MOTIUS

WE R&D.

eMil ML System

Motius GmbH

November 19, 2025 20:31 (ccd1cd0)

eMil - Machine Learning System on RISC-V

Government-Funded Research Project (2022 - Today)

> **BayVFP-Funded Research**
>
> Research project funded by the Bayerisches Verbundforschungsprogramm (BayVFP) des Freistaates Bayern; Förderlinie "Digitalisierung"

Project Overview

Funding: Bayerisches Verbundforschungsprogramm (BayVFP)

Partners: Perif and Motius (partially covering development costs)

Timeline: 2022 - Today (Ongoing)

Type: Research & Development

Objective: Design and build a machine learning system that is networked across different levels, from sensors to the cloud, and optimized as a whole.

Status: Functionality phase (FPGA) - did not proceed to manufacturing

---

Multi-Level ML Architecture

The project aimed to create an end-to-end ML processing pipeline optimized for edge computing:

```
Sensors → Edge Processing (RISC-V + FFT) → Network → Cloud
```

Edge ML Processing

Machine learning workloads running on RISC-V at the edge

FFT Accelerator

Custom hardware for sensor data preprocessing

System-Wide Optimization

End-to-end optimization across the entire pipeline

Real-Time Signal Processing

Hardware acceleration for radar applications

Development Workflow

We figured out what the project needed, built the architecture, and implemented it successfully through multiple phases:

