

TOPICAL WORKSHOP ON ELECTRONICS FOR PARTICLE PHYSICS  
RETHYMNO, CRETE, GREECE  
6–10 OCTOBER 2025

## Design and evaluation of TriglaV: a prototype SoC ASIC for particle physics applications realized with the SOCRATES platform

Anvesh Nookala,<sup>a,b,c,\*</sup> Marco Andorno,<sup>a</sup> Alessandro Caratelli,<sup>a</sup> Davide Ceresa,<sup>a</sup> Benoît Denkinger,<sup>a</sup> Georgios Nikolaidis,<sup>a</sup> Kostas Kloukinas,<sup>a</sup> Jeffrey Prinzie<sup>b</sup> and Axel Jantsch<sup>c</sup>

<sup>a</sup>*EP-ESE-ME, CERN,*

*Espl. des Particules 1, Meyrin, Switzerland*

<sup>b</sup>*Electronic Circuits and Systems, KU Leuven,  
Kleinhoefstraat 4, Geel, Belgium*

<sup>c</sup>*Institute of Computer Technology, TU Wien,  
Gufshausstraße 27-29, Vienna, Austria*

*E-mail:* [anvesh.nookala@cern.ch](mailto:anvesh.nookala@cern.ch)

**ABSTRACT.** As part of an ongoing R&D effort to introduce software programmability and general computing functionalities to ASICs for future particle physics experiments, the microelectronics section at CERN is developing fault-tolerant System-on-Chips (SoC) for applications in harsh environments such as LHC detectors. This work presents the TriglaV ASIC, a RISC-V-based fault-tolerant 32-bit SoC built using our SOCRATES framework for rapid and customizable SoC generation. Our aim with TriglaV is to validate the SOCRATES flow with a silicon prototype and to demonstrate this prototype for resilient operation of micro-controller workloads applicable to particle physics accelerators and detectors. The chip was fabricated in a commercial 28 nm CMOS technology targeting a clock frequency of 250 MHz. Its design integrates extensive protection against single-event effects combining triplication, error-correcting codes and special layout techniques. We have experimentally qualified TriglaV for TID robustness up to 1 Grad through high-dose rate x-ray irradiation. Single-event upset cross-section for corrected and uncorrected errors were measured up to an energy of 62 MeV cm<sup>2</sup>/mg. This effort has paved the way for future designs generated with SOCRATES and upcoming ASICs incorporating SoCs for operation in particle physics environments.

**KEYWORDS:** Digital electronic circuits; Radiation-hard electronics; VLSI circuits

\*Corresponding author.

---

## Contents

|          |                                       |          |
|----------|---------------------------------------|----------|
| <b>1</b> | <b>Introduction</b>                   | <b>1</b> |
| <b>2</b> | <b>Design description</b>             | <b>2</b> |
| 2.1      | System architecture                   | 2        |
| 2.2      | Upset and error observability         | 3        |
| 2.3      | Physical design                       | 3        |
| <b>3</b> | <b>Experimental validation</b>        | <b>3</b> |
| 3.1      | Power and area consumption            | 3        |
| 3.2      | Heavy-ion testing at UCLouvain        | 4        |
| 3.3      | X-ray irradiation at room temperature | 6        |
| <b>4</b> | <b>Conclusions</b>                    | <b>6</b> |

