

# Evaluation of HPC Storage Systems for HEP Analyses

Asal

Mehrabi

Supervisors: Jonas Mannfeld, Vincenzo Padulano

## Table of Contents

1

**Introduction to RNTuple**

2

**High-Performance Computing Systems**

3

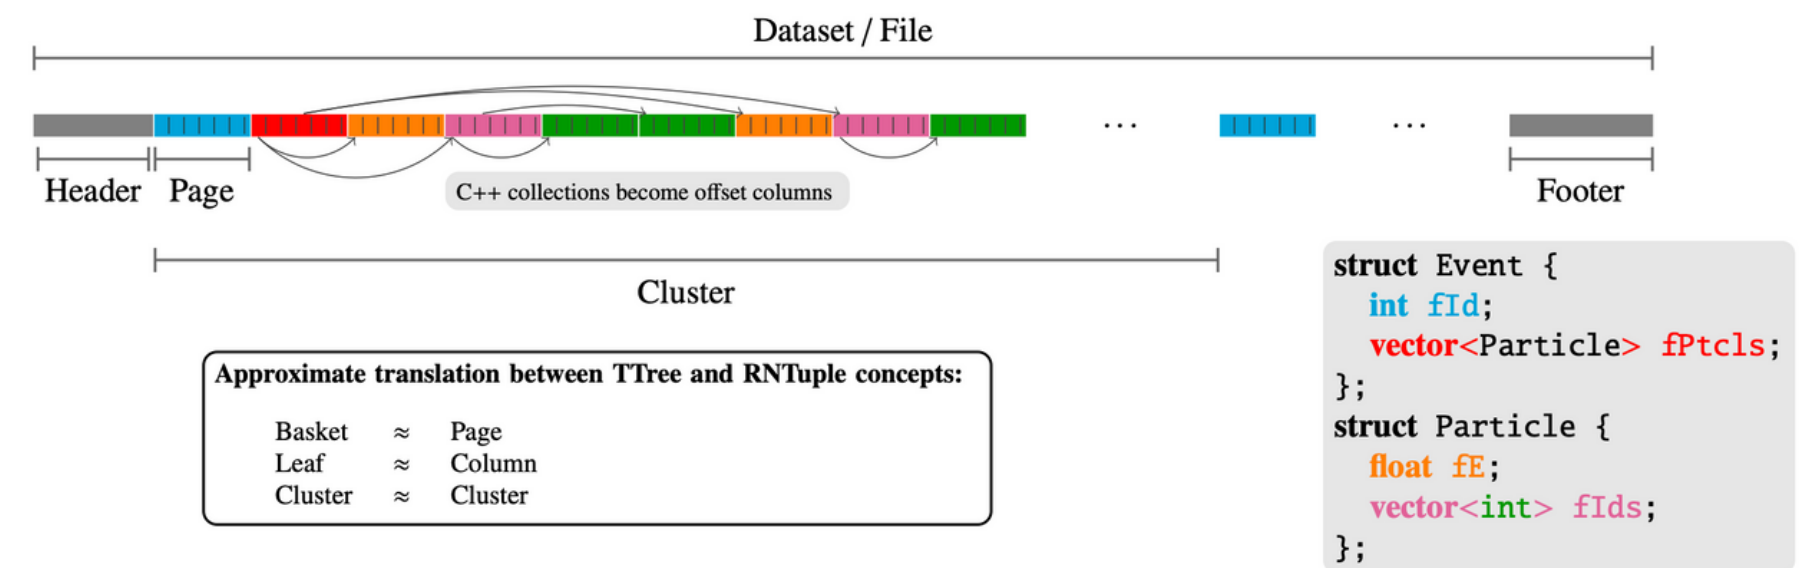
**Performance Analysis and Benchmarking**

4

**Conclusions and further work**

# Introduction to RNTuple

- RNTuple:
  - Data Storage format
  - Structure: Columnar data format
  - Fast access, optimized for modern hardware.
- Analysis Grand Challenge





