SPOCK: Enabling Cognitive Cloud Computing in Space for Low Latency Earth Observation

Leonardo Carrer
Aresys srl
Rome, Italy
Leonardo.carrer@aresys.it

Gabriele Meoni
(1) ESA Φ-lab, European Space Agency, ESRIN, Italy
gabriele.meoni@esa.int

Álvaro Morón Elorza
Indra Space
Madrid, Spain
amorone@indra.es

Antonio Gajardo Gonzalez, *Indra Space*
Madrid, Spain
agajardo@indra.es

Rohaan Ahmed
Indra Space
Madrid, Spain
roahmed@indracompany.com

Fabio Gerace
Aresys srl
Rome, Italy
fabio.gerace@aresys.it

Juan Carlos García-Taheño Hijes, *Indra Space*
Madrid, Spain
jcgarciat@indra.es

Karen Nicole Pirazan,
Indra Space
Puertollano, Spain
knpirazan@eservicios.indra.es

Pedro José López Rosillo, *Indra Space*
Madrid, Spain
pjlopezr@indra.es

Alessandro Princi,
Indra Space
Madrid, Spain
aprinci@indra.es

Marcos Quintana González, *Indra Space*
Madrid, Spain
mquintana@indra.es

Firoozeh Foroozan
Indra Space
Madrid, Spain
fforoozan@indra.es

Manuel Indaco
Indra Space
Madrid, Spain
mindaco@indra.es

Francisco Jesús López Honrubia, *Indra Space*
Madrid, Spain
fjlopezh@indra.es

Davide Giudici
Aresys srl
Milan, Italy
davide.giudici@aresys.it

***Abstract*—Earth Observation (EO) missions are facing increasing operational pressure from growing data volumes, stringent latency requirements and the need for timely and actionable information products. Traditional EO system architectures, which rely primarily on ground-based processing, introduce bandwidth bottlenecks and limit responsiveness for time-critical applications such as disaster monitoring and maritime surveillance. The SPOCK (Smart Processing Orbital Cloud Konstellation) initiative proposes a novel EO architecture to enable cognitive cloud computing in space. The proposed architecture integrates onboard processing, artificial intelligence and inter-satellite links to shift key elements of the EO information value chain from ground to orbit. SPOCK aims to significantly reduce latency, optimize bandwidth use and demonstrate scalable and federated EO on-board processing. This is achieved by combining edge AI, software-defined payloads and inter-satellite links. The initiative positions itself as a precursor to future orbital cloud infrastructures for next-generation EO missions.**

Keywords— Earth Observation, Cognitive Cloud Computing, On-Board Processing, Edge AI, Inter-Satellite Links, Software-Defined Satellites

I. INTRODUCTION

Earth Observation (EO) systems are undergoing a profound architectural transition driven by the rapid growth of sensor resolution, constellation size and user demand for timely and actionable information. While recent commercial and institutional missions have significantly expanded data acquisition capabilities, the underlying data handling remains largely constrained by traditional architectures that prioritize bulk data downlink and ground-based processing. This approach increasingly struggles to meet latency and scalability requirements. This is particularly true for time critical applications such as disaster response, wildfire monitoring and maritime surveillance.

In conventional EO missions, satellites function primarily as passive data collectors transferring raw or lightly processed data to Earth for downstream exploitation. As data volumes grow, this model introduces systemic bottlenecks such as downlink congestion, delayed information delivery and inefficient utilization of onboard resources. These limitations are motivating a shift toward on-board processing and edge computing for autonomous decision-making thus enabling satellites to extract semantic information directly in orbit and transmit only high value products rather than raw imagery thus also reducing latency.

Fig. 1. Conceptual illustration of the SPOCK architecture depicting a spaceborne cognitive processing node operating and coordinating with other satellites to to enable low-latency and autonomous Earth Observation services.

Early demonstrations of this paradigm have been realized through ESA's Φsat programme, particularly Φsat-1 [1], which successfully demonstrated onboard artificial intelligence (AI) and selective downlink of hyperspectral data. Φsat-1 provided the first in-orbit validation that AI-based inference can significantly improve EO efficiency by reducing unnecessary data transmission and accelerating product availability. The follow-on Φsat-2 [2] mission further extends this concept by supporting a software-defined AI platform capable of hosting multiple onboard applications by stressing on the modularity and in-flight reconfiguration capabilities for future EO systems.