---

## 1 Introduction

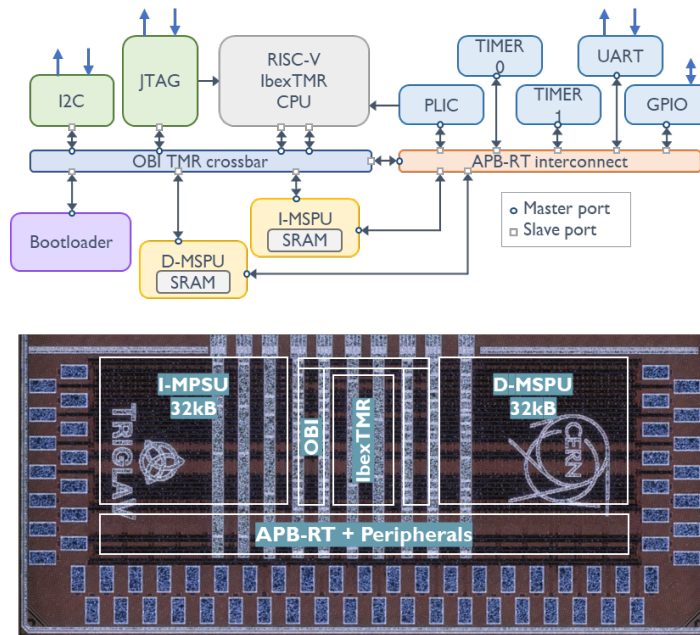
The complexity of ASICs used for instrumentation in future generations of High-Energy Physics (HEP) applications is expected to increase, driven by advances in semiconductor manufacturing technologies that improve performance, as well as by the trend toward greater programmability and logic capabilities in designs [1]. HEP environments pose challenging radiation, power and material constraints to electronics. Traditionally, ASICs deployed in such environments relied on customized digital circuit state machines with limited configurability [2]. In contrast, modern ASICs often adopt microprocessor-based modular System-on-Chips (SoC) architectures, integrating reusable IP blocks interconnected over standardized busses and supported by software programmability and general-purpose computing capabilities. Moving forward, ASIC designers are actively developing fault-tolerant SoCs for future applications in HEP accelerators and experiments.

In this work we present the design and experimental validation of the TriglaV ASIC, a radiation-hardened SoC designed to demonstrate reliable operation of 32-bit microcontroller workloads under demanding HEP radiation environments. TriglaV is the first ASIC built by the toolset of SOCRATES, which is a framework developed at CERN for automated generation of fault-tolerant SoCs [3]. Using the SoCMake build system, the full system is composed out of IP block sources and SystemRDL top-level descriptions of the architecture [4]. Our objectives with TriglaV are centered on demonstrating the viability of the SOCRATES platform to produce fault-tolerant SoC ASICs as well as providing a radiation effects study on SoC circuits manufactured in a commercial 28 nm bulk CMOS technology. The design of TriglaV prioritizes maximum protection against Total-Ionizing-Dose (TID) and Single-Event Effects (SEE), ensuring reliable operation in the intense radiation conditions found in the inner layers of High-Luminosity LHC detectors. We present hardness assurance results from heavy-ion testing as a means to measure SEE sensitivity alongside results from x-ray irradiation to validate TID robustness.

## 2 Design description

### 2.1 System architecture

The architecture of Triglav follows a general scheme for microcontroller SoCs as shown in figure 1. It includes an Ibex RISC-V processor and 64kB of SRAM interfaced through an on-chip low-latency interconnect following the Open Bus Interface (OBI) protocol [5, 6]. The OBI interconnect is implemented as a multi-master-and-slave N x M crossbar. I2C is the main communication interface, but interfaces implementing the Joint Test Action Group (JTAG) for debugging and Universal Asynchronous Receiver-Transmitter (UART) protocols are also included, all of which can be used for programming and configuration of the SoC. An additional peripheral interconnect, the APB-RT,<sup>1</sup> connects common peripheral modules such as GPIO, timers and interrupt control.



**Figure 1.** (Top) Triglav SoC architecture block diagram. (Bottom) Microscope image of the 1x2 mm<sup>2</sup> die with main SoC blocks overlaid.

In order to mitigate SEUs and allow for fine-grained measurements of upset rates, we applied Radiation-Hardening-by-Design (RHBD) to nearly all logic by protection with either Triple Modular Redundancy (TMR) or usage of Error Correcting Codes (ECC). All TMR logic was implemented at the Register Transfer Level (RTL) using TMRG, a tool to generate TMR RTL code out of non-triplicated input code [7]. We opted for a fully protected TMR scheme, where both combinatorial and sequential elements, as well as TMR majority voter circuits, are triplicated. Any voted output of TMR flip-flops is fed back to the input for immediate self-correction within one clock cycle. Any signals or memory content protected by ECC use byte-level Hamming (13,8) encoding.<sup>2</sup> This means up to four errors can be corrected within one 32-bit word, provided they occur in separate bytes.

<sup>1</sup>RT: Radiation-Tolerant.

<sup>2</sup>Hamming (13,8) denotes a code which has a total of 13 bits, of which 8 are data and 5 are parity bits.

