E/0904-611 – GSTP Element 1 “Develop”
ESA Contract No. 4000139825/22/NL/AS

SoC-HEALTH^2

Hierarchical Health Management
in Heterogeneous Systems

OCFM Framework for FPGA-Based Satellite Electronics
Executive Summary

2025

Notices

For information, please contact us by e-mail: info@testonica.com

This document is intended to inform interested stakeholders about the key achievements in SoC-HEALTH2 project (ESA Contract No. 4000139825/22/NL/AS) with respect to Health Management demonstration framework development.

Document version 1.0, date 23.01.2025

1 Overview

Ensuring the reliability and fault tolerance of satellite electronics is paramount for mission success in diverse space environments. Over the last decade, the space sector has seen significant diversification between traditional institutional missions led by agencies such as ESA and NASA and commercial New Space initiatives driven by private companies. Both domains have increasingly adopted high-performance commercial-off-the-shelf (COTS) components, such as FPGAs, accelerators, microprocessors, and memories, to meet growing demands for computational power and cost efficiency. COTS components offer orders of magnitude higher performance compared to radiation-hardened or radiation-tolerant alternatives. However, their susceptibility to radiation-induced Single-Event Effects (SEE) and Total Ionizing Dose (TID) radiation effects poses critical challenges.

For institutional missions, reliability and compliance with stringent standards remain paramount. These missions often use proven and qualified components, but efforts are underway to incorporate COTS devices alongside advanced fault mitigation strategies. Conversely, commercial New Space missions prioritize high performance and rapid deployment, leveraging COTS components with systematic testing and software-based fault mitigation to achieve cost and time efficiencies. CubeSat deployments and Low Earth Orbit (LEO) constellations often operate with minimal qualification processes, relying on shorter mission lifetimes and operational flexibility to accommodate risks.

The **On-Chip Fault Management (OCFM)** framework provides a unified solution to address the challenges of both institutional and commercial missions. The framework integrates real-time fault detection, isolation, and recovery capabilities into FPGA-based satellite electronics, including platforms like **Xilinx Versal ACAP**. By implementing fine-grained health monitoring and hierarchical cross-layer fault management, OCFM ensures autonomous and adaptive system operation even under adverse conditions.

OCFM’s proactive fault management enables graceful degradation, leveraging remaining healthy resources to sustain mission-critical functions. Its modular design supports scalability, from single SoC implementations to complex multi-board systems, ensuring applicability across diverse mission architectures. The framework also facilitates compliance with industry standards for institutional missions while offering the agility needed for New Space applications. This dual focus positions OCFM as an essential enabler of next-generation satellite systems, balancing performance, reliability, and cost efficiency.

A key advantage of this solution is its hierarchical cross-layer health management approach. Fault detection, data collection, and recovery mechanisms are distributed across the system’s hardware, OS software, and application layers. This architecture enables fault-aware task scheduling and adaptive resource utilization, ensuring the highest possible performance using remaining healthy resources. By leveraging the OCFM framework, satellite systems can achieve unparalleled resilience and operational efficiency, reducing downtime and mission risks while minimizing the reliance on full hardware redundancy. The ESA-funded GSTP project SoC-HEALTH2 expanded the OCFM framework to support advanced FPGA SoC platforms like Xilinx Versal ACAP (Adaptive Compute Acceleration Platform) to a significant extent.

2 The Changing Paradigm of Space Missions

In the evolving landscape of satellite technology, ensuring reliability and fault tolerance has become a critical challenge. Institutional space agencies like ESA and NASA prioritize stringent reliability and compliance standards, ensuring mission success over extended lifetimes. Conversely, New Space ventures, driven by commercial entities, emphasize rapid deployment, cost efficiency, and high performance. This divergence underscores the need for scalable and adaptable fault handling solutions capable of addressing the unique challenges of both domains.

The adoption of high-performance commercial-off-the-shelf (COTS) components, such as FPGAs, accelerators, microprocessors, and memories, has revolutionized the space sector by providing superior computational capabilities. However, these components bring inherent vulnerabilities to radiation-induced effects, such as Single Event Upsets (SEUs) and Total Ionizing Dose (TID), making robust fault management indispensable.

