The METASAT Hardware Platform: A High-Performance Multicore, AI SIMD and GPU RISC-V Platform for On-board Processing

Leonidas Kosmidis[1,2], Marc Sole[1,2], Jannis Wolf[1], Ivan Rodriguez[1,2] and Matina Maria Trompouki[1]

[1]Barcelona Supercomputing Center (BSC)
[2]Universitat Politècnica de Catalunya (UPC)

Abstract

The METASAT Horizon Europe project which is funded by the European Commission and started in January 2023, will enable model-based design methodologies in order to manage the complexity of upcoming hardware and software for space on-board processing. As a representative high performance platform for on-board processing, METASAT will design a multicore platform featuring accelerators prototyped on an FPGA. This includes both an AI SIMD accelerator tightly integrated with the CPU, as well as a GPU. All hardware components of the METASAT platform will be open source and based on the RISC-V open ISA. In this abstract, we provide an overview of the platform architecture as well as preliminary implementation decisions and the current development status.

Introduction

Current and upcoming space missions become increasingly complex, incorporating new functionalities and even the use of Artificial Intelligence (AI) for on-board processing. For this reason there is a trend to move towards more powerful hardware architectures which can provide the computational power required by this type of processing.

Multicores have been introduced into space platforms for more than a decade with the introduction of NGMP, which has been recently qualified. Next Front-Grade (ex-CAES) Gaisler's platforms are also based on multicores. In terms of RTOS support, RTEMS SMP has full support for multicores and its fully qualified for NGMP, with the pre-qualification toolkit for GR721RC and GR740 being openly available.

Despite the proliferation of multicores in space, currently their use is limited to the ability of executing single threaded tasks on different cores, but not real parallel processing. While this can increase the overall computation capacity of the on-board processing platform, the single thread performance provided by each core is not enough for the advanced functionalities mentioned earlier. This can be solved either with the introduction of more capable hardware such as accelerators or with more complex software, through the use of new parallel programming models like OpenMP. Hardware accelerators, especially the ones focused on AI processing require also complex programming models and software stacks.

Such an increased complexity of both hardware and software of future space platforms is hard to manage. For this reason, model-based engineering approaches are increasingly employed in the design of space systems. In particular, ESA has developed the open source TASTE framework[1] which is constantly under development with new functionalities.

The Horizon Europe project METASAT [1], which has started in January 2023 and it is funded by the European Commission, will develop model-based design approaches which will help to manage the complexity of programming such advanced high performance platforms, including AI accelerators and GPUs. For this reason, the METASAT reference platform is currently under development, which will serve as a target for the development of the aforementioned model based design methods.

The METASAT consortium consist of 5 partners and it is led by the Barcelona Supercomputing Center (BSC), which provides the high hardware platform described in this article, as well as two computationally intensive AI-based use cases, based on ESA's OBPMark-ML open source benchmarking suite [2][3]. The consortium includes fentISS which provides a hypervisor solution, IKERLAN and ALES/Collins Aerospace which provide expertise on the Model-Based Engineering and OHB which provides a real world satellite use case.

The METASAT hardware platform will be prototyped on an FPGA and will target an architecture that will be possible to be qualified and used in institutional missions in the future.

In this abstract, we provide an overview of the architectural design of the METASAT platform.

