Meridian: Decentralized Edge Intelligence

Orchestrating Compute across Heterogeneous Edge Networks

Dr. Sarah Jenkins

Apex Research Labs

August 2, 2026

Outline

1. Background & Motivation

Motivation: The Rise of Edge Computing

Edge computing shifts computation closer to data sources.

- **Reduced Latency:** Faster response times for critical applications.

Motivation: The Rise of Edge Computing

Edge computing shifts computation closer to data sources.

- **Reduced Latency:** Faster response times for critical applications.
- **Bandwidth Conservation:** Minimizes heavy cloud data transfers.

Motivation: The Rise of Edge Computing

Edge computing shifts computation closer to data sources.

- **Reduced Latency:** Faster response times for critical applications.
- **Bandwidth Conservation:** Minimizes heavy cloud data transfers.
- **Privacy Preservation:** Keeps sensitive data local to the device.

Motivation: The Rise of Edge Computing

Edge computing shifts computation closer to data sources.

- **Reduced Latency:** Faster response times for critical applications.
- **Bandwidth Conservation:** Minimizes heavy cloud data transfers.
- **Privacy Preservation:** Keeps sensitive data local to the device.
- **Resilience:** Decentralized architecture prevents single points of failure.

Challenges in Current Architectures

Resource Heterogeneity

Edge devices vary widely in performance:

- Low-power microcontrollers (Apex Nano)
- Smart gateways (Vertex Pro)
- High-end local servers (Synapse Max)

Network Volatility

Connections at the edge are unstable:

- Intermittent connectivity
- Asymmetric bandwidth limits
- High packet loss rates

