

Titles &  
Abstract

2024-2025



**ELYSIUMPRO**  
INSPIRING THE LEADING EDGE TECHNOLOGIES


Advanced Academic Final Year Projects

**VLSI**

**IEEE PROJECTS**

[elysiumpro.in](http://elysiumpro.in)





**EPRO-VLSI-001**

**Low Complexity VLSI Architecture for OTFS Transceiver under Multipath Fading Channel**

Orthogonal Time Frequency Space (OTFS) modulation has established itself as a dependable protocol for high-speed vehicular communication. This pioneering technique operates within a novel two-dimensional (2D) delay-Doppler domain waveform. When juxtaposed with conventional modulation methods like orthogonal frequency division multiplexing (OFDM), OTFS demonstrates superior performance enhancements in scenarios involving rapidly moving wireless channels. This paper begins by initially unveiling the input-output association of the OTFS signal within the delay-time domain. A comprehensive comparison with the established OFDM waveform highlights the potential of OTFS for achieving notably lower bit error rate (BER) under various conditions, which has been obtained by using the Minimum Mean Square Equalizer (MMSE) equalization technique. Finally, we have proposed novel and low complexity Very-large-scale integration (VLSI) architecture for the OTFS transmitter and the receiver by using the LU decomposition technique for the first time in state-of-the-art.


**High-Throughput Bilinear Pairing Processor for Server-Side FPGA Applications**

**EPRO-VLSI-002**

The goal of this work is to speed up server-side cryptographic pairing procedures on field-programmable gate arrays (FPGAs). Prior research on FPGA pairing implementations concentrated on embedded device efficiency efficiency, aiming to maximise performance with the least amount of circuit resources. For server-side applications, where the main goal is maximum performance after FPGA resources are depleted, these topologies are probably inefficient. Their low operation frequency and low digital signal processor (DSP) utilisation make their architectures inefficient. In this work, we fully use DSPs to offer a high-throughput pairing processor architecture for server-side FPGAs. First, we provide a server-side FPGA-compatible loop-unrolled modular multiplication method. When compared, the algorithm displays the maximum throughput and efficiency efficiency. Secondly, we create a pairing processor architecture that incorporates the suggested modular multiplier, enabling redundant adders and interspersed executions to sustain its high throughput. The findings of our test of the BN254 and BLS12\_381 pairs on the suggested processor architecture indicate that it achieves good throughput, which is roughly two and five times quicker than that of earlier studies, respectively.







**EPRO-VLSI-003**

## Low Complexity VLSI Architecture for OTFS Transceiver under Multipath Fading Channel


Orthogonal time frequency space (OTFS) modulation has established itself as a dependable protocol for high-speed vehicular communication. This pioneering technique operates within a novel 2-D delay-Doppler domain waveform. When compared with conventional modulation methods like orthogonal frequency-division multiplexing (OFDM), OTFS demonstrates superior performance enhancements in scenarios involving rapidly moving wireless channels. This article begins by initially unveiling the input-output association of the OTFS signal within the delay-time domain. A comprehensive comparison with the established OFDM waveform highlights the potential of OTFS for achieving a notably lower bit error rate (BER) under various conditions, which has been obtained by using the minimum mean square equalizer (MMSE) equalization technique. Finally, we have proposed a novel and low-complexity VLSI architecture for the OTFS transmitter and the receiver by using the lower-upper (LU) decomposition technique for the first time in the literature.

## Stream Processing Architectures for Continuous ECG Monitoring Using SubsamplingBased Classifiers

Monitoring of biomedical data, such as electrocardiogram (ECG) signals, requires accelerators, which can process data streams in a continuous manner. Especially, wearable monitoring systems require both ultralow power consumption and sufficiently complex deep neural network (DNN) classifiers to identify asymptomatic and critical health conditions, such as atrial fibrillation (AF). Such continuous data streams pose unique constraints on the processing pipeline for classification systems, which can be addressed in the design methodology of application-specific integrated circuits (ASICs). In this work, we identify specific constraints to define common operating conditions, which guide the design of ECG accelerators in an algorithm-hardware codesign methodology. In specific, we show that the input frame size and the number of classifications per time frame play a significant role for the computational complexity (CC) of the classifier, as well as the ECG accelerator executing the classifier in a continuous manner. As an example, the constraints are applied in a top-down algorithm-hardware codesign flow.