COTS components are pivotal in bridging performance gaps, enabling advanced capabilities for Earth observation, telecommunications, and scientific research. However, their susceptibility to radiation-induced faults poses a significant barrier to their widespread adoption in critical missions. Satellites operating in Low Earth Orbit (LEO) and beyond encounter harsh environments where robust fault mitigation strategies are non-negotiable.

The **On-Chip Fault Management (OCFM)** framework emerges as a cutting-edge solution, addressing the diverse needs of institutional and commercial New Space missions alike.

3 Radiation Environment and Challenges in Space

The radiation environment that poses significant challenges to satellite electronics primarily consists of:

- **Protons and Electrons**:
  - High-energy protons are abundant in the Van Allen radiation belts surrounding Earth, particularly affecting satellites in LEO.
  - Electrons with lower mass but significant energy levels also contribute to cumulative radiation exposure, causing degradation over time.
- **Cosmic Rays**:
  - Galactic Cosmic Rays (GCRs) are high-energy charged particles originating outside our solar system. They include protons, heavy ions, and nuclei, and are capable of penetrating satellite shielding, leading to ionization effects in sensitive components.
- **Solar Particle Events (SPEs)**:
  - During periods of heightened solar activity, solar flares and coronal mass ejections release bursts of high-energy protons and heavy ions. These events can cause sudden and severe radiation exposure to satellites.

4 Faults Induced by Radiation

Radiation exposure leads to various faults in satellite electronics, which can disrupt or degrade performance. The primary fault mechanisms include:

- **Single Event Effects (SEEs)**:
    - **Single Event Upsets (SEUs)**: Bit flips in memory or logic circuits caused by a single charged particle striking a sensitive node.
    - **Single Event Latchups (SELs)**: Persistent high-current states in CMOS devices that may lead to component failure.
    - **Single Event Transients (SETs)**: Temporary voltage spikes in analog or digital circuits caused by ionizing particles.
- **Total Ionizing Dose (TID)**:
    - Accumulation of ionizing radiation over time can degrade the performance of semiconductor materials, leading to parameter shifts or permanent damage in components.
- **Displacement Damage (DD)**:
    - Non-ionizing energy deposition displaces atoms in the crystal lattice of semiconductor materials, reducing performance in photodiodes, sensors, and transistors.

These faults necessitate robust fault mitigation strategies to ensure uninterrupted satellite operations and prevent mission-critical failures.

5 Typical Solutions and Their Limitations

Several fault mitigation solutions are typically employed in the space sector, primarily focusing on hardware redundancy, error correction codes (ECC), and Triple Modular Redundancy (TMR). While these methods have proven effective in mitigating faults, they exhibit several limitations:

- **Hardware Redundancy**:
    - Relies on duplicating or triplicating hardware components to ensure fault tolerance. This approach significantly increases weight, power consumption, and cost, making it impractical for many New Space missions.
    - Redundant hardware cannot dynamically adapt to evolving fault scenarios, leading to inefficient resource utilization.
- **Error Correction Codes (ECC)**:
    - While effective for memory integrity, ECC is limited to detecting and correcting single-bit or small-scale errors. Larger-scale radiation-induced faults may go undetected or require additional mechanisms for recovery.
- **Triple Modular Redundancy (TMR)**:
    - Ensures fault tolerance through majority voting but increases system complexity and resource requirements.
    - Vulnerable to common-mode failures caused by simultaneous faults in redundant modules, particularly in harsh radiation environments.
- **Software-Based Fault Management**:
    - Commonly used in New Space missions due to its cost efficiency and flexibility. However, software-only approaches are slower in fault detection and recovery, lacking the real-time responsiveness required for mission-critical operations.

6 SoC-HEALTH2 OCFM Framework: A Unified Approach

Ensuring the reliability and fault tolerance of satellite electronics is paramount for mission success in diverse space environments. The On-Chip Fault Management (OCFM) framework offers a holistic cross-layer fault management approach to fault detection, isolation, and recovery (FDIR) that spans from hardware (e.g., DSPs, NoC) to software (OS, application tasks), leveraging Versal's integration of ARM cores and programmable logic. In this manner, the OCFM ensures sustained operations for a diverse range of satellite architectures even in adverse conditions.