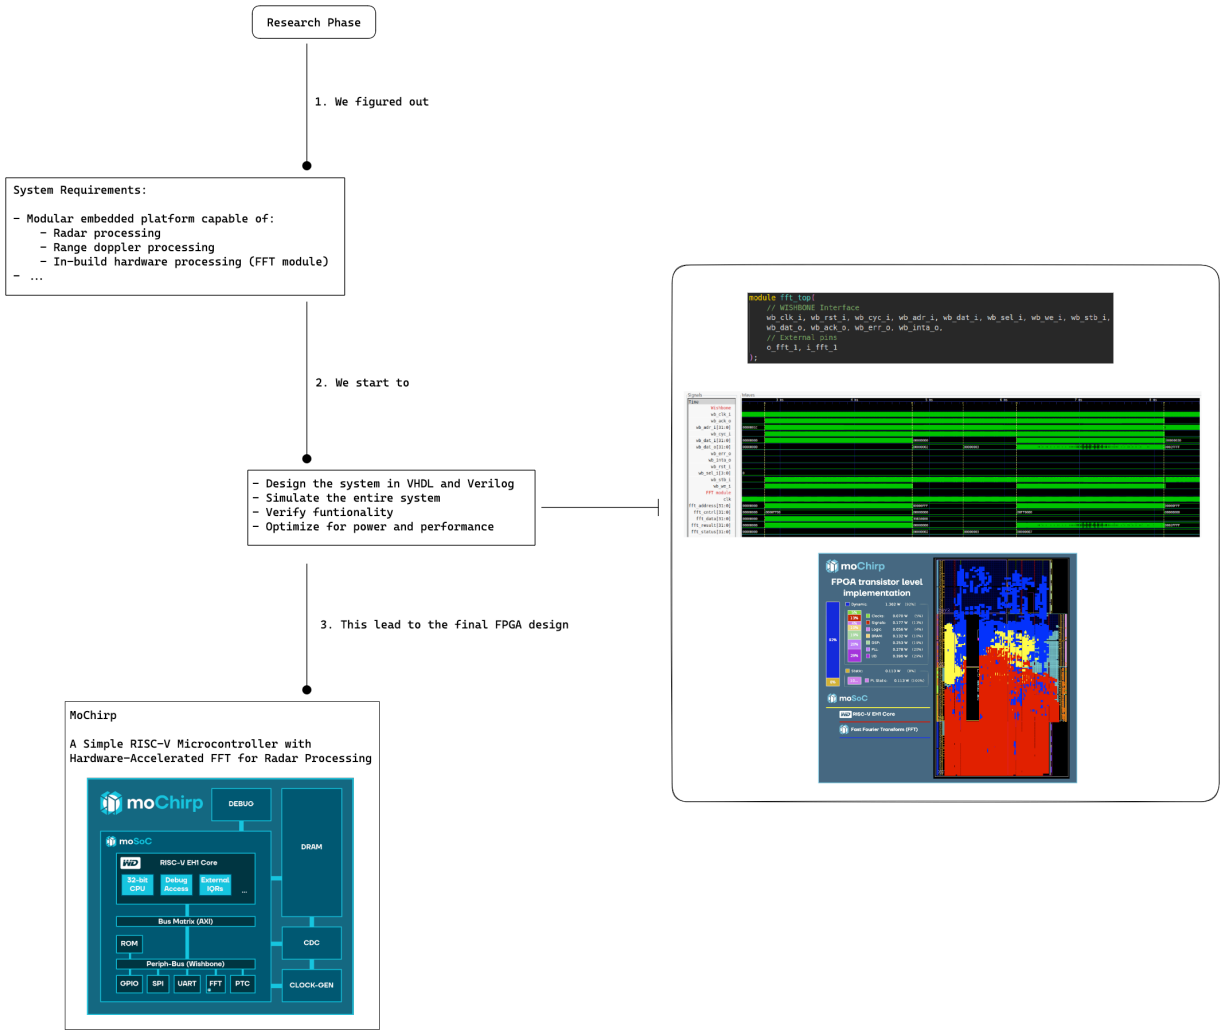

Development Phases: 1. **Requirements gathering** - Analyzed embedded platform capabilities for radar processing 2. **System design** - Designed the system in HDL and verified functionality 3. **Hardware implementation** - Built and optimized the design for FPGA 4. **Validation** - Tested performance and validated ML workload processing

---

System Architecture: moSoC (Motius System-on-Chip)

