P3-26: Emerging Systems

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

SHREC Annual Workshop (SAW25-26)

University of Pittsburgh
BYU BRIGHAM YOUNG UNIVERSITY
VIRGINIA TECH
UF UNIVERSITY of FLORIDA

January 13-14, 2026

Dr. Alan George
Mickle Chair Professor of ECE
University of Pittsburgh

Dr. Rajkumar Kubendran
Assistant Professor of ECE
University of Pittsburgh

Linus Silbernagel
Diego Wildenstein
Jakob Bindas
Myles Fernau
Nischal Kharel
Research Students
University of Pittsburgh

Number of requested memberships ≥ 5

Goals, Motivations, & Challenges

Goals

- Evaluate **next-gen processing** architectures, sensors and algorithms
- Investigate use of **neuromorphic systems** for space applications
- Benchmark **PIM architecture** and test radiation resiliency

Motivations

- Novel architectures are needed to **combat SWaP-C** constraints of on-board space processing
- Neuromorphic systems and PIM architectures offer **power- and memory-efficient** computation

Challenges

- Next-generation architecture requires **unique design considerations**
- **Lack of software maturity** for new architectures poses challenges when designing novel solutions

PIM: Processing-in-memory
SWaP-C: Size, weight, power, and cost

Proposed Tasks for 2025

T1 — Event-Based Object Detection and Tracking

Develop and improve algorithms for space-related object detection and tracking using EBSs

Task Members: Linus Silbernagel, Jakob Bindas

T2 — ML Acceleration Using RISC-V

Investigate acceleration of SNNs using RISC-V Vector Extension

Task Member: Myles Fernau

T3 — Processing-in-Memory Architecture

Analyze performance and determine reliability of second-generation Gemini-II APU devices

Task Members: Diego Wildenstein, Nischal Kharel

EBS: Event-based sensor
SNN: Spiking neural networks
APU: Associative Processing Unit

T1 Event-Based Object Detection and Tracking

Develop and improve algorithms for space-related object detection and tracking using EBSs

Task Leader: Linus Silbernagel

Linus Silbernagel – linussilbernagel@pitt.edu
Jakob Bindas – jmb442@pitt.edu

EBSs: Event-based sensors

T1: Background

Event-Based Sensors

- Event-based vision sensors produce **asynchronous events** and offer unique characteristics such as **high temporal resolution** and **high dynamic range**
- Uses **minimal power** during operations due to asynchronous nature of sensor

