

VLSI

Python | IEEE Projects

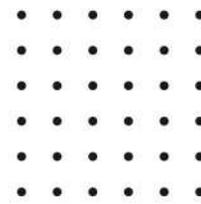


Advanced Academic Final Year Projects

Titles & Abstract
2025 - 2026

elysiumpro.in





Reverse-Engineering Optimization Techniques of High-Level Synthesis: Simulation Insights into Accelerating Applications with AMD-Xilinx Vitis

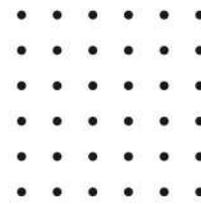
PROJ24027

Modern AI applications often include computationally intensive components that benefit from hardware acceleration. Tools like AMD Vitis HLS enable the creation of custom hardware designs through high-level synthesis (HLS) and hardware optimizations. However, mastering these tools is challenging due to the complexity of optimization techniques and limited understanding of their interactions and hardware impact. This article presents a quantitative analysis of Vitis HLS optimization directives by reverse engineering their behavior. Over 150 experiments were conducted to study three main objectives: evaluating pragma behavior and application rules, modeling latency estimates provided by Vitis HLS, and assessing the effects of optimizations on design space exploration, particularly in terms of area and latency. Experiments explored various combinations and placements of optimizations within loop and function hierarchies in the test bench. The results provide practical guidance for effective use of Vitis pragmas and highlight configurations that optimize latency and area, aiding developers in navigating Vitis HLS optimizations more effectively.

High Efficiency Multiply-Accumulator Using Ternary Logic and Ternary Approximate Algorithm

PROJ24047

The research presents a novel ternary multiply-accumulator (MAC) unit designed to enhance computational efficiency, particularly in applications like neural networks. This work addresses the limited progress in ternary-based vector processing despite ternary logic's higher information density compared to binary systems. The unit incorporates ternary approximate algorithms that reduce power consumption by 30% with only a 2% computation error, along with sophisticated circuits achieving a remarkable 74-80% lower power-delay-product (PDP). Evaluated using both CNTFET and 180 nm CMOS processes, the ternary MAC unit demonstrates significant advantages over its binary counterpart, including a 45% reduction in area and a 30% decrease in power consumption. These findings highlight the strong potential of ternary logic for developing more compact and energy-efficient hardware.



Optimization of Chirality Variation in Carbon Nanotube Field Effect Transistor Spiking Neurons

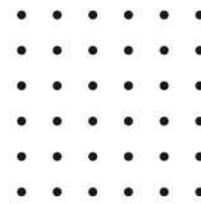
PROJ24115

Carbon Nanotube Field Effect Transistors (CNFETs) are emerging as a potential successor to traditional silicon-based CMOS technology, which is reaching its scaling limits due to Short-Channel Effects (SCEs). These effects arise when CMOS device dimensions shrink to the point where the channel length is comparable to the depletion layer widths, hindering further performance improvements. CNFETs, with their superior electrical properties like ballistic transport, offer a promising solution to overcome these challenges. This research presents the first known study investigating chirality variation in CNFETs for optimizing spiking neurons. Specifically, the study uses the Penta-Transistor Integrate & Fire (PTIF) architecture to achieve two distinct goals: a 6.89x increase in spiking frequency and an 87.43% energy saving, demonstrating CNFETs' remarkable potential for high-performance and low-power neuromorphic computing applications.

Tunable Energy-Efficient Approximate Circuits for Self-Powered AI and Autonomous Edge Computing Systems

PROJ24116

To meet the demands of deploying computationally intensive AI/ML models on energy-constrained edge devices, this research proposes approximate compressors for designing energy-efficient Multiply and Accumulate (MAC) units in Deep and Convolutional Neural Networks (DNN/CNN). The core idea is to reduce hardware complexity by trading off a small amount of computational accuracy. A key innovation is the introduction of tunable compressors and MAC units, which can switch between two approximation modes to allow for runtime adjustments of energy efficiency and accuracy. Validated at 7nm and 55nm technology nodes, the proposed tunable compressor shows an average 49% reduction in energy consumption and a 30% reduction in delay compared to state-of-the-art designs. The complete MAC unit also demonstrates significant gains, with a 36% average energy reduction and 18% delay reduction. The minimal accuracy trade-off is confirmed by a low MRED (mean relative error distance) of 0.19. When integrated into a simple CNN, the solution achieves 15% lower power consumption and 7% lower area overhead without any loss of functionality.



