EMSA5: A RISC-V Processor System for Enhanced Functional Safety in Embedded Applications

Michael Faulwasser[1], Martin Zimmerling[1]

[1] Department DCC, Fraunhofer Institute for Photonic Microsystems (IPMS), Dresden, Germany

E-mail:
michael.faulwasser@ipms.fraunhofer.de

Abstract

The RISC-V RV32 processor system EMSA5 was originally designed to meet the stringent requirements of functional safety as specified by ISO26262. It consists of a redundant multi-core architecture that supports advanced safety mechanisms such as Dual-Mode-Redundancy with Lockstep (DMR-L) and Triple-Mode-Redundancy (TMR), ensuring high reliability in critical applications. However, functional safety is also highly relevant in industrial, avionics and space applications. A comprehensive RISC-V system requires more than just a core; the EMSA5 includes a complete ecosystem tailored for software development, as well as essential peripherals for both processing and communication. To facilitate different communication interface requirements, EMSA5 covers both classic CAN and modern Ethernet-based network architectures, demonstrating its versatility in various automotive and industrial applications.

Furthermore, the incorporation of RISC-V processor extension significantly improves performance in a trade-off with resource utilization, enabling more efficient data processing and computational capabilities.

This work shows the EMSA5 processor for embedded sensor signal processing in a network environment. The algorithms are based on neural networks, and it is evaluated in performance with and without Vector Extension Zve32x for quantized data and Floating Point Extension for floating point data. In addition to performance, the system is analyzed in terms of logic resources. The EMSA5 processor is compatible with leading FPGA platforms such as Microchip and AMD Xilinx, and its silicon has been proven by foundries like Global Foundries. This underlines the EMSA5's contribution to the development of safe, efficient, and versatile processing solutions in the ever-evolving landscape of embedded systems.

Keywords: RISC-V, redundancy, TMR, tracing, AI

Introduction

The EMSA5 is an advanced embedded processor IP core designed for applications that require fail-safe systems, such as space environments. It features a RV32 RISC-V processor with a 5-stage pipeline and provides redundancy options including DMR-L and TMR to ensure reliability. The processor supports Privilege Modes M+U and includes PMP for memory access protection, along with non-maskable interrupt (NMI) for handling hardware faults. Notably, it offers optional vector extensions Zvex32 and an optional floating-point unit (FPU) capable of single and double precision calculations, making it suitable for embedded AI applications requiring enhanced computational capabilities.

The EMSA5 is designed with low-power and low-footprint considerations and includes an extensive AMBA-based infrastructure featuring AHB-lite [1] and APB [2] buses with protection mechanisms such as parity and ECC (error correcting code), alongside ECC-protected memories for memory protection. Its application areas include serving as a safe, reliable real-time controller and communication processor, as well as enabling AI-powered system monitoring and AI edge computing.

In terms of hardware validated design, the EMSA5 is compatible with FPGA (Field-Programmable Gate Array)

TRISTAN Project has received funding from the Chips Joint Undertaking (Chips-JU) under the grant agreement nr. 101095947. Chips-JU receives support from the European Union's Horizon Europe's research and innovation programme and Austria, Belgium, Bulgaria, Croatia, Cyprus, Czechia, Germany, Denmark, Estonia, Greece, Spain, Finland, France, Hungary, Ireland, Israel, Iceland, Italy, Lithuania, Luxembourg, Latvia, Malta, Netherlands, Norway, Poland, Portugal, Romania, Sweden, Slovenia, Slovakia and Turkey.

© 2025 Fraunhofer IPMS

platforms such as Xilinx, Microchip, and Intel, as well as ASIC implementations. The processor's reliability is bolstered by its adherence to ISO26262 functional safety standards, with features like processor subsystem redundancy (DMR-L, TMR), watchdogs, error detection and correction (EDC) scheme in bus logic, ECC-protected memories, and observability through a trace interface. Additionally, a Safety Test Library (STL) is available to further ensure system reliability.

