International Astronautical Congress, Sydney, Australia, 29-September - 03 October 2025.

IAC-25,D1,IP,4,x95717

Feasibility of a COTS GPU-Based On-board AI Platform for Robotic Space Missions

Marco Sewtz[a]*, **Roberto Lampariello**[a]

[a] *German Aerospace Center (DLR), Institute for Robotics and Mechatronics, Wessling, Germany*
E-mail: Marco.Sewtz@dlr.de

* Corresponding author

Abstract

The feasibility of deploying a high-performance COTS GPU platform for robotic space missions is assessed, focusing on NVIDIA Jetson modules integrated into the ScOSA architecture. The design couples the Jetson with a radiation-hardened Co-CPU for active fault detection, error correction, and controlled shutdown during high-radiation events, while exploiting the GPU's AI capabilities for real-time perception and manipulation tasks. Radiation simulations for a representative LEO demonstrate that aluminum shielding reduces the TID to well below the Jetson's tested tolerance. SEEs remain a concern and are addressed through system-level mitigation. Thermal modeling shows that a PCM buffer significantly extends operational run-time compared to passive heatsinking alone, enabling longer AI inference periods without active cooling.

1. Introduction

Modern and future space missions increasingly require advanced on-board data processing for tasks such as on-orbit servicing and assembly, satellite rendezvous, and robotic arm operations, as illustrated in Figure 1. Similar processing demands also arise in prospective rover-type missions. These scenarios require real-time analysis of complex sensor data and the execution of Artificial Intelligence (AI) models directly in space.

This work investigates the feasibility of a commercial off-the-shelf (COTS) Graphics Processing Unit (GPU)-based on-board AI platform for robotic space applications. The proposed platform integrates into the Scalable On-board Computing for Space Avionics (Scalable On-board Computing for Space Avionics (ScOSA)) framework to realize reliable On-Board Computer (OBC) platforms that blend radiation-tolerant nodes with high-performance COTS components [1, 2]. While alternatives such as Central Processing Unit (CPU)- or Field-Programmable Gate Array (FPGA)-based solutions are viable, they often entail limited parallelism, specialized development flows, and longer deployment times. In contrast, NVIDIA Jetson modules provide state-of-the-art AI processing capabilities with a widely adopted toolchain. However, these modules are not inherently designed for the space environment, which necessitates additional measures for reliability under radiation and thermal constraints [3–6].

The design addresses two main challenges: radiation tolerance and thermal management. For radiation, the approach combines passive shielding with active fault detection and recovery using a radiation-hardened Co-CPU. This Co-CPU detects and corrects faults and can temporarily shut down the Jetson during passages through high-radiation regions such as the South Atlantic Anomaly (SAA). For thermal control, the concept considers a Phase Change Material (PCM) buffer to extend operational run time without resorting to complex active cooling.

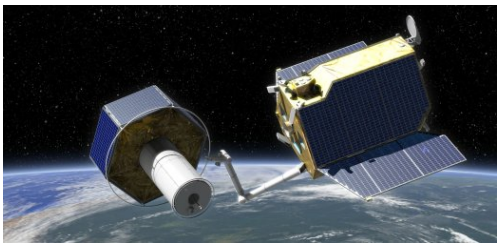

Fig. 1: On-Orbit Servicing mission of a target satellite.

The software architecture and AI pipeline streamline the development and deployment of neural networks by using the Open Neural Network Exchange (Open Neural Network Exchange (ONNX)) format and NVIDIA TensorRT for inference optimization [7, 8]. This reduces qualification overhead and accelerates time-to-orbit. The platform will be validated through two missions: a technology demonstration cubesat in 2027 and an on-orbit servicing mission in 2028.