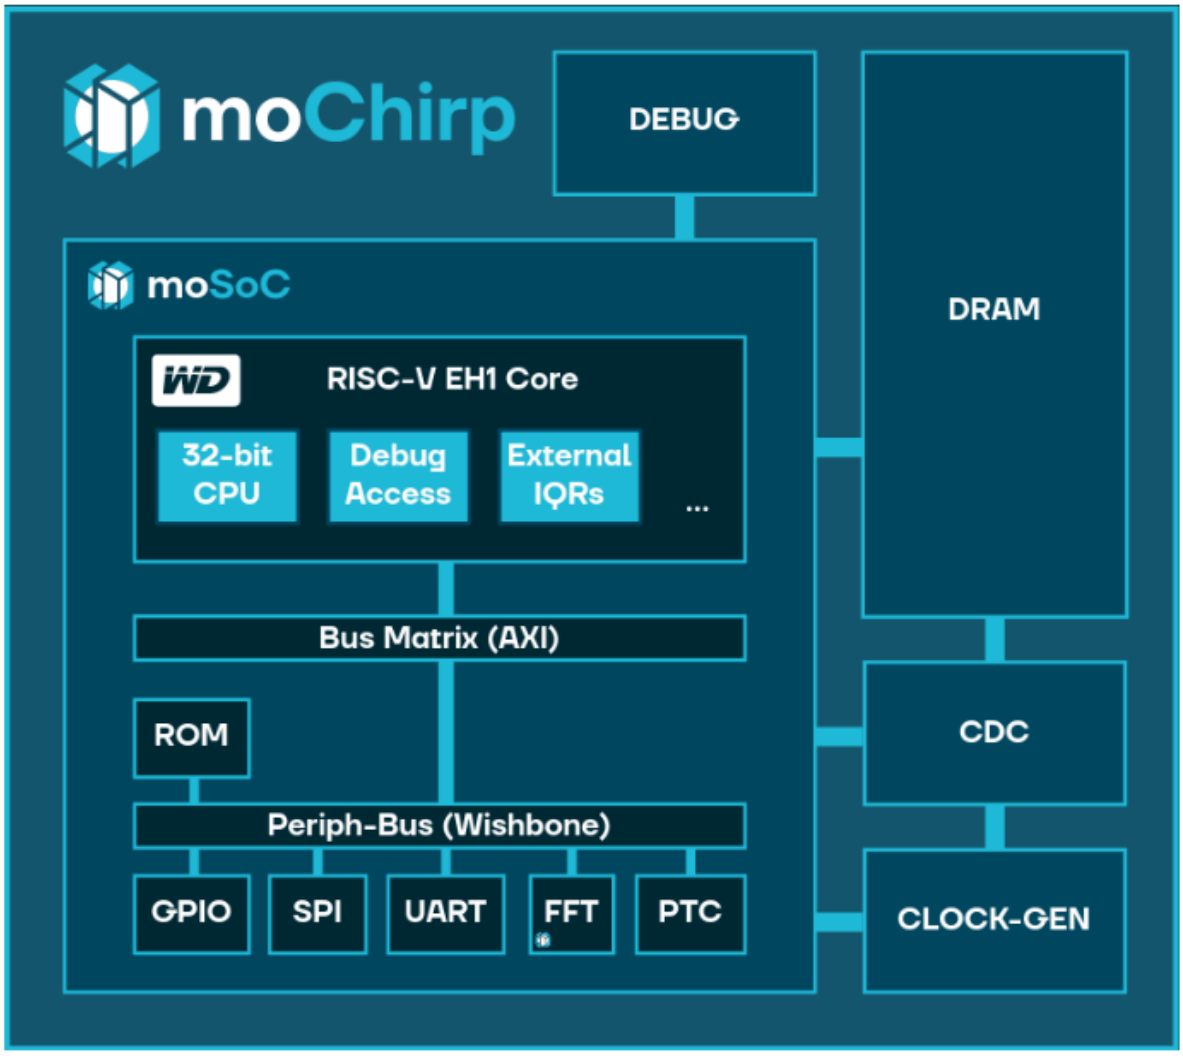

Figure 2.1.: MoSoC system architecture (Bello 2022)

moChirp System Components:

Western Digital RISC-V EH1 Core

32-bit RISC-V CPU with debug access and external I/Os

Memory Hierarchy

ROM, DRAM controller for high-bandwidth data processing

Dual Bus Architecture

AXI bus matrix for high-speed transfers, Wishbone peripheral bus

Custom FFT Accelerator

Hardware-accelerated Fast Fourier Transform designed by Motius

Communication Interfaces

UART, SPI, GPIO for sensor and debug connectivity

Support Modules

PTC (Programmable Timer Counter), CDC (Clock Domain Crossing), CLOCK-GEN

Custom FFT Accelerator

Purpose: Hardware-accelerated Fast Fourier Transform for sensor signal processing

FFT Architecture

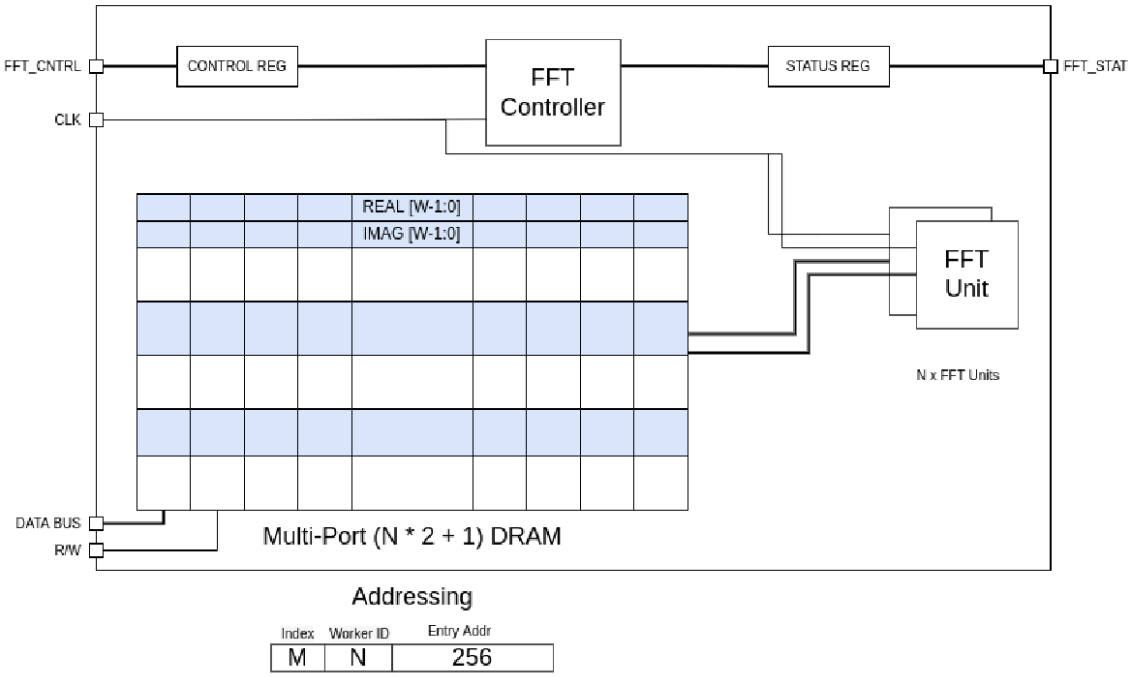

Figure 2.7.: Architecture diagram of the whole FFT surroundings in the moSoC (Meléndez 2022)

