National Aeronautics and Space Administration

NASA's High Performance Spaceflight Computer

I. Introduction

As humans and robotic space missions continue to explore the universe, scientific observation and analysis of remote locations require advanced levels of autonomy to successfully achieve NASA's science goals. Spacecraft computational power, connected systems, power efficiency and resilience in dynamic, harsh environments are critical to improving the quantity and quality of science return.

NASA's Science & Technology Mission Directorate (STMD) has embarked on a Game Changing Development (GCD) effort to significantly improve the capabilities of space-qualified computing technology and funded the *High Performance Spaceflight Computing* (HPSC) program to deliver this technology.

HPSC is a fault tolerant, rad-hard-by-design (RHBD), modern cache-coherent multicore System-On-Chip (SoC) 64-bit microprocessor with unparalleled end-to-end sensor data ingestion and edge processing capabilities facilitated by a built-in 240Gbps enterprise grade TSN Ethernet switch and High Performance Computing (HPC) features. HPSC effectively combines modern computing architectures, advanced SoC design techniques, fault tolerance and recovery mechanisms with a power-aware RHBD approach to achieve over a 100x improvement in performance per watt over NASA's current space-qualified computers.

HPSC is a groundbreaking partnership with Microchip Technology and its industry partners. It's the culmination of Microchip's over 60 years of space flight heritage stretching back to the Apollo program and NASA's experience and expertise in autonomous systems and fault tolerant computing technologies, ushering in a new era of possibilities for space exploration.

HPSC was designed for autonomy and AI at the edge of space and the demanding fault-tolerance needs of Human Exploration and Operations (HEO) missions. This whitepaper contains a high level overview of the features and benefits of HPSC.

white paper

2024

II. HPSC is the future of AI science return

Spacecraft autonomy is the key to exploring dynamic environments in remote locations [1,2,3]. As opportunities to study significant locations and events are farther from Earth, ground-in-the-loop (GITL) based data processing and decision making become inefficient and time consuming, wasting precious spacecraft resources. Furthermore, longer communication light times reduce the ability of mission control, and therefore the spacecraft, to respond to rapidly changing conditions in harsh environments, opportunities for science discovery may be tragically lost or simply overlooked.

The Deep Space Network (DSN), an infrastructure used to communicate with spacecraft, has limited bandwidth to send massive amounts of data that will be generated by future advanced sensors and on-board science instruments and the cost-per-byte to send, receive and process this data has risen. Moreover, any future space-qualified processor must possess features that allow it to recognize and recover from faults, maintain operational readiness and adapt to changing mission timelines and science objectives. Recent space missions have proven to be particularly long lived and it's not uncommon to realize unexpected and profitable science return and results during extended mission phases and in the pursuit of ad-hoc science goals.

"The objective of NASA space exploration missions is quantified in terms of science return, and achievement of those objectives is represented as a set of mission operational functions" [4].

Based on a comprehensive study of past and future spaceflight missions and use-cases, modern computer architectures and science software applications, HPSC was designed with an innovative set of operational features and functions that give scientists and mission designers the ability to achieve unprecedented levels of autonomy and science data processing far from the historied confines of mission control. Figure 1 depicts a high-level block diagram of HPSC and its major sub-systems, where all the sub-systems are attached to a *Network-on-Chip* (NoC) fabric as the primary organizing principle.

These key capabilities are:

- High Performance Computing (HPC)
- Artificial Intelligence At the Edge (AIAE)
- Radiation Hardening and Fault Tolerance
- Ethernet Connected Spacecraft
- Mixed Criticality and Power Efficiency
- Security and Mission Assurance

1. High Performance Computing

HPC is used by scientists to study astronomical phenomenon, analyze large sets of data, perform science simulations and understand the fundamental nature of the universe and our place in it. Typically, space-borne sensors such as telescopes and radars, remote sensing orbiters and rovers conducting orbital and *surface* science send *raw* data back to Earth, which can be analyzed using high performance computers. One such HPC system is the Pleiades supercomputer at NASA Ames' High End Computing Capability (HECC) center.

HPSC brings a similar capability to spacecraft that enables *in situ* science data processing and analysis. Furthermore, by enabling HPSC to use the same software, tools and HPC processing techniques as terrestrial supercomputers, we have coined the notion of *Desktop to Spaceflight Software Development*. Essentially, this means modern science software, built upon existing HPC software libraries and methodologies, is available to NASA scientists supporting the migration of datacenter based HPC tasks to spaceborne HPC tasks.

