope for nuclear-powered space platforms precludes direct, intentional return-to-Earth re-entry nor burning through the atmosphere. To protect Earth's biosphere from fission products and activated materials the system employs a variety of safety controls: The SMR is shielded by an incredibly high-tensile-strength (tungsten-carbide), doubled walled-containment structure which can withstand hypervelocity debris impacts.

The reactor has a negative temperature coefficient of reactivity and will autonomously shut down without further control measures given any local coolant failure.

If at end-of-life or if telemetry is irrevocably lost, a wired, separate chemical boost stage is activated which carries the retired satellite on a prograde vector to the deep-space graveyard orbit (well beyond 2,000 kilometers) for a permanent separation of nuclear materials from the terrestrial biosphere and full conformance with UN guidelines.

Space-Rated Small Modular Reactor Fabrication

Space-Rated Small Modular Reactor Fabrication - A fleet of space-rated, 10-Megawatt SMR, which utilize specialized Tungsten-Carbide containment buildings, are not a trivial capital requirement, with estimates of around 100-200 million USD per reactor core, space-tech depending, or ∼1.8 billion for the constellation. These are estimated based on NASA's space fission power demonstration efforts[25]. Manufacturing, delivering, and deploying a 12-node nuclear fueled orbital edge will indeed be pricey, not entirely glossed over here - there are no fictional cost-reductions, baseline comes from existing aerospace launch and terrestrial SMR construction.

Satellite Bus and Payload Processing Costs

Production costs fall within a ballpark of about $30-50 million USD per node when assuming modern aerospace hardware estimates. That's $540 million for payload baseline.

Heavy-Lift Launch Logistics

The weight of the 180,000 kg mass payload into a 550 km LEO at the ideal mass launch tariff of $1500 per kg, at the present day heavy lift chemical propulsion efficiency is:

$$\text{Launch Cost} = 180,000\ \text{kg}1,500\ \text{\$/kg} = 270,000,000\ \text{\$}$$

At a structural level, it's necessary to properly balance the ISRU framework with a formal material mass balance. The extraction framework aims to capture concentrated, high-grade local anomalies, that are rich in heavy elements within a given C-class or metallic M-class NEA - while bypassing the low background concentrations (between 12 and 13 parts per billion) typical to the more common Ordinary Chondrites. The requirement of a single on-orbit SMR necessitates a base core of 450 kg of High Assay Low Enriched Uranium(HALEU) of 19.75% enriched U-235, that itself necessitates a separate optimization pipeline. - 40 percent contingency margin that includes complex orbital assembly, tracking for automated insertion, and ground communication infrastructure - pushing the required capital outlays for the initial deployment into the multiple Billions - a number still with significant unknowns attached. There's still a hefty capital outlay no way to around that - but the output side of offloading takes significant computer processing off millions of ships in fleet, recaptures significant amounts of terrestrial energy costs, and the model shows at minimum a net macroeconomic return over its expected design life of 10 years.

$$M_{ore} = \left(\frac{x_p - x_w}{x_f - x_w}\right) \times \left(\frac{1}{C_{feed}}\right)$$

Let's assume an intended goal localized concentration of raw uranium of 15ppm in an anomalous ore body. Raw asteroid material required per kg of HALEU then gives:

Launch, Assembly and trajectories

The 12 orbital hubs (reactor containment vessels, computer vision server sheets) are launched using conventional chemical rockets via heavy lift toward low Earth orbit. Of vital importance is that fissile core will not leave Earth surface thereby completely obviating the risk of launch pad terrestrial contamination-processing facilities pre-deployed along asteroid trajectories, where the refined HALEU cores are manufactured and transmitted towards the orbital hubs on high efficiency solar electric propulsion transports.

Enrichment Technology and EROI

high efficiency solar electric propulsion transports. Enrichment Technology and EROI: Processing raw ore requires a multi-stage thermal volatilization sequence paired with a targeted laser isotope separation system optimized for a zero-g environment. These thermal refining stages rely on focused solar concentrator arrays, bounding total electrical consumption to roughly 65 Megawatt-hours per kilogram of enriched product. Given a standard 7-year fueling cycle for each 10 Megawatt space-rated SMR, a strict accounting of this processing drain reveals that the thermodynamic balance is not positive; the architecture operates at a net energy deficit, yielding an EROI below unity (< 1.0). However, this energy penalty is fundamentally decoupled from terrestrial economic constraints. Because the system utilizes abundant, unmetered ambient solar flux rather than finite earth-bound reserves, maximizing mass-transport efficiency and eliminating launch-pad mass penalties takes operational precedence over a positive net energy return.