# High Performance Computing Systems

LUMI

**CPU Nodes: 128 cores per node**  
**Memory: Up to 512 GB per node**  
**Storage: Parallel file system (Lustre)**




ALPS

**CPU Nodes: 128 cores per node**  
**Memory: Up to 256 GB per node**  
**Storage: Parallel file system (Lustre)**



**Job scheduling and Resource allocation**

# Workflow



**NanoAOD dataset  
Analysis Grand challenge**

**HPC Systems**

**Benchmarking Storage with  
Performance Metrics**

**LUMI**



## Challenges



**Data Handling and Storage**



**Different HPC System Configurations**



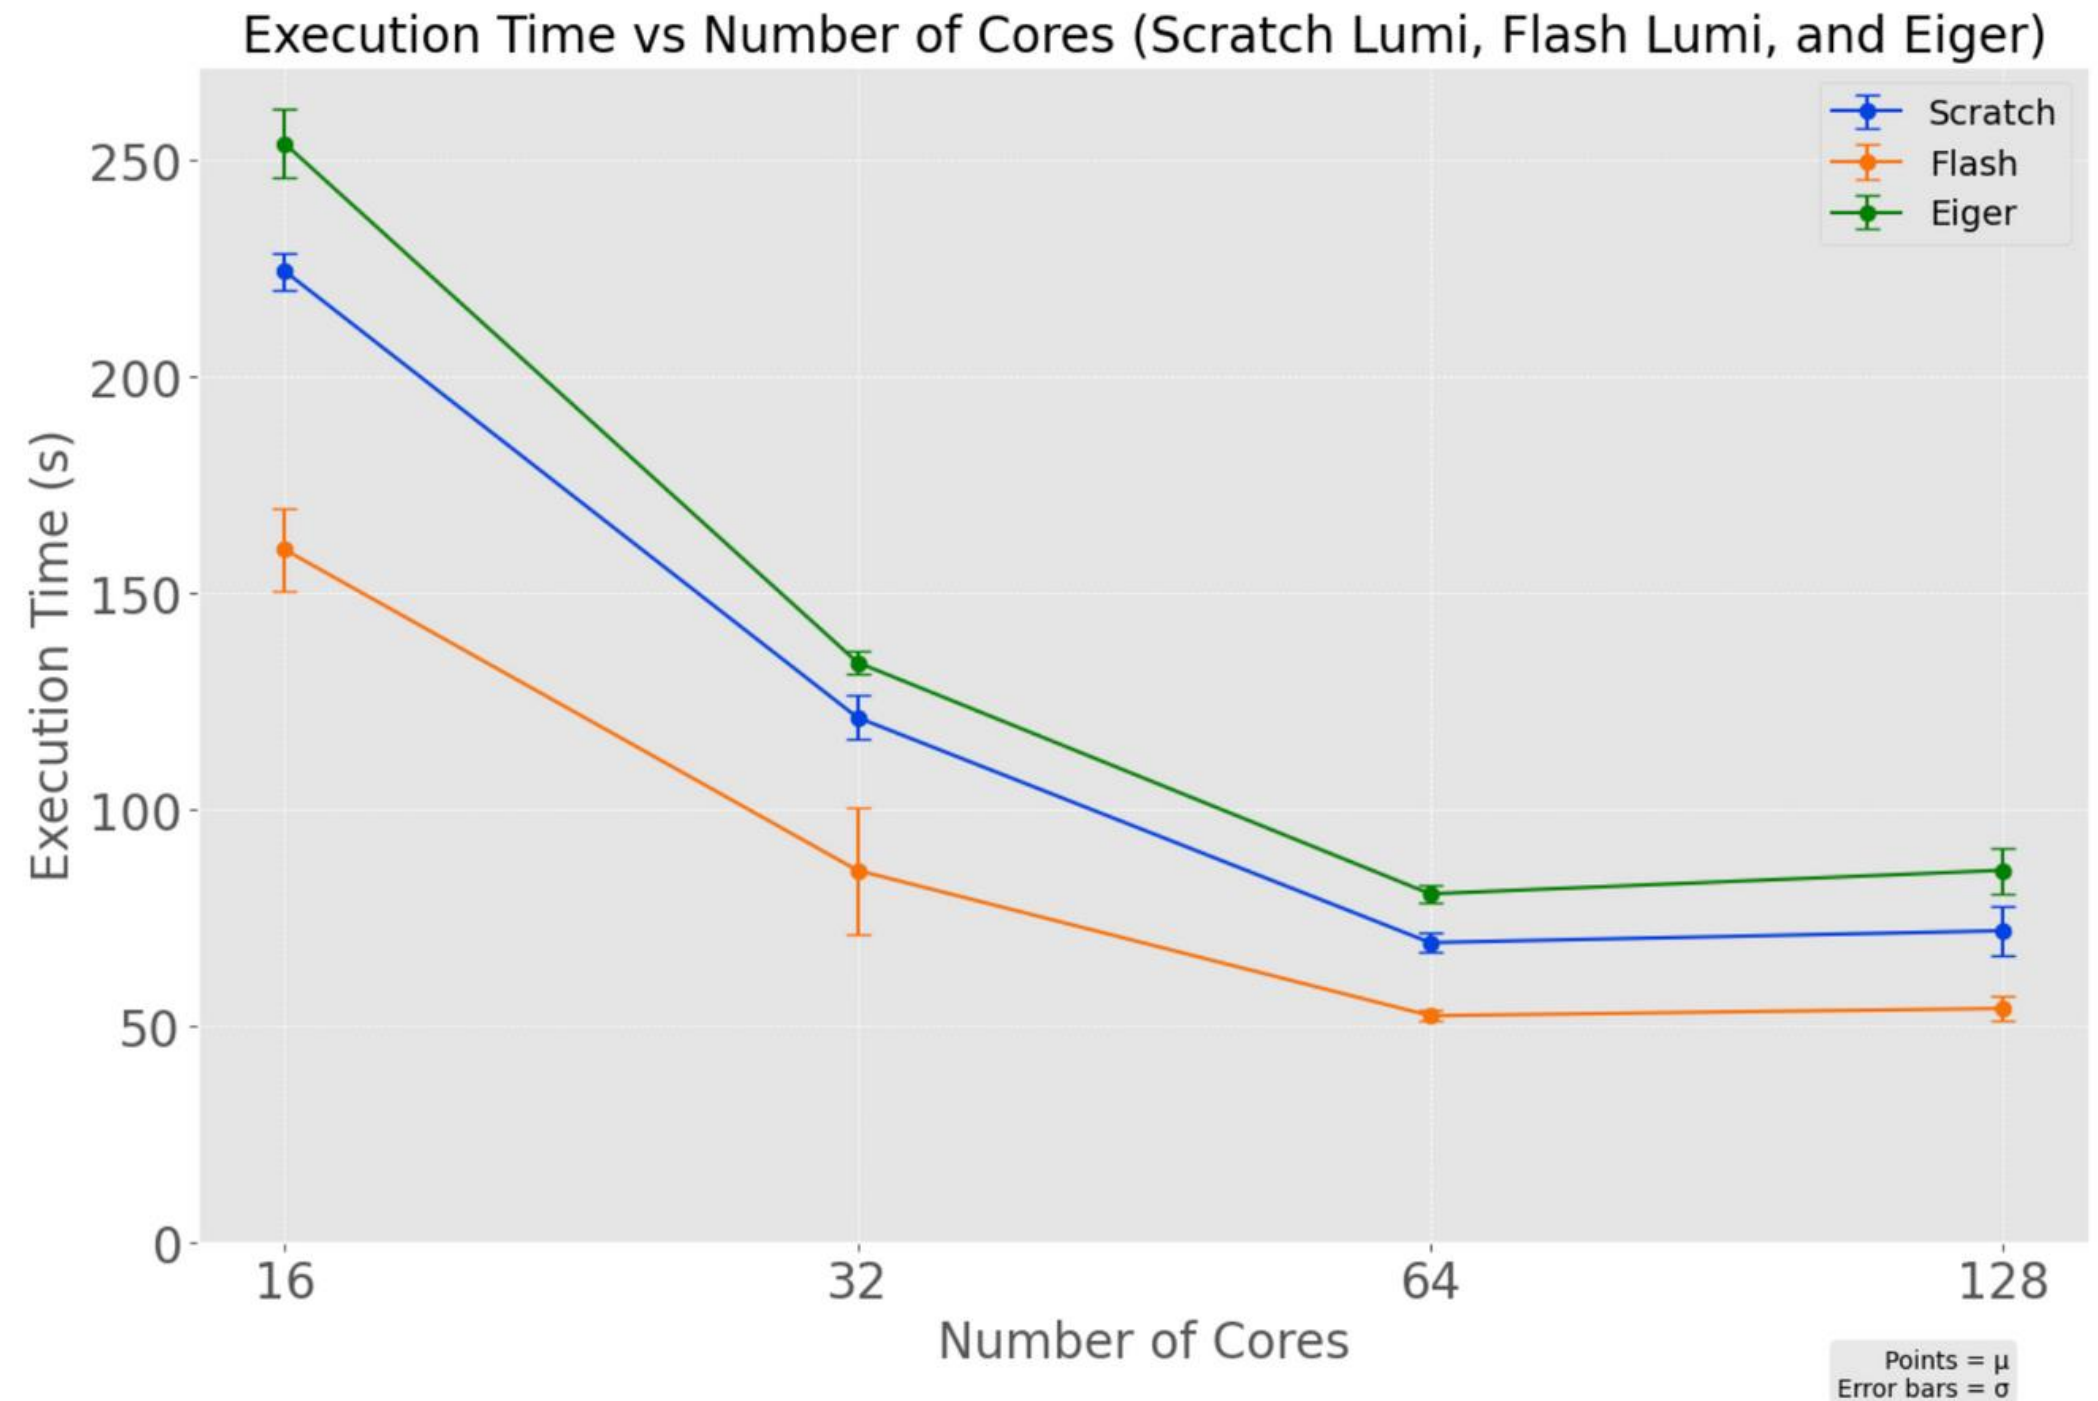
**Performance Variability**



**Software and Tool Compatibility**

# Impact of Storage on Execution Time

- **Scratch Storage: Spinning disks, Lustre**