Design Features:

Multi-Port DRAM

Efficient data access architecture with multiple ports

Configurable Controller

FFT controller with adjustable parameters

Complex Data Support

REAL and IMAGINARY data channel processing

Pipelined Computation

Parallel FFT computation units for throughput

Optimized Addressing

Smart addressing (Index, Worker ID, Entry Address)

Implementation:

Designed custom FFT IP core in HDL from scratch. Integrated with RISC-V CPU via Wishbone bus. Optimized for real-time edge processing with memory-efficient multi-port DRAM architecture.

FPGA Implementation

Resource Utilization

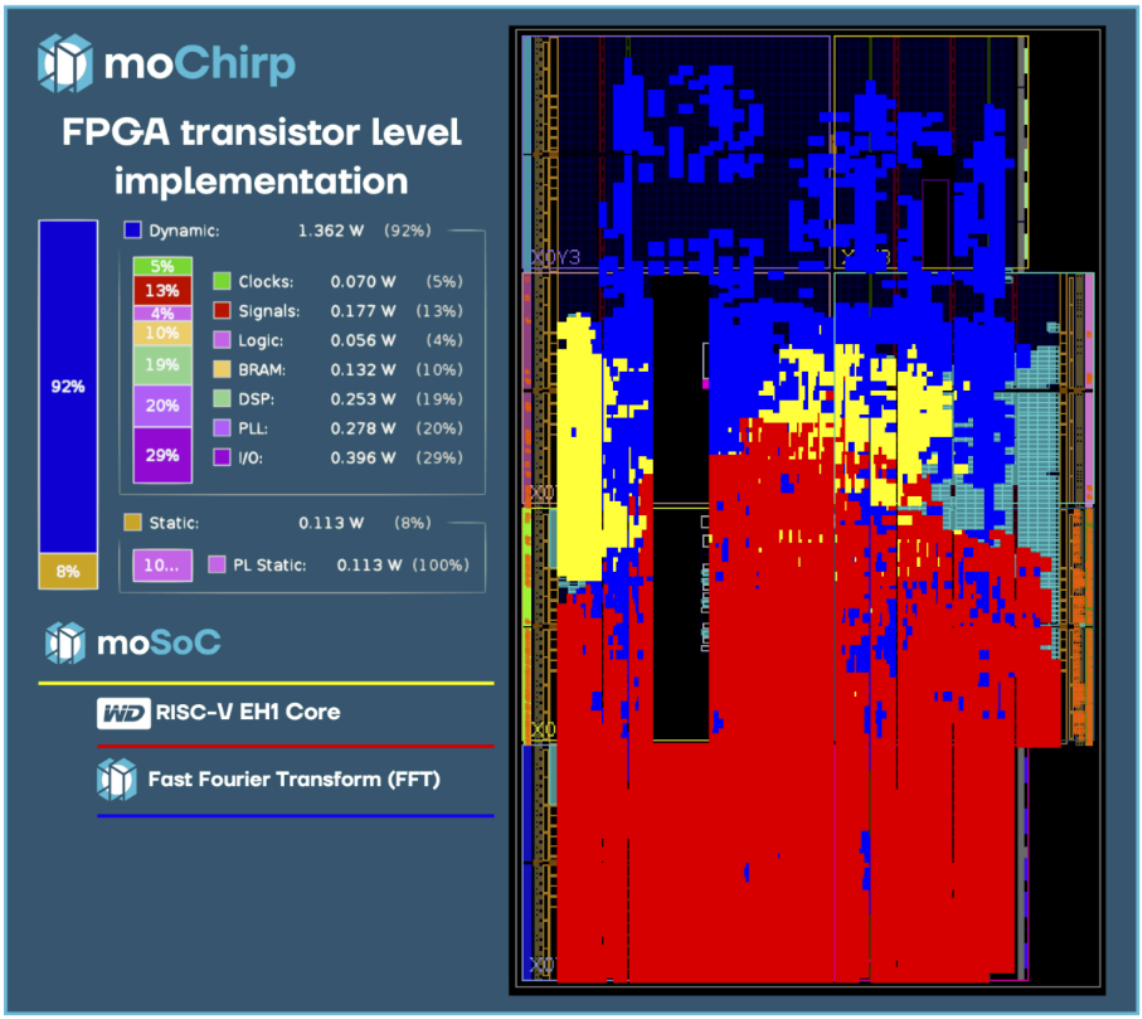

Figure 3.3.: Resource utilization of the Artix-7 100T on the Nexys A7 board from Xilinx