6.1 Key Features

- **Real-Time Fault Management**: Integration of fault monitors and sensors for detecting and isolating hardware faults across processing, memory, and interconnects, while the Asynchronous Fault Propagation Network (AFPN) ensures low-latency fault signalling.
- **Health-Aware Scheduling**: Dynamic task allocation based on real-time health maps optimizes resource utilization and operational continuity. Tasks are continuously reassessed against the system's evolving health status, prioritizing critical operations on the healthiest available resources. In degraded states, non-essential tasks are postponed or redistributed to prevent further system strain. This adaptive mechanism ensures mission-critical functionality while preserving system longevity.
- **Configuration Memory Protection**: Dedicated mechanisms such as scrubbing, partial reconfiguration, and redundant bitstream storage for safeguarding FPGA configuration memory against radiation-induced faults, ensuring operational integrity.
- **Radiation-Hardened Design**: Enhanced resilience against radiation-induced faults, including SEUs and SEFIs, safeguarding critical operations in space environments.
- **Scalability**: Modular design for seamless integration with diverse satellite architectures, allowing for customization to meet specific mission requirements.
- **AI-Driven Fault Prediction**: Machine learning models predict potential hardware failures, enabling proactive fault mitigation and preemptive recovery strategies.

6.2 Architecture Enabling Resilience and Scalability

At the heart of the OCFM framework lies the **Instrument Manager (IM)** and the **Health Manager (HM)**, each serving distinct yet complementary roles in ensuring fault resilience and operational reliability:

- **Instrument Manager (IM)**:
    - Acts as the centralized hub for fault aggregation and diagnostics at the hardware level. The IM interfaces with embedded sensors, fault monitors, and Reconfigurable Scan Networks (RSN) to collect and process real-time health data.
    - Employs redundancy mechanisms such as Triple Modular Redundancy (TMR) and Error Correction Codes (ECC) to ensure fault-free operation even under radiation-induced soft errors.
    - The IM ensures swift responses to urgent fault conditions using latency-optimized communication pathways via the Asynchronous Fault Propagation Network (AFPN).

  - Operates autonomously within the hardware layer to isolate, analyze, and propagate fault data, ensuring the integrity of the system's physical components.