The processor is based on the openly-available Ibex RV32IMC design, which we have modified for triplication with TMRG. Of particular concern is the protection of the memories due to their size on the chip and the higher SEU sensitivity of SRAM cells. For this purpose, we developed the Memory Scrubbing and Protection Unit (MSPU) to handle ECC encoding of memory accesses and periodic scrubbing of the SRAM content. Scrubbing is required to prevent the accumulation of SEUs that can lead to uncorrectable multi-bit upsets within one encoded byte. The scrubber reads one 32-bit word per cycle and requires two cycles for a correction. At 250 MHz, one 32 kB memory can be traversed by its scrubber in 33  $\mu$ s. Due to the use of dual-port SRAMs, memory accesses through the OBI can occur concurrently with scrubbing accesses. Regarding the protection of the interconnects, the OBI interface is fully triplicated, whereas in the APB-RT its control signals are triplicated and data and address use ECC. Design verification was conducted with SEU and SET fault simulation on post-layout netlist representation, as well through emulation on FPGA development board. Furthermore, we performed verification of the triplication applied to the Ibex core using formal proofing alongside fault injection [8].

## 2.2 Upset and error observability

The amount of SEUs occurring in TriglaV is monitored using TMR and ECC correction counters. Upset counts are distinguishable between modules due to separate TMR and ECC counters for each distinct module. Uncorrectable ECC errors, that is two bit-flips occurring within a single encoded byte, are counted and identified as Double Errors (DE) and are externally detectable through a dedicated signal. Successful or failed program execution is indicated through a dedicated program return signal and readout of a return value register. Any observed uncorrected upsets leading to failures, such as a program terminating with a failure exit code or timing out, are referred to as Single-Event Functional Interrupts (SEFIs).


## 2.3 Physical design

TriglaV was implemented in a 28 nm bulk CMOS technology. This technology has previously been shown to be robust to the TID levels required for HEP experiments and is therefore considered as one of the target technologies for foreseeable upcoming ASIC projects [9]. The target clock frequency of TriglaV is 250 MHz, which includes an additional margin of 10% on the clock period uncertainty, set to preempt a potential degradation of the delay of the circuit due to the TID-induced threshold voltage shift. Flip-flops belonging to one TMR set were physically spaced by at least 10  $\mu$ m in order to prevent potential uncorrectable Multi-Cell Upsets (MCU) caused by a single particle hit. This spacing was also applied to buffer cells within the triplicated reset tree.


# 3 Experimental validation

## 3.1 Power and area consumption

At the nominal clock frequency of 250 MHz and supply voltage of 0.9 V, the total average power draw of TriglaV is approximately 69 mW or 0.28 mW/MHz. Of this, about 15 mW are consumed by scrubbing of two memories at the maximum rate, while 6 mW are attributable to activity of the triplicated Ibex core. The chip consumes 9.4 mW in the reset state without a running clock, representing the total leakage power of the circuit. The total silicon area of TriglaV is approximately



**Figure 2.** Chip area breakdown for Triglav showing large fractions occupied by SRAM blocks and I/O pad ring.



**Figure 3.** Logic area breakdown showing almost half the area taken up by the two MSPUs and the other half occupied by Ibex CPU, interconnects and periphery combined.

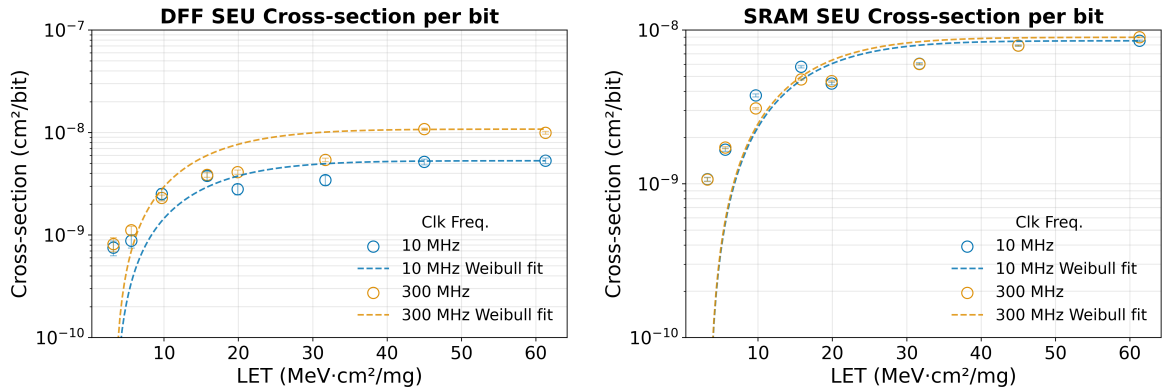
1.6 mm<sup>2</sup>, of which SRAM, I/O pads, and logic respectively occupy about a third of this area as shown in figure 2. As shown in 3, out of the logic area almost half of it is taken up by the two MSPUs with the Ibex core occupying around 20%.