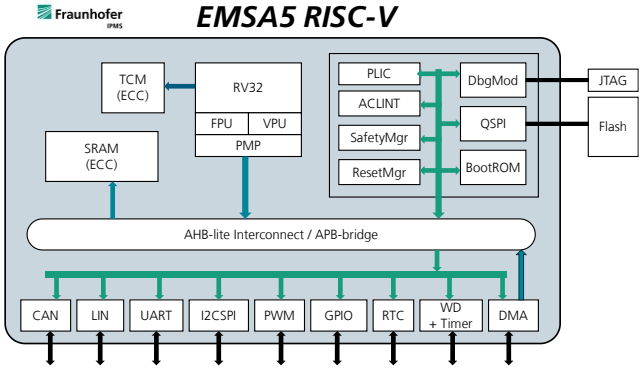

Figure 1: Reliable System by Redundancy and Error Detection and Correction

The EMSA5 integrates seamlessly with a range of other communication IPs, allowing for the creation of custom systems tailored to specific needs. Furthermore, it supports a comprehensive ecosystem including the Gnu Compiler Collection (GCC), IAR Embedded Workbench, Lauterbach uTrace, and FreeRTOS, providing developers with robust tools for efficient system design and implementation.

Table 1: Resource Numbers for a Minimal RV32EC Configuration

	Technology	Resources
EMSA5-GP	FPGA[1]	LUTs: 4385, FFs: 2142
EMSA5-FS[2]	FPGA[1]	LUTs: 15949, FFs: 7717
EMSA5-GP	180 nm	0.217 mm², 21.6k ND2gates
EMSA5-FS[2]	180 nm	1.003 mm², 100k ND2gates

[1] Xilinx Artix7

[2] TMR, ECC enabled

The design of the EMSA5 processor has been thoroughly tested on both FPGA and ASIC platforms, demonstrating its versatility and reliability across different hardware implementations. Despite its advanced capabilities, the EMSA5 maintains a low resource count, making it an efficient choice for systems where resource optimization is critical. Detailed numbers regarding the hardware resource utilization are provided in a table within this paper, highlighting the processor's minimal general-purpose (GP) IP design and its fail-safe TMR (FS) design. Specific hardware resources are presented for the FPGA Xilinx Artix7 device and a 180 nm ASIC technology, offering insights into its adaptable and efficient architecture.

Robustness by Observability

The EMSA5 processor incorporates a sophisticated trace interface that plays a pivotal role in enhancing system robustness. This interface is crucial for ensuring reliable software verification, efficient fault detection, and comprehensive firmware validation. By providing detailed observability of the system state, the trace interface allows engineers to monitor and analyze operations in real-time, facilitating the prompt identification of anomalies and ensuring system integrity.

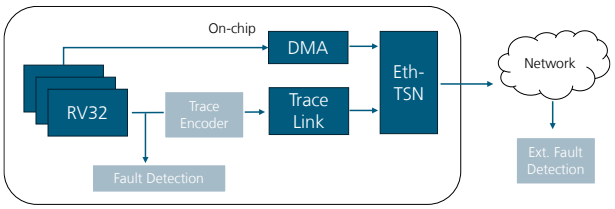

Figure 2: TSN TraceLink for Observability

A key feature of the EMSA5 processor is the integration of TraceLink over TSN Ethernet. This setup allows trace data and application data to share a common interface using deterministic Time-Sensitive Networking (TSN) capabilities. Such integration ensures seamless and reliable data transfer, which is vital for maintaining the usability and efficiency of embedded systems in aviation and space applications. The deterministic properties of TSN enable precise control of data flow, ensuring that both trace and application processes operate harmoniously within the system architecture.

The combination of robust observability through the trace interface and the efficient data transfer facilitated by TSN Ethernet positions the EMSA5 processor as a key component in advancing the reliability of embedded systems in harsh environments.