- **Flash Storage: Based on SSDs, Faster access times.**
- **Eiger Storage**



## Further Work



**Analysis on HPC systems**



**Benchmarking the  
Distributed RDF**





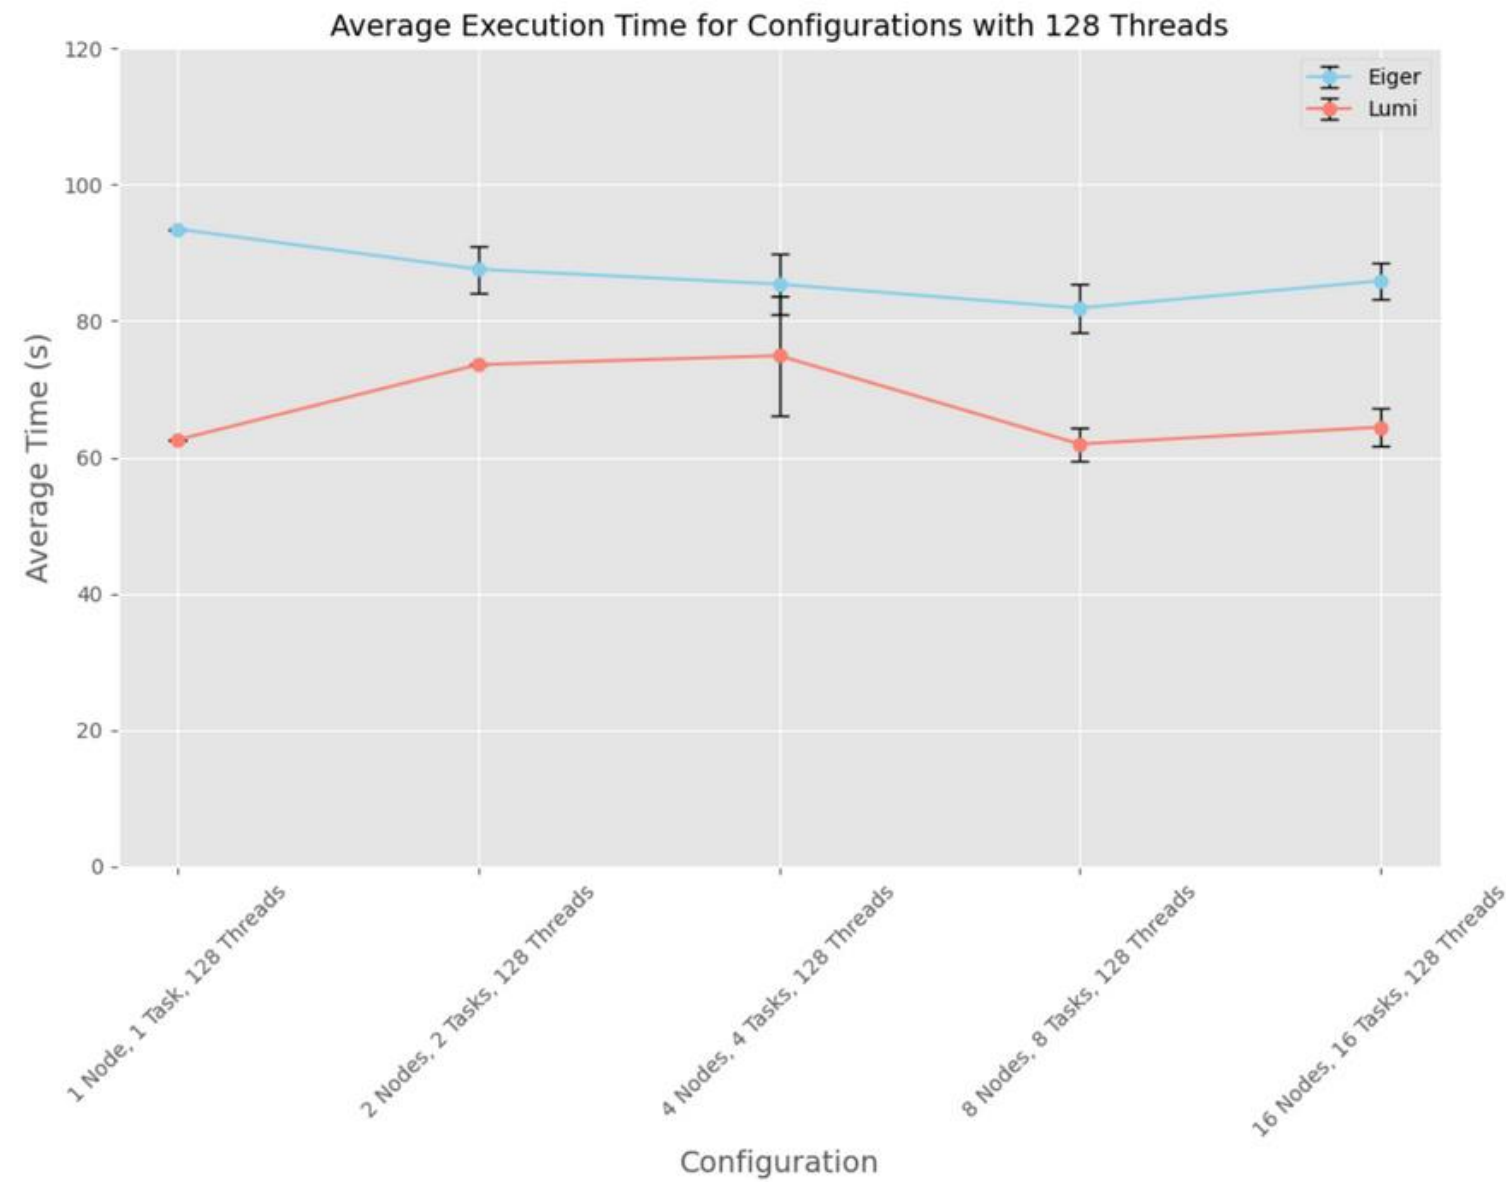
# Thank you!