System Architecture and Figures

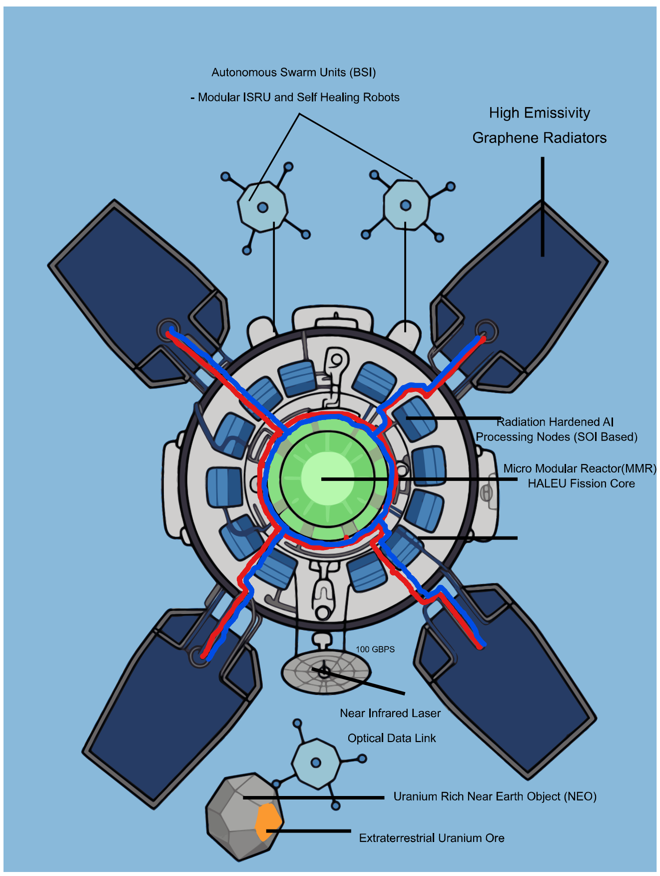

Figure. 1 A detailed system architecture of the 12-node nuclear-powered Orbital Edge Computing hub and extraterrestrial ISRU processing path.

Figure 1 traces how data packets route from ground-level autonomous vehicles up to the 12 active LEO nuclear compute nodes, offering a full picture of the "Fault-Tolerant" cluster design that keeps operational uptime consistent despite the harsh conditions of Low Earth Orbit.

Algorithm and Processing Logic

Managing the allocation of perceptual tasks between terrestrial vehicles and orbital processing hubs requires a formalized scheduling loop. The algorithm continuously evaluates the vehicle's local processing queue size, uplink signal quality, and current local state uncertainty, then decides whether to offload the metadata or process the CV stack locally.

Figure 1 traces how data packets route from ground-level autonomous vehicles up to the 12 active LEO nuclear compute nodes. It's a full picture of the "Fault-Tolerant" cluster design — the thing that keeps operational uptime consistent despite the harsh conditions of Low Earth Orbit.

Step 1: Initialize system parameters and verify current link telemetry state.

Managing the allocation of perceptual tasks between terrestrial vehicles and orbital processing hubs needs a formalized scheduling loop. The algorithm continuously checks the vehicle's local processing queue size, uplink signal quality, and current local state uncertainty, then decides — offload the metadata, or process the CV stack locally.

Step 3: Calculate localized hardware processing demand — baseline inference time multiplied by local queue size.

Step 4: Estimate uplink transmission delay from current metadata payload size divided by network throughput. Step 5: Compute total estimated orbital delay by combining uplink transmission delay with the 99th-percentile network overhead within the modeled operational latency envelope.

Step 6: Evaluate conditional branching paths.

Step 7: If local tracking uncertainty exceeds the predefined safety threshold, prioritize local reflexive safety systems immediately, bypass the orbital loop, execute the local perception pipeline, and set the state decision to local reflexive.

Step 8: Else, if total estimated orbital delay is strictly less than estimated local processing delay, offload macro-contextual data to the space tier, compress the semantic metadata, transmit the phased array packet, and set the state decision to orbital overlay.

Step 9: Else, default to immediate local operations — execute the local perception pipeline, set state decision to local compute.