2. Background

The growing demand for high-performance on-board data processing in space missions such as on-orbit servicing, active debris removal, and autonomous planetary exploration is pushing the limits of traditional radiation-hardened (RH) processors [1]. Established RH architectures, such as the LEON family or the BAE RAD750, offer proven reliability in harsh environments, but their computational performance, energy efficiency, and AI capability lag far behind modern COTS devices. Current RH processors often deliver performance comparable to early-2000s commercial CPUs, whereas platforms like the NVIDIA Jetson family integrate multi-core ARM CPUs, advanced GPUs, and dedicated AI accelerators on a single module.

The use of embedded GPUs in space has already been demonstrated, for example on NASA's Mars Helicopter [4], and European projects such as GPU4S [9, 10] have identified them as suitable candidates for high-performance payload processing, provided their radiation response is well understood. Recent studies have performed both Total Ionizing Dose (TID) and Single Event Effect (SEE) testing on multiple Jetson variants, including Jetson Nano [3], Xavier NX [4], Orin NX [6], and Orin AGX [5]. Reported TID tolerances range from approximately 16 krad(Si) for early Nano boards up to more than 36 krad(Si) for Orin NX, with Error-Correcting Code (ECC)-protected industrial models showing improved resilience against Single Event Upset (SEU)s and Single Event Functional Interrupt (SEFI)s.

However, COTS modules are not inherently designed to survive the harsh radiation and thermal conditions of space. Radiation-induced failures may include SEUs, SEFIs, and destructive latch-ups, which must be mitigated through passive shielding, active fault detection, and system-level redundancy [5, 11]. Similarly, thermal management is critical, as high-load AI workloads can drive power consumption beyond 15–50 W depending on the Jetson model, requiring conduction paths, thermal transfer plates, or PCM buffers to maintain operational temperatures [12].

A combined GPU + Co-CPU architecture addresses these challenges by pairing high-performance computation with robust fault tolerance. The GPU excels at massively parallel workloads such as image processing, object detection, and deep neural network inference—tasks essential for real-time perception and manipulation in robotic missions. The Co-CPU, implemented with RH technology, continuously monitors the GPU's health via watchdogs and heartbeat signals, isolates faults, and can reset or reflash the GPU's storage from a protected Programmable Read-Only Memory (PROM) in case of data corruption caused by SEEs. This division of responsibilities allows the GPU to deliver cutting-edge AI performance while the Co-CPU ensures mission-critical reliability.

A promising system-level realization of this approach is the ScOSA architecture [1, 2], which combines RH Reliable Computing Node (RCN)s with High-Performance Node (HPN)s in a distributed, reconfigurable framework. The middleware manages fault detection, task migration, and workload distribution transparently, enabling developers to port terrestrial applications with minimal modification. This heterogeneous architecture gives mission designers the flexibility to balance performance, reliability, and cost, and to dynamically adapt computational resources to different mission phases.

In summary, recent research and mission heritage indicate that with appropriate shielding, active mitigation, and thermal design, COTS GPU-based System-on-Module (SoM)s embedded into hybrid architectures like ScOSA can be viable for short- to medium-duration Low Earth Orbit (LEO) missions. This enables advanced AI capabilities in space while significantly reducing time-to-orbit compared to custom RH designs.

3. Mission

Two space missions will serve to progressively validate the AI platform, starting with a technology demonstration and culminating in operational use. Both flights are designed to assess performance, reliability, and fault tolerance in representative LEO conditions, providing critical data for future deployments.

The first mission, planned for launch in 2027, is a rideshare payload on-board a technology demonstration cubesat dedicated to testing and qualifying new space technologies. The COTS GPU-based platform will be integrated as an OBC experiment to evaluate computational performance, reliability, and resilience in orbit. During this mission, another on-board experiment will release a small object into space. Although this release experiment is entirely independent of our work, it creates an opportunity to test our perception pipeline in a realistic orbital scenario. By using the experiment's camera feed, the platform will perform real-time detection, tracking, and pose estimation of the released object, thereby validating AI algorithms under true space conditions. In addition to functional performance, the mission will capture detailed telemetry on radiation-induced faults and thermal behavior. While the exact orbital parameters are not yet fixed, the mission is expected to operate for one to two years. Data from this flight will directly inform hardware and software refinements for operational deployment.