**EPRO-VLSI-004**





**EPRO-VLSI-005**

**FELIX: FPGA-Based Scalable and Lightweight Accelerator for Large Integer Extended GCD**

The extended greatest common divisor (XGCD) computation is a critical component in various cryptographic applications and algorithms, including both pre- and postquantum cryptosystems. In addition to computing the greatest common divisor (GCD) of two integers, the XGCD also produces Bézout coefficients  $ba$  and  $bb$  which satisfy  $GCD(a, b) = a * ba + b * bb$ . In particular, computing the XGCD for large integers is of significant interest. Most recently, XGCD computation between 6479-bit integers is required for solving Nth-degree truncated polynomial ring unit (NTRU) trapdoors in FALCON, a National Institute of Standards and Technology (NIST)-selected postquantum digital signature scheme. To this point, existing literature has primarily focused on exploring software-based implementations for XGCD. The few existing high-performance hardware architectures require significant hardware resources and may not be desirable for practical usage, and the lightweight architectures suffer from poor performance.


**Enabling HW-Based Task Scheduling in Large Multicore Architectures**

Dynamic Task Scheduling is an enticing programming model aiming to ease the development of parallel programs with intrinsically irregular or data-dependent parallelism. The performance of such solutions relies on the ability of the Task Scheduling HW/SW stack to efficiently evaluate dependencies at runtime and schedule work to available cores. Traditional SW-only systems implicate scheduling overheads of around 30K processor cycles per task, which severely limit the ( core count , task granularity ) combinations that they might adequately handle. Previous work on HW-accelerated Task Scheduling has shown that such systems might support high performance scheduling on processors with up to eight cores, but questions remained regarding the viability of such solutions to support the greater number of cores now frequently found in high-end SMP systems. The present work presents an FPGA-proven, tightly-integrated, Linux-capable, 30-core RISC-V system with hardware accelerated Task Scheduling. We use this implementation to show that HW Task Scheduling can still offer competitive performance at such high core count, and describe how this organization includes hardware and software optimizations that make it even more scalable than previous solutions.

**EPRO-VLSI-006**







**EPRO-VLSI-007**

## Noise Analysis and Design Methodology of Chopper Amplifiers With Analog


Biopotential acquisition chopper instrumentation amplifiers require a dc-servo loop (DSL) in order to filter electrode dc offsets. However, the noise performance degradation due to the addition of the DSL is often overlooked despite that it can be very detrimental at the frequencies of interest. This article presents an in-depth noise analysis of biopotential acquisition chopper instrumentation amplifiers with analog DSLs. Analytical expressions that predict the noise of different DSL implementations are found and a design flow to minimize their noise contribution is proposed. The design methodology is demonstrated with example circuits targeting biopotential recording systems. These circuits are implemented using a standard 180 nm CMOS technology, and their performance is verified through postlayout simulations. The findings of this work provide a comprehensive understanding of the noise characteristics of a DSL, its impact on noise performance, and design strategies for noise optimization.

## An Open-Source Tool to Model and Explore Complex Routing Architecture for FPGA

—As the benefits of Moore’s Law diminish, computing performance and efficiency gains are increasingly achieved through specializing hardware to a domain of computation. However, this limits the hardware’s generality and flexibility. Field Programmable Gate Arrays (FPGAs), microchips which can be re-programmed to implement arbitrary digital circuits, enable the benefits of specialization while remaining flexible. A challenge to using FPGAs is the complex computer aided design flow required to efficiently map a computation onto an FPGA. Traditionally these design flows are closed-source and highly specialized to a particular vendor’s devices. We propose an alternate data-driven approach which uses highly adaptable and re-targetable open-source tools to target both commercial and research FPGA architectures. While challenges remain, we believe this approach makes the development of novel and commercial FPGA architectures faster and more accessible. Furthermore, it provides a path forward for industry, academia, and the open-source community to collaborate and combine their resources to advance FPGA technology.