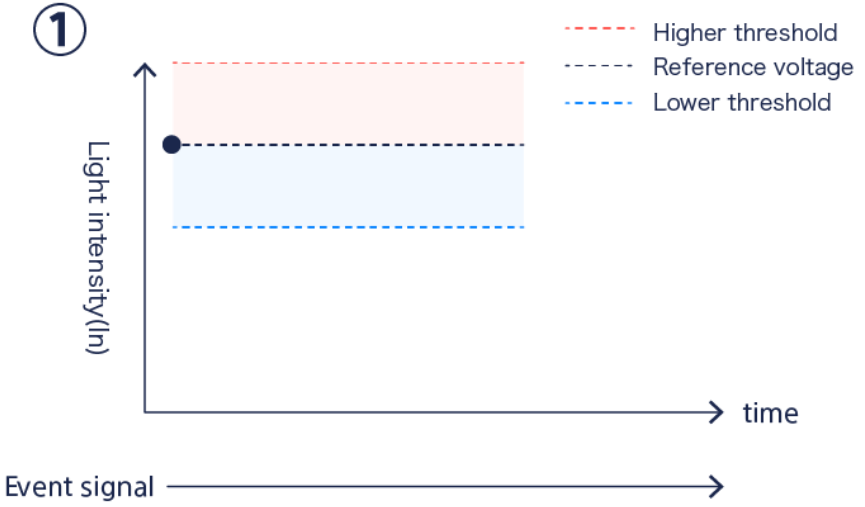

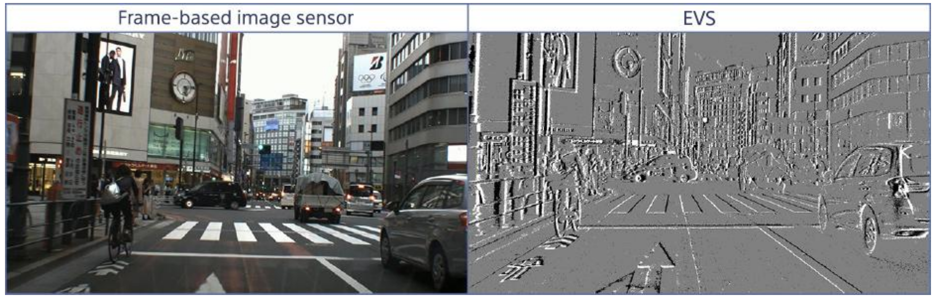

Applications

- Attractive properties of neuromorphic systems can be leveraged for **on-board space applications**
- Event-based data enables **small object tracking** with **high data efficiency**

EVS: Event-based vision sensor

T1: Event-Based Object Detection and Tracking

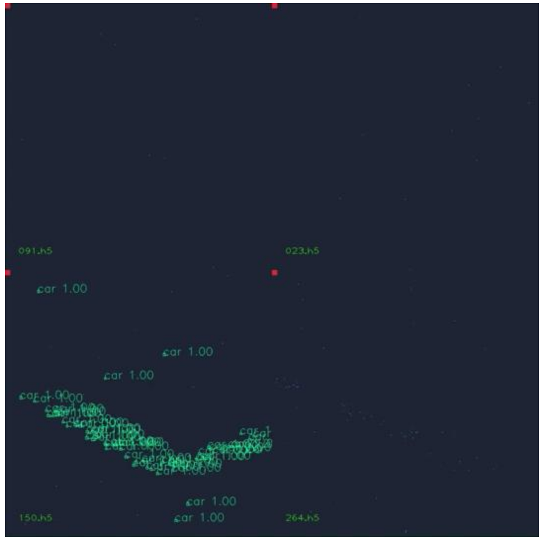

Object Detection in Satellite Imagery

- Expand upon research in **event-based object detection** in large-scale **satellite imagery**
- Evaluate performance of novel architectures such as **spike-based** and **sparse transformers**

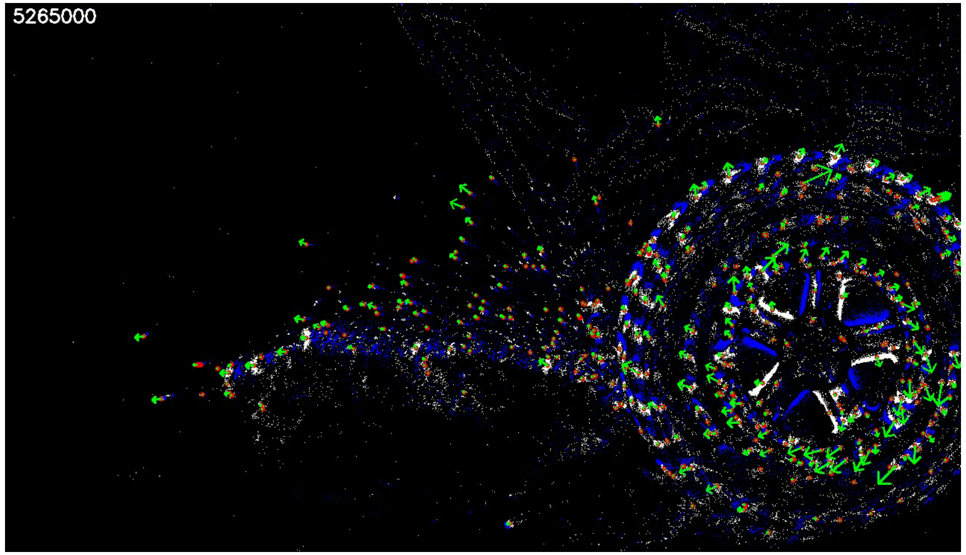

Small Object Tracking

- Transform existing cluster-based tracking pipeline into **lightweight application** for **online tracking**
- Compare performance between **lightweight and original** to **analyze tradeoffs** between versions

T2 ML Acceleration Using RISC-V

Investigate acceleration of SNNs using RISC-V Vector Extension

Task Leader: Myles Fernau

Myles Fernau – mjf140@pitt.edu

SNN: Spiking neural networks

T2: Background

Event-Based Sensors and Algorithms

- SNNs are **powerful** and **efficient**, especially when paired with event-based sensor data
- Sparse, **asynchronous** computation enables **low-power** edge deployment