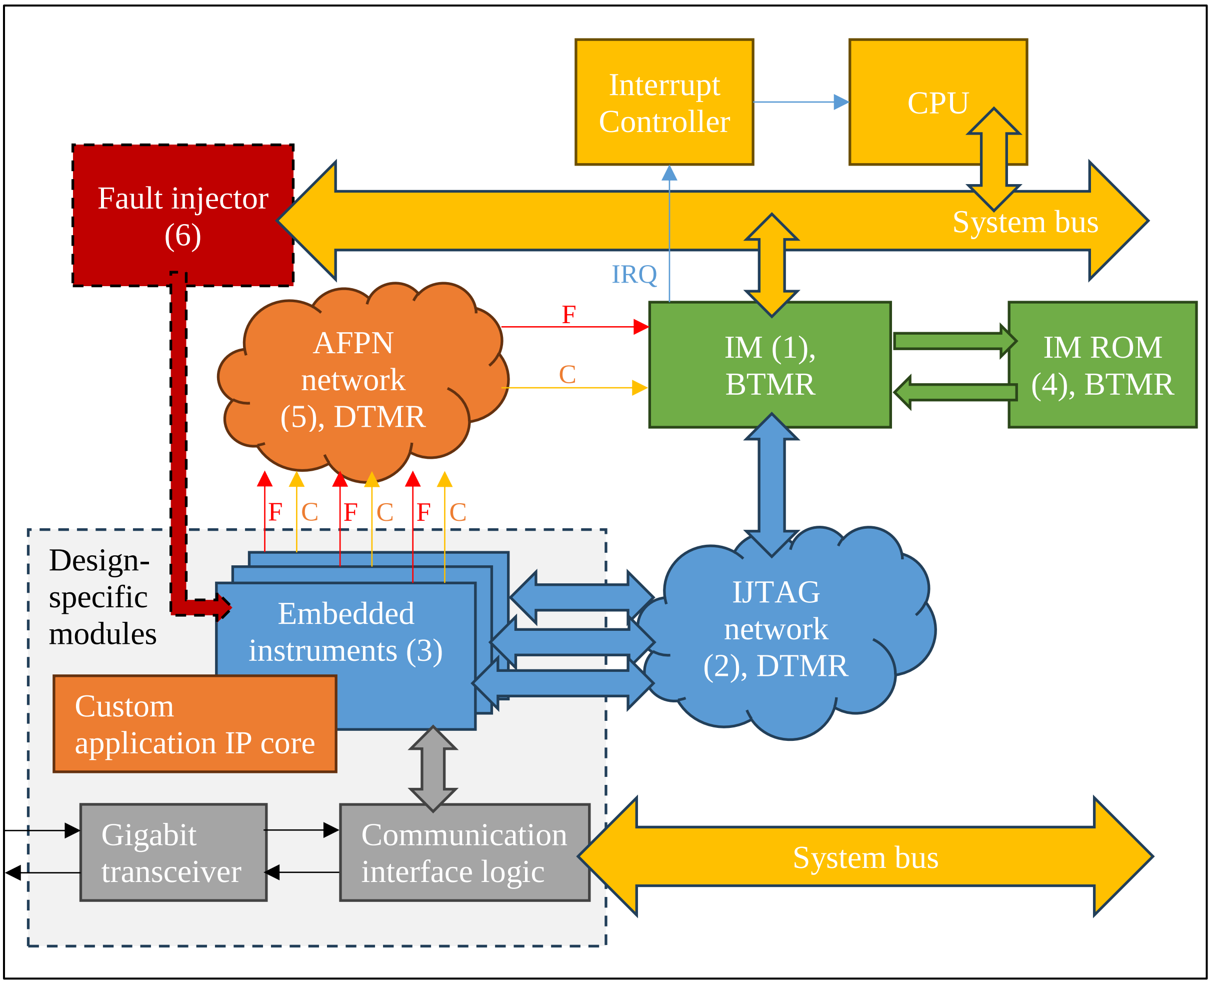

Figure 1. HM hardware block diagram

- **Health Manager (HM)**:
  - Operates at the software layer, leveraging the data provided by the IM to build dynamic health profiles of system components.
  - Maintains the **Health Map (HM)** and **Resource Map (RM)**, which provide real-time insights into system health and resource availability.
  - Facilitates health-aware scheduling and resource allocation, dynamically adjusting system operations to optimize performance and mitigate risks during degraded states.
  - Implements recovery protocols for high-level system operations, coordinating with the IM to initiate hardware-level fault mitigation actions as required.

Together, the IM and HM form a hierarchical fault management ecosystem, where the IM ensures low-level fault detection and immediate response, while the HM oversees system-wide resource management and recovery planning. This collaboration ensures uninterrupted fault management even under severe radiation-induced disruptions, providing a robust foundation for satellite electronics. Other important OCFM framework components are:

- **Fault Monitors**:
    - Sensors embedded in FPGA pipelines, memory, and interconnects provide granular health data and fault detection capabilities.
    - Integration with IJTAG-based networks enables comprehensive fault coverage and efficient instrumentation control.
- **Asynchronous Fault Propagation Network (AFPN)**:
    - Enables low-latency fault signaling across the system, ensuring rapid detection and isolation of critical faults.
    - Designed for scalability to handle large multi-board systems, maintaining efficiency under high fault loads.
- **Health Map (HM) and Resource Map (RM)**:
    - HM dynamically updates the health status of system components, enabling fault-aware decision-making.
    - RM tracks the availability and utilization of system resources, facilitating adaptive scheduling and resource allocation.
- **Self-Protection Mechanisms**:
    - The OCFM framework itself is safeguarded against soft errors through redundancy in critical fault management paths, including the IM and AFPN subsystems. Error Correction Codes (ECC) are implemented in memory blocks to ensure the integrity of stored health data.
    - Periodic self-check routines monitor the integrity of the OCFM hardware and software layers, with a fallback mechanism to recover from transient faults.
    - Triple Modular Redundancy (TMR) is employed in critical components such as the IM, ensuring fault-free operation even under radiation-induced soft errors.