HPSC is a modern cache-coherent shared memory multicore microprocessor with eight application processing cores implemented using the open standard 64-bit RISC-V Instruction Set Architecture (ISA) [5]. HPSC integrates two SiFive X288 core complexes, with each complex consisting of 4x X280 RISC-V cores. The X280 cores were designed with an advanced feature called a *vector unit*, which are compliant with the RISC-V Vector Extension (RVV) standard. The vector unit features a 512-bit vector register length and variable vector length computations are supported, up to 4096-bits wide.

RISC-V vectors are a powerful and super-efficient extension that features compact code size, high performance capability and the on-die SoC structures consume a limited area compared to Single Instruction Multiple Data (SIMD) architecture approaches favored by other ISAs. Furthermore, RVV can utilize different vector lengths within the same software, enabling scalability, flexibility and future compatibility.

Due to the unique properties of the X280 RISC-V vector unit design and implementation, it boasts a high ratio of FLOPS/watt. Furthermore, the X288 scalar and vector pipelines are decoupled, allowing the vector unit to run behind the scalar pipeline, which enables memory loads to commit early. This improves machine-level parallelism and tolerates latency.

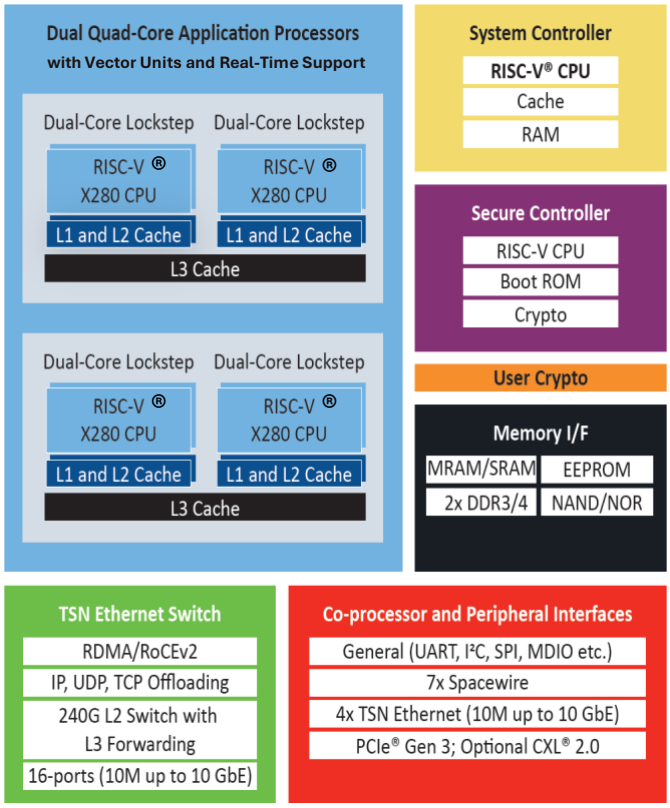

Figure 1 – HPSC Architecture *(courtesy Microchip Technology)*

Additionally, as RVV possesses similar qualities as Digital Signal Processors (DSPs), Graphics Processing Units (GPUs) and Tensor Processing Units (TPUs), it eliminates the need for multiple on-die processing units, their associated power consumption and additional software tool chains resulting in decreased power consumption and reduced software complexity compared to devices with multiple heterogeneous computing elements utilizing different ISAs. Furthermore, HPSC can simultaneously process scalar, vector and AI/ML instruction streams generated from a single optimizing compiler toolchain. Our compiler suite (LLVM) is the same one used on a significant number of terrestrial supercomputers in the Top500 list, which is a ranked measure of a given computers ability to solve a series of linear equations, A x = b, using a dense random matrix A.

To enable this extreme leap in space-borne data processing capability, HPSC features two high-performance DDR4 ports that have a combined bandwidth of 51.2 GBytes/second. This helps to satisfy the high rates of memory communication required to sustain computations by the 8 vector units in the application cores. Additionally, HPSC has two DMA engines that can be used to offload the movement and transposition of data inside the SoC with minimal involvement of the cores. These DMA engines can overlap data transfer while computation is occurring, increasing performance and minimizing latencies.