The second mission, scheduled for 2028, is an on-orbit servicing flight where the GPU hardware will serve as the execution platform for real-time perception and AI-based manipulation. In this operational scenario, the system will process high-rate vision data to detect mission-specific objects and estimate their pose relative to a manipulator arm, ensuring precise grasping without collision risk. Mission-critical reliability is essential, as perception errors could lead to damage to the target satellite or the servicing spacecraft. Lessons from the first mission will be applied to refine shielding strategies, optimize active mitigation via the Co-CPU, and implement operational procedures such as suspending GPU activity during SAA passages. This mission is also planned for a LEO orbit with up to one year of operational use, and will serve as the first end-to-end demonstration of the platform in a time-critical robotic servicing task.

4. Hardware

The hardware architecture is designed to integrate a high-performance COTS computing platform into a reliable, space-qualified system. It builds on the ScOSA framework [1, 2], which addresses the fundamental trade-off between the computational power of modern electronics and the proven robustness of traditional space-qualified hardware.

In the ScOSA concept, highly reliable RH RCNs are combined with high-performance COTS HPNs in a single distributed system. This architecture is fault-tolerant and reconfigurable, allowing faulty nodes to be masked while maintaining mission reliability. Its distributed nature also enables scalability and flexible resource allocation across mission phases.

At the core of the system is a middleware layer that abstracts this heterogeneous architecture for application developers, presenting a unified execution environment. It provides fault-tolerance features such as a voter service for Triple Modular Redundancy (TMR) [13] and a checkpointing service for restoring state after a task or node failure. The network supports heterogeneous interfaces, including Ethernet and SpaceWire, and can run different operating systems such as Real-Time Executive for Multiprocessor Systems (RTEMS) and Linux.

The implementation used in this work follows a two-module design:

1. **Rad-Hard module** – based on space-qualified components, including a Co-CPU and a PROM. The Co-CPU, currently a GR740, monitors the GPU node via heartbeat signals and a watchdog. If the GPU fails to respond, the Co-CPU can reset it or, if necessary, re-flash the Non-Volatile Memory express (NVMe) storage from the protected PROM to correct radiation-induced data corruption [4, 5]. It also handles all communications with other spacecraft subsystems.

2. **COTS module** – based on an NVIDIA Jetson SoM [12], mounted on a custom, rad-hard carrier board to interface with the Co-CPU and the spacecraft bus. The Jetson Orin AGX Industrial variant is under consideration due to its ECC memory, which increases resilience against SEUs in a radiation environment [5, 6]. However, it has higher mass, volume, and thermal dissipation requirements. If integration constraints on size or power budget cannot be met, the smaller Jetson Orin NX 8 GB will be used as a fallback, offering lower thermal load but without industrial-grade ECC protection.

Thermal management is a key design factor. For example, the Jetson Orin NX 8 GB has a worst-case total power draw of 15 W [12] when using the power-saving mode, and components must remain below 80 °C. Aluminum structures may be used both as mechanical mounting and as part of the thermal path. Shielding concepts are being evaluated to provide both radiation protection and heat spreading. The final thermal solution—potentially integrating a PCM buffer as discussed in Section 6.2—will be defined based on detailed thermal simulations.

The current hardware configuration is designed for a single-node evaluation in orbit, but the architecture remains compatible with multi-node or TMR setups for future missions. This aligns with the ScOSA philosophy of scaling processing capabilities by distributing workloads across heterogeneous computing resources.

5. Software

The software architecture integrates RH and COTS components within the ScOSA framework to provide a modular, fault-tolerant execution platform. An abstraction layer based on Core Flight Software (cFS) [14] or DLR Outpost [15] decouples flight applications from hardware and operating system (OS) details, enabling reuse and simplifying qualification. The middleware masks node failures, supports task migration, and exposes a unified interface so that applications can be deployed without knowing which node executes them.