Custom Systems Enabled by a Diverse IP core Portfolio

In the realm of embedded processing for space and other environments demanding functional safety, robust and reliable communication is paramount. The EMSA5 ecosystem provides a comprehensive suite of communication interfaces, catering to both high-speed and low-speed applications. This versatility is crucial for ensuring effective data exchange in diverse settings, from automotive to industrial applications.

EMSA5 supports widely adopted communication protocols such as UART, I2C, and SPI for low-speed needs, alongside classic CAN and LIN protocols. Notably, the CAN protocol includes redundancy options, enhancing reliability in mission-critical scenarios. For high-speed applications, the EMSA5 portfolio features LLEMAC-1G, a low-latency Ethernet MAC

with redundancy capabilities, and 10G Ethernet for high-speed communication demands.

A significant advantage of the EMSA5 system is its adherence to Time-Sensitive Networking (TSN) standards over Ethernet, which ensures deterministic communication. The support for standards like 802.1AS enables sub-microsecond time synchronization, vital for coordinating time-sensitive operations. Additionally, 802.1Qav and 802.1Qbv facilitate traffic shaping and scheduling, optimizing the flow of data to meet stringent timing requirements.

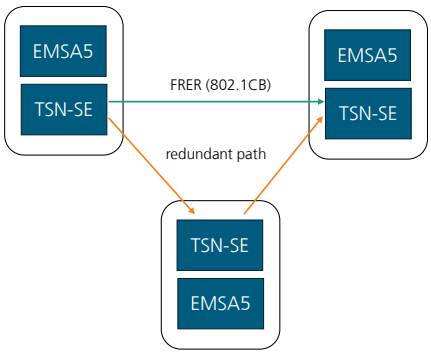

Figure 3: Redundant Data Communication Using FRER

Of particular importance is the implementation of 802.1CB, which enables frame replication and elimination for redundant communication paths (FRER). This feature is essential in environments where reliability and fault tolerance are critical, allowing data to traverse multiple paths and ensuring continuous communication even in the event of a fault.

In conclusion, the EMSA5 communication IP portfolio not only meets the diverse interface requirements for embedded processing in functional safety-critical environments but also emphasizes the importance of deterministic communication through well-established protocols and TSN standards. This ensures that systems can operate reliably and efficiently, even in the most challenging conditions.

Embedded AI with EMMI

Embedded AI is a burgeoning field that focuses on deploying artificial intelligence models on embedded systems, which are typically constrained by limited computational resources and power. As described by the TinyML foundation, Embedded AI involves hardware, algorithms, and software capable of performing on-device sensor data analytics at extremely low power. This requires a specific workflow for developing embedded AI applications, including optimization techniques such as approximation, batch normalization folding, and quantization.

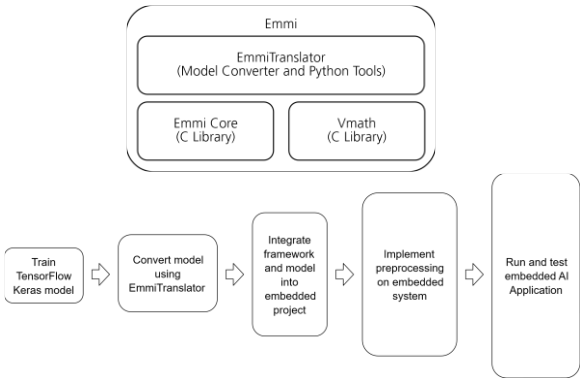

Figure 4: AI Framework Emmi Components and Workflow

The Emmi framework is an inference-only AI framework specifically designed to run neural networks on embedded systems. It is written entirely in C and supports both floating-point and integer operations with quantized weights, biases, and feature maps [4]. Models for Emmi are converted from TensorFlow Keras, supporting both sequential and functional architectures. The framework comprises three components: a model converter that transforms trained TensorFlow Keras models into C code, the Emmi Core which implements supported activation functions and layers, and the vmath library that performs mathematical operations for both floating-point and quantized data types. A notable feature of Emmi is its optimization for automatic vectorization by the compiler, resulting in faster execution times while maintaining compatibility across various platforms.