HPSC has software support for OpenCL, a programming language and open standard for exploiting the differing strengths of multiple connected parallel processors. This standard allows the composition of parallel computing tasks independent of the underlying ISA or architecture. OpenCL in HPSC can target both the on-chip RISC-V application cores and off-chip attached GPUs, accelerators and FPGAs. With the ability to connect disparate, specialized processing elements via dual x8 or bifurcated quad x4 PCIe Gen 3 interfaces and target them with a unified programming language, HPSC can scale out to handle the most demanding and specialized workloads.

Furthermore, HPSC features Compute eXpress Link (CXL) [6], a protocol that rides on top of its PCIe interfaces and allows CXL enabled devices to participate in HPSC's memory coherency mechanism. This eliminates the need for specialized data handling due to the "host to device" memory transfer paradigm.

HPSC's HPC capability via RVV, its high-performance memory subsystems, standards-based software compatibility and custom data movement features allow very efficient processing of large datasets and is well suited to scientific computing workloads utilized by planetary scientists, astronomers and instrument designers.

Finally, utilizing HPSCs massive external I/O capability (discussed below), multiple HPSCs can be combined to realize HPC clusters of rad-hard computing services for use on larger spacecraft, orbiting satellites, gateways, space stations and HABs.

2. AIAE

Artificial Intelligence At the Edge (AIAE) features allow HPSC to use neural network computational techniques to make sense of its environment, plan and execute missions based on changing real-time priorities and objectives and resolve critical science questions.

Combined with HPSC's unprecedented end-to-end system throughput and carefully realized system balance (the ratio of computations to communications), the RVV execution pipeline in each application processing core also enables HPSC to perform high-performance AI/ML tasks. This is further enhanced by specialized extensions to the RVV specification called *SiFive Intelligence Extensions*. These custom vector instructions speed up critical computations necessary for efficient execution of Convolutional Neural Networks (CNNs) with mixed datatypes and complex data reductions.

HPSC can process AI workloads, data flow management, object detection and recognition, recommendation

processing, mission planning and assist flight and mission software with decision making. These AI tasks can execute while simultaneously ingesting sensor data at speed.

Finally, the HPSC software suite uses the same libraries and AI/ML API interfaces (Tensorflow, XNNPACK and OpenXLA) as terrestrial applications. Models can easily be developed and validated on the desktop and in the datacenter and then directly used inside HPSC.

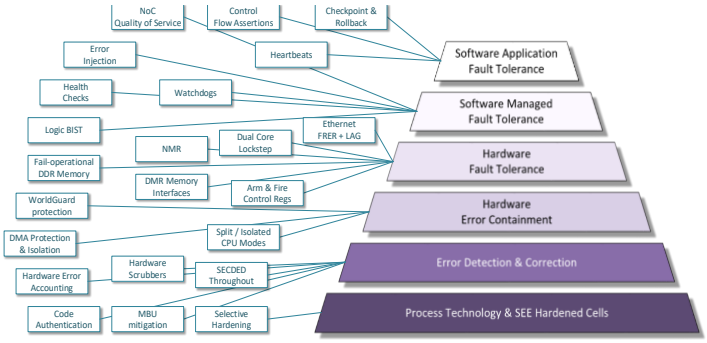

Figure 2 – Fault Tolerance *(courtesy Microchip Technology)*

3. Radiation Hardening and Fault Tolerance

HPSC is developed using an RHBD approach that blends multiple radiation hardened design techniques. This approach allows for high processor clocks rates, while simultaneously protecting HPSC from harsh radiation environments. Logic structures on the device are protected from radiation damage and radiation-induced operational failures using a variety of radiation hardened flip flops and gates. These logic elements are distributed within HPSC according to functional criticality, power, timing, and die area tradeoffs. This ensures high-performance processing, lower power consumption and a balance of radiation protection for deep space missions.

A. Radiation Hardening Design Approach

HPSC is fabricated on GlobalFoundries 12LP+ (12nm Leading-Performance Plus) process node to fabricate HPSC chips. GlobalFoundries is an onshore high-volume, trusted supply chain provider to the Department of Defense (DoD) with low defect rates and excellent quality control. HPSC is available in several packages including a space-qualified organic QML Class-Y package. 12LP+ was chosen to enhance radiation tolerance in the following ways:

Radiation-Hardened Design

Total Ionizing Dose (TID) performance is enhanced by using advanced FINFET materials and structures specifically designed and tested to resist TID radiation damage.