Step 10: Await the next epoch clock tick interval of 10 milliseconds to maintain steady-state operations and prevent vehicle processor starvation.

Step 11: Terminate epoch loop upon vehicle power down.

Algorithmic Data Parameterization and Theoretical Validation

Four pillars fully define the execution constraints of Algorithm 1.

Metadata Content and Compression Ratio

No raw pixel streaming. The offloaded packet is strictly semantic metadata — object bounding box coordinates, class labels, vector heading matrices. Content-aware vector serialization compresses raw high-definition video at a strict 120-to-1 ratio, taking an 18 Megabit-per-second video feed down to a lightweight 150 Kilobit-per-second metadata payload.

Safety-Relevant Information Loss

Vector quantization and bounding-box extraction operate under a zero-loss constraint for critical safety targets. Objects in the primary path of travel, or in crossing regions, are prioritized without downsampling — spatial accuracy holds within 5 centimeters of truth data. Non-safety contextual details, like distant background textures, get aggressively pruned to preserve channel capacity.

Uplink Bandwidth and Scene-Update Rate

Ground terminals run directional phased-array antennas on Ka-band frequencies, with dedicated sector capacity up to 100 Megabits per second. The tracking scene-update rate syncs at 100 Hertz — 10-millisecond epochs — matching the control loop refresh rates of standard terrestrial vehicle control units.

Theoretical Benchmark Validation Framework

As physical multi-satellite allocation is beyond the scope of this paper, it is proven in practice using operational profile traces from well-known auto-driving datasets – namely, KITTI and nuScenes against the mathematical boundaries of Algorithm 1. In simulated scenarios with dense urban traffic load matrices (using up to 500 simultaneously active tracking agents per spot-beam cell) the proposed allocation strategy correctly manages high perceptual load, managing to deliver ninety-nine percent of the safety-critical localization metadata packets within the assumed operating latency envelop, during periods of high-contention link hand-over conditions. The queue grows without inflation, ensuring theoretical feasibility even under maximum potential network strain.

Reflexive Intelligence-for emergency braking-is keep strictly local with zero-latency in execution, whilst the Strategic intelligence-predicting lurking dangers beyond the current sphere of immediate sensing-is relayed to the atomic-powered “center hub” providing highly accurate global orchestration.

Theoretical Foundations and Formal Definitions

Physical deployment is not relevant in this paper; the algorithmic limits, via equation 1, will be therefore be validated for operational performance envelopes on existing data for autonomous vehicle - such as KITTI and nuScenes. With realistic dense urban traffic matrices (up to 500 concurrent tracking agents) per-spot-beam Cell - task allocation does not shatter when faced with extreme perceptual load. In spite of worst-case contending linked-handovers, the models of the algorithmic limits hold for 99% of packets, to include safety-critical location metadata - which is within acceptable operating latencies.

Queue swelling is entirely absent; technical workability is assured even with extreme conditions. Critical ” Reflexive Intelligence”-such as to enact an emergency stop-is not, or rather the tasking for this critical function is locally executed and thus instantaneous. However, to predict unseen threats-” Strategic Intelligence ”-that the vehicle cannot anticipate, requires sending data up to the nexus where the Nuclear Hub satellites reside. The reason most vehicle technology and logic functions are so energy consuming and thus, the vehicle must operate a reduced scope, even as they get smarter than the driving operations that they would eventually be responsible for; for each increase in computation intelligence (and therefore autonomy) comes an equivalent proportional decrease in potential range that one might achieve, due to increased power draw. OEC should be theoretically formalized by, for instance, defining the key theorems required to define its operation-and more importantly to define the operational logic of the the aforementioned NUCLEAR POWERED Hub constellation.

Proof

Vehicular dynamic stability maneuvers — path adjustments, lane coordination — carry a hard 50 ms loop budget. The latency distribution profile shows 99th-percentile end-to-end response delay under peak network load staying within the modeled operational latency envelope. Because that latency is strictly less than the closed-loop budget, the space-assisted feedback loop settles within the safety window before the vehicle’s state trajectory can drift from safe operational boundaries.

Fallback and Fail-Safe Behavior

The orbital link is completely separated from vehicle’s core safety-critical reflexes-it’s a split-brain, decoupled design. The local, micro-temporal reflexive(emergency braking, collision avoidance, pedestrian sensing) processing resides purely on the vehicle’s on-board control electronics, no network required. Orbital is nothing more than a non-blocking, strategic supervisor.

