

Ubaid Bakhtiar

📍 Maryland, USA 📧 ubaidb@umd.edu 📞 (301)768-5256 🌐 [/in/ubaid-bakhtiar](https://in.ubaid-bakhtiar) 🌐 [/ubaidb](https://ubaidb)

RESEARCH SUMMARY

My research focuses on optimizing resource utilization to accelerate the throughput and latency in various computing architectures, including CPUs, GPUs, and FPGAs, while maintaining area and timing constraints. Utilizing dynamic reconfiguration, hardware/software co-design, and novel architectural support, my research encompasses applications across diverse domains, from sparse algebra to scientific computing and large language models (LLMs).

EDUCATION

Ph.D. Electrical and Computer Engineering | University of Maryland-College Park | Fall, 2022- Present

BSc. Electrical Engineering | Lahore University of Management Sciences (LUMS), Pakistan | 2018-2022

PUBLICATIONS

Chasón: Supporting Cross HBM Channel Data Migration to Enable Efficient Sparse Algebraic Acceleration | Ubaid Bakhtiar, Amirmahdi Namjoo Bahar Asgari | MICRO-2025 | Acceptance Rate ~21%

- A novel HBM-based sparse algebraic accelerator that improve the resource utilization by migrating the non-zero value across HBM channels. Chasón achieves upto $\sim 6\times$, $\sim 20\times$, $\sim 11\times$ and $\sim 3\times$ speedup over state-of-the-art SpMV accelerator Serpens, Nvidia RTX 4090, Nvidia RTX 6000Ada and Intel Core-i9-11980HK respectively.

Acamar: A Dynamically Reconfigurable Scientific Computing Accelerator for Robust Convergence and Minimal Resource Underutilization | Ubaid Bakhtiar, Helya Hosseini, Bahar Asgari | MICRO-2024 | Acceptance Rate ~22.7%

- Acamar is the first scientific computing accelerator that leverages the partial dynamic reconfiguration of FPGA. It provides robust convergence across a wide range of workloads while enhancing resource utilization for sparse kernels—features that were lacking in the previous scientific computing accelerator designs.

Pipirima: Predicting Patterns in Sparsity to Accelerate Matrix Algebra | Ubaid Bakhtiar, Donghyeon Joo, Bahar Asgari | DAC-2025 | Acceptance Rate ~23%

- Pipirima predicts the matrix sparsity pattern using lightweight predictor and leverage it to accelerate SpMM and SpMV kernels. It shows latency speed up of upto $\sim 20\times$ over state -of-the-art accelerators; ExTensor and Tensaurus.

Segin: Synergistically Enabling Fine-Grained Multi-Tenant and Resource Optimized SpMV | Helya Hosseini, Ubaid Bakhtiar, Donghyeon Joo, Bahar Asgari | IEEE-CAL 2025 | Acceptance Rate ~20%

- Segin leverages a novel fine-grained multi-tenancy approach to allow multiple SpMV operations to be executed simultaneously on a single hardware with minimal modifications, enhancing resource utilization and improving throughput by $1.92\times$.

RELEVANT EXPERIENCE

Graduate Research Assistant | Computer Architecture and Systems Lab | University of Maryland | May, 2023-Present

- Research Advisor: Dr. [Bahar Asgari](#)
- Research Area: Computer Architecture and Domain Specific Designs
- Crafting domain-specific architecture designs to tackle computational challenges and devising methods to enhance their performance as well as simulation and prototyping on modern architectures, i.e. CPUs, GPUs, FPGAs

COURSEWORK

Programming Languages and Computer Architecture | Domain Specific Architecture | Digital Computer Design | Compilers and Optimizations

University of Maryland-College Park

Computer Architecture | Digital System Design | Embedded Systems | VLSI Design | Machine Learning

Lahore University of Management Sciences

SKILL

Topics	HW/SW Co-design, CPUs, GPUs, Simulators, Data Structures, AI/ML
Programming Languages	C/C++, Python, RTL (Verilog), OpenCL, MIPS/RISC Assembly, MATLAB
Hardware Platforms	AMD Xilinx FPGAs and ZYNQ SoC, RISC
Software	TAPA, Rapidstream, Xilinx Vitis, High-level Synthesis, CACTI, Synopsys DC, GPGPU-SIM, Nvidia Nsight Compute