An Ultra-Low-Power Fully-Static Contention-Free Single-Phase-Clock Flip-Flop With Low Area

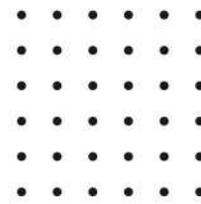
PROJ24117

This paper introduces a novel Fully-Static, Contention-Free Single-Phase-Clock Flip-Flop specifically designed for IoT devices. The architecture's key features—the elimination of invalid toggling, floating nodes, and contention paths—significantly enhance its robustness and minimize power consumption. This is further optimized by using logic merging and topology compression to achieve a smaller area. Implemented in a 28nm process, the design shows remarkable power savings. Post-simulation results reveal that at 0.9V and 1GHz, it achieves up to an 86.95% dynamic power reduction compared to traditional MSFF. It also delivers substantial static power savings, with up to a 98.91% reduction. Furthermore, a 10K Monte Carlo simulation confirmed 100% functional integrity from 0.4V to 0.9V, highlighting its reliability and efficiency for ultra-low-power applications.

Special Issue on Defect and Fault Tolerance in VLSI and Nanotechnology Systems (DFTS 2023) in the IEEE Transactions on Device and Materials Reliability

PROJ24118

The ten articles in this special issue present innovative research in the field of defect and fault tolerance in VLSI and nanotechnology systems and provide readers with valuable insights into the latest advances and future trends in these challenging research areas. The focus of these articles is on the reliability in the design, technology and testing of electronic devices and systems, integrated circuits, printed modules, as well as methodologies and tools used for reliability and security prediction, verification and design validation.



A 65-nm CMOS Downconverter-Less Clock Generator Architecture Using Voltage Stacking of Oscillator and Frequency Dividers for Scaling-Friendly IoTs

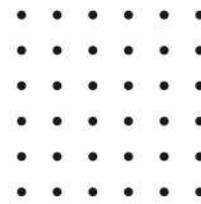
PROJ24119

This study introduces a novel CMOS clock generator designed to meet the high energy efficiency demands of IoT devices by eliminating the need for power-hungry, scaling-unfriendly step-down converters. The prototype chip, fabricated in a 65-nm CMOS process, employs a unique architecture based on voltage stacking and charge recycling. This is achieved through a stacked oscillator and a series of frequency dividers. Two distinct configurations were developed. The first prioritizes low power (LP) consumption, generating a 2.09 Hz clock with an exceptionally low power of 0.22 nW, which is the lowest reported power for a sub-10-Hz clock at a nominal voltage. The second configuration is optimized for ultra-low frequencies, achieving an output of 0.079 Hz. This architecture offers a scaling-friendly solution with significant potential for advanced technology nodes.

Rule-Based Reinforcement Learning on FPGA for QoS-Aware Dynamic Frequency Scaling

PROJ24120

Modern processors in multiprocessor system-on-chips (MPSoCs) have built-in hardware features to handle submillisecond load variations, but traditional software-based dynamic frequency scaling (DFS) governors are too slow to leverage them effectively. To address this, this work proposes a novel hardware reinforcement learning (RL) agent, augmented with preemptive shielding and eligibility traces, designed to optimize the execution of deadline-bound quality-of-service (QoS) tasks in mixed-critical environments. The algorithm's features are demonstrated through a hardware-in-the-loop simulation using LLVM's single-source benchmarks on SparcV8 processors. The research also presents a practical field-programmable gate array (FPGA) implementation, which achieves optimized resource usage and timing performance through quantization and approximation, offering a significant advancement in real-time system performance optimization.



