

Meridian Research Group

Onboarding & Lab Orientation

Research Group: Dr. Marcus Vance (Advisor)

Affiliation: Department of Electrical Engineering and Computer Science
Meridian Institute of Technology

Academic Year 2026–2027

Outline

Lab Overview

Who We Are

Tooling & Infrastructure

Lab Safety & Protocols

Expectations & Culture

First Steps

Who We Are

Meridian Lab Orientation Deck

Welcome to the Meridian Lab

Mission & Core Values

Our Core Mission


- We design high-throughput machine learning architectures for distributed, heterogeneous edge environments.
- We focus on bridging the gap between algorithmic complexity and actual hardware deployment constraints.

Core Values

- **Methodological Rigor:** Code and math must be verified and fully reproducible.
- **Collaboration:** Open communication and mutual help.

Primary Objective

To build robust intelligence that operates seamlessly within constrained environments (see VertexFlow architecture [1] and Synapse systems [2]).



Research Focus Areas

Where Your Work Fits



Edge Inference

- Model execution pipelines on Jetson and Edge TPU.
- Real-time scheduling algorithms.



Decentralized Opt

- Federated learning under bandwidth limits.
- Asynchronous consensus networks.



Efficient Hardware

- Custom FPGA accelerators.
- Quantization and compilation frameworks.

Tooling & Infrastructure

Meridian Lab Orientation Deck

Core Compute Resources

The Apex GPU Cluster

Cluster Architecture


- **GPU Nodes:** 16x DGX H100 nodes, each with 8x 80GB H100 GPUs.
- **CPU Nodes:** 32x dual-socket EPYC nodes for preprocessing.
- **Storage Pool:** 2PB high-speed Ceph storage.

Scheduling Policies

- Always use Slurm to request sessions.
- Keep interactive jobs under 2 hours.

Compute Allocations

Compute queues are allocated weekly. Post requests in the `#meridian-compute` Slack channel by Sunday night.



Slurm Pool

Software Stack & Reproducibility

Standard Dev Stack

Reproducibility Standards

- **Strict Formatting:** Use Black and Ruff for Python code.
- **Containers:** Apptainer/Singularity for cluster runs.
- **Dependencies:** Manage environments using Pixi or Poetry.

Resource	Platform	Address
Shared Code	GitLab	git.nimbus.io
Models	Registry	reg.nimbus.io
Compute	Slurm	apex-cluster
Wiki	Nimbus Wiki	wiki.nimbus.io

Note: Accounts for GitLab and the cluster are created during your first onboarding week. Reach out to the admin.

Lab Safety & Protocols

Meridian Lab Orientation Deck

Lab Safety & Protocols

Physical Lab Guidelines

Physical Access & Hardware Safety

- **ESD Protection:** Always wear Electrostatic Discharge (ESD) wrist straps when handling custom boards and FPGAs.
- **Cleanroom Rules:** No food or open drinks are allowed in the server closet or custom testing benches.
- **Emergency Shutdown:** Locate the red power-off buttons on the side walls in case of hardware malfunction.

Warning: Laser Safety

Our communication testbed uses Class 3B lasers. Never operate the testbed without wearing laser safety goggles (located in Drawer 4).

Incident Reporting

Report equipment failures or safety hazards in the #lab-hardware Slack channel.

Data Safety & Ethics

Protecting Code & Datasets

Data Backup Policies

- All raw data must be mirrored to the Nimbus storage cluster (which undergoes nightly backups).
- Do not store critical experimental data solely on local lab computer hard drives.

Responsible Code Sharing

- Always keep private repositories private until publication.
- Never commit credentials, server IPs, or API tokens to Git.

Ethical Publication Rules

- Follow IEEE guidelines for authorship and credit.
- Ensure all open-source models published on Vertex Registry are sanitized.

Expectations & Culture

Meridian Lab Orientation Deck

Expectations & Culture

Weekly Research Rhythms

Weekly Meetings

- **Lab Meeting:** Mondays 10:00 AM.
Attendance is mandatory. We discuss high-level project statuses and administrative updates.
- **Reading Group:** Wednesdays 4:00 PM.
We present and critique recent SOTA papers on edge machine learning.

Individual Synchronizations

- Weekly 1-on-1 with your direct mentor/advisor. Schedule via Nimbus Wiki calendar.

Working Hours

We support flexible schedules, but core collaboration hours are between **10:00 AM and 4:00 PM** (when team members are expected to be responsive on Slack and email).

Onboarding Milestones

First Year Timeline

Phase 1: Getting Up to Speed

- **Month 1:** System access complete, literature review draft ready, first baseline replication successful.
- **Month 3:** Core project proposal approved, initial experiment data collection finished.

Phase 2: Execution & Writing

- **Month 6:** First-author manuscript draft written and reviewed.
- **Month 12:** Paper submission to leading venue (e.g., NeurIPS, MLSys) and presentation.

First Steps

Meridian Lab Orientation Deck

Your First Project

Replicating VertexFlow


Step-by-Step Task List

1. **Access GitLab:** Clone the baseline repository:

```
git clone git@git.nimbus.io:vt/flow.git
```
2. **Start Slurm Session:** Request a GPU allocation:

```
srun -p gpu -gpus=1 -pty bash
```
3. **Run Replication:** Run the onboarding test script:

```
python run_onboard.py -config cfg.json
```
4. **Publish Results:** Paste performance metrics on the Nimbus Wiki.



References I

- [1] Alistair Sterling and Elena Rostova. “VertexFlow: Active Stream Scheduling for Distributed Heterogeneous Edge Devices”. In: *Proceedings of the Meridian Conference on Edge Intelligence (MCEI)*. 2025, pp. 102–114.
- [2] Marcus Vance and Sarah Jenkins. “Decentralized Optimizations in Constrained Environments: A Survey”. In: *Synapse Journal of Robotics and Automation* 12.4 (2024), pp. 245–258.