

# Framework for Lifecycle enrichment of HPC Applications on Exascale Heterogeneous Architecture

Karan Sapra

Advisor: Melissa Smith



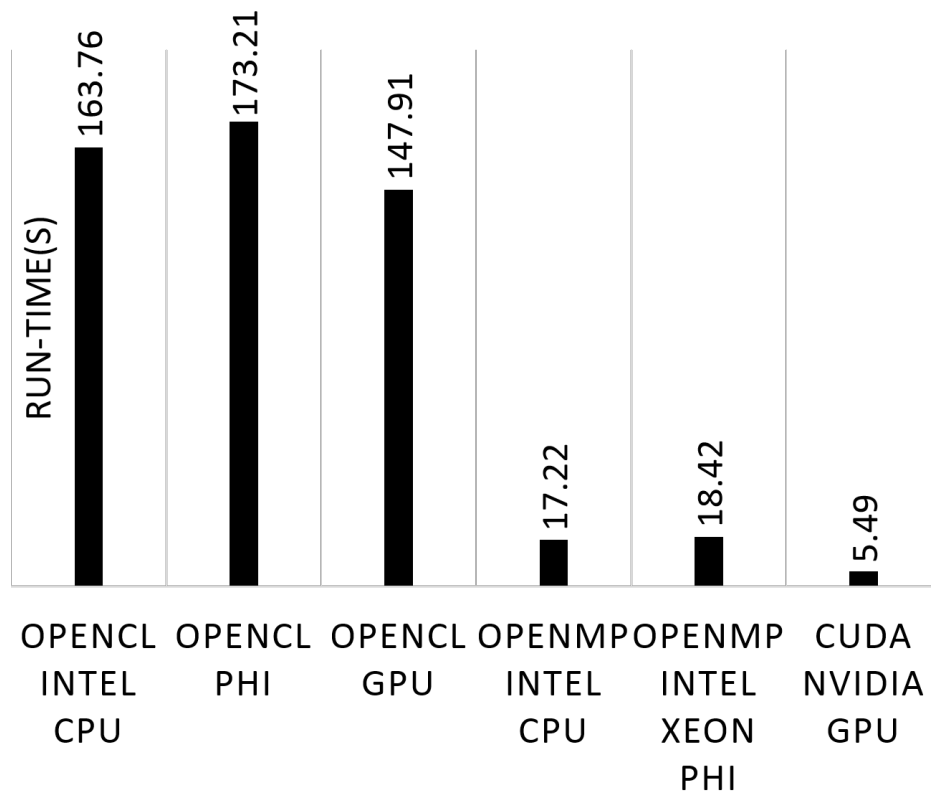
# Overview

- A study and analysis for optimal mapping of application-to-architecture(A2A) mapping.
- A framework for optimal performance of scientific applications on heterogeneous supercomputers.

# Motivation

- New performance milestones achieved
  - GPGPUs
  - Intel's Xeon-Phi coprocessors
  - FPGA
- Application performance depends on multiple factors
  - Architecture
  - Programming model
  - Implementation Detail

# Motivation



Comparison of Performance of K-means for different programming models and different architectures

# Primary Contributions

- Qualitative A2A framework: Application to Architecture Mapping
- Function Partitioning (FP) Framework for run-time maintenance of Application and Subtask in heterogeneous architectures
- A2A with FP: optimal mapping of application task at runtime for optimal performance

# A2A Mapping

- Various performance factor based on algorithmic implementation of application
  - Memory Accesses
  - Algorithmic Complexity
  - Floating Point Operations(FLOPs)
  - I/O
  - Network Latency and Network Bandwidth
  - Network Congestion
  - Communication Ratio
  - Etc.

# Application Performance

Characteristics	Description
Computation Ratio	The amount of time required to perform computation

✓ Time Complexity

Characteristics	Description
FLOPs	The amount of floating point computation instructions
Non-FLOPS	The amount of Non-floating point instructions