The middleware and flight-software framework provide: (i) process isolation and health monitoring for HPN tasks, (ii) command/telemetry, file and parameter services, (iii) time synchronization and checkpoint/restore, and (iv) op-

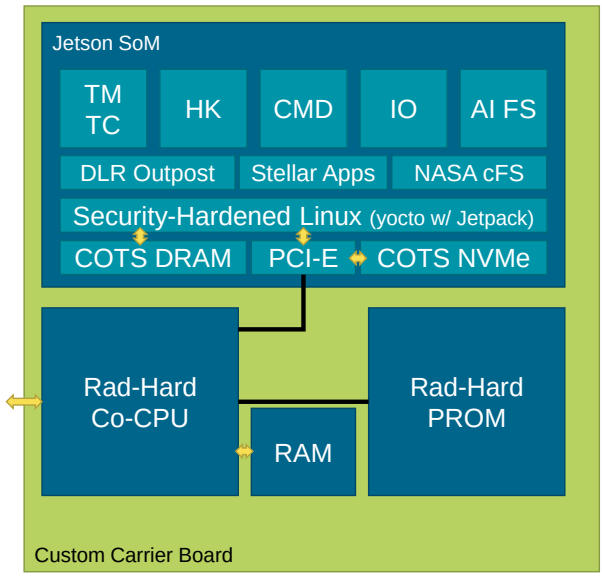

Fig. 2: Proposed stack for the whole architecture: RH control via Co-CPU and PROM, supervised HPN on Jetson SoM with security-hardened Linux, and services.

tional voter services for TMR. Figure 2 illustrates how these services supervise AI workloads on the GPU node while the Co-CPU enforces deterministic recovery via watchdogs and heartbeats.

The learning and deployment flow is split into ground optimization and on-board execution (cf. Figure 3):

1. **Ground development and optimization.** Networks are designed and trained using standard frameworks (e.g., TensorFlow, PyTorch) [16, 17]. Trained models are exported to ONNX [7] and optimized for the target SoM with NVIDIA TensorRT [8] to reduce latency and memory footprint within mission accuracy limits.

2. **On-board deployment and execution.** The optimized model becomes a configuration item, is versioned, and deployed to the HPN. Inference runs inside a monitored sandbox; the middleware supervises liveness and performance and can pause or terminate the task on policy triggers (e.g., SAA passage), while the Co-CPU executes resets or fallback if required.

For transient faults (SEU/SEFI), the system escalates recovery: restart the affected AI task; if unsuccessful, reset the HPN process/container; if still unsuccessful, perform node reset. For suspected data corruption, the Co-CPU reimages the NVMe from PROM. Operational procedures include pre-planned suspension of GPU workloads during high-radiation segments and automatic resume afterwards.

Perception pipelines that feed guidance and manipulation are bounded by cycle time. The runtime pins critical threads, constrains memory growth, and uses pre-allocated buffers. Interface nodes (camera/IO) can run on RCNs when timing requires it, forwarding frames to the HPN with back-pressure and drop policies defined in the middleware.

The initial flights integrate a single GPU node to reduce complexity. The same interfaces allow later upgrades to multi-node or TMR deployments without refactoring applications: tasks can be replicated across HPNs, voted on the RCN, or migrated based on thermal or radiation constraints.

6. Evaluation

The Evaluation section examines the key challenges of using a COTS GPU platform for robotic space missions, focusing on the radiation environment and thermal management. The analysis assesses both passive and active mitigation strategies to ensure the platform's reliability and resilience in a LEO environment. A detailed thermal analysis is also presented, exploring solutions like a PCM buffer to extend the operational period of the COTS module.

6.1 Radiation Environment Analysis for LEO Mission

The radiation environment for the proposed mission was simulated using the European Space Agency's Space Environment Information System (SPENVIS) [18], utilizing the AP-8 MIN (AP8MIN) model for trapped protons [19] and the ESA-SEE1 model for trapped electrons [20]. The simulation was configured for a 365-day mission in a 500 km circular orbit with a 40° inclination. This orbit is characteristic of LEO and includes passes through the SAA, where the majority of the trapped particle radiation is concentrated. The following is an analysis of the unshielded, orbit-averaged particle fluxes.

