The METASAT Space Platform: High Performance On-Board Processing for Institutional Missions Using Multicores, GPU and AI Accelerator

Leonidas Kosmidis[1,2] Marc Solé[2,1] Aridane Álvarez Suárez[3] Jannis Wolf[1]
Matina Maria Trompouki[1] Eckart Göhler[4] Alfred Hönle[4]

[1]Barcelona Supercomputing Center (BSC) [2]Universitat Politècnica de Catalunya (UPC)
[3]Fent Innovative Software Solutions (FENTISS) [4]OHB

***Abstract*—METASAT is a Horizon Europe project [3] focusing on enabling the high-performance processing capabilities of next generation European Space Missions, including advanced robotics and Artificial Intelligence (AI). In order to achieve this goal, an open source, high performance RISC-V platform which includes a multicore space processor enhanced with the SPARROW AI short vector unit [2] and the RISC-V Vortex GPU has been developed and prototyped in an FPGA. The platform supports mixed criticality workloads and has a fully qualified software stack which makes it appropriate for institutional missions. The platform has been demonstrated with space-related use cases, achieving significant speedups over the sequential versions and increased performance compared to GR740.**

***Index Terms*—component, formatting, style, styling, insert**

I. Introduction

New space missions need high performance processing in order to operate autonomously and process an increased amount of data. Such performance level cannot be provided by existing space processing technologies.

Graphics Processing Units (GPUs) and Artificial Intelligence (AI) accelerators are potential solutions for providing this level of performance, since previous studies have shown very promising results with space relevant software. However, such Commercial-off the shelf (COTS) devices cannot be used in high cost Institutional Missions (e.g. ESA Mission Classes I, II and III) which can only accept a very low risk. The reason for this is that at hardware level, COTS devices do not have the required reliability to withstand the harsh space environment, with radiation and extreme temperatures. Moreover, their COTS software stack depends on non-qualifiable components such as as Linux. In particular, all commercial GPUs currently only have drivers for Linux or Android.

In order to overcome this limitation, in the METASAT project [3] we have developed a RISC-V based platform which could be implemented in a radiation hardened FPGA and in the long term fabricated in a radiation-hardened process.

The platform is created around FrontGrade Gaisler's NOEL-V multicore processor, which has been integrated with the SPARROW AI accelerator unit [2], and a Vortex GPU [11], which is enhanced with desired properties for space. The platform has mixed criticality capabilities, allowing the co-existence of both high criticality and low criticality tasks in the same hardware. This is achieved thanks to virtualisation.

On the software side, METASAT has a fully qualifiable software stack, which consists of software with space heritage. In particular, it supports the XtratuM Next Generation (XNG) hypervisor, which is qualified up to ECSS Category B and provides robust partitioning.

Inside each partition, bare metal, Rust, Ada/SPARK or the widely used in space RTEMS SMP real-time operating system can be used. Regardless of the particular operating system or runtime in each partition, the SPARROW AI accelerator and the GPU can be used and shared among partitions.

The METASAT project has been completed in December 2024, achieving a fully functional status of its hardware and software. Moreover, it has been demonstrated with several use cases, running in parallel.

II. Hardware Overview

The METASAT platform is prototyped on Xilinx VCU118 FPGA. The basis of the METASAT system-on-chip is the GPL versions of NOEL-V and GRLIB from FrontGrade Gaisler, which provides the baseline multicore processor and its peripherals. We are using the highest performance configuration of the processor, which includes a dual-issue pipeline, memory management unit (MMU), and support for the RISC-V hypervisor extension. Since the GPL version is used, the

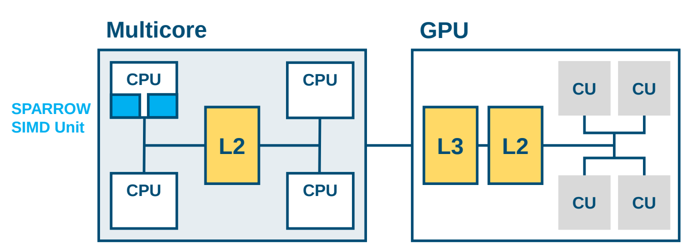

Fig. 1. METASAT Platform computing architecture.

core contains a non-pipelined, area-optimised floating point unit (nanoFPU-lite).

Each NOEL-V core is integrated with two SPARROW AI accelerator units – one for each of the two pipelines – which has been extended to work with the 64-bit registers of the core.

The SPARROW accelerator is a short vector / Single Instruction Multiple Data (SIMD) unit, specifically designed for the acceleration of the most common operations found in AI software, using 8 bit operations. SPARROW reuses the integer register file of the NOEL-V processor similar to the RISC-V Packed Extension and requires minimal hardware resources. Moreover, it requires minimal modifications of the core, without impact on qualification efforts. In fact, despite being developed with the GPL version of NOEL-V, FrontGrade Gaisler's feedback indicates that it integrates in a clean way with the commercial NOEL-V version, too. Moreover, BSC has been granted permission by Gaisler for the relicensing of SPARROW with a permissive Solderpad license, which opens the door for commercial adoption.