✓ FLOPs to Non-FLOPs Ratio

# Application Performance

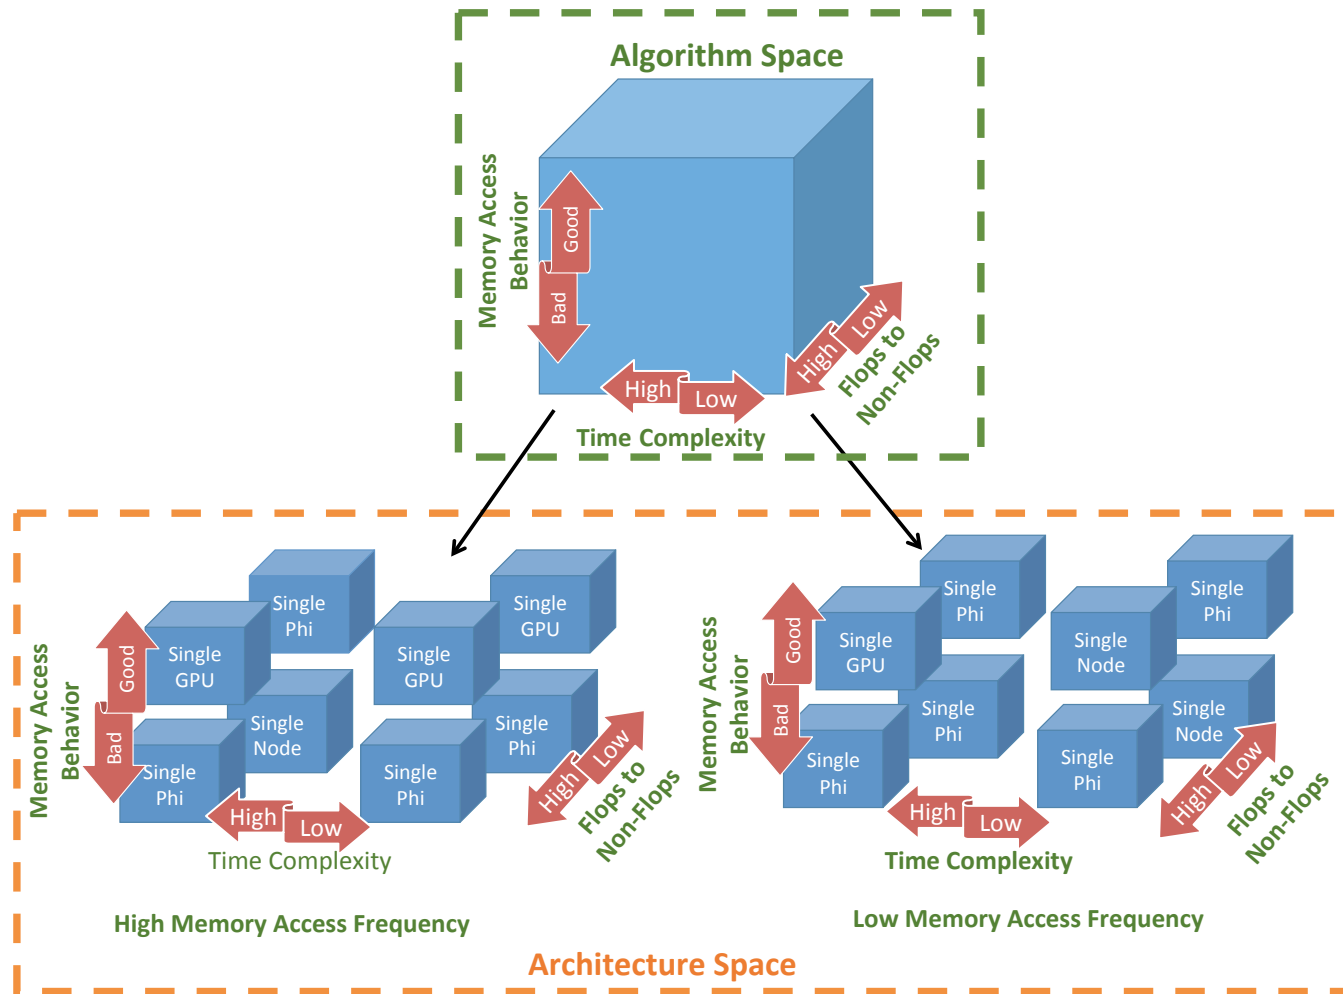
Characteristics	Description
L1 cache Accesses	L1 cache Hit and Miss Ratio
L2 cache Accesses	L2 cache Hit and Miss Ratio
L3 / shared cache	L3 cache Hit and Miss Ratio
Characteristics	Description
Main Memory Accesses and Latency	Time required to access main memory
Non-uniform Memory Accesses NUMA	The amount of time required for NUMA