Unshielded Particle Flux Characterization

The SPENVIS output provides the omnidirectional integral and differential fluxes for both protons and electrons as a function of energy. The raw data indicates a significant flux of low-energy particles. For instance, the orbit-averaged integral flux for protons with energies greater than 1 MeV is approximately $1.21{\times}10^2$ protons/cm^2/s. For electrons, the integral flux for energies greater than 0.04 MeV is even higher, at approximately 7.61×10^4 electrons/cm^2/s. These unshielded values represent the maximum possible exposure and are dominated by the proton and electron pop-

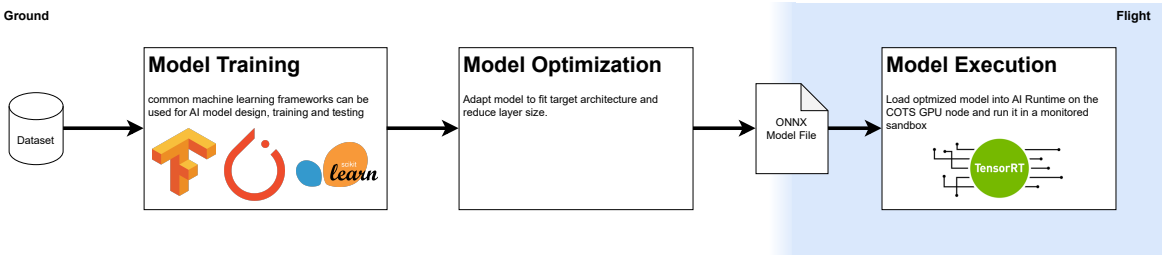

Fig. 3: AI pipeline workflow. Ground training and optimization (export to ONNX, optimization with TensorRT) and on-board execution on the COTS GPU node in a monitored sandbox.

ulations trapped in the Van Allen belts. This high flux of low-energy particles is the primary contributor to the TID that would be accumulated by an unshielded component.

Evaluation of Shielding and Particle Flux Attenuation

To mitigate the radiation effects on the NVIDIA Jetson platform, a passive shielding strategy was simulated using a 10 mm aluminum block and the MFLUX (MFLUX) tool within SPENVIS [18]. The MFLUX simulation calculated the orbit-averaged particle fluxes at the center of a solid sphere of aluminum. This new data confirms the shielding's effectiveness by showing a drastic reduction in the flux of low-energy particles.

- **Proton Flux:** The unshielded orbit-averaged integral proton flux at 1 MeV was found to be approximately 1.21×10^2 protons/cm^2/s. After passing through the 10 mm aluminum shield, the simulated flux at the same energy is reduced to 2.92×10^1 protons/cm^2/s. This demonstrates a significant attenuation of low-energy protons.

- **Electron Flux:** The effect is even more pronounced for electrons. The unshielded integral flux at 0.04 MeV was approximately 7.61×10^4 electrons/cm^2/s. The shielded flux at this energy is reduced to a much lower value of 1.24×10^0 electrons/cm^2/s, highlighting the shield's high efficiency in stopping low-energy electrons, which are the main cause of TID.

- **Neutron Flux:** The simulation also identified a new radiation component: secondary neutrons produced by the interaction of high-energy primary particles with the shielding material. The shielded neutron integral flux at 0.04 MeV is approximately 2.23×10^0 neutrons/cm^2/s. While this is a lower flux compared to the primary particles, it is a new consideration for radiation effects.

The absorbed TID was quantitatively evaluated using the SHIELDOSE-2 (SHIELDOSE2) tool [21]. The simulation showed that for a shielding thickness of 10 mm of aluminum, the mission-averaged TID is 184.29 rad(Si). This value is well below the published TID tolerance thresholds for the NVIDIA Jetson platform [5, 6], which range from 19 to 36 krad(Si). This confirms that the passive shielding strategy is highly effective in protecting the Jetson from TID effects over the course of the one-year mission.