While these initiatives represent significant advances, they predominantly address stand alone missions or single satellite intelligence. The future goal is scaling onboard cognition across distributed, heterogeneous constellations and enabling cooperation and federated processing among multiple space assets. This challenge has been explicitly addressed by ESA within the 3CS4EO (Cognitive Cloud Computing in Space for Earth Observation) project which defined the concept of satellite constellations behaving as distributed computational infrastructures rather than isolated platforms. 3CS4EO explicitly introduces the notion of technological constellations acting as space-based cloud infrastructures, where onboard processing, AI, inter-satellite links and distributed decision-making replace traditional ground-centric pipelines. Within this framework, satellites collectively provide sensing, processing, storage and decision-making capabilities analogous to terrestrial cloud systems but under the unique constraints of space environments.

To this extent and independently, Wang et al. [3] present a study of space computing constellations introducing an integrated architecture that tightly couples onboard computing, intersatellite laser links and AI inference in low-Earth orbit. The paper reports the design and in-orbit deployment of the *3-body computing constellation*. This constellation is a 12 satellite experimental platform equipped with high-speed laser intersatellite links and heterogeneous spaceborne computers enabling multi satellite distributed computing in orbit. Experimental results demonstrate that cooperative task execution across satellites significantly reduces end-to-end latency achieving up to 90% response-time reduction for remote-sensing tasks.

The SPOCK (Smart Processing Orbital Cloud Konstellation) initiative, funded by the European Space Agency and conducted by Indra and Aresys, extends the state of the art by explicitly targeting Cognitive Cloud Computing in Space (3CS) for Earth Observation (see Fig.1). The project investigates how onboard artificial intelligence, inter-satellite links and software defined payloads can be integrated into scalable mission architectures capable of supporting low-latency and event-driven EO services. SPOCK adopts a system level perspective in which satellites operate as cooperative computing nodes rather than focusing on the optimization of a single satellite or mission. One of the goals of SPOCK is to translate long term ESA vision studies on the matter into implementable architectural and operational concepts. The initiative does not introduce a fixed satellite or payload design, but instead explores how constellations can be architected and operated as distributed computing infrastructures in orbit by integrating sensing, processing and coordination functions within the space segment.

SPOCK also examines two complementary mission configurations for demonstration to support the aforementioned objectives. The first is a multi-satellite demonstrator enabling autonomous tip-and-cue operations among heterogeneous sensors, designed to validate distributed cognition and inter-satellite coordination. The second is a single high-capability satellite acting as a space-based edge computing and coordination hub, supporting sensor fusion, dynamic tasking and federation with external assets. These configurations enable a systematic assessment of architectural trade-offs between distributed autonomy and centralized in-orbit coordination, providing a practical pathway toward future federated Earth Observation systems aligned with emerging multi-operator and multi-mission ecosystems.

The rest of the paper is organized as follows. Section II introduces the SPOCK cloud computing capabilities. Section III the SPOCK system concept mission configurations and requirements. Section IV the key enabling technologies and Section V the conclusions of this paper.

Fig. 2. Novel Capabilities and Mission Paradigms of 3CS4EO

II. COGNITIVE CLOUD COMPUTING FOR EARTH OBSERVATION CAPABILITIES

Cognitive Cloud Computing for Earth Observation represents a transformation at EO system level in which elements such as sensing, on-board processing, decision-making and coordination are tightly integrated to enable adaptive information driven missions (see Fig. 2). Rather than a single technological innovation, the system is realized from the coordinated combination of onboard processing, cloud-native software principles, distributed learning and autonomous coordination across heterogeneous sensing assets. In this context, cognition refers to the system's ability to interpret data, adapt its behavior and optimize resource usage across the full *observation-processing-delivery* chain. This paradigm shifts EO missions away from rigid, pre-planned acquisition strategies and ground centric processing pipelines. The EO data is not treated as a uniform stream of pixels but as a dynamic source of semantic content whose value depends on context, timeliness and relevance. Realizing this vision requires the introduction of novel cloud computing capabilities for EO missions with increasing complexity (see Fig. 3).