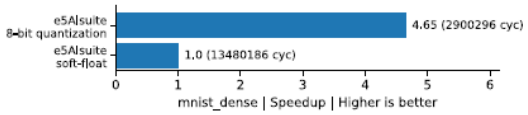

Figure 5: Speedup of the quantized network mnist_dense in the emmi AI framework [1]

The results of the thesis demonstrate significant enhancements to Emmi, including the implementation of 1D convolutional layers, batch normalization, concatenation, and global pooling. These improvements contribute to lowering the memory requirements for quantized neural networks. The EmmiTranslator now supports batch normalization folding through the integration of the Python package tf-batchnorm-fold, which has been adapted to work with Emmi's full feature set. Rigorous testing and benchmarking have validated these enhancements, showcasing Emmi's capability to execute neural networks efficiently on embedded systems.

A practical example of Emmi's application is illustrated through a conveyor demonstrator for predictive maintenance. Conveyors are automated mechanical devices with belts stretched over pulleys, essential for transporting goods and materials. Optimal belt tension is crucial for performance. The embedded AI application developed analyzes sensor data from accelerometers, gyrometers, and magnetometers to predict whether the belt is too loose or too tight. This real-world

application underscores Emmi's effectiveness in sensor data analysis for predictive maintenance tasks.

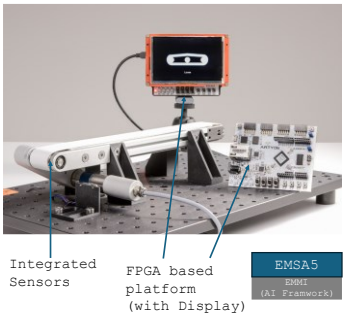

Figure 6: Demo setup of the conveyor predictive maintenance application

Conclusion

The EMSA5 processor system provides a robust solution for functional safety in embedded applications, leveraging RISC-V architecture with advanced redundancy mechanisms such as DMR-L and TMR. Its comprehensive ecosystem supports both software development and multiple communication interfaces, making it versatile for automotive, industrial, avionics, and space applications. Integration of RISC-V extensions enhances performance while maintaining efficient resource utilization. Compatibility with leading FPGA and ASIC platforms further underscores its reliability and adaptability. The EMSA5 system's focus on deterministic communication through TSN standards and its embedded AI capabilities with the Emmi framework demonstrate its potential to advance safe and efficient processing solutions in challenging environments, reinforcing its role in predictive maintenance and other critical applications.

Funding

TRISTAN Project has received funding from the Chips Joint Undertaking (Chips-JU) under the grant agreement nr. 101095947. Chips-JU receives support from the European Union’s Horizon Europe’s research and innovation programme and Austria, Belgium, Bulgaria, Croatia, Cyprus, Czechia, Germany, Denmark, Estonia, Greece, Spain, Finland, France, Hungary, Ireland, Israel, Iceland, Italy, Lithuania, Luxembourg, Latvia, Malta, Netherlands, Norway, Poland, Portugal, Romania, Sweden, Slovenia, Slovakia and Turkey.

References

[1] ARM AMBA 3 AHB-Lite Protocol Specification (IHI0033A), 06.06.2006, developer.arm.com

[2] ARM AMBA APB Protocol Specification v2.0 (IHI0024C), 13.04.2010, developer.arm.com

[3] Pfeiffer, P. 2023 Master *Thesis Development of a Machine Learning Framework for Quantized Neural Networks on Embedded RISC-V Systems*

[4] Pfeiffer, P. 2024 Master *Thesis Extension of an Embedded AI Framework for RISC-V Systems*