Single Event Effect and Mitigation Strategy

Despite the strong attenuation of low-energy particles, the MFLUX simulation shows that a non-negligible flux of high-energy protons and electrons persists after passing through the shield. The presence of these high-energy particles, along with the newly identified secondary neutrons, means that the risk of SEEs remains. SEEs, which can cause non-destructive bit flips or destructive latch-up events [22, 23], are not fully mitigated by passive shielding. Therefore, passive shielding alone is insufficient to guarantee mission success. The critical components of the mission's radiation hardness assurance are the active mitigation strategies, including the use of a rad-hard Co-CPU [24] for error detection and correction and the procedure to temporarily shut down the Jetson during passes through the most intense regions of the SAA [23]. This combination of passive shielding and active mitigation is necessary to ensure the reliable operation of the COTS Jetson platform in the LEO radiation environment.

6.2 Thermal Management with PCM Buffer

The thermal management of the COTS module is a critical aspect of ensuring its reliable operation in a space environment. As a baseline, a simple aluminum block with dimensions $90 \times 90 \times 10\,\mathrm{mm}$ was considered as a passive heatsink. With a heat dissipation of 15 W, this aluminum block alone provides a sensible heat buffer, al-

lowing the system to operate for approximately 13 minutes before reaching its maximum operational temperature of 80 °C from a starting temperature of 20 °C. The time (t) for this temperature change can be calculated using the following equation, which considers the mass of the aluminum (m_{Al}), its specific heat capacity ($c_{p,Al}$), the temperature change (ΔT), and the power input (P):

$$t = \frac{m_{Al} c_{p,Al} \Delta T}{P} \tag{1}$$

While this provides some thermal mass, the temperature rises continuously, limiting the stable operational period. To significantly extend this window, a more effective thermal storage solution is required. This led to the investigation of a Phase Change Material (PCM) buffer integrated within a hollow aluminum block.

To achieve a prolonged thermal buffer, a PCM with a high latent heat of fusion is required. Hard Fischer–Tropsch wax was selected for this purpose, with a melting temperature of 77 °C-80 °C and a latent heat of fusion (L_f) of approximately 240J/g [25]. This material effectively absorbs a large amount of energy during its phase change from solid to liquid, providing a temperature plateau and a stable operational state. The total heat energy absorbed (Q) by the system over a given time (t) is composed of the sensible heat absorbed by the aluminum block and the wax, as well as the latent heat absorbed by the wax:

$$Q = P \cdot t = m_{Al} c_{p,Al} \Delta T + m_{wax} c_{p,wax} \Delta T + m_{wax} L_f \tag{2}$$

By integrating the wax into a hollow aluminum block with a 2 mm wall thickness, the operational time to reach 80 °C is extended to 22.6 minutes, with a critical 10.6-minute period where the system temperature remains stable at 77 °C as the wax melts.

The time to reach the target temperature can be calculated by considering the total heat absorbed and the power input. For the solid aluminum block, the time (t_{Al}) is determined by the sensible heat absorption of the aluminum:

$$t_{Al} = \frac{m_{Al} c_{p,Al} \Delta T}{P} \tag{3}$$

For the hollow block with the PCM, the total time (t_{PCM}) is the sum of two phases. The first is the time to heat the system to the wax's melting point ($t_{sensible}$), and the second is the time it takes for the wax to fully melt (t_{latent}):

$$\begin{aligned} t_{PCM} &= t_{sensible} + t_{latent} \\ &= \frac{m_{Al} c_{p,Al} \Delta T_1 + m_{wax} c_{p,wax} \Delta T_1}{P} + \frac{m_{wax} L_f}{P} \end{aligned} \tag{4}$$

where ΔT_1 is the temperature change from the initial state to the wax melting point.