### 3.2 Heavy-ion testing at UCLouvain

In order to measure the SEU sensitivity and test the reliability of Triglav, a heavy-ion testing campaign was conducted at the Heavy-Ion Facility (HIF) of UCLouvain, Belgium. The chip was irradiated with various ion types representing a Linear Energy Transfer (LET) range from 3.2 to 61.3 MeV cm<sup>2</sup>/mg. A maximum flux of  $1.5 \times 10^4 \text{ s}^{-1} \text{ cm}^{-2}$  was used for all ions alongside a minimum particle fluence of  $5 \times 10^6 / \text{cm}^2$  per ion type. The corrected SEU cross-section accumulated across all attempted ions and normalized across the total flip-flop count in Triglav is shown in figure 4(a). This cross-section result is of the same order of magnitude as previously reported results obtained for flip-flops of the same 28 nm technology [10], indicating that, for this circuit, Single-Event Transients (SETs) occurring throughout combinatorial logic are not contributing significantly to the overall rate of SEEs. Figure 4(b) shows the cross-section for corrected ECC bytes in the SRAM memory. It was previously shown that for the SRAM cells the chance for MCU during heavy-ion testing goes above 50% at an LET of 5 MeV cm<sup>2</sup>/mg [10]. For the SRAM memories of Triglav, the MSPU counts corrected upsets during scrubbing without distinguishing between individual particle hits.

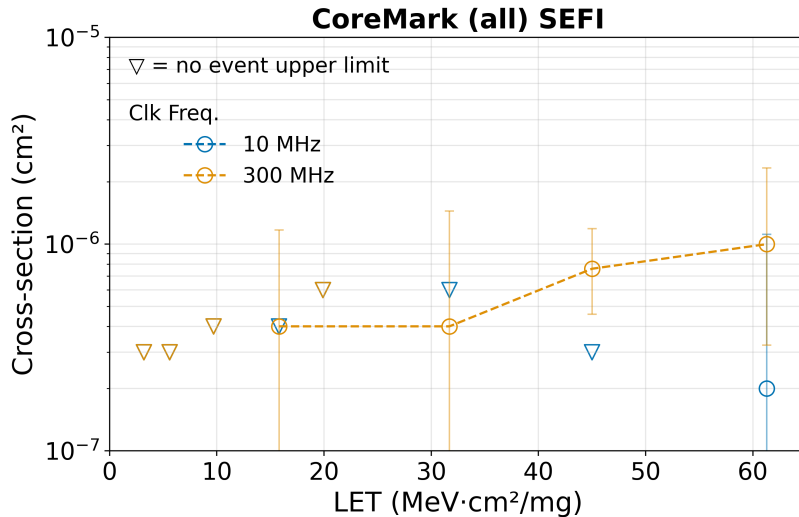
Despite the extensive protections against SEUs and design verification applied prior to the tapeout, uncorrected upsets leading to SEFIs, i.e., program timeouts, memory corruptions or exceptions caught by the CPU, were also observed. The SEFI cross-section measured during continuous execution of the CoreMark benchmark is depicted in figure 5 [12]. Out of 47 total recorded SEFIs, 11 are caused by unintentional resets, in 8 cases recovery was possible through software exception handling, and at least 9 required a reboot following a memory corruption. In the 19 remaining cases, a reset was needed to resume normal execution.

We are currently investigating if the cause for these SEFIs can be attributed to several known vulnerabilities in the layout of the design: firstly, the foundry dual-port SRAM blocks are not internally



(a) Per-bit DFF cross-section shows a small dependency on the clock frequency. (b) Per-bit SRAM cross-section with no observable frequency dependency.

**Figure 4.** Cross-section of corrected SEUs measured with TMR error and ECC counters during heavy-ion irradiation.



**Figure 5.** SEFI cross-section during CoreMark benchmark (47 total SEFIs). CoreMark execution time set to 5 s for all frequencies. 95% confidence intervals are shown and upper limit for zero events is marked by ∇. A generally higher cross-section for SEFIs was observed at higher LET and 300 MHz operation.

hardened and an increased SEFI rate was observed when both ports are active simultaneously compared to just one port being accessed. Secondly, closely spaced combinatorial logic belonging to one TMR set could be upset by a single particle hit causing a Single-Event Multi-Transients (SEMT) and thereby bypassing the TMR majority voting. Several such locations were identified within the layout. Lastly, an SET sensitivity in the IO ring power-on-control cell is known to cause unintended assertions of the reset signal.