Although SPARROW is a custom ISA extension with full software support, as detailed in the following, BSC is also involved in the RISC-V Soft Core standardisation effort [9], which will facilitate its adoption within FPGA implementations or solutions like GR740-MINI which combine a hard processor with an FPGA.

The multicore cluster of the METASAT platform contains a shared Level 2 cache among the cores. Since we are using the NOEL-V GPL version, the L2-lite module of GRLIB is used, which offers less configurations than the proprietary version.

The CPU cluster is connected through AXI to a Vortex GPU. Vortex has been ported from PCIe accelerator FPGAs (i.e. Intel Arria 10 and Xilinx Alveo) to the target FPGA. Moreover, its driver and software stack has been ported from x86 Linux to RISC-V. Unlike COTS GPUs platforms which rely on Linux, the METASAT platform allows the GPU to be used either in bare-metal, RTEMS or through the XNG hypervisor. This opens the door for the implementation of a qualifiable space cloud approach which could be used in institutional missions, as opposed to Linux and docker existing approaches [5]. The platform includes a UART and an ethernet device, GRETH from GRLIB.

III. Software Stack

In terms of software support, the XtratuM NG hypervisor has been ported to the platform, offering time and space partitioning. It allows complete separation of different partitions, which can have different criticality and host a different operating system.

The hypervisor allows each partition to access the ethernet device, the UARTS, the SPARROW AI accelerator and the GPU, from whatever operating system they are using, including bare metal.

Access to GRETH is provided with an XNG driver, and an I/O ethernet server, which provides network connectivity to any partition requiring this functionality.

UART devices can also be virtualised as well as access to SPARROW. Finally, a GPU Server has been implemented, which provides exclusive access to each partition that needs to use the GPU. In this way, the GPU can be timeshared.

Under the RTEMS SMP real-time operating system, which has been modified to support SPARROW in the toolchain and in its context switch mechanism, multiple cores can be used for homogeneous parallel processing using OpenMP [8] and can also be combined with SPARROW code.

Although OpenMP is not included in the pre-qualified version of RTEMS SMP, it is a feature that is going to be included soon, thanks to the increased interest from the space community. Parallel processing allows to meet performance requirements which are not possible with a single core. Interestingly, OpenMP under RTEMS SMP achieves a remarkably linear scaling [8] for space- relevant algorithms [6] [10] [1].

The platform includes a complete software support for SPARROW, with supported added both in GCC and LLVM, as well as to Gailer's NCC, and RTEMS gcc toolchain. SPARROW can be programmed either in assembly, or in intrinsics similar to ARM NEON's, thanks to a high level library [7].

The Vortex GPU supports a low-level programming model, the Vortex Driver API, as well as OpenCL. Moreover, an OpenGL SC 2.0 driver is under development, which provides compliance with Khronos' Open API specification for safety critical workloads. Finally, a Vortex backed for the qualifiable CUDA-like high-level programming model Brook Auto [12] with the ISO 26262 qualifiable BRASIL GPU compiler [13] is under development.

In terms of AI support, the METASAT software stack supports TensorFlow Micro, with a SPARROW and a Vortex backend for the neural networks of the OBPMark-ML applications. Similar to the rest of the software, can be used in baremetal, RTEMS and XNG. Intergration with the GPU Server allows seamless sharing of the GPU of any AI application, without application code modifications.

IV. Use Cases

In order to demonstrate the maturity of the hardware and software platform, several space relevant use cases have been implemented. This includes an instrument control software from OHB, an AI-based FDIR (fault detection, isolation and recovery) application developed real housekeeping data from DLR's EnMAP satellite developed by OHB, and two high performance use cases, from ESA's OBPMark-ML benchmarking suite. These applications include a cloud screening and a ship detection application.

The use cases have reached a high level of integration and are possible to be executed in parallel on the METASAT platform. The METASAT prototype is demonstrated in a space representative lab environment, interfaced with EGSE and fed with EnMAP house keeping data and telecommands.

V. Performance Evaluation

Performance results comparing the parallel performance of our 100 MHz FPGA prototype against an ASIC implementation of FrontGrade Gaisler's GR740 space processor, provided very interesting and promising results. In fact, the general purpose performance of the two is comparable, despite the fact that they use different technologies (FPGA vs ASIC) and that the METASAT platform uses the GPL version of GRLIB's component which provide inferior performance than Gaisler's commercial components [8].

However, in terms of 8-bit AI inference performance, the METASAT platform is faster than GR740, thanks to the performance boost of its Single Instruction Multiple Data (SIMD) SPARROW AI accelerator [4].

In terms of the use cases, the METASAT platform is capable of providing the required performance of all applications, both the higher criticality ones (AI-based FDIR) and the payload processing ones (AI-based Cloud Screening and Ship Detection). The acceleration features provide a speedup of 2-3$\times$ compared to the scalar version of the applications, which can be further increased in the future since the priority of the project so far has been functionality over performance. Moreover, a future implementation of the METASAT platform with Gaisler's proprietary version instead of the GPL, would provide an additional performance boost.