Additionally, the framework's integration of configuration memory protection mechanisms ensures system integrity. Periodic scrubbing corrects transient errors, while ECC and TMR enhance resilience against persistent faults. These features are particularly critical for FPGAs in space, where configuration memory integrity is paramount for mission success.

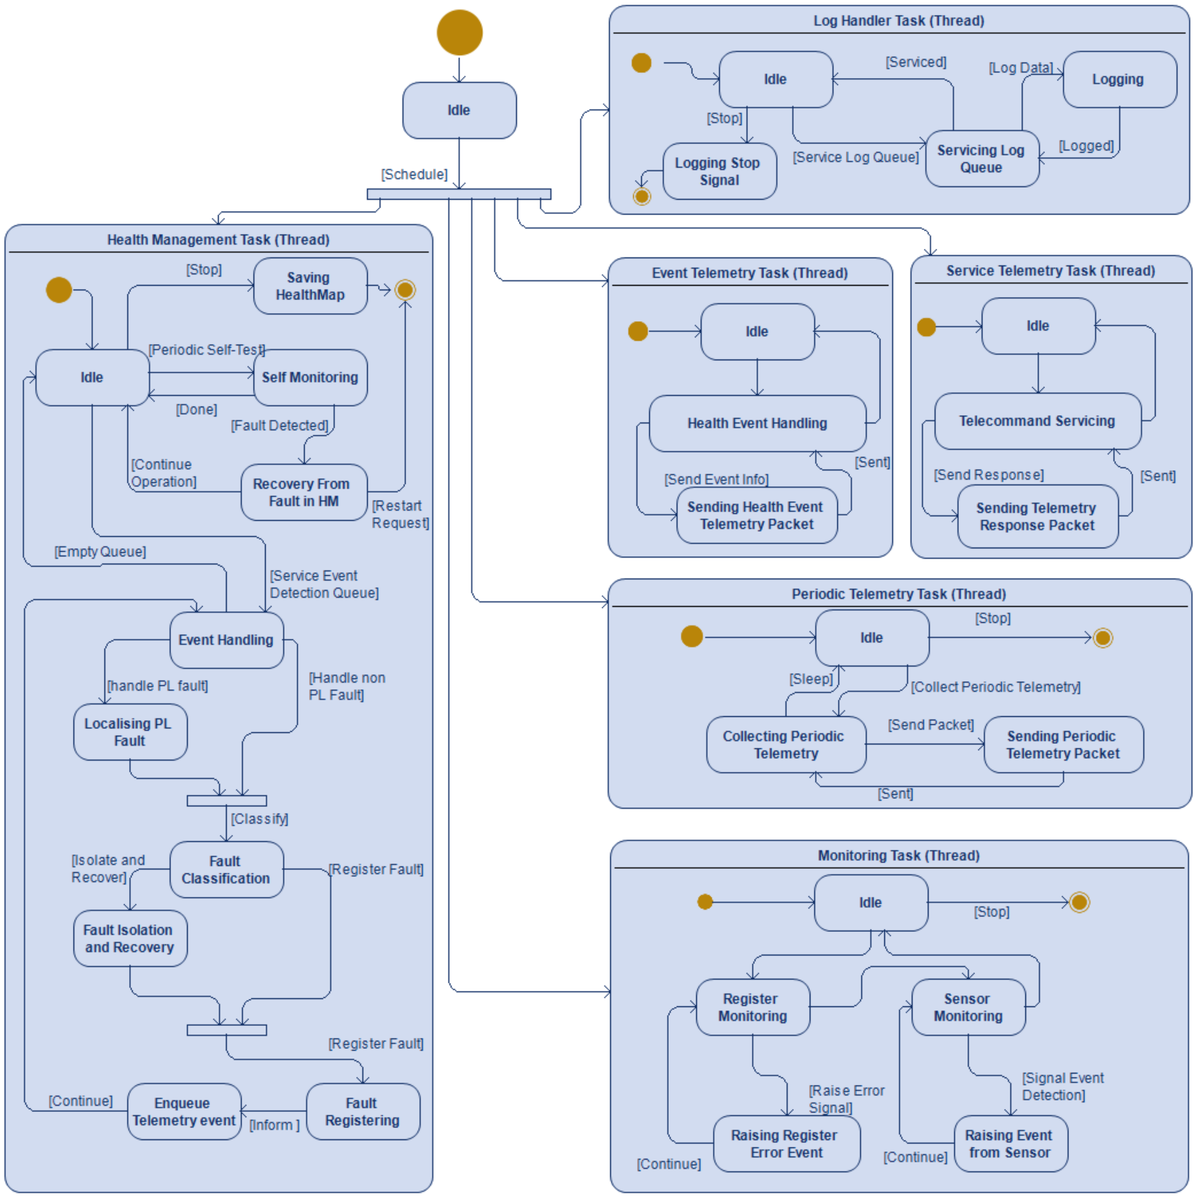

Figure 2. HM software state machine diagram

6.3 Health Management Software Stack

- **Real-Time Fault Management**:
    - Integrated into the operating system to handle fault detection, reporting, and recovery processes autonomously.
    - It applies pre-configured rules and adaptive algorithms to detect anomalies and trigger fault-handling mechanisms.
    - It is built with a hierarchical architecture to manage scalability across complex systems, including multi-board and heterogeneous configurations.
    - Supports multi-layer health monitoring across hardware, OS, and application levels.
- **APIs and Extensions**:
    - HealthMap API provides a standardized interface for health data collection and analysis.

    - Instrument Manager API enables custom configurations for system-specific fault monitoring needs.
- **AI and Predictive Models**:
    - Machine learning algorithms analyze historical health data to identify trends and predict potential hardware failures.
    - Integration with adaptive schedulers ensures proactive reallocation of resources based on fault predictions.

6.4 FPGA-Specific Enhancements

The ESA-funded GSTP project SoC-HEALTH2 expanded the OCFM framework to support advanced FPGA SoC platforms like **Xilinx Versal ACAP** (Adaptive Compute Acceleration Platform) to a significant extent. Versal ACAP introduces a heterogeneous and highly flexible architecture, including scalar engines (ARM CPUs), adaptable hardware (programmable logic), and intelligent engines (AI cores) as well as some integrated fault-tolerance infrastructure. These vendor-provided fault-tolerance mechanisms are, however, fragmented, leading to inefficiencies and unprotected areas. The OCFM extension to Versal ACAP aims at closing these gaps by offering advancements in several areas.