HPSC's radiation-induced Single Event Effects (SEE), Single Event Transient (SET), Single Event Latchup (SEL), Single Event Upset (SEU) and Single Event Functional Interrupt (SEFI) performance is enhanced throughout the SoC by utilizing radiation-aware design techniques in hardware and software. This includes deploying extensive EDAC/ECC functions, memory scrubbing, n-modular redundant logic, multi-level architectural error detection, reporting and recovery functions in hardware and software and deploying radiation hardened (SEE resistant) logic and radiation hardened design rules and practices throughout the design.

Furthermore, all HPSC memories are protected using methods to help detect and correct data corruption, external non-volatile random access memories (NVRAM) have redundant access ports and significant attention was paid to memory error reporting and fault handing capabilities.

B. Fault Tolerance

HPSC has multilayered and comprehensive support for fault tolerance based on several major categories. Broadly outlined in Figure 2, they consist of process technology and SEE hardened cells, error detection and correction, hardware error containment, hardware fault tolerance, software managed fault tolerance and software application level fault tolerance. This paper does not cover every one of these categories, instead we choose several major functional blocks, which provide insight into the comprehensive design of HPSC.

System Controller

HPSC features an always-on processing complex known as the *System Controller* (SysC). In terms of radiation tolerance, SysC is the hardest part of the device. This RISC-V complex, which has its own separate memory, I/O and network-on-chip (NoC) fabric for managing HPSC, is primarily responsible for the health, safety, error reporting and handling, power and clock mode configuration, fault tolerance features and overall management of the SoC. The SysC is also responsible for configuring the SoC, boot operations and externalizing the state of HPSC for consumption by outside facilities. Also, HPSC has emergency modes that allow its recovery by various means in worst-case scenarios.

The SysC is constantly monitoring the state of the SoC and tracks the status of all systems and sub-systems. A historical record of events and their significance can be queried and analyzed via various means. Furthermore, SysC can be configured using built-in fault response

heuristics to handle hundreds of possible failure and near-failure event scenarios and recover from them utilizing various means.

SysC software has multiple Built-In Self-Test (BIST) procedures that are executed during device boot up and can be called periodically or on demand depending on their capability and scope. If errors or anomalies are detected in sub-systems, SysC can disable them, or section them off, as mandated by its configuration and the event's risk profile to maintain the health and safety of device.

HPSC supports the ability to switch between software loads for different mission phases such as Cruise, Entry, Descent and Landing (EDL), Roving, Planning, Science Data Acquisition and Science Data Processing. Transitions to different software loads (mission phase) or configurations can be requested via external pins or a call to the SysC via Flight Software (FSW) after certain fail safe conditions are met.

Finally, the SysC also runs the software for the built in TSN ethernet switch. Because SysC runs the critical functions of the Ethernet switch, HPSC can be put into a configuration that completely disables the application cores and HPSC can be used solely as a TSN Ethernet switch, thereby solving the need for the funding and production of a separate TSN-capable Ethernet switch Application Specific Integrated Circuit (ASIC).

Dual Core Lockstep (DCLS)

The application cores on HPSC can be paired off into a configuration known as DCLS. This configuration increases the fault-tolerance capability of the HPSC by providing built-in redundancy of software execution at the processor core's microarchitectural[1] level.

DCLS or *split lock* replicates portions of an application processor core's execution pipeline and two cores are combined into a *lockstep domain*. When two cores are configured in DCLS, they are referred to as the *main core* and *shadow core*. All outputs of the two cores are compared at each clock cycle of processor operation. The *main core* outputs are the ones selected as the lockstep domain outputs, while the *shadow core* values are discarded after comparison, provided the two outputs match.

This comparison logic provides a direct, externally (to the lockstep domain) visible notification of differences in the outputs. Any difference is flagged immediately and is handled according to the lockstep register settings, which typically involves preventing writes to memory and halting the operation of both the main and shadow cores. At this point, SysC can decide, in conjunction with FSW, how to best handle the fault while ensuring the integrity of SoC operation in the context of the mission.