Hardware Platform: AMD Xilinx Artix-7 (Nexys A7)

System Integration:

Complete SoC with RISC-V CPU, custom FFT IP, and peripherals running at 100MHz on FPGA. Pipelined FFT computation enables low-latency sensor data processing.

Technical Validation

FFT Waveform Analysis

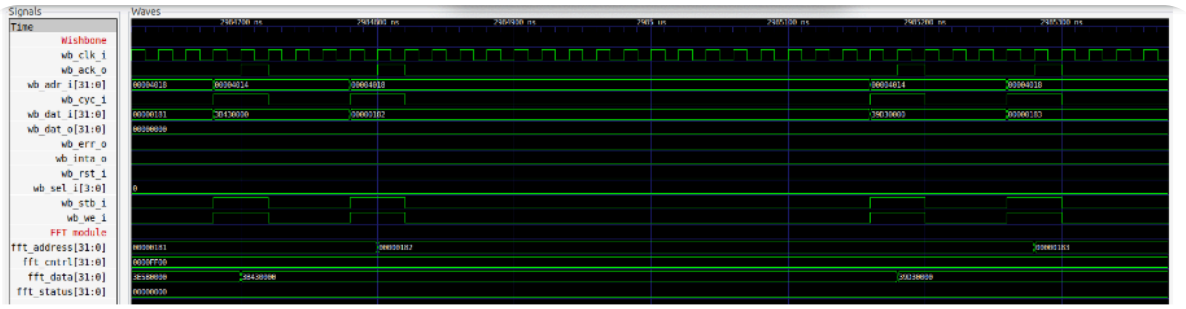

Figure A.1.: Waveform of two typical read cycles on the Wishbone bus (from memory to the FFT ROM)