The first capability is on-board multi-sensor data fusion, which allows information from heterogeneous payloads, such as SAR and optical, to be combined directly in orbit. Fusion improves scene understanding by compensating for the limitations of each sensor thus enabling more robust object detection, classification and environmental assessment. Building on this enhanced situational awareness, cognitive mission planning replaces static acquisition timelines with adaptive strategies that are responsive to detected events, evolving environmental conditions and user priorities. Within a constellation, this enables tip-and-cue operations, dynamic task redistribution and context-aware scheduling. The tasking decisions are increasingly driven by onboard analytics rather than ground intervention allowing the system to react on timescales compatible with time-critical EO use cases. The practical implementation of such adaptive behavior depends on agile, cloud-native software architectures onboard spacecraft. By adopting containerized applications, modular services and standardized interfaces the EO platforms decouple mission functionality from hardware constraints. This allows algorithms, processing chains and decision logic to be updated or reconfigured in orbit.

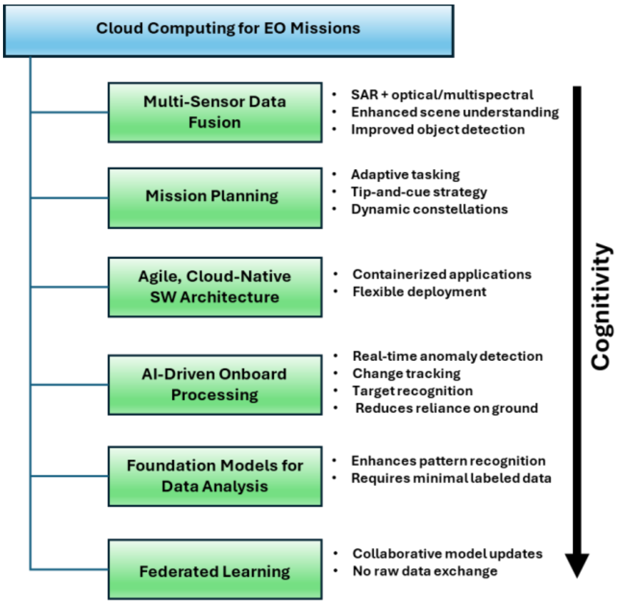

Fig. 3. Novel Cloud Computing Capabilities for EO Missions.

From a system perspective, the satellite becomes an execution environment for EO applications rather than a fixed function payload. This approach mirrors established cloud computing principles while operating under space specific constraints. Within this software defined environment, machine learning models perform real-time inference tasks directly on acquired data. The system reduces downlink volumes and shortens the time between observation and actionable insight by extracting semantic information in orbit. As EO systems scale in complexity and application diversity, foundation models emerge as a key architectural enabler [4]. These models provide strong generalization across sensors and use cases while requiring minimal task specific labeling. This enables rapid deployment of new analytics and improved robustness to changing observation conditions. Their adoption reduces the need for mission specific model development and facilitates interoperability across missions and operators. Finally, federated learning [5] enables continuous improvement of onboard intelligence without exchanging raw data. Models are locally updated using newly acquired observations while only parameters or gradients are shared across the system and with ground infrastructures. This preserves bandwidth, respects data sovereignty constraints and allows learning to scale across constellations.

The above elements define Cognitive Cloud Computing in Space as a paradigm where (i) multi-sensor fusion enhances perception, (ii) cognitive planning enables adaptive behavior, (iii) cloud-native software provides flexibility, (iv) onboard AI delivers low-latency intelligence, (v) foundation models ensure scalability and generalization and (vi) federated learning sustains long-term evolution. In SPOCK, this paradigm serves as basis for the architectural framework that guides the selection of mission configurations, technologies, and operational concepts for low-latency, scalable and federated EO systems.