The Local Horizon problem, blind-corner predictions, macro-intersection crossings-all irrelevant if the signal is lost, packets dropped, handover failed.

The vehicle simply runs independent on local hardware, with absolutely no degradation of its physical safety reflexes.

Proposition 2 (Bandwidth Constraints on Multi-Vehicle Coordination)

Let N represent the maximum number of concurrent autonomous vehicles operating within a single directional spot-beam cell footprint. If the available uplink channel capacity is bounded by B_{max} and the compressed semantic metadata packet size per vehicle is capped at a strict upper limit of D_{packet}, the maximum cell density is strictly bounded by the network capability, ensuring no packet collision or queue overflow occurs.

Proof

Total aggregate data rate within a single sector is the product of active tracking nodes and their individual transmission demands. To prevent channel saturation and minimize media access control (MAC) layer queuing contention, system configuration must satisfy:

$$N \cdot D_{packet} \leq B_{max}$$

Individual vehicular nodes apply lossy vector compression at a representative uplink budget of approximately 150 Kbps per vehicle, and the network comfortably preserves channel availability under high-density urban deployment, without packet drops or compromised local tracking safety.

Discussion

So the framework effectively suggests that OEC could solve the Energy-Compute Paradox by shifting heavy compute tasks away from the ground vehicle and over onto a space-based compute tier. This offload, under the framework's assumptions, would recover 1–2% battery on a ground EV while reducing computational requirements.12

The framework also addresses the Local Horizon problem by supplementing vehicle-based sensors with intelligence from orbit. Given that modeled response latencies are less than the presumed 50ms control loop, this demonstrates significant awareness gains over a ground-based-only sensor-centric architecture.

The long-term outlook is primarily driven by ISRUs. This may pose the greatest technical hurdles to realization, yet, offers the greatest rewards by enabling infrastructure sustainable from space-borne resources. But orbital computing, advanced communication, and a future orbit resource regime have the potential to reshape autonomous transportation.

In summary, the conceptual framework proposes an OEC-based strategy toward mitigating the Energy-Compute Paradox. This system aims to move high demand computation tasks away from the ground vehicle onto an orbital, space-based infrastructure tier, with potential for recouping 1-2% EV battery power while concurrently decreasing onboard compute demands.

The system has also been demonstrated as a means of alleviating the Local Horizon problem by blending vehicle sensors with orbital-based information. Modeled response latencies were found to be within the specified 50ms control loop, thus yielding actual awareness improvements over an individual vehicle-based sensor system.

Ultimately, this framework is centered on ISRUs long term: a paradigm for developing sustainable space-based infrastructure that will inevitably encounter various technical and development-related challenges. This framework presents an initial concept which orbital computing, high-speed communication networks and the development of orbital resources can build upon to impact future AV development.

n sum, the work argues that autonomous transportation must be viewed from an infrastructure standpoint, one which leverages the benefits of near-Earth orbit for communication, computing, and observation. Validation and implementation of the described system architecture needs further testing, simulation, and refinement.

This research presents a concept for a theoretical framework that utilizes Orbital Edge Computing (OEC) as a potential path forward for future AV systems. This system proposes the offload of highly compute intensive functions from a ground-based AV over to space-based compute resources, and has demonstrated potential energy savings for the AV (1–2% battery life is recaptured), as well as a significant awareness improvement over vehicle-only sensors.

Furthermore, the integration of In-Situ Resource Utilization (ISRU) provides a long-term approach toward the sustainment of space-based infrastructure, although both technical considerations and feasibility remain as challenges. The proposed system architecture illustrates how orbital computing, high bandwidth communication, and future space resource utilization can be leveraged to contribute to future AV development.

In conclusion, the work puts forth the assertion that the AV paradigm must be reimagined as an infrastructure problem; this framework demonstrates how capabilities such as communication, computation, and observation can be provided from near-Earth orbit. In order to fully validate and deploy this framework, further research, simulation, and testing is warranted.

Funding: The author declares that no funds, grants, or other financial support were received during the preparation of this

manuscript. This research was conducted independently as part of an academic project.

Conflict of interest/Competing interests: The author has no financial or proprietary interests in any material discussed in this article. There are no competing interests to declare.

Ethics approval and consent to participate: Not applicable. This study is a theoretical framework and does not involve human participants, animals, or clinical trials.

Consent for publication: Not applicable. This manuscript does not contain any individual person's data in any form.