Emerging Technologies in CMOS Integrated Sensing System-On-Chip

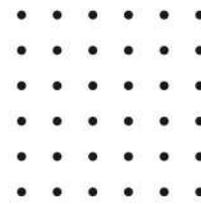
PROJ24121

This paper provides a comprehensive review of the evolution of sensing technologies toward CMOS integrated sensing System-on-Chip (SoC), which consolidate sensors, signal processing, and communication onto a single chip. These systems offer compact, low-power, and cost-effective solutions for various applications, including wearable health monitoring and industrial automation. The review details the architectural components of these SoCs, such as sensing modalities and signal conditioning circuits, while also examining key design challenges like power efficiency and miniaturization. The paper also explores current trends, including the integration of wireless communication, energy harvesting, and on-chip AI. Finally, it discusses future directions, such as advances in semiconductor manufacturing and the use of sustainable materials, emphasizing the importance of collaborative innovation to drive the next generation of intelligent, adaptable, and energy-efficient sensing solutions.

Hardware-Software Stitching Algorithm in Lightweight Q-Learning System on Chip (SoC) for Shortest Path Optimization

PROJ24122

This paper presents a novel hardware-software co-design to accelerate Q-learning algorithms using a RISC-V-based System-on-Chip (SoC). A key innovation is the maze-stitching algorithm, which decomposes large, complex mazes into smaller sub-mazes, allowing for efficient computation on a low-complexity hardware accelerator. The proposed system, combining a 64-bit RISC-V core with a Q-Learning accelerator on an Arty A7-100T FPGA operating at 50MHz, demonstrates significant performance improvements. The hardware-accelerated approach achieves speedups ranging from 84× to 233× over traditional software, while the algorithm alone provides a 13× to 36× speedup. This maze-stitching technique ensures the solution is scalable to larger problem sizes while maintaining hardware efficiency, making it a viable and low-footprint solution for resource-constrained edge computing environments.



A Novel Algorithm for Aspect Ratio Estimation in SRAM Design to Achieve High SNM, High Speed, and Low Leakage Power

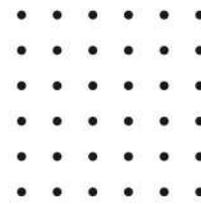
PROJ24123

This paper introduces a novel algorithm for optimizing transistor sizing in static random-access memory (SRAM) to enhance speed, improve Static Noise Margin (SNM), and reduce leakage power consumption. The SRAM, designed with 45 nm technology and operating at a 1.2 V supply, was validated through extensive Monte Carlo simulations. The results demonstrate impressive speed, with read and write access times as low as 9.97 ps and 12.00 ps, respectively. The design also exhibits robust stability, with SNM values of 328.2 mV (read), 453.7 mV (write), and 452.3 mV (hold). A compact layout, achieved by including precharge and write driver circuits, results in a total area of 9.79 μm^2 , with the cell occupying 4.1 μm^2 . Finally, the proposed design achieves an exceptionally low leakage power of 1.64 pW, confirming the efficiency and performance benefits of the optimized sizing approach.

Ultra-low power and 1.5 bit/cell ternary-SRAM stability modeling for always-on applications

PROJ24124

We present an ultra-low power ternary SRAM (T-SRAM) with a storage capacity of 1.5 bit/cell, using a commercial 110-nm CMOS foundry for always-on applications, along with an analysis of its stability. By designing T-CMOS with SPICE compact model parameters, which are body-effect coefficient (m), peak electric field coefficient (CEP), and gate width (W), band-to-band tunneling current (IBTBT) can be reduced to hundreds of fA range and it allows VDD to scale down to 0.55 V. Finally, we experimentally demonstrate T-SRAM cell which static and dynamic powers are decreased to 4.5×10^{-2} and 1.3×10^{-7} , respectively.