4. Ethernet Connected Spacecraft

HPSC has a built in 240Gbps 16-Port TSN Ethernet switch with an attached Remote Direct Memory Access (RDMA) feature implemented using the RDMA over Converged Ethernet (RoCE) version 2 standard. This sea change advancement in capability over existing space-qualified microprocessors uniquely positions HPSC as the heart of an Ethernet connected spacecraft.

HPSC's collection of TSN protocols are designed to ensure reliable, deterministic delivery of time-sensitive traffic across Ethernet networks. TSN protocols handle time propagation and synchronization, traffic scheduling, and network redundancy.

HPSC is conformant with and implements a superset of the IEEE P802.1DP TSN for Aerospace profile standard and NASA is participating in the IEEE TSN standards body to influence the IEEE TSN standard for NASA's needs.

HPSC's Ethernet Switch has significant hardware and software capabilities and is built from existing deployed Microchip switch IP, which helps to ensure network interoperability, reliability, compatibility and flexibility. Some of the main features of the switch are:

- 240G non-blocking switch with 16 external ports (up to 10G each) and 2 internal ports
- 80G of Remote Direct Memory Access (RDMA, using RoCEv2 protocol) and general packet offload
- Virtual Local Area Network (VLAN) and Quality of Service (QoS) support
- TCAM-based packet inspection and classification
- Layer 2 and layer 3 forwarding
- Operations, administration, maintenance (OAM) and performance monitoring
- Link Aggregation (LAG)
- Time Distribution with IEEE 1588v2 Precision Time Protocol (PTP) and IEEE 802.1AS
- Highly integrated Time Sensitive Networking (TSN) support

[1] *Microarchitecture* is a term of art that refers to the connected structures in a microprocessor that implement the ISA.

HPSC also features four stand-alone TSN-enabled Ethernet endpoints with similar features to the switch.

Using HPSC's RDMA capability, it's possible to sink 80Gbps of the switch bandwidth into DDR4 memory with zero-touch by the application cores. The applications cores can consume these memory blocks at the appropriate time using the standard *libibverbs* software library. *libibverbs* is part of the OpenFabrics Enterprise Distribution (OFED) is open-source software for RDMA and kernel bypass applications. This feature puts HPSC network and RDMA connectivity on-par with terrestrial HPC systems.

HPSC also has Spacewire router with seven endpoints for connection to Spacewire enabled devices and systems. Furthermore, HPSC has support for deep introspection, trace and debug capabilities, both on the ground and in space.

By utilizing the combination of RDMA and TSN features, HPSC can ingest high-bandwidth sensor data and support the creation of HPSC purpose built or *ad hoc* processing clusters to work on larger science problems. HPSC has hardware and software support to scale its capabilities up to very large space HPC systems.

5. Mixed Criticality and Power Efficiency

HPSC has several operating modes and capabilities that enable mixed criticality processing, hardware separation to reduce interference, an advanced notion of time and hardware support for ARINC-653 time and space partitioning.

A. Mixed Criticality

HPSC has multiple major modes that split up the architecture in various ways using its dual application core cluster topology and RISC-V WorldGuard. WorldGuard is a fine-grain security model for isolated code execution and data protection. WorldGuard offers SoC-level information control with advanced isolation control, based on multiple levels of privilege per world, and 32 worlds. WorldGuard offers multi-domain security and data protection for core, cache, interconnect, peripheral, and memory. WorldGuard works in conjunction with RISC-V Physical Memory Protection (PMP) integrated into HPSC. In the HPC implementation, WorldGuard has low level support for limiting the access to/from and improving the QoS for sub-systems attached to the NoC. These different modes, many of which utilize DCLS (explained above), allow HPSC to attach NoC endpoints (e.g. DDR4, PCIe, DMA, etc.) to each *split* for exclusive use. By partitioning HPSC into different portions with hardware level separation, the various splits have reduced or no practical interference channels with one another.

Each type of split mode has different capabilities. In core-isolated split mode for example, application cores can turn on or off features to enable modes that help with deterministic latency of real-time processing. Figure 3 represents the split modes possible for configuring HPSC.