- **Embedded Instrumentation**:
    - Custom IP blocks deployed in FPGA fabric for monitoring buses, interfaces, and accelerators.
    - Protects flip-flops and distributed memory not covered by existing vendor ECC mechanisms.
    - Leveraging programmable logic to implement fault-detection logic tailored to mission-specific requirements.
    - Monitoring Versal’s AI matrix multipliers and DSP cores for faults in tensor and vector processing units within AI engines.
    - Monitoring Versal’s PMC and NoC to gather detailed health data and optimize task scheduling across the NoC-connected resources.

- **Integration with Heterogeneous Components**:
    - The modularity of OCFM allows it to monitor scalar (CPU cores), adaptable (programmable logic), and intelligent (AI cores) components of Versal ACAP.
    - Flexible interfaces for external health monitoring modules to expand coverage beyond the SoC.
    - the Resource Map (RM) includes specialized components like DSPs, AI cores, and adaptable logic blocks, alongside CPUs and memory units.
    - Enhance fault classification algorithms support distinct fault models across these heterogeneous elements.
- **Dynamic Reconfiguration Support**:
    - The OCFM can benefit from Versal’s ability to dynamically reconfigure hardware at runtime to isolate faulty components and redirect workloads to healthy ones.
    - Interface APIs to Versal’s dynamic reconfiguration controllers
- **Advanced Scheduling**:
    - The health-aware task scheduler can integrate with Versal’s hardware and software orchestration, ensuring resource allocation considers real-time fault and health data.