A Review on Sub-0.21-V Ultra-Low-Supply-Voltage Analog-to-Digital Converters

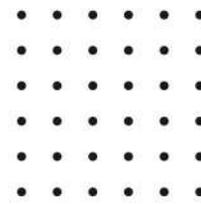
PROJ24125

Ultra-low-supply-voltage (ULV) analog-to-digital converters (ADCs) operating at 0.21 V or lower are attractive for Internet-of-Things (IoT) and embedded applications due to their extremely low power consumption. This paper surveys state-of-the-art ULV ADCs to evaluate current trends and design strategies. Architectures, circuit implementations, and calibration techniques are analyzed and key trends are identified. Based on the observations, the paper provides recommendations for the circuit designer to make judicious design choices to obtain the desired performance for ULV ADCs. This paper further explores the VCO-based architecture and proposes a new topology to achieve high resolution for ULV ADCs.

Biologically-Inspired, Ultra-Low Power, and High-Speed Integrate-and-Fire Neuron Circuit With Stochastic Behavior Using Nanoscale Side-Contacted Field Effect Diode Technology

PROJ24126

This research introduces a novel, biologically inspired neuron circuit designed to advance neuromorphic computing. The circuit, a leaky integrate-and-fire (LIF) neuron, leverages nanoscale side-contacted field effect diodes (S-FEDs) to achieve key biological characteristics. A primary innovation is its ability to mimic the stochastic behavior of biological neurons in a tunable manner. This design demonstrates significant performance improvements over existing designs, boasting ultra-low power consumption and high-speed operation. The circuit minimizes the energy per spike, making it exceptionally efficient. By combining these attributes, the research provides a powerful and practical foundation for creating more efficient and realistic brain-inspired computing systems, potentially leading to more advanced artificial intelligence hardware that can process information with greater fidelity to biological processes.



ApproXAI: Energy-Efficient Hardware Acceleration of Explainable AI using Approximate Computing

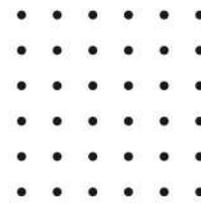
PROJ24127

This paper introduces XALedge, a novel framework designed to overcome the energy inefficiency of real-time Explainable AI (XAI). While traditional XAI methods are computationally intensive and slow, existing hardware accelerators fail to fully address the power consumption issue on edge devices. XALedge solves this by leveraging approximate computing techniques applied to core XAI algorithms like integrated gradients and Shapley analysis. It translates these processes into approximate matrix computations, exploiting the synergy between convolution, Fourier transform, and approximation paradigms. This approach facilitates efficient hardware acceleration on TPU-based edge devices. The comprehensive evaluation shows that XALedge achieves a 2x improvement in energy efficiency compared to current accurate XAI hardware accelerators, all while maintaining comparable accuracy. This advancement holds significant potential for deploying explainable AI in energy-constrained, real-time applications.

Very-Large-Scale Integration (VLSI) Implementation and Performance Comparison of Multiplier Topologies for Fixed- and Floating-Point Numbers

PROJ24128

This study performs a VLSI design and performance comparison of three multiplier types—array, Wallace tree, and radix-4 Booth—to address the need for energy-efficient, high-performance multipliers in portable devices. The circuits were designed using Alliance open-source tools in a 350 nm process technology, with an emphasis on single- and double-precision floating-point numbers. The findings indicate that the array multiplier consistently had the highest delay and largest area across all configurations. The Wallace multiplier emerged as the most efficient, exhibiting the lowest delay for single-precision mantissa multiplication and utilizing the smallest area for both single- and double-precision floating-point designs. While its delay advantage was less pronounced in double-precision numbers, the Wallace multiplier's overall performance and compact footprint make it the superior choice for designing high-performance arithmetic circuits.



VLSI Architecture for FIR Filter using Radix-4 Booth Multiplier and CBL Adder

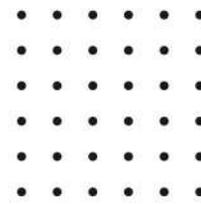
PROJ24129

This research focuses on designing high-speed architectures for 16-tap and 32-tap FIR filters, which are crucial for audio and video signal processing. The primary objective is to improve speed by utilizing a radix-4 Booth multiplier (BM) based on a partition multiplier. The Booth algorithm, originally developed for multiplying binary numbers in two's complement notation through repeated addition, is used here in its more efficient radix-4 form. By leveraging this modified algorithm, the paper aims to create an optimized FIR filter design. The proposed architecture is implemented using the Vertex device family software from Xilinx, demonstrating a practical approach to enhancing the performance of these essential digital filters. This method provides a clear solution for improving computation speed by addressing the multiplication bottlenecks in filter coefficient processing.