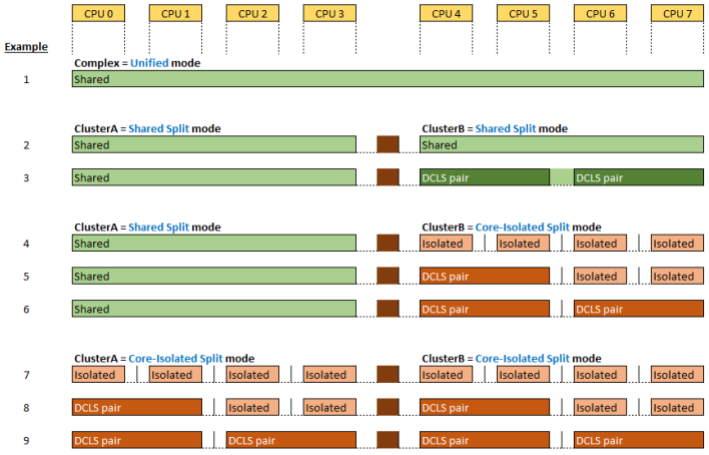

Figure 3 – HPSC Split Modes *(courtesy Microchip Technology)*

HPSC split modes allow the execution of multiple, logically isolated operating systems, including multiple copies of the same operating system on different sides of a split. For example, reference mode 6 in Figure 3, Cluster A can run SMP Linux RT, which is ingesting and organizing data via RDMA and performing AI tasks, while the Cluster B has two identical copies of Windriver VxWorks 7 running time critical FSW tasks, each on its own DCLS pair. These redundant copies of VxWorks can be used as high availability instances with FSW managing failover during a mission phase such as EDL. Or alternately, can be running different vehicle control software for less critical but more complex operations such as simultaneous sensing and roving. In either case, using a WorldGuard reserved memory window between Clusters A and B, FSW on Cluster B can ask for data and results from software daemons handling data ingestion and AI tasks on Cluster A.

HPSC will support several operating systems at launch including Linux with PREMPT_RT, VxWorks 7, RTEMS, bare metal execution using a low level SDK and the Xen Hypervisor. The depth and breadth of this software offering (including CSPs and BSP support for the evaluation board design of all operating systems) is one of the hallmark features of HPSC and underscores NASA and Microchip's commitment to widespread support of many scenarios, use-cases and mission profiles.

B. Time and ARINC-653

HPSC has an advanced notion of time with several onboard time keeping systems for various purposes, including redundant gPTP timers inside the Ethernet switch. Every HPSC timer has sub-nanosecond resolution with a range of bit widths and capabilities, some of which can be automatically synchronized amongst themselves, network time and external time sources. Furthermore, the timer subsystem is connected throughout the device and provides the basis for other important timing systems such as the ARINC-653 [7] scheduler block.

One example is the high precision timer (HP_TIMER). It's a monotonic timer that is available with uncontended access locally to each core using a native RISC-V instruction and is guaranteed to be the same value on every core at the same instant, even across split modes. The value of HP_TIMER in the core is warm reset invariant, so even after a core is warm reset and restarted, the current HP_TIMER value will be reloaded from the main HP_TIMER housed in the SysC.

The ARINC-653 scheduler block allows hardware timing of major, minor and sub-minor frames based on HP_TIMER ticks and can be used as the basis of ARINC software support for time partitioning. Furthermore, the ARINC scheduler block is directly connected to the DMA engine. This connection and associated logic can track many transfer and byte limit schedules while halting errant DMA transactions that occur outside or run over their assigned major, minor and sub-minor timeframes.

These advanced features, usable across disparate operating systems and application software allow independent execution of mixed criticality workloads, minimize time and space interference, coordinate precise timing and execution of FSW tasks and support complex mission and science processing requirements during critical mission phases.

C. Power Efficiency

HPSC can adjust its power usage based on the unique requirements of the mission at any given point in the timeline. This is accomplished through several key technologies: Power Islands, Clock Control, Clock Gating and Power Dial.

HPSC has hardware support and software control over 70+ power islands that can be turned on or off depending on how the SoC will be used and for what purpose. Additionally, a particular mission profile can adjust the clock speed of the processor before boot time to reduce power consumption. Entire portions of HPSC can be put into low power modes that use very little quiescent current.

The application core power modes include normal run mode with the Power Dial option and multiple idle modes with increasing depths of sleep, including a full power down mode. Power Dial is a method of scaling down power consumption in a processor core by restricting cycles allowed to advance instructions in the processor execution pipeline. Wait-for-interrupt (WFI) wake up mode puts the processor into a state that waits for an interrupt to resume processing. When the processor is in the WFI state, it consumes less power.