SNN Acceleration with RV

- Current SNN acceleration requires custom hardware designs with **limited portability**
- RISC-V is being adopted for **next-generation** spaceflight computing, such as HPSC, creating demand for portable solutions to **onboard** ML acceleration

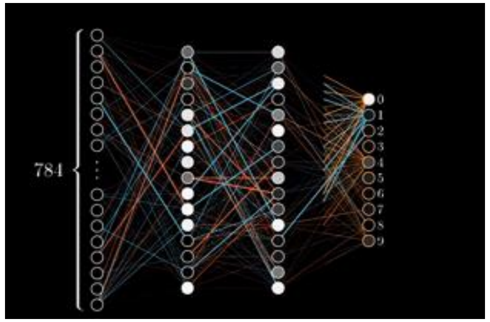

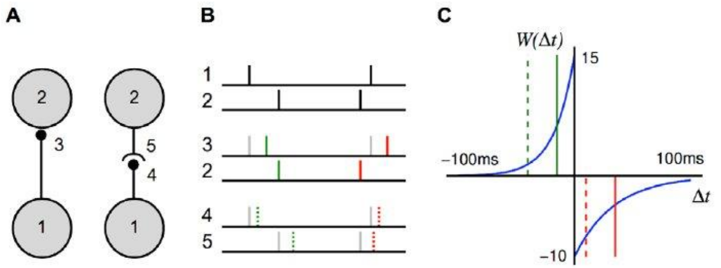

SNN: Spiking Neural Network
RVV: RISC-V Vector Extension

T2: RISC-V SNN Acceleration Studies

Vectorizing SNN Inference

- Write core **SNN** kernels using **RVV** intrinsics
- **VLEN-agnostic** design enables portability across RISC-V hardware implementations

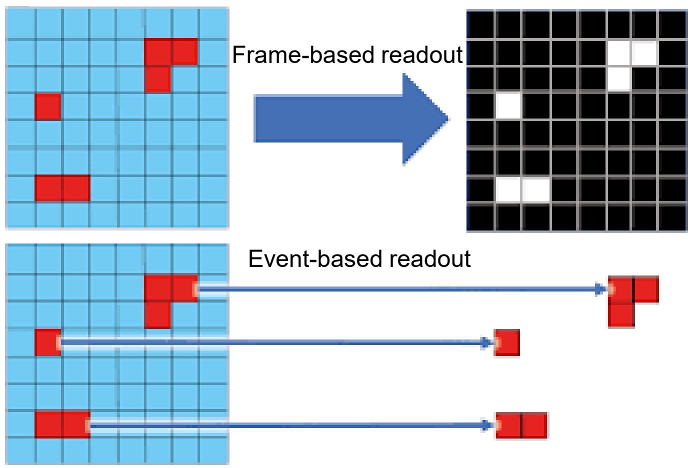

What Will We Measure?

- Inference **throughput** and **speedup** vs. scalar RISC-V baseline
- Relative benefit of **vectorization** across SNN operations

SNN: Spiking Neural Network
RVV: RISC-V Vector Extension
VLEN: Vector Length

T3 Processing-in-Memory Architecture

Analyze performance and determine reliability of second-generation Gemini-II APU devices

Task Leader: Diego Wildenstein

Diego Wildenstein – diw29@pitt.edu
Nischal Kharel – nik96@pitt.edu

APU: Associative Processing Unit

T3: Background

In-Memory Processing

- Architectures with processing logic **directly integrated** into top-level cache memory cells
- **Improves runtime latency** by reducing need for frequent memory transfers