Efficient Low-Power VLSI Design of a High-Speed Baugh-Wooley Multiplier Using Verilog HDL

PROJ24130

The rising demand for portable systems and ultra-high-density chips has made low-power optimization a critical concern in high-performance digital systems. Since multiplication is a fundamental operation in many digital signal processing (DSP) algorithms, developing power-efficient multipliers is crucial for VLSI system design. This project focuses on the implementation of a high-speed, low-power multiplier using the shift-and-add method of the Baugh-Wooley multiplier. The study delves into the design and implementation of these multipliers, specifically introducing a Modified Baugh-Wooley architecture aimed at achieving minimal area, power, and delay. The design was developed using Verilog HDL and was simulated and synthesized using the Xilinx ISE tool, demonstrating a practical approach to creating efficient arithmetic circuits for modern computing applications.



Finite Impulse Response using Different Multiplier and Adder

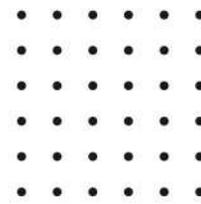
PROJ24131

In the design of DSP processors, the primary objectives are area optimization and power reduction. The fundamental component for these processors is the Finite Impulse Response (FIR) Filter, which is composed of adders, flip-flops, and multipliers. The performance of the FIR filter is most significantly impacted by the multiplier, which is the slowest of these modules. To address this bottleneck, this paper studies different types of adders and focuses on the Booth's multiplication algorithm. The Booth algorithm is particularly useful for computer architecture as it efficiently multiplies signed binary numbers in two's complement notation without requiring a separate correction step. By analyzing different implementations of adders and the Booth multiplier, the paper aims to provide insights for designing more efficient and high-performance FIR filters, thereby improving the overall DSP processor.

Acceleration of Timing-Aware Gate-Level Logic Simulation Through One-Pass GPU Parallelism

PROJ24132

This paper introduces a novel approach for accelerating VLSI circuit simulation using GPU parallelism, addressing the limitations of conventional parallel strategies. The research focuses on 4-value logic timing-aware gate-level circuits, proposing a waveform-based GPU parallelism method. The key innovation is an algorithm that manages task dependencies in combinational circuits to significantly reduce the need for CPU-GPU synchronization. This enables one-pass parallelism, requiring only a single data transfer between the CPU and GPU. To support this strategy, the authors developed optimized data structures for dynamic allocation of outputs with uncertain scales. Experimental results on industrial-scale benchmarks demonstrate that this approach offers substantial performance gains over existing state-of-the-art baselines, highlighting its potential for meeting the demands of modern chip design complexity.



Analog VLSI Implementation of Subthreshold Spiking Neural Networks and Its Application to Reservoir Computing

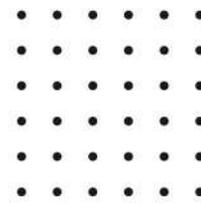
PROJ24133

To enhance the energy efficiency of neuromorphic computing, this study introduces a fully analog two-variable spiking neuron and spiking neural network (SNN) circuit. Taking advantage of the physical properties of transistors in their subthreshold region, the design achieves a remarkably low energy consumption of just 22.7 fJ/spike. The circuits were implemented on an analog VLSI chip using a 0.18 μm CMOS process, with measurements confirming their ability to exhibit complex spike dynamics at low power. The research successfully applied the generated spike sequence to a spoken digit recognition task using a reservoir computing framework, achieving an impressive efficiency of 14.4 fJ/SOP. These results underscore the potential of this SNN-based hardware to significantly advance edge AI applications by providing a highly energy-efficient solution for processing time-series data.