III. SPOCK SYSTEM CONCEPT, MISSION CONFIGURATIONS AND REQUIREMENTS

The SPOCK system architecture is driven by a set of mission-level and system-level requirements derived from user needs for low latency, scalable and information driven Earth Observation services. Accordingly, the following key areas of constellation architecture requirements have been identified (see Table 1):

- Observation capabilities in terms of coverage, payloads and revisit time,
- On-Board processing and Latency,
- Data quality and consistency (e.g., geolocation accuracy),
- Scalability and modularity,
- Flexibility.

While the final objective of SPOCK is the definition of a scalable constellation architecture enabling cognitive cloud computing in space, the project adopts an incremental approach to system development. The rationale is that given the technological novelty and architectural complexity associated with onboard AI processing, inter-satellite links and software-defined operation, SPOCK introduces a set of representative mission configurations explicitly conceived as technology and concept demonstrators that can be further expanded in the future.

In this context, SPOCK investigates multiple reference configurations that progressively exercise the building blocks required for a future operational constellation. By separating long-term architectural objectives from near-term demonstrator implementations, the project enables early validation of critical capabilities such as autonomous event detection, tip-and-cue tasking, distributed processing and in-orbit coordination while maintaining a clear traceability toward the ultimate constellation based vision. The following subsections first introduce a tradeoff analysis on the constellation geometry and then the demonstrator configurations and discuss how each contributes to assessing the feasibility and performance of Cognitive Cloud Computing in Space architectures for Earth Observation.

A. *Constellation geometry and connectivity*

Understanding the number of satellites required for a given constellation architecture is a critical step in the design of space-based communication systems. The satellite count directly affects not only the mission total cost, including manufacturing, launch, and maintenance, but also influences system complexity, deployment time, coverage performance, and network reliability. In optical ISL constellations, where link range and geometry play a fundamental role in connectivity, having a precise estimate of the number of required satellites is essential for ensuring continuous and redundant communication paths while maintaining a feasible system architecture. The analysis begins by considering a classical Walker star configuration, which distributes satellites evenly across orbital planes with consistent spacing in both right ascension and mean anomaly. To estimate the number of satellites needed within a single orbital plane, several

TABLE I. IDENTIFIED CONSTELLATION REQUIREMENTS DERIVED FROM THE USER/CUSTOMER REQUIREMENTS

Title	**Description**
Wide-Area Coverage	The constellation shall provide wide-area observation capability enabling monitoring of extended regions (maritime zones, floods, wildfires) within a single pass, with constellation sizing to achieve regional coverage goals.
Multi-Sensor Payload Support	Space and platform resources shall support heterogeneous payloads (e.g., optical, SAR, thermal) to enable all-weather/day-night monitoring and cross-sensor fusion.
Revisit Time for Crisis Monitoring	The constellation shall ensure a worst-case revisit time of ≤12 hours over priority AOIs for flood and wildfire scenarios via orbital design and dynamic scheduling.
On-board Processing Latency	On-board AI/Computer Vision pipelines shall deliver event detections and derived products within about 1 minute per imagery product.
Automated Alert Latency	The system shall generate and distribute automated alerts within 5 minutes after on-board event detection, subject to ground segment availability.
Geolocation Accuracy	The system shall meet geolocation accuracy thresholds for products: vessels ≤±10 m; flood and fire products ≤±30 m under nominal conditions.
Inter-Satellite Link	The system shall incorporate inter-satellite communication capability to enable data relay and coordination between constellation elements.
Modular & Reconfigurable HW/SW Design	The constellation shall adopt an open systems paradigm with modular and reconfigurable hardware/software assets to enable in-flight experimentation and low-latency acquisition tasking .
Modular & Expandable Constellation	The constellation shall be modular and expandable to include future satellites and enable IOD/IOV of novel 3CS4EO-enabled capabilities
Open Interface Standards	Satellite-to-satellite and satellite-to-ground interfaces shall use open, modular standards

altitudes and link ranges are evaluated. The aim is to assess how many satellites are necessary to ensure in-plane ISLs, providing continuous connectivity along the orbit path.