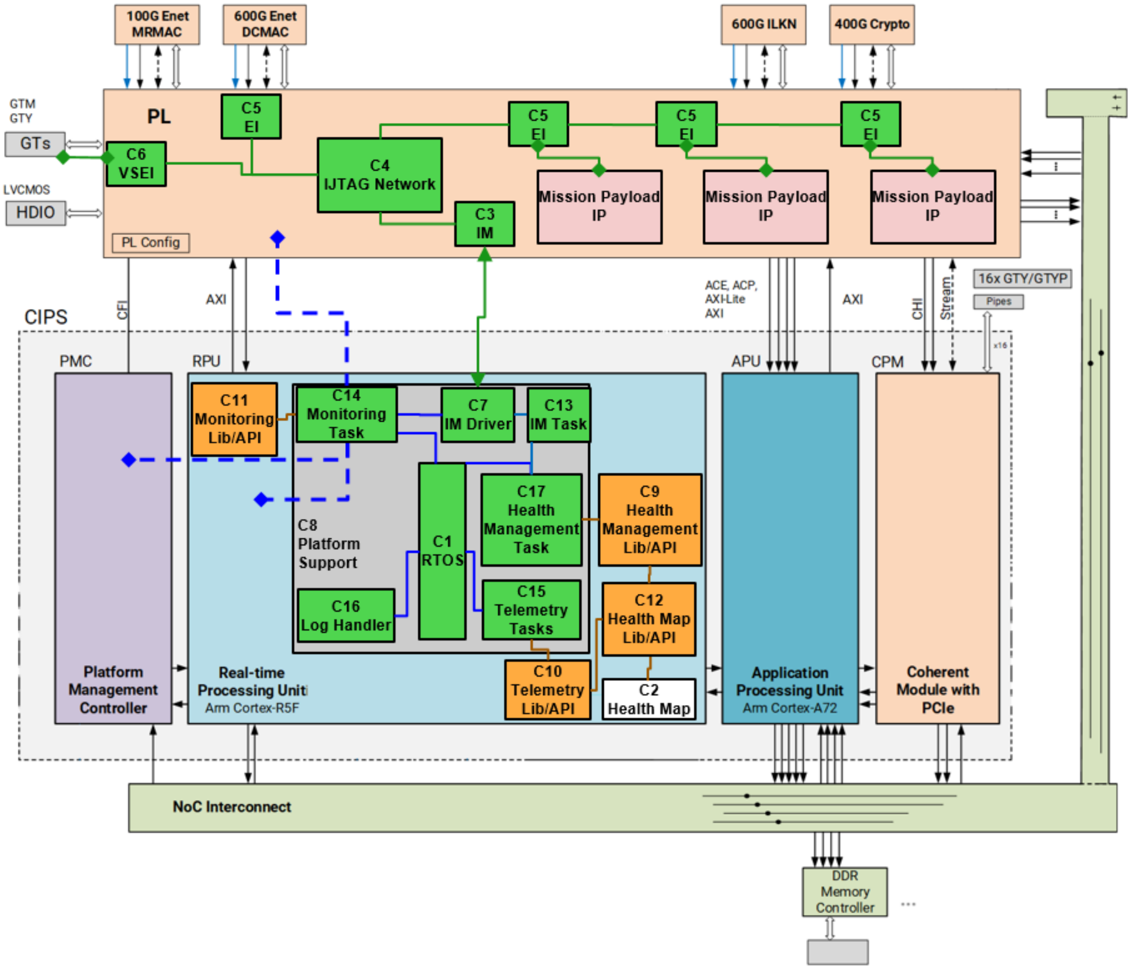

Figure 3. HM demo components location in Xilinx Versal ACAP platform

Incorporating a cross-layer reliability model integrating hardware health data with OS-level fault handling and recovery protocols based on a unified health map that correlates faults across hardware and software layers, the OCFM framework provides a strong foundation for fault handling and health monitoring, supporting a platform as complex as Versal ACAP.

7 Benefits and Differentiators

Adapting the OCFM framework to Versal ACAP allows to leverage its complexity, reconfigurability, and heterogeneity towards advancements in dynamic reconfiguration, intelligent fault prediction, and cross-layer integration. OCFM maximizes the reliability and efficiency of Versal-based systems in the following ways.