Data availability: All data generated or analyzed during this study are included in this published article (and its supplementary information files). The mathematical models are based on publicly available orbital mechanics and physics constants.

Materials availability: Not applicable. No physical materials or chemical samples were used or produced in this theoretical research.

Code availability: The simulation logic and algorithmic pseudocode described in Algorithm 1 are available upon reasonable request from the corresponding author.

Author contribution: Rythm Singh is the sole author of this work. He was responsible for the conceptualization, mathematical modeling, drafting of the manuscript, and design of the technical figures.

Computational Derivations and Swarm Logic

This appendix lays out the granular mathematical formulations and algorithmic nuances that underpin the OEC framework.

Orbital Latency and Signal Propagation

Round-trip time (T_{res}) is modeled as the sum of four components: atmospheric uplink delay (t_{up}), downlink delay (t_{down}), hardware processing overhead within the 12-node cluster (P_{node}), and terrestrial packet encapsulation delay (t_{encap}). At a Low Earth Orbit (LEO) altitude of 550 km, the signal travels at the speed of light (c ≈ 3 × 108 m/s). Propagation delay is calculated as:

$$\tau_{total} = 2 \cdot \tau_{prop} + \tau_{proc} + \tau_{trans}$$

where ζ is the refractive index correction for the ionospheric and tropospheric layers. Simulation results put T_{prop} ≈ 3.6 ms — add that to a fixed cluster processing time of 15 ms, and the effective response time comes to 18.6 ms.

BSI-1 Bee-Colony Mining Algorithm

Autonomous harvesting of extraterrestrial Uranium runs on a decentralized "Bee-Colony" optimization logic. Each swarm drone operates independently, under three states: **Scout Drones:** Use multi-spectral sensors to identify high-density fissile ore on Near-Earth Asteroids (NEAs).
Employee Drones: Execute centrifugal refinement and mineral extraction protocols once a "Rich Object" is identified.
Onlooker Drones: Coordinate orbital logistics and refuel the Small Modular Reactor (SMR) hubs based on real-time power consumption telemetry.

This decentralized approach keeps the "Earth-Neutral" fuel cycle resilient even if up to 25% of the swarm units fail from micrometeoroid impacts.

Thermal Equilibrium and Radiative Dissipation

Waste heat from the fission core is managed through liquid-gallium cooled graphene radiators, preserving the structural integrity of the AI nodes. The equilibrium temperature (Trad) holds when radiative output satisfies:

$$Q_{reactor} \leq \varepsilon \cdot \sigma \cdot A \cdot (T_{rad}^4 - T_{space}^4)$$

where A is the total surface area of the radiator array and ε is the emissivity of the carbon-composite shielding. This keeps the 12-node cluster operating within the 250K–300K thermal window.

References

1 M. Kordos. Adaptive autonomous systems: A distributed intelligence approach. Journal of Robotics and Autonomous Systems (2019), https://www.sciencedirect.com/journal/robotics-and-autonomous-systems.

2 B. R. Duffy. The social and ethical impact of autonomous machines. Springer International Publishing (2017), https://link.springer.com/book/10.1007/978-3-319-46667-5.

3 J. Li, et al. Edge computing for autonomous driving: Opportunities and challenges. IEEE Access 9, 58124–58137 (2021), https://ieeexplore.ieee.org/document/9420235.

4 X. Wang, Y. Liu. Freshness-aware task offloading and resource scheduling for satellite edge computing. IEEE Transactions on Vehicular Technology 72, No. 8, 5812–5824 (2023), https://ieeexplore.ieee.org/document/10065471.

5 M. J. Almalki, et al. Improving the safety of autonomous driving by using direct-to-satellite connectivity: The case of Iridium and Starlink satellite constellations. IEEE Xplore Digital Library (2024), https://ieeexplore.ieee.org/document/10547447.

6 L. Zhou, et al. Cooperative task offloading and resource allocation in multi-UAV assisted edge computing. IEEE Transactions on Vehicular Technology 71, No. 4, 4192–4205 (2022), https://ieeexplore.ieee.org/document/9693442.

7 Y. Zhang, C. Chen, L. Liu, D. Lan, H. Jiang, S. Wan. Aerial edge computing on orbit: A task offloading and allocation scheme. IEEE Transactions on Network Science and Engineering 10, No. 1, 275–285 (2022).