asal.mehrabi@cern.ch

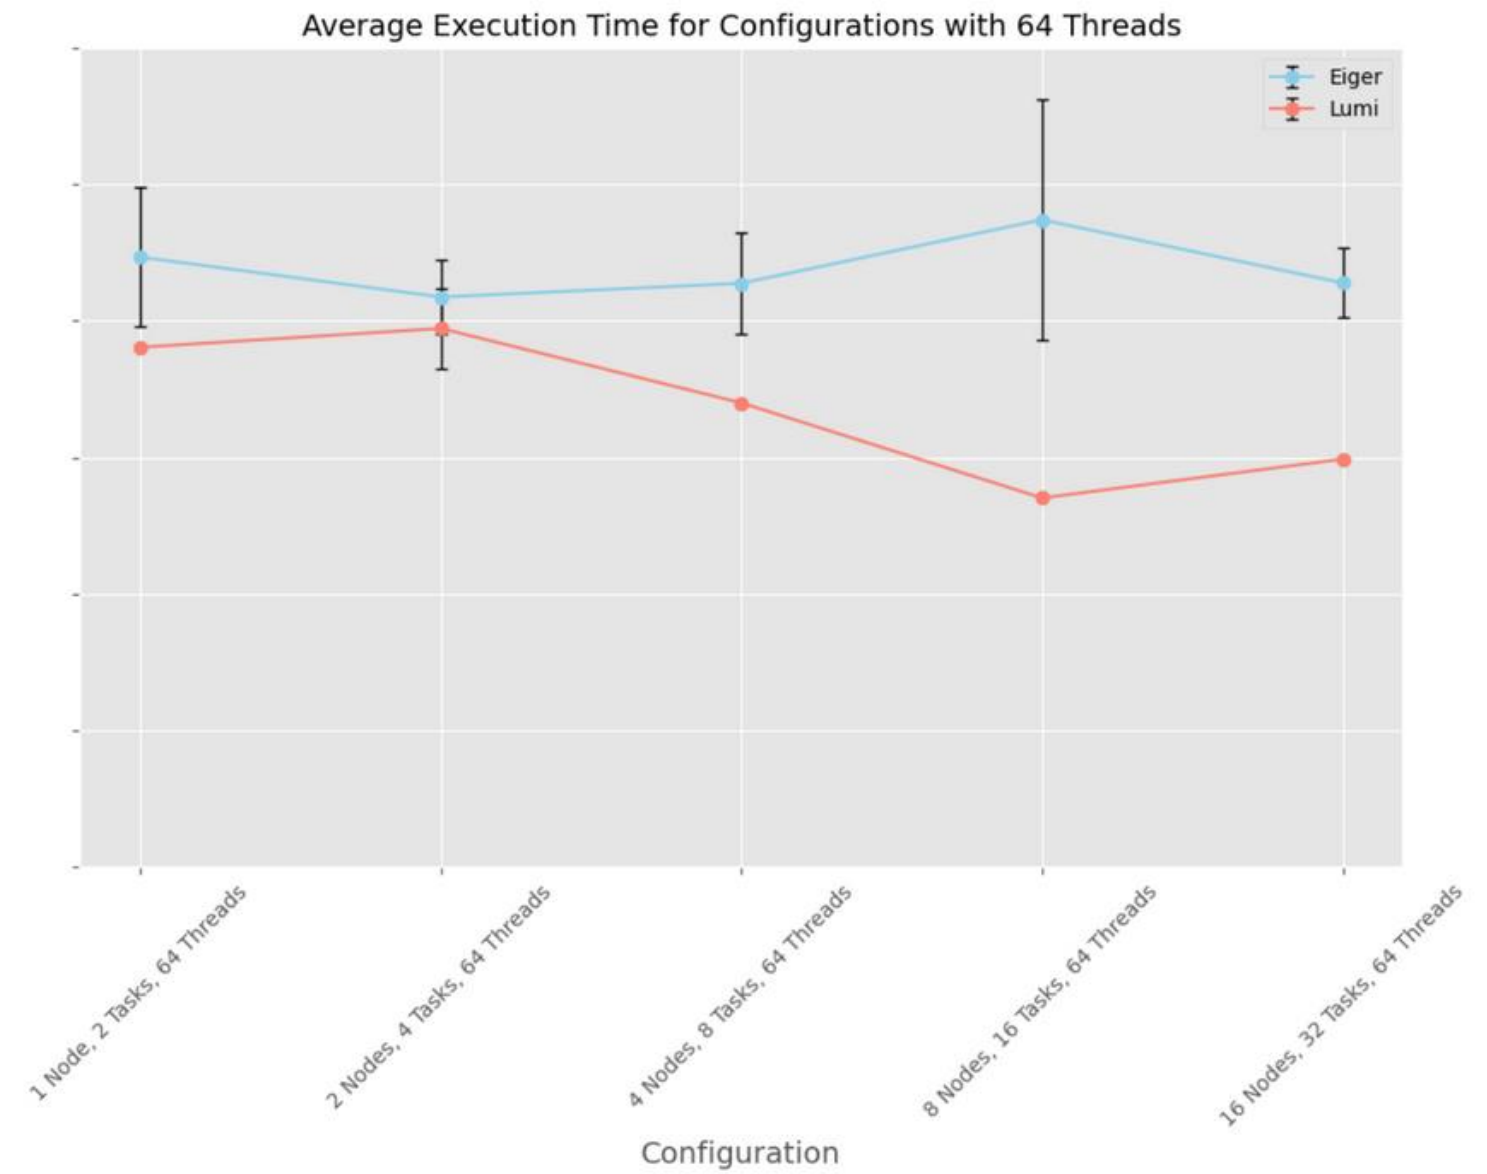
[linkedin.com/in/asal-mehrabi](https://www.linkedin.com/in/asal-mehrabi)



# Multithreading



Full Node



Half Node

# Multithread RDF vs Distributed RDF