Using the assumed Walker star configuration, the number of satellites per plane is computed by determining the maximum angular separation between satellites that can be supported by a given optical link range at different altitudes. For each scenario, the satellite spacing is derived such that the distance between adjacent satellites does not exceed the maximum allowable link range, ensuring uninterrupted communication within each orbital plane. A circular LEO orbit is considered, and a variety of link range and orbit altitudes are used to evaluate the number of satellites to guarantee continuous along-track link in-plane. From geometry considerations, the anomaly angle subtended by two spacecrafts on the same orbital plane with inter-satellite link distance d_{ISL} is:

TABLE II. TOTAL NUMBER OF SATELLITES, IN WALKER-STAR CONFIGURATION, REQUIRED TO GUARANTEE IN-PLANE AND CROSS-PLANE LINKS PER ORBITAL ALTITUDE [KM] AND MAXIMUM RANGE LINK [KM].

Total satellites		**Orbit altitude [km]**					
		500	600	700	800	900	1000
Maximum ISL range [km]	1000	968	968	1035	1058	1058	1128
	2000	242	242	276	276	276	288
	3000	120	120	120	120	128	128
	4000	66	66	66	72	72	72
	5000	45	45	45	45	45	50

$$\theta_{ISL} = 2\, arcsin\, \frac{d_{ISL}}{2(R_E + h)} \quad (1)$$

From this angle, it is possible to estimate the number of satellites needed to cover the orbital plane as:

$$\mathrm{n_{SAT-IP}} = \left\lceil \frac{2\pi}{\theta_{ISL}} \right\rceil \quad (2)$$

After evaluating intra-plane connectivity, the next step is to determine the number of orbital planes required to achieve consistent cross-plane (out-of-plane) ISLs. Similar to the in-plane case, the cross-track angular spacing between planes is constrained by the optical terminal's link range and the geometry imposed by the orbital altitude. This ensures that each satellite maintains one or more stable cross-links to satellites in adjacent planes, thereby enhancing network resilience and enabling low-latency routing across the constellation. By evaluating the across-track distance at the equator between adjacent planes, the maximum allowed longitudinal distance between each orbital plane to guarantee cross-link is found as:

$$\theta_{LONG} = \frac{d_{ISL}}{(R_E + h)\, sin\, i} \quad (3)$$

where R_E is the radius of the Earth, h the satellite orbit height and i the orbit inclinations. Assuming Walker-Star configuration, and assuming a SSO orbit at 97°, the half-longitude globe shall be covered. This means that the number of orbital planes to be set for the Walker-Star configuration is found as:

$$n_{planes} = \left\lceil \frac{\pi}{\theta_{LONG}} \right\rceil \quad (4)$$

Finally, the total number of satellites required for each configuration is calculated by multiplying the number of satellites per plane by the number of orbital planes. This estimation provides a clear overview of the constellation scale necessary to support full in-plane and cross-plane connectivity, as a function of orbital altitude and ISL range constraints. The total number of satellites is approximated as (see Table 2):

$$n_{WS} = n_{SAT-IP} \cdot n_{planes} \quad (5)$$

As expected, the total number of satellites increases when increasing the altitude of the orbit, due to the higher distances to be covered and decreases with increasing capability of maximum inter-satellite link range.

B. Four Satellites Demonstrator Mission Concept

A practical approach to validate the 3CS4EO vision is through a phased deployment strategy. An idea for the initial phase is to launch four satellites along the same orbit and separated by a certain along track distance dictated by the maximum distance allowed by the intersatellite link. This approach serves as a precursor and tests the full constellation. Moreover, these satellites will act as a technology demonstrator, allowing them to test critical components such as inter-satellite laser communication, onboard routing and edge computing processing and thermal management.

In addition to computing experiments, this type of mission based on four satellites can demonstrate tip and cue operations, a feature which is a critical capability for reducing the time gap between data collection and actionable insights. The tip and cue approach involves one satellite detecting an event or anomaly and cueing the other satellite for a more detailed observation. An example of a potential synergistic and complementary tip and cue architecture in terms of payload is SAR paired with three optical sensors.