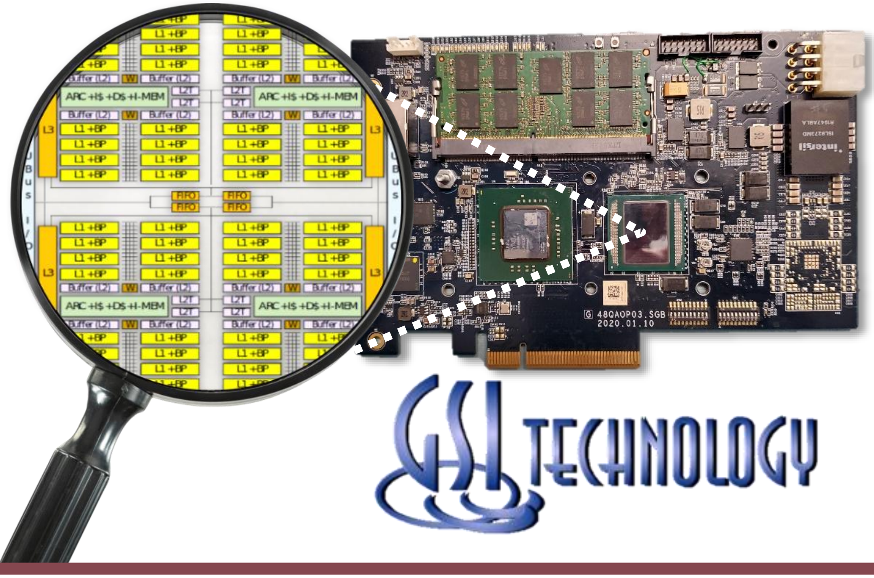

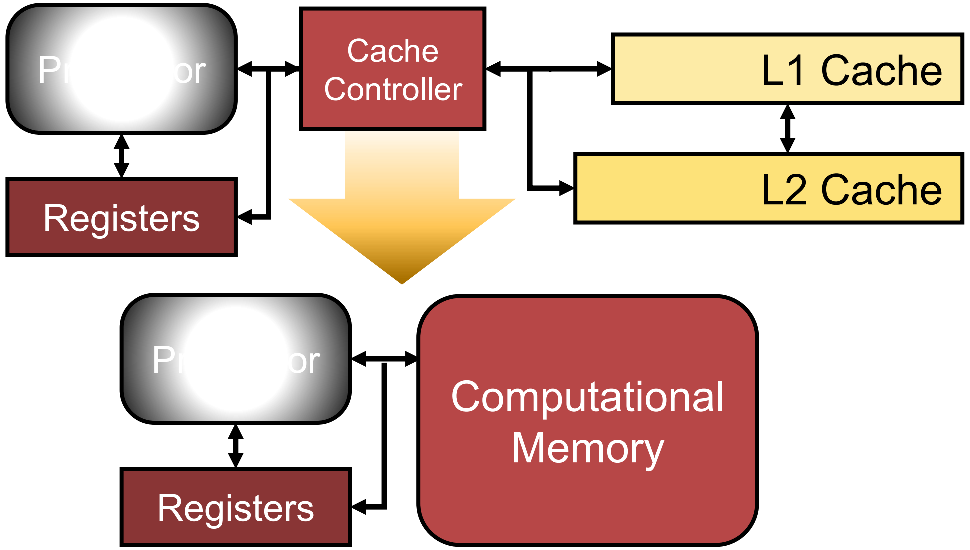

Gemini-II APU

- New generation of state-of-the-art PIM devices with over **2 million bit-processor** memory cells
- Device exhibits **low power usage** profile and **consistent, fast performance** for memory-bound apps

APU: Associative Processing Unit
PIM: Processing-in-memory

T3: Processing-in-Memory Studies

P3

SpaceBench

$$\begin{bmatrix} A_{11} & A_{12} & A_{13} \\ A_{21} & A_{22} & A_{23} \\ A_{31} & A_{32} & A_{33} \\ A_{41} & A_{42} & A_{43} \\ A_{51} & A_{52} & A_{53} \\ A_{61} & A_{62} & A_{63} \end{bmatrix} \times \begin{bmatrix} B_{11} \\ B_{21} \\ B_{31} \end{bmatrix} = \begin{bmatrix} C_{11} \\ C_{21} \\ C_{31} \\ C_{41} \\ C_{51} \\ C_{61} \end{bmatrix}$$

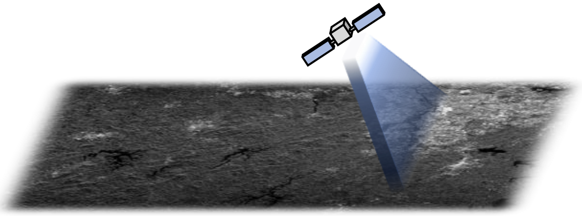

APU Performance on Compute Kernels

- **Benchmark G2 APU** device with compute kernels commonly used in machine-learning apps
- **Compare performance** of G2 APU with other upcoming **spaceflight CPU and GPU** devices