[1] `https://taste.tools/`

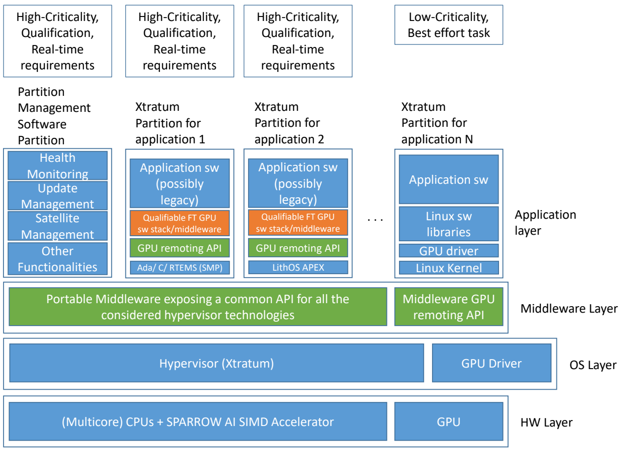

Figure 1: *METASAT's Hardware and Software Overview*

The METASAT Hardware Platform Overview

The METASAT hardware platform will be a mixed-criticality platform, allowing the deployment of software of different criticality on the same hardware. In order to achieve this, it employs the concept of virtualisation. In particular, it uses the Xtratum Hypervisor [4]. Figure 1 shows a high level hardware and software architecture of the platform. In order to provide a high-performance design, in which partitions of different criticality will be assigned in separate CPU cores, METASAT employs a multicore version of Front-Grade Gaisler NOEL-V [5] RISC-V CPU, which is enhanced with AI processing capabilities, through integration with the SPARROW open source AI SIMD (Single Instruction Multiple Data) unit [6]. This will satisfy the AI needs of applications with moderate acceleration needs, with low latency requirements and with the need of high criticality, qualifiable software.

For the acceleration of applications with much higher performance needs, the METASAT platform includes an open source RISC-V based GPU [7][8], which will be extended with real-time capabilities such as hardware features that allow the computation of Worst Case Execution Time (WCET) of GPU tasks and reduction of interference between multiple GPU tasks executed concurrently on the GPU, and reliability features which are required for use in space. The GPU is fully configurable in the number of shader cores, number of threads, presence and size of shared L2 and L3 caches etc.

One of the current limitations of the use of GPUs in institutional missions of high criticality is that most GPUs require device drivers and user space libraries for non-qualifiable operating systems like Linux and Android [9]. Moreover, their closed source nature prevents their porting to qualifiable, Real-Time operating systems used in space like RTEMS. In METASAT, we will overcome this limitation by adapting Vortex's bare metal open source GPU driver, and develop a portable method for the use of GPUs among multiple partitions, no matter whether they will be running on bare-metal, RTEMS, Xtratum Runtime Environment or a full linux partition for low criticality software.

Current Status

The selected FPGA platform that will host the hardware prototype is the Xilinx VC118. A preliminary synthesis of the METASAT platform shows that this FPGA is enough to include 8 64-bit configurations of NOEL-V with SPARROW AI accelerators with a 48% utilisation. This leaves half of the FPGA for the implementation of a multicore Vortex GPU consisting of 4 64-bit shader cores and a 64KB L2 cache. Currently the integration between the CPU and GPU is on-going. Then a design space exploration will be performed to decide the exact CPU and GPU configuration according to the needs of the project use cases. It is worth noting that the platform will be released as open source.

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] BSC, IKERLAN, OHB, fentISS, ALES. *METASAT: Modular model-based design and testing for applications in satellites*. `https://metasat-project.eu/`.

[2] David Steenari et al. In: *2nd European Workshop on On-Board Data Processing (OBDP2021)*. `https://doi.org/10.5281/zenodo.5638577`. 2021.

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

[4] fentISS. *XtratuM hypervisor*. `https://www.fentiss.com/xtratum/`.

[5] Gaisler. *NOEL-V Processor*. `https://www.gaisler.com/index.php/products/processors/noel-v`.

[6] 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. DOI: `10.23919/DATE54114.2022.9774730`.

[7] Georgia Tech. *Vortex GPU*. `https://vortex.cc.gatech.edu/`.

[8] Blaise Tine et al. "Vortex: Extending the RISC-V ISA for GPGPU and 3D-Graphics". In: *MICRO-54: 54th Annual IEEE/ACM International Symposium on Microarchitecture*. MICRO '21. 2021. ISBN: 9781450385572.

[9] Leonidas Kosmidis et al. "GPU4S: Major Project Outcomes, Lessons Learnt and Way Forward". In: *2021 Design, Automation Test in Europe Conference Exhibition (DATE)*. 2021.