- **Increased Reliability**: Continuous operation of critical systems under adverse conditions ensures mission success and operational stability.
- **Reduced Costs**: Minimization of hardware redundancy requirements and associated weight and launch costs provides significant financial and payload efficiency.
- **Extended Mission Lifespan**: Fault prediction, coupled with graceful degradation mechanisms, extends system longevity and reduces the need for frequent interventions.
- **Operational Flexibility**: Adaptable to heterogeneous FPGA-based systems and future mission needs, ensuring compatibility with a variety of satellite platforms and payload configurations.
- **Standards Compliance**: Fully aligned with ECSS requirements for space electronics, ensuring rigorous quality and reliability.
- **Improved Sustainability**: Proactive fault management reduces system downtime, enabling more efficient utilization of satellite resources throughout the mission lifecycle.

7.1 Applicability to payload and OBC subsystems

The OCFM framework is arguably better suited for the on-board computer (OBC), given its centralized role in satellite operations and the need for system-wide fault resilience. However, implementing OCFM in payload subsystems provides a value for missions where payload data is the mission's primary output. For maximum mission reliability, a hybrid approach could integrate OCFM into both subsystems, leveraging its adaptability to address the specific needs of each.

7.1.1 Applicability to payload subsystems:

- The payload often contains highly specialized instruments and sensors critical to the mission's primary objectives (e.g., imaging, communications, scientific experiments).
- OCFM's fine-grained health monitoring and fault-tolerant mechanisms ensure sustained operation of these sensitive components.
- Health-aware scheduling can prioritize mission-critical tasks during degraded states, ensuring payload functionality is maintained even when resources are constrained.
- Configuration memory protection and partial reconfiguration mechanisms safeguard the payload's programmable logic and FPGA subsystems against radiation-induced faults.
- While ensuring payload resilience is vital, the complexity of implementing OCFM across diverse and specialized payload designs may require additional customization.

7.1.2 Applicability to OBC subsystems:

- The OBC is the brain of the satellite, responsible for managing mission operations, processing data, and maintaining communication with ground stations.
- OCFM is inherently well-suited for the OBC, as its hierarchical fault management integrates seamlessly with the processing cores, memory units, and interconnects typically found in OBC architectures.

- AI-driven fault prediction can enhance the OBC's ability to preemptively address potential failures, ensuring uninterrupted command and data handling (C&DH) operations.
- The framework's ability to dynamically reallocate resources ensures continued operation of critical OBC functions, even in partially degraded states.
- The OBC's centralized role in mission success means fault management here is non-negotiable, and OCFM's robust design can maximize reliability in this subsystem.

7.2 Added Value for Users

While traditional approaches depend heavily on full hardware redundancy, increasing cost and power consumption, the OCFM approach balances the workload only being activated for detecting and mitigating faults where it’s needed, thus reducing redundancy needs while maintaining high reliability. Designed for FPGA-based platforms, including Xilinx Versal ACAP, enabling scalability across diverse satellite architectures and use cases. AI-driven fault prediction ensures adaptability to evolving mission requirements and emerging technologies, making the system robust against unforeseen challenges.

- **Mission Success Assurance**: Real-time fault management and predictive maintenance ensure uninterrupted satellite operations and maximize uptime.
- **Cost Efficiency**: Reduced dependency on hardware redundancy translates to lighter payloads, lower launch costs, and better utilization of mission budgets.
- **Enhanced Reliability**: Comprehensive fault coverage mitigates the risks of radiation-induced faults and other environmental challenges, ensuring long-term operational stability.
- **Long-Term Support**: Post-deployment services, including updates, optimizations, and technical support, safeguard mission longevity and adaptability.
- **Future Adaptability**: Modular and AI-driven design ensures readiness for future technologies and mission requirements, offering flexibility for evolving operational needs.

8 Conclusion

The proposed OCFM implementation for FPGA-based satellite electronics offers a cutting-edge solution for improving system reliability, fault tolerance, and cost efficiency. By leveraging advanced fault detection, isolation, and AI-driven prediction, this framework addresses the unique challenges of operating in extreme space environments.

The ESA-funded GSTP project SoC-HEALTH2 expanded the OCFM framework to support advanced FPGA SoC platforms like Xilinx Versal ACAP (Adaptive Compute Acceleration Platform) to a significant extent.

Our team is committed to delivering a customized and scalable solution tailored to target satellite mission needs. With a focus on long-term support, adaptability, and mission-critical reliability, the OCFM framework is the ideal choice for next-generation satellite electronics.

For further discussions and a detailed proposal, please contact us at info@testonica.com.