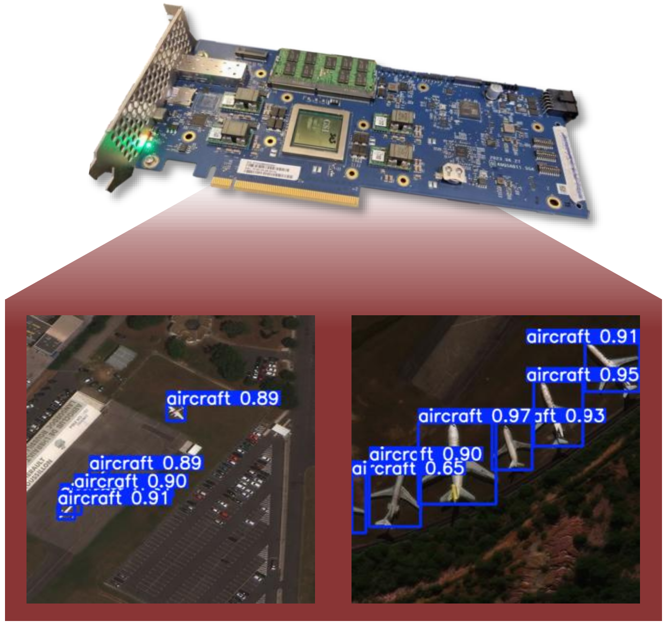

APU Performance on ML Models

- Accelerate additional **YOLO-based object detection** models using G2 APU hardware
- Explore feasibility of Generative AI on G2 APU

G2 APU: Gemini-II Associative Processing Unit

Milestones, Deliverables, & Budget

MILESTONES

- SMW26 (Summer 2026): Showcase preliminary results on all project tasks
- SAW26-27 (January 2027): Completion of all project tasks

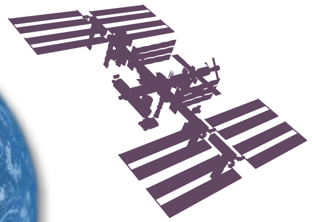

DELIVERABLES

- Monthly progress reports from all projects
- Midyear and end-of-year full reports from all projects
- 1-2 conference or journal publications per task

BUDGET

- Minimum recommended: Five (5) memberships (250 Votes)

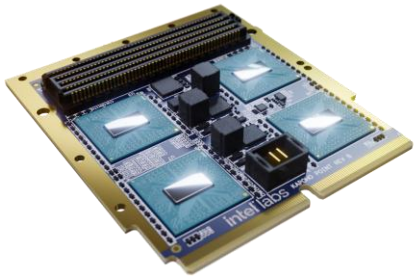

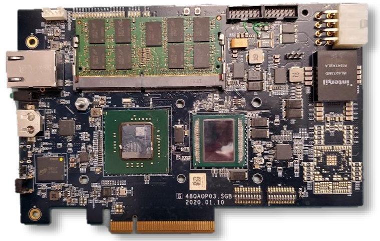

Conclusions & Member Benefits

Conclusions

- Analyze performance of novel methods such as **transformers** for **event-based object detection** in satellite imagery
- Develop **lightweight algorithm** for **detecting and tracking** small objects on board rovers and satellites
- Accelerate SNN inference using **RISC-V** vector intrinsics and evaluate performance scaling across **vector-length** configurations
- Characterize performance and reliability of second-generation **in-memory processing** devices

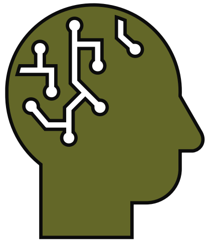

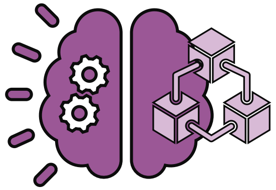

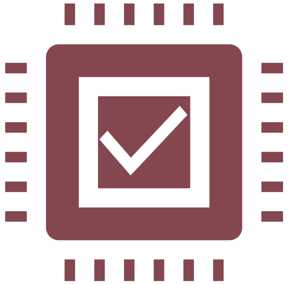

Member Benefits

- Direct influence over processors and frameworks studied
- Direct influence over apps and datasets studied
- Direct benefit from new methods, data, code, models, and insights from metrics, benchmarks, and emulations

SNN: Spiking neural network