To meet the mission requirements for extended operational periods, the dimensions of the wax volume were optimized. To achieve a thermal buffer capable of supporting the 15 W heat load for a total of 30 minutes, the wax volume would need a height of approximately 10.6 mm. For a longer operational period of 45 minutes, a more substantial wax volume with a height of 15.9 mm would be necessary. These calculations demonstrate that a compact and weight-efficient PCM solution can provide a robust thermal buffer, significantly extending the mission's operational capabilities without relying on complex active cooling systems.

7. Conclusion

The presented work demonstrates the feasibility of integrating a high-performance COTS GPU platform, specifically NVIDIA Jetson modules, into a fault-tolerant, space-qualified architecture for robotic space missions. By embedding the Jetson into the ScOSA framework alongside a radiation-hardened Co-CPU, the design balances computational capability with mission reliability.

Radiation environment simulations for a representative LEO orbit confirm that a 10 mm aluminum shield can effectively reduce TID to well below the Jetson's tested tolerance limits. Nevertheless, SEEs remain a concern, requiring active mitigation through system-level fault detection, error correction, and controlled shutdown during high-radiation passes such as those through the SAA.

Thermal analyses indicate that while passive aluminum heatsinking offers limited operational windows, integrating a PCM buffer can extend stable run-times significantly without adding complex active cooling. This supports longer AI inference periods and more flexible mission operations.

The two upcoming missions—a technology demonstration cubesat in 2027 and an operational on-orbit servicing mission in 2028—will provide progressive in-orbit validation of the architecture, its mitigation strategies, and its AI processing pipeline. The approach offers a reusable, rapid-deployment platform for spaceborne AI applications, reducing qualification overhead while enabling advanced perception and autonomy capabilities for future robotic missions.

Appendix

The calculations presented here use the following constants:

- Power input (P): 15 W
- Initial temperature (T_i): 20°C
- Target temperature (T_f): 80°C
- Temperature change (ΔT): $60K$
- Temperature change to wax melting point (ΔT_1): $57K$
- Density of Aluminum (ρ_{Al}): 2700kg/m^3
- Specific heat of Aluminum ($c_{p,Al}$): 900J/kg $\cdot$ K
- Density of Hard Fischer–Tropsch wax (ρ_{wax}): 900kg/m^3 (0.9g/cm^3)
- Specific heat of Hard Fischer–Tropsch wax ($c_{p,wax}$): 2500J/kg $\cdot$ K
- Latent heat of fusion of Hard Fischer–Tropsch wax (L_f): 240J/g (240, 000J/kg)

References

[1] Kurt Schwenk, Moritz Ulmer, and Ting Peng. Scosa: application development for a high-performance space qualified onboard computing platform. In *Proc. SPIE 10792, High-Performance Computing in Geoscience and Remote Sensing VIII*, 2018.

[2] Moritz Ulmer, Kurt Schwenk, Ting Peng, and Michael Gerndt. Scosa: A scalable on-board computing platform for space avionics. In *2020 IEEE Aerospace Conference*, 2020.

[3] Windy S. Slater, Nayana P. Tiwari, Tyler M. Lovelly, and Jesse K. Mee. Total ionizing dose radiation testing of nvidia jetson nano gpus. In *2020 IEEE Radiation Effects Data Workshop (REDW)*, 2020.

[4] Ivan Rodriguez-Ferrandez, Maris Tali, Leonidas Kosmidis, Marta Rovituso, and David Steenari. Sources of single event effects in the nvidia xavier soc family under proton irradiation. In *2022 IEEE International Symposium on On-Line Testing and Robust System Design (IOLTS)*, 2022.

[5] Windy S. Slater, Benjamin B.W. Rutherford, Jesse K. Mee, Ryan E. Pinson, Matthew Gruber, Daniel Sabogal, and Ian A. Troxel. Single event effects and total ionizing dose radiation testing of nvidia jetson orin agx system on module. In *2023 IEEE Radiation Effects Data Workshop (REDW) in conjunction with NSREC*, 2023.