**EPRO-VLSI-008**





**EPRO-VLSI-009**

## MaliGNNoma: GNN-Based Malicious Circuit Classifier for Secure Cloud FPGAs

Detecting such threats before loading onto the FPGA is crucial, but existing methods face difficulty identifying sophisticated attacks. We present MaliGNNoma, a machine learning-based solution that accurately identifies malicious FPGA configurations. Serving as a netlist scanning mechanism, it can be employed by cloud service providers as an initial security layer within a necessary multi-tiered security system. By leveraging the inherent graph representation of FPGA netlists, MaliGNNoma employs a graph neural network (GNN) to learn distinctive malicious features, surpassing current approaches. To enhance transparency, MaliGNNoma utilizes a parameterized explainer for the GNN, labeling the FPGA configuration and pinpointing the sub-circuit responsible for the malicious classification. Through extensive experimentation on the ZCU102 board with a Xilinx UltraScale+ FPGA, we validate the effectiveness of MaliGNNoma in detecting malicious configurations, including sophisticated attacks, such as those based on benign modules, like cryptography accelerators.


## HUB Meets Posit: Arithmetic Units Implementation

The posit™ format was introduced in 2017 as an alternative to replacing the widespread IEEE 754. Posit arithmetic provides reproducible results across platforms and possesses tapered accuracy, among other improvements. Nevertheless, despite the advantages provided by such a format, their functional units are not as competitive as the IEEE 754 ones yet. The HUB approach was presented in 2016 to reduce the hardware cost of floating-point units. In this brief, we present HUB posit, a new format to mitigate the hardware overhead of posit units. Results show that it is possible to reach up to 15% and 12% in terms of area-delay product for adders and multipliers, respectively, while maintaining a similar level of accuracy. In addition, synthesis results show that HUB posit units are able to reach higher frequencies than conventional ones.

**EPRO-VLSI-010**







## MaliGNNoma: GNN-Based Malicious Circuit Classifier for Secure Cloud FPGAs

Transport triggered architectures (TTAs) follow the static programming model of very long instruction word (VLIW) processors but expose additional information of the processor datapath in the programming interface, which enables low-level code optimizations but results in lower code density. Multi-instruction-set architectures add flexibility via their ability to switch instruction sets during execution. The added flexibility is interesting for VLIW-style processors because it enables reducing the large instruction stream energy footprint by using an instruction set with enhanced code density in regions with limited opportunities for exploitation of instruction level parallelism. In this article, we introduce a dual instruction-set architecture, "Dual-IS", that implements both RISC-V and TTA instruction sets with shared datapath resources by means of a lightweight microcode unit. In order to utilize the flexible architecture automatically, we introduce a compilation method that is able to independently target code for both instruction sets based on static code analysis and a microarchitectural model of the processor. Compared to a single-ISA TTA processor, we were able to lower the instruction stream energy consumption 45% on average in the best design point, which resulted in a total energy consumption reduction of 26% and a 0.4% lower run time.





**2.1L+** Customer Satisfied | **27+** Specialized Domains | **107+** Campus Tie-ups | **24/7** Support Center | **57+** Countries Covered | **2600+** Journal Access

📞 99447 93398

✉ info@elysiumpro.in

🌐 elysiumpro.in

📍 229, First Floor, A Block,  
Elysium Campus,  
Church Rd, Anna Nagar,  
Madurai, Tamil Nadu -  
625020.

**Delivery Centres**

Chennai | Coimbatore | Tirunelveli | Virudhunagar | Trichy