The tip satellite equipped with the SAR payload continuously monitors with wide swath coverage but at a lower resolution with respect to an optical system. A SAR system can detect several phenomena such as detecting changes such as flooding or infrastructure damage. SAR could detect these changes even under cloud cover or at night. Upon detecting an anomaly through on-board AI processing (i.e., edge computer), the tip satellite equipped with the SAR sends a cue signal to the other satellite with the optical sensor via inter satellite link. The optical satellites then autonomously reorient their optical payload to capture imagery of the same area under daylight conditions, providing detailed visual confirmation and further on-board edge computing. This strategy will help in validating low latency inter satellite communication, autonomous tasking, and AI driven prioritization, which are essential for future large scale constellations supporting disaster response, security, and environmental monitoring.

C. One Satellite Demonstrator Mission Concept

A practical approach to validate the 3CS4EO vision with reduced complexity in terms of mission operations is through a single, highly capable but larger satellite acting as an integrated technology demonstrator. This satellite hosts multiple payloads (optional) and advanced computing capabilities serving as a precursor to future modular constellations. Unlike the multi satellite tip and cue architecture this concept focuses on (i) sensor fusion and high-performance on-board computing (ii) onboard autonomy within a single platform and (iii) platform as a service and interoperability with other constellations and ground systems. Also, this mission concept is suitable for disaster Monitoring (wildfire and floods) and Maritime Domain Awareness aligning the user level requirements. A Sun-synchronous orbit (SSO) is selected for this mission for a number of reasons such as continuous solar power availability and predictable thermal conditions, which are essential for powering advanced computing hardware and maintaining thermal stability. These features not only support energy intensive edge computing but also simplify thermal management making SSO the preferred choice for cloud computing in space. Furthermore, since most EO payloads operate in SSO, this orbit facilitates seamless data exchange between the proposed cloud computing satellite and other sensors enhancing interoperability and collaborative processing. The satellite will integrate:

TABLE III. HIGH-LEVEL REQUIREMENTS FOR SINGLE SATELLITE MISSION CONCEPT ENABLING THE 3CS4EO CONCEPT.

Category	Requirement
Orbit	Sun-synchronous orbit (SSO) at 400–700 km altitude. Expandable according to Walker principle.
Payloads	Integrated SAR (wide swath, medium resolution) and optical camera (high resolution). These are optional.
Onboard Computing	≥500 TOPS performance with parallel GPUs (clustered architecture) and FPGA acceleration. Modular and reconfigurable. Fault-tolerant design.
Modularity and Reconfigurability	Open systems paradigm enabling HW/SW reconfiguration (platform-as-a-service).
Data Handling	Onboard storage, real-time prioritization and compression for downlink.
Autonomy	Full self-tasking and dynamic scheduling. Health monitoring without ground intervention for ≥24 hours.
Security	End-to-end encryption for all links
Ground Integration	Cloud API compatibility for orchestration and federated learning experiments.
Interoperability	Open interface standards for future constellation expansion, ground systems and third-party missions' interaction.
External Connectivity	Capability to interface external communication constellations for high throughput communication and global coverage.

- A high-resolution optical camera for detailed imaging (optional). High-resolution optical imagery complements SAR by providing detailed visual confirmation of anomalies detected by radar. This is important for wildfire and floods mapping and maritime vessel identification.

- A Synthetic Aperture Radar (SAR) for all-weather, day and night observation (optional). SAR acquisitions are critical for disaster monitoring during floods or wildfires, especially under cloud cover or at night. SAR can detect water extent, infrastructure damage, and changes in terrain even when optical sensors are obstructed.

- Advanced edge computing hardware featuring parallel processing units combined with accelerators enabling real time AI inference and dynamic tasking.

- Modular and reconfigurable HW/SW design supporting in-flight experimentation and future upgrades (platform-as-a-service).

- Open interface standards for interoperability with ground systems and potential future satellite EO constellations

To further enhance data relay and global accessibility, the one satellite concept shall integrate with external communication constellations. This capability enables high throughput communication, reduces latency for time sensitive applications, and provides resilient connectivity even in scenarios where direct ground station access is limited. Leveraging commercial broadband constellations ensures easy integration with terrestrial cloud infrastructures and supports real time orchestration of AI driven tasks across distributed networks. In addition to processing its own sensor data, the satellite can ingest information from ground stations and other satellites enabling collaborative processing and enriching onboard analytics with external datasets. This combination of high-performance computing, multi-source data integration and low latency connectivity positions the satellite as a space edge computing hub paving the way for future orbital data centers. The proposed mission concepts allow simultaneous acquisition of complementary datasets (optical and SAR), onboard fusion with both on-board data and data from third party and low latency product generation. Moreover, it will validate critical capabilities such as autonomous anomaly detection, adaptive sensing and cloud integrated orchestration. Table 3 summarizes the high-level requirements for this one satellite mission concept.