In addition to the RTL (register transfer level) implementation, a QEMU model of the platform is available, allowing faster software development and debugging [4].

VI. Conclusion and Future Work

The METASAT hardware platform provides high performance for future space systems. It includes a multicore CPU, a SIMD accelerator and a GPU, a deep cache memory hierarchy, an Ethernet and UART. It features a completely qualifiable software stack, which makes it appropriate for institutional missions.

As a future work, the GPU will be enhanced with real-time and reliability features. In this way the Worst Case Execution Time (WCET) of GPU tasks will be possible to be computed. The design could be a good fit for space thanks to its built-in reliability features.

The METASAT platform and most of each software will be open-sourced in the following months, likely by the time of the event.

VII. Acknowledgments

The research presented throughout this paper has received funding by the European Community's Horizon Europe programme under the METASAT project (grant agreement 101082622).

References

[1] OBPMark (On-Board Processing Benchmarks). https://obpmark.github.io/.

[2] Marc Solé Bonet and Leonidas Kosmidis. SPARROW: A Low-Cost Hardware/Software Co-designed SIMD Microarchitecture for AI Operations in Space Processors. In *DATE*, 2022.

[3] Leonidas Kosmidis, Alejandro J. Calderón, Aridane Álvarez Suárez, Stefano Sinisi, Eckart Göhler, Paco Gómez Molinero, Alfred Hönle, Álvaro Jover Álvarez, Lorenzo Lazzara, Miguel Masmano Tello, Peio Onaindia, Tomaso Poggi, Iván Rodríguez Ferrández, Marc Solé Bonet, Giulia Stazi, Matina Maria Trompouki, Alessandro Ulisse, Valerio Di Valerio, Jannis Wolf, and Irune Yarza. METASAT: Modular Model-Based Design and Testing for Applications in Satellites. In *SAMOS*, 2023.

[4] Leonidas Kosmidis, Marc Solé, Ivan Rodriguez, Jannis Wolf, and Matina Maria Trompouki. The metasat hardware platform: A high-performance multicore, ai simd and gpu risc-v platform for on-board processing. In *European Data Handling Data Processing Conference (EDHPC)*, 2023.

[5] Leonidas Kosmidis, Matina Maria Trompouki, Marc Solé, Aridane Álvarez Suárez, and Jannis Wolf. Towards a qualifiable space cloud approach. In *2024 IEEE 10th International Conference on Space Mission Challenges for Information Technology (SMC-IT)*, pages 109–114, 2024.

[6] Ivan Rodriguez, Leonidas Kosmidis, Jerome Lachaize, Olivier Notebaert, and David Steenari. GPU4S Bench: Design and Implementation of an Open GPU Benchmarking Suite for Space On-board Processing. Technical Report UPC-DAC-RR-CAP-2019-1, Universitat Politècnica de Catalunya. https://www.ac.upc.edu/app/research-reports/public/html/research_center_index-CAP-2019,en.html.

[7] Marc Solé and Leonidas Kosmidis. Compiler support for an ai-oriented simd extension of a space processor. *Ada Lett.*, 42(1):95–99, December 2022.

[8] Marc Solé, Jannis Wolf, Ivan Rodriguez, Alvaro Jover, Matina Maria Trompouki, and Leonidas Kosmidis. Evaluation of the Multicore Performance Capabilities of the Next Generation Flight Computers. In *Digital Avionics Systems Conference (DASC)*, 2023.

[9] Marc Solé and Leonidas Kosmidis. Sparrow: Small and portable ai acceleration for risc-v softcores. In *Soft RISC-V Systems Workshop 2024, RISC-V International*, 2024.

[10] David Steenari, Leonidas Kosmidis, Ivan Rodríguez-Ferrández, Álvaro Jover-Álvarez, and Kyra Förster. OBPMark (On-Board Processing Benchmarks) - Open Source Computational Performance Benchmarks for Space Applications. In *2nd European Workshop on On-Board Data Processing (OBDP)*, 2021.

[11] Blaise Tine, Krishna Praveen Yalamarthy, Fares Elsabbagh, and Kim Hyesoon. Vortex: Extending the RISC-V ISA for GPGPU and 3D-Graphics. In *International Symposium on Microarchitecture (MICRO)*, 2021.

[12] Matina Maria Trompouki and Leonidas Kosmidis. Brook auto: high-level certification-friendly programming for gpu-powered automotive systems. In *Proceedings of the 55th Annual Design Automation Conference, DAC 2018, San Francisco, CA, USA, June 24-29, 2018*, pages 100:1–100:6. ACM, 2018.

[13] Matina Maria Trompouki and Leonidas Kosmidis. BRASIL: A high-integrity GPGPU toolchain for automotive systems. In *37th IEEE International Conference on Computer Design, ICCD 2019, Abu Dhabi, United Arab Emirates, November 17-20, 2019*, pages 660–663. IEEE, 2019.