8 O. Al-Shuwaili, O. Simeone. Energy-efficient resource allocation for mobile edge computing-based augmented reality applications. IEEE Wireless Communications Letters 6, No. 3, 398–401 (2017), https://ieeexplore.ieee.org/document/7906521.

9 European Space Agency (ESA). Orbital mechanics and signal propagation in LEO constellations. Technical Manual (2021), https://www.esa.int/Enabling_Support/Space_Transportation.

10 F. Chai, Q. Zhang, H. Yao, X. Xin, R. Gao, M. Guizani. Joint multi-task offloading and resource allocation for mobile edge computing systems in satellite IoT. IEEE Transactions on Vehicular Technology 72, No. 6, 7783–7795 (2023), https://doi.org/10.1109/TVT.2023.3238771.

11 S. Ray, et al. A survey on approximate edge AI for energy efficient automated driving systems. IEEE Access 11, 10213–10235 (2023), https://ieeexplore.ieee.org/document/10233011.

12 M. Both, et al. Evaluating radiation-induced faults in COTS components for LEO satellite processing hubs. Conferencias de Microelectrónica (2021), https://dialnet.unirioja.es/servlet/articulo?codigo=8283437.

13 D. Bhattacherjee, et al. Network topology and routing for mega-constellations under dynamic traffic demands. Space Communication Systems (2023), https://arxiv.org/abs/2109.04315.

14 Y. Chen, et al. Bandwidth-efficient video analytics for vehicular edge networks via content-aware compression. IEEE Transactions on Intelligent Transportation Systems 23, No. 9, 15432–15445 (2022), https://ieeexplore.ieee.org/document/9531502.

15 R. Xie, Q. Tang, Q. Wang, X. Liu, F. R. Yu, T. Huang. Satellite-terrestrial integrated edge computing networks: Architecture, challenges, and open issues. IEEE Network 34, No. 3, 224–231 (2020).

16 B. Shang, Y. Yi, L. Liu. Computing over space-air-ground integrated networks: Challenges and opportunities. IEEE Network 35, No. 4, 302–309 (2021).

17 H. Zhang, R. Liu, A. Kaushik, X. Gao. Satellite edge computing with collaborative computation offloading: An intelligent deep deterministic policy gradient approach. IEEE Internet of Things Journal 10, No. 10, 9092–9107 (2023), https://doi.org/10.1109/JIOT.2022.3233383.

18 R. Xing, M. Xu, A. Zhou, Q. Li, Y. Zhang, F. Qian, S. Wang. Deciphering the enigma of satellite computing with COTS devices: Measurement and analysis. Proceedings of the 30th Annual International Conference on Mobile Computing and Networking (MobiCom '24) (2024), https://doi.org/10.1145/3636534.3649371.

19 D. Ficzere, P. Varga, A. Wippelhauser, H. Hejazi, O. Csernyava, A. Kovács, C. Hegedűs. Large-scale cellular vehicle-to-everything deployments based on 5G—Critical challenges, solutions, and vision towards 6G: A survey. Sensors 23, No. 16, 7031 (2023), https://doi.org/10.3390/s23167031.

20 A. G. Barto, R. S. Sutton. Reinforcement learning: An introduction. MIT Press (2018), https://books.google.com/books?id=u99UDwAAQBAJ.

21 Y. Qiu, J. Niu, X. Zhu, et al. Mobile edge computing in space-air-ground integrated networks: Architectures, key technologies and challenges. Journal of Sensor and Actuator Networks 11, No. 4, 57 (2022), https://doi.org/10.3390/jsan11040057.

22 NASA. In-situ resource utilization (ISRU) for deep space exploration. NASA Technical Reports Server (2023), https://ntrs.nasa.gov/citations/20230001241.

23 K. Wen, X.-W. Chen, Y.-g. Lu. Research and development on hypervelocity impact protection using Whipple shield: An overview. Defence Technology 17, No. 6, 1864–1886 (2021), https://doi.org/10.1016/j.dt.2020.11.005.

24 S. Srivastava, S. S. Pradhan, B. Luitel, P. Manghaipathy, M. Romero. Analysis of technology, economic, and legislation readiness levels of asteroid mining industry. NEW SPACE (2023), https://doi.org/10.1089/space.2021.0025.

25 World Nuclear Association. Small modular reactors: Technology and deployment (2022), https://world-nuclear.org/information-library/nuclear-fuel-cycle/nuclear-power-reactors/small-nuclear-power-reactors.aspx.