IV. KEY ENABLING TECHNOLOGIES

The SPOCK concept relies on the enabling technologies reported in Fig. 4 namely (i) edge computing hardware, (ii) high-speed router and (iii) intersatellite link whose joint maturity makes cloud computing in space technically feasible within the next decade. While none of these technologies is entirely novel in isolation, their coordinated application within a spaceborne, distributed and autonomous EO architecture represents a fundamental difference from traditional mission designs. This section discusses the most critical technological enablers for SPOCK with emphasis on their system-level role rather than individual component specifications.

A. Edge Computing

SPOCK strongly relies on the availability of high-performance edge computing platforms capable of executing complex EO processing and AI inference directly in orbit. Historically, onboard computers were designed primarily for telemetry handling and payload control with limited computational headroom. Recent advances, however, have introduced space-qualified or space-proven heterogeneous computing platforms that combine CPUs, GPUs and FPGAs within compact and power-efficient architectures. GPUs provide the raw computational throughput required for AI-driven workloads such as image segmentation, object detection, multimodal fusion and foundation-model inference. Demonstrated in-orbit platforms based on COTS GPUs already achieve on the order of 100–250 TOPS [3] while maintaining acceptable power consumption for small- to medium-class satellites. These capabilities enable near-real-time processing of SAR and optical data streams, transforming the satellite into an active processing node rather than a passive sensor. FPGAs complement GPU processing by supporting deterministic, low-latency and power-efficient computation. FPGA accelerators are particularly well suited for tasks such as SAR focusing, signal preprocessing, data compression and encryption. Offloading these functions reduces GPU load and improves overall system efficiency and enhances fault tolerance. The resulting heterogeneous architecture allows computational tasks to be matched to the most appropriate hardware resource, a key requirement for sustaining autonomous operations under constrained power and thermal budgets. From a system perspective, computational performance alone is insufficient. Edge computing platforms must also deliver high efficiency in terms of performance per watt, per kilogram, and per unit volume. These metrics directly affect payload accommodation, mission lifetime, and launch cost, and therefore represent primary design drivers rather than secondary optimizations.

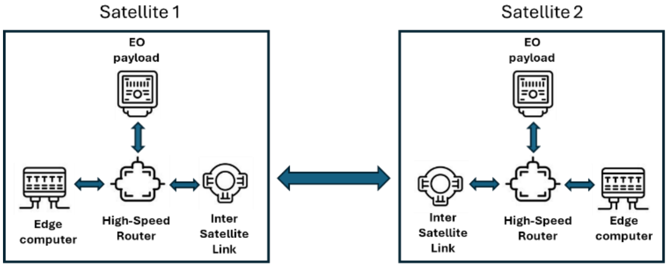

Fig. 4. Key components for space computing are processing, routing and transmission.

B. Distributed Computing and Onboard Networking

Cognitive EO operations at constellation scale require efficient data routing and coordination not only within a satellite but also across multiple spacecraft. SPOCK therefore relies on distributed computing enabled by onboard networking elements such as space-qualified Ethernet switches and high-throughput inter-processor links. These components form the digital backbone of the spaceborne cloud, enabling data exchange between payloads, edge computers, and communication subsystems with predictable latency. Modern space-qualified switches already support multi-gigabit aggregate throughput with sub-millisecond latency, making them suitable for coordinating AI inference pipelines, routing ISL traffic and supporting fault-tolerant system architectures. Importantly, these switches enable modular payload and processing architectures in which subsystems can be added, removed or reconfigured without redesigning the entire onboard data handling chain. At the constellation level, distributed computing enables workloads to be partitioned across multiple satellites, either to parallelize processing or to exploit spatial and temporal diversity in observations. While fully dynamic workload migration in orbit remains a long-term objective, SPOCK adopts a pragmatic approach in which cooperative processing and data sharing are initially limited to well-defined tasks such as tip-and-cue coordination, alert propagation and federated learning updates.