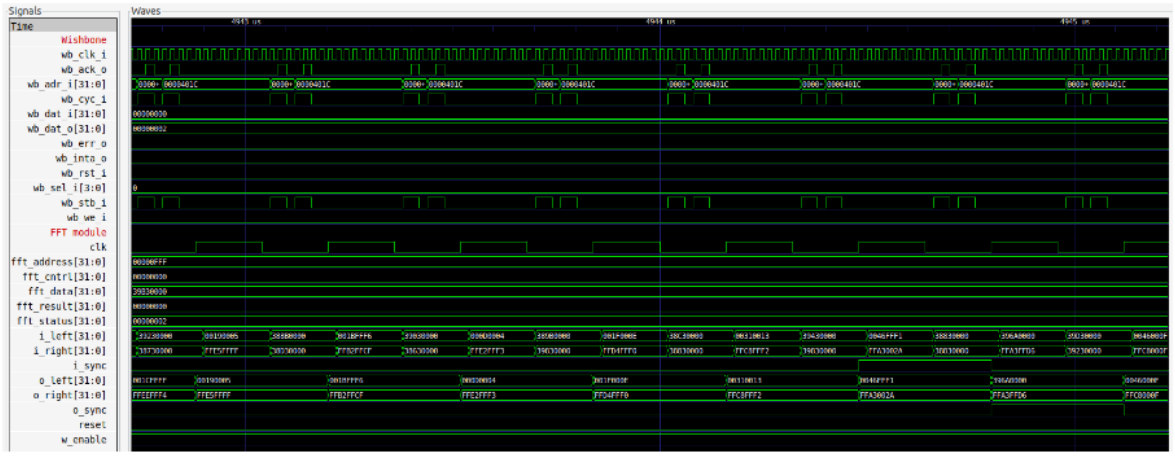

Figure A.2.: Waveform of some input cycles of the FFT pipeline on the moSoC

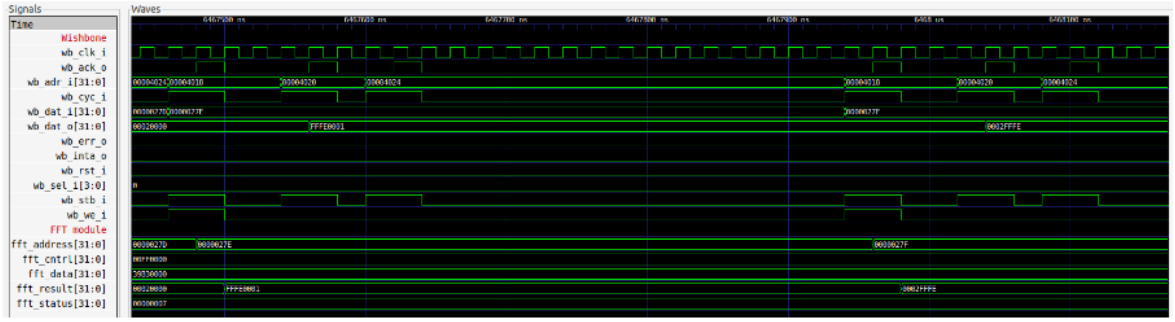

Figure A.3.: Waveform of two typical write cycles on the Wishbone bus (from FFT ROM to memory)

Testing and Verification:

Validated typical read cycles on Wishbone bus (memory to FFT ROM), input cycles of FFT pipeline on moSoC, and write cycles (FFT ROM to memory). Complete functional validation of the FFT accelerator

confirmed 100MHz operation on FPGA with real-time signal processing and successful ML workload execution.

Technologies & Tools

Category	Technology
FPGA	AMD Xilinx Artix-7 (Nexys A7 board)
CPU Core	Western Digital RISC-V EH1 (32-bit)
Custom IP	FFT accelerator (designed by Motius)
Bus Architecture	AXI bus matrix, Wishbone peripheral bus
Peripherals	UART, SPI, GPIO, DRAM controller, PTC
Memory	ROM, DRAM with multi-port architecture
Debug	JTAG debug interface, external I/Os

Key Achievements

RISC-V ML System

Built complete ML processing system on RISC-V, demonstrating edge ML feasibility with custom acceleration. Integrated multiple open-source IP modules and achieved target performance on FPGA.

∿ Custom FFT Accelerator

Designed and verified FFT IP core from scratch for hardware-accelerated sensor signal processing. Optimized for low-latency edge computing and successfully integrated with RISC-V CPU.

System Integration

Complete SoC architecture with CPU, custom IP, and peripherals. Multi-level bus architecture (AXI + Wishbone), memory hierarchy (ROM, DRAM), and comprehensive peripheral set.

Research Insights

Edge ML on RISC-V

RISC-V is viable for edge ML workloads when paired with custom acceleration. The open-source IP ecosystem is mature enough for research, and FPGA prototyping effectively validates concepts before silicon.

Custom IP Development

HDL design for ML acceleration requires careful optimization. Bus interface standards (AXI, Wishbone) enable modularity. Simulation and verification are critical, as integration testing reveals timing and resource challenges.

System-Wide Optimization

Optimize across the sensor → edge → cloud pipeline. Balance CPU processing with hardware acceleration. Memory bandwidth is often the bottleneck, and power efficiency requires early architectural decisions.

Project Status

Current Phase: FPGA functionality phase (ongoing)

Achieved: RISC-V system functional on FPGA, custom FFT accelerator validated, ML workloads demonstrated, and system-wide integration complete.

Not Pursued: Manufacturing and tape-out were not pursued. This research project focused on validation and feasibility rather than production manufacturing.

Documentation & Resources

Internal Resources

GitLab Repository: gitlab.motius.de/motius-silicon/Motius-radar-processing

Documentation: docs-motius-silicon.apps.motius.ci

Impact on Motius Capabilities

RISC-V Experience

Proven ability to integrate Western Digital RISC-V cores with deep experience in the open-source RISC-V ecosystem and understanding of RISC-V for ML workloads.

Custom IP Development

FFT accelerator design expertise and ML acceleration hardware design. Experience integrating custom IP with standard buses (AXI, Wishbone).

Research Capabilities

Government-funded project experience with academic and research partnerships. Strong publication and technical documentation skills.

← Commercial Platform

Next: NEORV32 →