WFI Tile Clock Gate mode detects when operations have finished on a core and initiates a coarse clock gating mode to save even more power. This feature can also be applied to any of the eight cores in the device and alternately power gated idle mode puts the entire core complex into the deepest sleep possible, while the rest of the SoC remains alive.

This multiplicity of power control features allows HPSC to scale up and down the mission power curve and be used in the smallest of helicopters, small sats and robots to the largest of HEO missions.

6. Security

HPSC has a comprehensive security solution with comprehensive defense-in-depth, tamper-resistant, future-proof security features. This solution revolves around several key elements:

Security Coprocessor

Also known as the User Crypto Accelerator, this is a separate peripheral block attached to the NoC that is responsible for encryption and decryption tasks. It's a side-channel resistant post-quantum enabled cryptography accelerator for use by the application cores. It appears as a native device in the operating systems. All currently used cryptography algorithms by the DoD are present.

Secure Controller

Integrated secure enclave for secure boot, encryption, key and secret handling and tamper detection. This controller is a RISC-V processor with special purview of the security implementation.

Supply Chain Assurance

HPSC is manufactured using secure manufacturing techniques, which include the use of Hardware Security Modules (HSMs) during fabrication process to ensure authenticity of wafer, die and packaged parts. Every HPSC device has a factory inserted, unique X.509 certificates traceable to a trusted certificate authority (CA).

II. Reducing the cost per byte of science return

NASA's HPSC project is working with industry partners to develop and support the HPSC ecosystem including participation in the new Sensor Open Systems Architecture (SOSA) for Space working group, which aims to standardize the interfaces and specifications for avionics board and chassis level electronics. Based on previous safety/mission critical efforts by the Telecom industry and DoD, standards crafted by this group will be a focal point for industry innovation and commoditization of space flight electronics.

HPSC gives scientists the ability to process data and conduct autonomous missions in deep space. HPSC can make decisions at the point of observation to best serve the mission and its goals.

HPSC is energy efficient, conserving precious resources and extending mission life. HPSC has significant fault tolerant handling and radiation mitigation capabilities ensuring it will operate in the harsh, dynamic environments of deep space.

HPSC's advanced architecture supports multiple configurations and mission phases for radically different science and spaceflight computing priorities.

HPSC's robust implementation and long life-cycle survival characteristics allows the extension of expensive spaceflight missions to support exigent and directed follow-on science goals.

Finally, HPSC reduces the cost-per-byte of science return by delivering a comprehensive software, hardware and reference design solution.

III. Conclusion

HPSC was built with AI, future spacecraft connectivity, HPC, resilience, efficiency and security as the cornerstones of its design and development.

HPSC represents a massive leap forward in space-qualified computing technology and will markedly improve the quality and quantity of science data return.

HPSC's unique high reliability and industry-standard architecture, takes a forward looking approach to the needs of future science missions. HPSC makes more advanced missions possible with lower cost, shorter time-to-value, higher reliability and with a greater regard to answering important and far ranging questions about the nature of the universe.

References

[1] N. A. of Sciences, D. on Engineering, P. Sciences, S. S. Board, and C. on A. S. S. for the S. for L. in the Universe, *An astrobiology strategy for the search for life in the universe*. National Academies Press, 2019.
[2] E. National Academies of Sciences, "Origins, worlds, and life: a decadal strategy for planetary science and astrobiology 2023-2032," 2022.
[3] N.R. Council, D. on Engineering, P. Sciences, S. S. Board, and C. on the P. S. D. Survey, *Vision and voyages for planetary science in the decade 2013-2022*. National Academies press, 2012.
[4] A. Howard and G. Rodriguez, "Validating mission relevance of autonomy technologies through increased science return," in *20th International Conference on Machine Learning*, Washington DC USA, Aug. 2003.
[5] A. Waterman, K. Asanovic, and C. Division, "The RISC-V Instruction Set Manual Volume I: unprivileged ISA." 2019.
[6] CXL – Compute Express Link, https://computeexpresslink.org/resources/
[7] SAE Industry Technologies Consortia, Avionics Application Software Standard Interface ARINC-653, https://www.sae.org/standards/content/arinc653p0-1/