[6] Michael A. Félix, Windy S. Slater, Derrek C. Landauer, Ryan E. Pinson, and Benjamin B.W. Rutherford. Total ionizing dose radiation testing of nvidia jetson orin nx system on module. In *2024 IEEE Space Computing Conference (SCC)*, 2024.

[7] Junjie Bai, Fang Lu, Ke Zhang, et al. Onnx: Open neural network exchange. https://github.com/onnx/onnx, 2019. Zugriff am 08. August 2025.

[8] NVIDIA Corporation. NVIDIA TensorRT. https://developer.nvidia.com/tensorrt, 2025. Accessed: 2025-08-08.

[9] Iván Rodríguez, Leonidas Kosmidis, Jérôme Lachaize, Olivier Notebaert, and David Steenari. Gpu4s bench: Design and implementation of an open gpu benchmarking suite for space on-board processing. Technical report, Universitat Politècnica de Catalunya, 2019. Technical Report UPC-DAC-RR-CAP-2019-1.

[10] Leonidas Kosmidis, Iván Rodríguez, Álvaro Jover, Sergi Alcaide, Jérôme Lachaize, Jaume Abella, Olivier Notebaert, Francisco J. Cazorla, and David Steenari. Gpu4s: Embedded gpus in space – latest project updates. *arXiv preprint arXiv:2109.11074*, 2021. Accessed via arXiv.

[11] European Cooperation for Space Standardization (ECSS). Ecss-e-st-10-04c rev.1: Space environment, 2020. Space environment standard.

[12] NVIDIA Corporation. *Jetson Orin NX Series and Jetson Orin Nano Series Modules Thermal Design Guide, Version 1.5*, May 2025. Document TDG-11127-001_v1.5.

[13] Israel Koren and C. Mani Krishna. *Fault-Tolerant Systems*. Morgan Kaufmann, Burlington, MA, 2 edition, 2010.

[14] NASA's Goddard Space Flight Center. core Flight System (cFS). Software, 2025. Abgerufen von https://cfs.gsfc.nasa.gov/.

[15] German Aerospace Center (DLR), Institute for Space Systems. Outpost core, 2025. Accessed: 2025-08-08.

[16] Martín Abadi et al. Tensorflow: Large-scale machine learning on heterogeneous systems, 2015.

[17] Adam Paszke et al. Pytorch: An imperative style, high-performance deep learning library. In *Advances*

in Neural Information Processing Systems 32, pages 8024–8035, 2019.

[18] ESA Space Environments and Effects Section. SPENVIS: Space Environment Information System, 2025. European Space Agency. Available at: https://www.spenvis.oma.be/.

[19] D. M. Sawyer and J. I. Vette. Ap-8 trapped proton environment for solar minimum and solar maximum. Technical Report NASA-TM-X-72605, NASA, Greenbelt, MD, USA, 1976. NASA Goddard Space Flight Center.

[20] European Space Agency. SEE1: Standard Electron Environment Model, version 1. ESA SP-1254, published as part of ESA Space Environment Standards, 2001. Available via ESA; commonly used within SPENVIS for trapped-electron environments.

[21] Stephen M. Seltzer. Updated calculations for electron and proton penetration and energy deposition in aluminum shielding – shieldose-2. Technical Report NASA-TP-3360, NASA Goddard Space Flight Center, Greenbelt, MD, USA, 1994.

[22] James R. Wertz and Wiley J. Larson. *Space Mission Analysis and Design*. Microcosm Press and Kluwer Academic Publishers, El Segundo, CA, USA, 3 edition, 1999.

[23] European Cooperation for Space Standardization. Space environment. Technical Report ECSS-E-ST-10-04C Rev.1, ECSS, Noordwijk, The Netherlands, 2020.

[24] Cobham Gaisler. GR740 Quad-Core LEON4FT SPARC V8 Processor, 2025. Available at: https://www.gaisler.com/products/gr740.

[25] Syntop Chemical Co Ltd. Fischer Tropsch Wax Ft Wax Synthetic Wax Xtft-70 Data Sheet, 2025.