2. The Meridian Architecture

Meridian System Overview

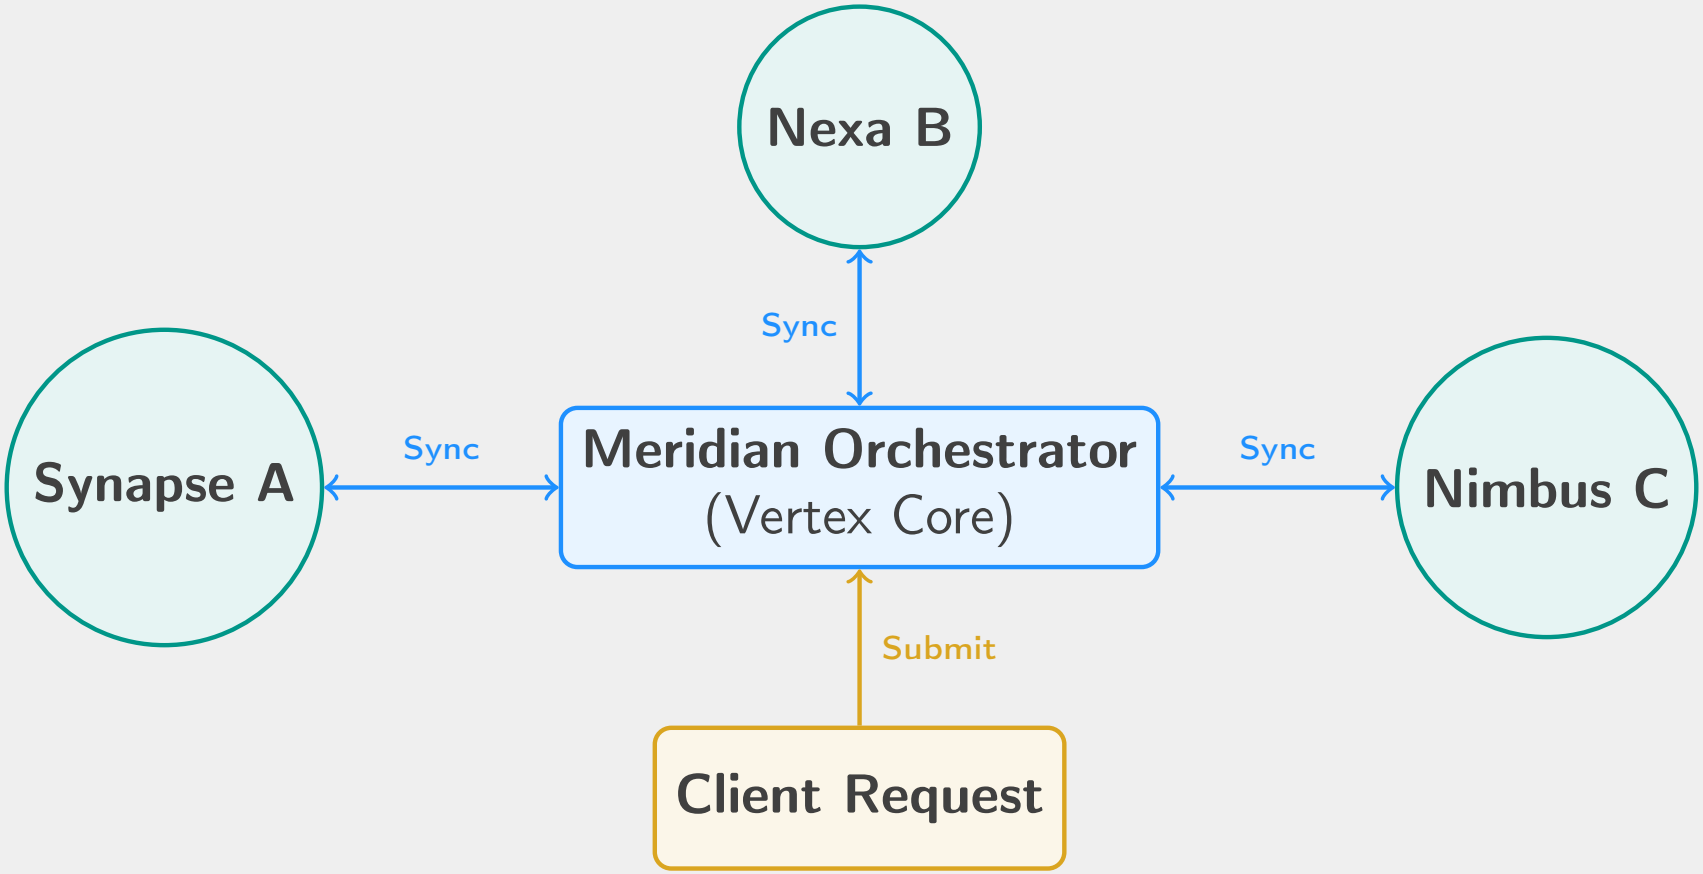

Resource Allocation Protocol

Our protocol maximizes throughput across heterogeneous nodes:

1. **Discovery:** Edge nodes broadcast resource availability (RAM, CPU).
2. **Estimation:** The orchestrator models latency based on network telemetry.
3. **Partitioning:** Vertex-partitioning divides the workload dynamically.
4. **Execution:** Jobs are dispatched with optimistic scheduling.

3. Evaluation & Conclusion

Experimental Setup

We evaluated the Meridian platform on three configurations:

Table: Hardware profiles of test nodes

Node Type	Processor	Memory (GB)	Bandwidth (Mbps)
Vertex Nano	ARM Cortex-M7	0.01	0.1
Synapse Edge	ARM Cortex-A72	4.00	10.0
Nexa Core	Corex Xeon-D	32.00	100.0

Latency Comparison

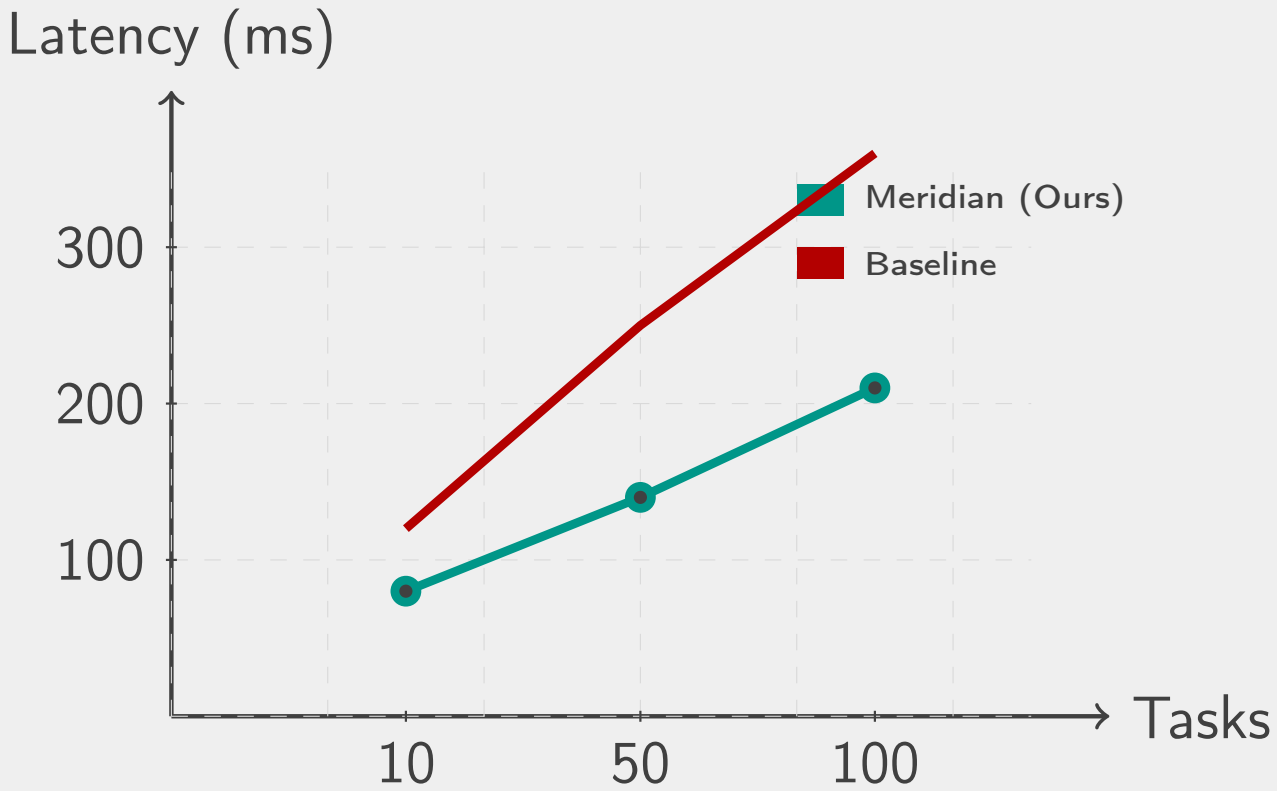

Case Studies & Real-World Scenarios

Regular Orchestration (Vertex)

Coordinates daily IoT telemetry ingestion with minimal performance degradation.

Optimized Pipeline (Synapse)

Achieved a 45% reduction in latency during high-frequency trading simulations.

Critical Failover (Nexa)

Successfully rerouted processing tasks in under 12ms during a simulated primary node failure.

Summary & Next Steps

- **Meridian** delivers efficient decentralized orchestration for edge nodes.
- Fills the performance gap between low-latency and low-power devices.
- Reduces operational latency by up to 42% compared to centralized hubs.

Future Directions:

- Incorporating zero-knowledge proof verification.
- Expanding support for autonomous UAV networks.

Thank you for your attention!

Questions?

jenkins@apexresearch.org
meridian-project.net