C. Inter-Satellite Communications

Inter-satellite links provide the low-latency and high-capacity communication channels required for constellation level cognition. Whereas traditional EO missions relied almost exclusively on ground relay, SPOCK leverages ISLs to decouple coordination and processing from ground visibility constraints. Optical inter-satellite links are identified as the primary candidate technology due to their superior data rates, inherent resistance to interference and reduced regulatory burden compared to RF systems. State-of-the-art optical terminals today demonstrate data rates in the multi-gigabit range over thousands of kilometers, with ongoing developments targeting even higher throughput and further miniaturization. These capabilities are sufficient to support the exchange of processed data products, semantic events, and coordination messages between satellites in LEO constellations.

The adoption of optical ISLs introduces new system-level challenges, most notably in terms of pointing, acquisition, and tracking. As a result, SPOCK treats ISLs as an integrated spacecraft subsystem rather than a standalone payload. Close coupling between attitude determination and control, terminal pointing mechanisms, and onboard scheduling is required to ensure robust link availability. Hybrid architectures combining optical ISLs with RF links or external relay constellations are therefore considered as a means to balance performance, availability and operational complexity.

V. CONCLUSIONS

This paper presented the SPOCK initiative as a system level approach toward enabling cognitive cloud computing in space for Earth Observation. The ultimate objective of SPOCK is the definition of a scalable constellation architecture in which satellites operate as cooperative computing nodes, capable of delivering low-latency, event-driven EO services through onboard processing and inter-satellite coordination. A core architectural principle for this vision is the adoption of an open, modular and reconfigurable hardware and software design. This enables in-flight experimentation, rapid reconfiguration of onboard functions and responsive acquisition tasking.

To progress toward this long-term goal, the paper introduced a set of reference mission configurations proposed as demonstrators. These demonstrators are conceived to validate key enabling technologies, operational concepts and system behaviors under realistic constraints. This also serves as intermediate steps for assessing feasibility, performance and architectural tradeoffs ahead of the full constellation deployment. The modular structure of both the space segment and the mission configurations is intentionally designed to support incremental expansion. This enables future satellites to be added to the system and allowing In Orbit Demonstration (IOD) and In-Orbit Validation (IOV) of novel 3CS4EO enabled capabilities without redesigning the overall architecture.

SPOCK establishes key design principles for future Earth Observation systems in which multiple satellites and operators can cooperate within a shared, intelligent space infrastructure by enabling interoperability, reconfigurability and scalable growth.

ACKNOWLEDGMENT

The authors would like to acknowledge the contributions of Paolo Falcone from Aresys. This work has been performed under ESA Contract No. 4000149090/25/I-DT.

REFERENCES

[1] Giuffrida, G., Fanucci, L., Meoni, G., Batič, M., Buckley, L., Dunne, A., ... & Aschbacher, J. (2021). The Φ-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.

[2] Marin, A., Coelho, C., Deconinck, F., Babkina, I., Longepe, N., & Pastena, M. (2021). Phi-Sat-2: Onboard AI apps for earth observation. Proc. Space Artif. Intell, 6.

[3] Wang, H., Yao, K., Gong, L., Jin, Y., Zhang, P., Liu, Y., ... & Zhao, Z. (2025). Space Computing Constellation: System Architecture, Implementations and Challenges. IEEE Internet of Things Journal.

[4] Guo, X., Lao, J., Dang, B., Zhang, Y., Yu, L., Ru, L., ... & Li, Y. (2024). Skysense: A multi-modal remote sensing foundation model towards universal interpretation for earth observation imagery. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (pp. 27672-27683).

[5] Moreno-Alvarez, S., Paoletti, M. E., Sanchez-Fernandez, A. J., Rico-Gallego, J. A., Han, L., & Haut, J. M. (2024). Federated learning meets remote sensing. Expert Systems with Applications, 255, 124583.