



P1-25: Reconfigurable Systems



Mission-Critical Computing

NSF CENTER FOR SPACE, HIGH-PERFORMANCE,
AND RESILIENT COMPUTING (SHREC)

SHREC Annual Workshop (SAW24-25)



January 14-15, 2025

Dr. Alan George

Mickle Chair Professor of ECE
University of Pittsburgh

Dr. Samuel Dickerson

Associate Professor of ECE
University of Pittsburgh

James Bickerstaff

Owen Lucas

Joseph Black

Peter Drum

Peri Hassanzadeh

Yasser Morsy

Nicholas Schmidt

Research Students

University of Pittsburgh

Number of requested memberships ≥ 5

Goals, Motivations, Challenges


Goals

- Develop and evaluate **scalable architectures** for FPGA-based apps
- Compare performance of **spatial and fixed-logic architectures**
- Evaluate FPGA **design tools** for high-level synthesis




Motivations

- FPGAs realize custom datapaths for **efficient processing**
- New processing paradigms required for **increased performance**
- Potential **productivity boost** from high-level design tools



Challenges


- High-level design tools often limit **optimization granularity**
- **SWaP-C constraints** can limit parallelism due to lack of resources
- Potential **resource overhead** produced from high-level design tools



Proposed 2025 Tasks

T1: High-Performance FPGA Architectures

- Accelerate **suite of graph operations** & compare HBM vs DRAM throughput
- Enhance OPIR processing with **heterogenous architecture** to boost performance




T2: Machine-Learning Acceleration

- Evaluate and improve machine-learning accelerators on **Versal devices**
- Accelerate **SIREN on FPGA** leveraging HLS tools and collect speedup results



T3: On-Orbit FPGA Accelerators

- Compare accelerated **multi-target apps** on FPGA devices
- Employ **dependability strategies** on PolarFire FPGA & evaluate reliability/utilization



T4: Distributed-Memory Parallel Computing

- Investigate acceleration of **memory-bound apps** using NVSHMEM and ISHMEM
- Compare performance across **GPU-based** implementations



T1: High-Performance FPGA Architectures

James Bickerstaff and Owen Lucas

T1: High-Performance FPGA Architectures

Background

- **HBM** & **efficient streaming techniques** enable rapid processing of large-scale graphs
- Low-latency sensors often require custom hardware acceleration in order to achieve **real-time performance** while maintaining **high accuracy**

High-Throughput Application Acceleration

- **Expand graph processing accelerator** to speed up suite of operations
- Optimize partitioning and PEs to **maximize parallelism** and **throughput**
- Investigate same architecture across **different memory technologies**

Multi-Sensor OPIR Object Tracking

- Accelerate ground station with multiple **parallel object tracking** methods for processing and combining object-tracking data from multiple OPIR sensors
- Take advantage of **heterogenous processing** to reduce latency and increase throughput



T2: Machine-Learning Acceleration

Joseph Black and Peter Drum

T2: Machine-Learning Acceleration

Background


- SIREN MLPs can **model continuous functions effectively** as well as having applications such as image compression and 3D model reconstruction
- Versal's unique hardware presents new opportunities to **accelerate ML models**

SIREN MLP Acceleration

- Accelerate training of **SIREN MLPs** on FPGA devices
- Leverage Vitis HLS tools to develop and **optimize training kernel performance** and examine speedup results

Machine Learning for Versal

- Evaluate **current and future machine-learning accelerators** for Versal devices such as FINN, DPU, NPU
- Measure and improve performance and **resiliency** of these accelerators



T3: On-Orbit FPGA Accelerators

Peri Hassanzadeh and
Nicholas Schmidt



T3: On-Orbit FPGA Accelerators

Background

- Multi-target **tracking** applications widely used in space benefit from **acceleration**
- FPGAs are widely used in space applications, but **SEU** are of major concern with potential to disrupt function and loss of information

Multi-target Tracking Acceleration

- **Evaluate** and **compare** methods to accelerate multi-target tracking applications leveraging FPGA- and GPU-enabled systems
- Measure **speedup**, **parallel efficiency**, and **resource utilization** for each design

Fault-Tolerant FPGAs

- Explore **dependability techniques** to be utilized on PolarFire FPGA mitigating SEUs
- Develop resilient image processing application and analyze **reliability**, **performance**, and **resource utilization** of different resilient configurations



T4: Distributed-Memory Parallel Computing

Yasser Morsy

T4: Distributed-Memory Parallel Computing

Background

- **NVSHMEM** and **ISHMEM** provide ability to perform **in-kernel, GPU-GPU communication** with symmetric memory model
- Unlike GeMM, **memory-bound apps** like **BFS** are constrained by interconnect speed

Distributed-Memory Graph Processing

- Evaluate **communication aggregation** method for BFS to effectively utilize interconnect bandwidth
- Evaluate GeMM & BFS using **ISHMEM** on Intel GPUs to compare performance

Experiment Hardware

- **NVSHMEM**
 - **Single node** cluster of 8 A100 GPUs with NVSwitch (Pitt CRC)
 - **Multi Node** cluster with 8 H100 GPUs with fat tree InfiniBand interconnect (TAMU)
- **ISHMEM**
 - **Single node** cluster of 4 Intel Max 1100 GPUs (TAMU)



Milestones, Deliverables, Budget

- **Milestones**

- SMW (June 2025): Showcase midterm results on all projects
- SAW (Jan 2026): Demonstrate completion of all projects



- **Deliverables**

- Monthly progress reports from all projects
- Midyear and end-of-year full reports from all projects
- 4 conference/journal papers (1 per task)
- Direct access to completed codes and architectures



- **Budget**

- **5+** memberships (**250+** votes) for all tasks

Conclusions & Member Benefits

Conclusions

- Develop **novel FPGA architectures** and systems capable of processing large amounts of data at a rapid rate while leveraging **state-of-the-art devices and memory technologies**
- Investigate acceleration of various **machine-learning methods and architectures** as well as measuring their resiliency for use in mission-critical applications
- Evaluate **computer-vision tasks** on FPGA and GPU devices and conduct fault-injection experiments to **uncover vulnerabilities within PolarFire devices** and discover possible mitigation techniques
- Leverage NVSHMEM and ISHMEM libraries for in-kernel communication to **overcome interconnect limitations** and accelerate memory-bound applications

Member Benefits

- **Direct influence** over research direction and projects with frequent meetings
- **Direct benefit** from accelerator designs and tools exploration
- **Direct insights** from research developments, analyses, and publications