A Very Large-Scale Integration (VLSI) Chip Design for Abnormal Heartbeat Detection Using a Data-Shifting Neural Network (DSNN)

PROJ24134

This paper introduces a data-shifting neural network (DSNN) for the highly accurate detection of six types of abnormal heartbeats from ECG signals. To enhance detection accuracy, the DSNN doubles the input signal using a data shifting scheme, effectively providing more information for training. While this approach doubles the computational time, it substantially improves the detection rate. The proposed DSNN chip was implemented using the TSMC 0.18- μm CMOS process, operating at 20 MHz with a compact chip area of 0.619 mm^2 and a maximum power dissipation of 0.75 mW. When tested on the MIT-BIH arrhythmia database, the chip achieved a high detection rate of 97.17%. The combination of high accuracy and a small chip footprint suggests that this DSNN is well-suited for integration into wearable or portable healthcare devices.



Low-Complexity VLSI Architecture for OTFS Transceiver Under Multipath Fading Channel

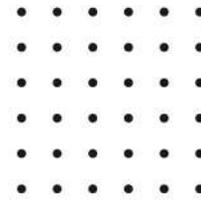
PROJ24135

This paper presents a novel and low-complexity VLSI architecture for the Orthogonal Time Frequency Space (OTFS) transmitter and receiver, a modulation technique superior to OFDM for high-speed vehicular communication. Operating in a 2-D delay-Doppler domain, OTFS demonstrates a significantly lower bit error rate (BER) using an MMSE equalization technique. The proposed architecture leverages the LU decomposition technique for the first time. Performance comparisons show that the new transmitter design is 7.394% faster than existing work while using 89.354% fewer LUTs and 79.984% fewer FFs, indicating a highly optimized design in both latency and resource utilization. Furthermore, this study is the first to propose an architecture for the OTFS receiver, providing its resource utilization and timing analysis. This work represents a significant advancement in the practical implementation of OTFS technology.

Design and Analysis of a High-Gain, Low-Noise, and Low-Power Analog Front End for Electrocardiogram Acquisition in 45 nm Technology Using gm/ID Method

PROJ24136

This work presents a novel analog front-end (AFE) circuit for electrocardiogram (ECG) detection systems, designed and investigated in a 45 nm technology node using Cadence. The AFE is composed of an instrumentation amplifier, a Butterworth band-pass filter, and a notch filter, all built using two-stage, Miller-compensated operational transconductance amplifiers (OTAs). The post-layout simulation results demonstrate high-performance metrics, including a bandwidth of 239 Hz, a variable gain between 44 and 58 dB, and a low total power consumption of just 10.88 μ W with a ± 0.6 V supply. The circuit also exhibits a deep notch of -56.4 dB at 50.1 Hz, a low total harmonic distortion (THD) of -59.65 dB, and a compact layout area of 0.00628 mm².



AI Approaches to Investigate EEG Signal Classification for Cognitive Performance Assessment

PROJ24137

This research proposes a new method for analyzing Electroencephalogram (EEG) brain waves to evaluate cognitive processes and motor imagery. The study addresses the need for effective feature extraction and classification by utilizing a novel approach. To validate its efficacy, the method was tested on two benchmark datasets, EEGMAT and EEGMMIDB. The analysis employed various machine learning classifiers, including LSTM, 1D-CNN, and DNN, to evaluate accuracy both before and after feature selection. The results show that the proposed process of feature extraction, selection, and classification outperforms current state-of-the-art models. It achieved a high accuracy of 93.36% for the EEGMAT dataset and an exceptional 98.65% for the EEGMMIDB dataset. These findings highlight a significant advancement in EEG data classification, demonstrating the potential for more accurate analysis of brain activity.



ELYSIUMPRO
INSPIRING THE LEADING EDGE TECHNOLOGIES



2.1L+ Customer Satisfied | **27+** Specialized Domains
107+ Campus Tie-ups | **24/7** Support Center | **72+** Countries

📞 99447 93398

✉ info@elysiumpro.in

🌐 elysiumpro.in

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