Following a cross-section measurement, a rate of occurrence for a given HEP radiation environment can be predicted [11]. We estimate the upper-bound for the SEFI rate in TriglaV to be less than 1 per day for a worst-case proton flux environment such as inner layers of HL-LHC detectors. Most applications in HEP will experience a less severe flux rate however.

### 3.3 X-ray irradiation at room temperature

We performed a TID characterization of TriglaV at a high-dose rate of 10 Mrad/h and room temperature. This irradiation was performed using the x-ray irradiation facility at CERN EP-ESE, Geneva. The final deposited dose was 1 Grad, with no loss of functionality observed across operating voltages of 0.8, 0.9, and 1 V and clock frequencies up to 300 MHz. The leakage power was measured to almost double at 1 Grad, increasing from 9.4 mW to 18 mW. While no failures of the chip were observed during this test, for the 28 nm node, dose-rate dependent effects may lead to underestimation of transistor degradation and therefore elevated temperature testing is recommended for high-dose rate irradiation [13].

## 4 Conclusions

In this work, we presented the design and demonstrated the functionality of TriglaV following a 28 nm design flow driven by the SOCRATES platform. The efficacy of applied protection measures against SEU and TID effects were validated through experimental irradiation. TID testing has shown robustness beyond 1 Grad, which is sufficient for most long-term LHC applications. Further TID irradiation tests at elevated temperatures need to be conducted in view of dose-rate dependent effects. While a majority of induced upsets were corrected during heavy-ion testing, a small number of SEFIs recorded show that there are still vulnerabilities left in the design. The root causes for these need to be further validated, e.g. by SEMT simulations and laser testing. Nevertheless, it is estimated that the observed SEFIs would occur rarely, even for the most severe radiation environments.

## Acknowledgments

This activity has received funding from the European Union's 2020 research and innovation programme under grant agreement No 101008126, corresponding to the RADNEXT project.

## References

- [1] ECFA Detectors R&D Roadmap Process Group, *The 2021 ECFA detector research and development roadmap*, [CERN-ESU-017](#) (2021).
- [2] F. Faccio, *ASIC survival in the radiation environment of the LHC experiments: 30 years of struggle and still tantalizing*, *Nucl. Instrum. Meth. A* **1045** (2023) 167569.
- [3] M. Andorno et al., *SOCRATES: a radiation-tolerant SoC generator framework*, [2025 JINST 20 C02017](#).
- [4] R. Pejašinović et al., *Hardware and software build flow with SoCMake*, [arXiv:2502.02065](#).
- [5] lowRISC, *Ibex RV32 Core*, <https://github.com/lowRISC/ibex?tab=readme-ov-file>.
- [6] OpenHW Group, *OpenBus Interface v1.6.0*, <https://github.com/openhwgroup/obi/blob/main/OBI-v1.6.0.pdf>.
- [7] S. Kulis, *Single Event Effects mitigation with TMRG tool*, [2017 JINST 12 C01082](#).
- [8] A. Pulli and M. Lupi, *A simulation methodology for verification of transient fault tolerance of ASICs designed for high-energy physics experiments*, [2023 JINST 18 C01038](#).
- [9] G. Borghello et al., *Total ionizing dose effects on ring-oscillators and SRAMs in a commercial 28 nm CMOS technology*, [2023 JINST 18 C02003](#).

- [10] G. Borghello et al, *Single Event Effects characterization of a commercial 28 nm CMOS technology*, in the proceedings of the *Topical Workshop on Electronics for Particle Physics*, Geremeas, Sardinia, Italy, 01–06 October 2023 [<https://indico.cern.ch/event/1255624/contributions/5443894/>].
- [11] M. Huhtinen and F. Faccio, *Computational method to estimate single event upset rates in an accelerator environment*, *Nucl. Instrum. Meth. A* **450** (2000) 155.
- [12] EEMBC, *CoreMark*, <https://www.eembc.org/coremark/>.
- [13] G. Borghello et al., *ELDRS in a Commercial 28-nm CMOS Technology*, *IEEE Trans. Nucl. Sci.* **72** (2025) 2276.