✓ Memory Access Behavior

+

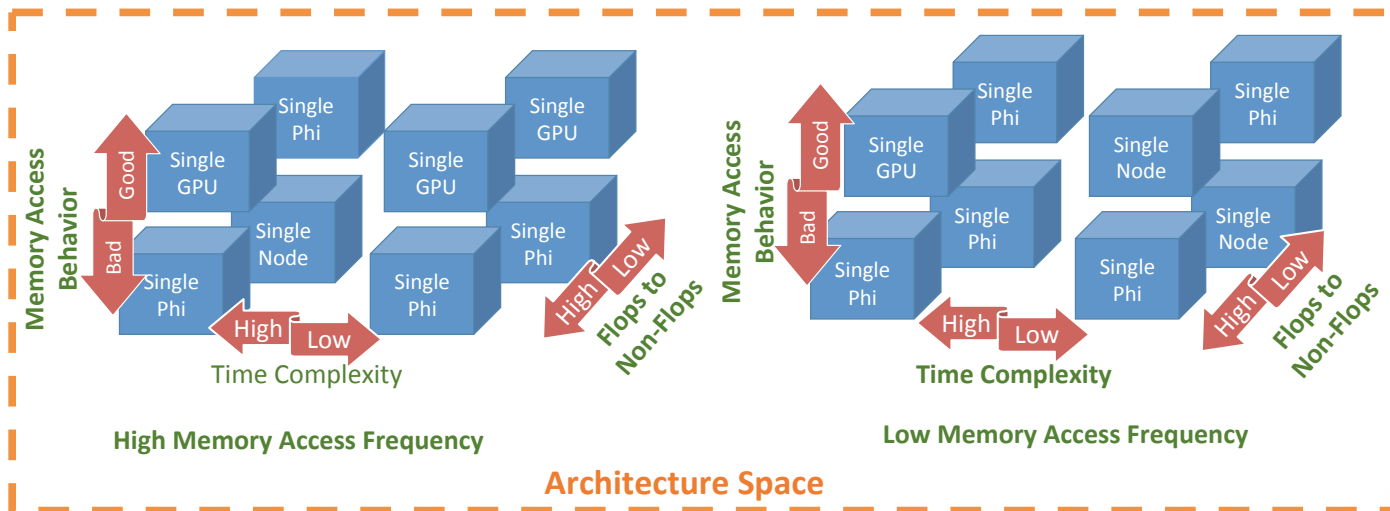
✓ Memory Access Frequency

# A2A Framework: Tesseract



# Application Performance Factors

- Time Complexity
- Floating-point operations (FLOPs) to Non-FLOPs ratio
- Memory Access Behavior
- Memory Access Frequency



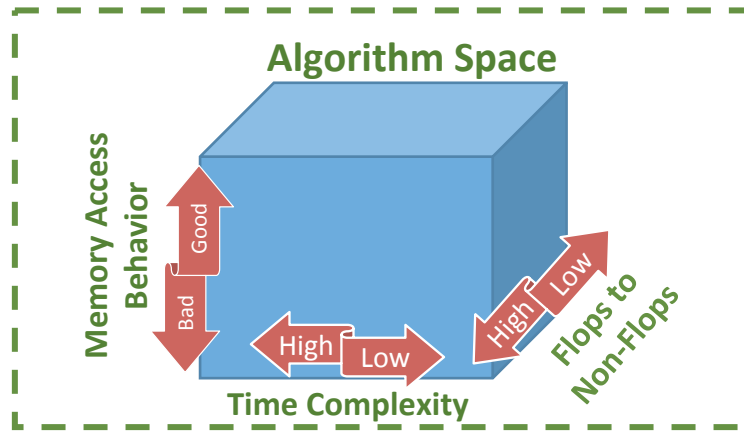
# Application and Micro-benchmarks

- OpenDwarfs Benchmark
- Rodinia Benchmark
- Scalable Heterogeneous Computing Benchmark Suite (SHOC)
- Linear Algebra, Divide and Conquer, Grid, N-Body, Sort, Graph Traversal, Logic, Dynamic Programming , Branch and Bound, Graph Models

Algorithm	Time Complexity	FLOPs-Non-FLOPs Ratio	Memory Access Behaviors	Memory Access Frequency	Result based Mapping	Tesseract based Mapping
GEM	High	High	Bad	High	$((P, N)toG)$	$(P)$
LAVAMD	High	High	Good	Low	$((G, P)toG)$	$(G)$
SWAT	High	Low	Bad	High	$(N)$	$(N)$
NW	High	Low	Bad	High	$(N)$	$(N)$
SRAD	High	Low	Good	High	$(G)$	$(G)$
CFD	Low	High	Bad	High	$(G, N)$	$(P)$
HMM (Vary Observation)	High	Low	Good to Bad	High	$((G, N)toP)$	$(GtoP)$
K-means (Vary Data Size)	High	High	Good	High	$((G, P)toG)$	$(G)$
K-means (Vary Max Cluster)	High	High	Good	High	$((G, P)toG)$	$(G)$
K-means (Vary Dimension)	High	High	Bad	High	$(G, P)$	$(P)$
LUD	High	Low	Good	High	$(G, N)$	$(G)$
SPMV (Vary Data Size)	High	High	Good	Low	$((G, N)toG)$	$(G)$
Triad	Low	High	Good	High	$(G, P)$	$(G)$
Reduction	Low	High to Low	Bad	Low	$(G)$	$(P)$
FFT	High	High	Good	Low	$((G, P)toG)$	$(G)$
SCAN	High	Low	Good	High	$((G, P)toG)$	$(G)$
Stencil2D	Low	High	Good	High	$(G)$	$(G)$
TDM	High	Low	Bad	High	$(N)$	$(N)$
CRC	Low	Low	Good	Low	$(N)$	$(N)$
Radix Sort	High	High	Bad	High	$(G)$	$(P)$
BFS	High	Low	Bad	High	$((G, N)toG)$	$(N)$
A*	High	Low	Bad	High	$((N, P), N)$	$(N)$

Table V: Algorithm Classification and Mapping to Single (G)PU, Single (P)hi and Single (N)ode.

# TESSERACT

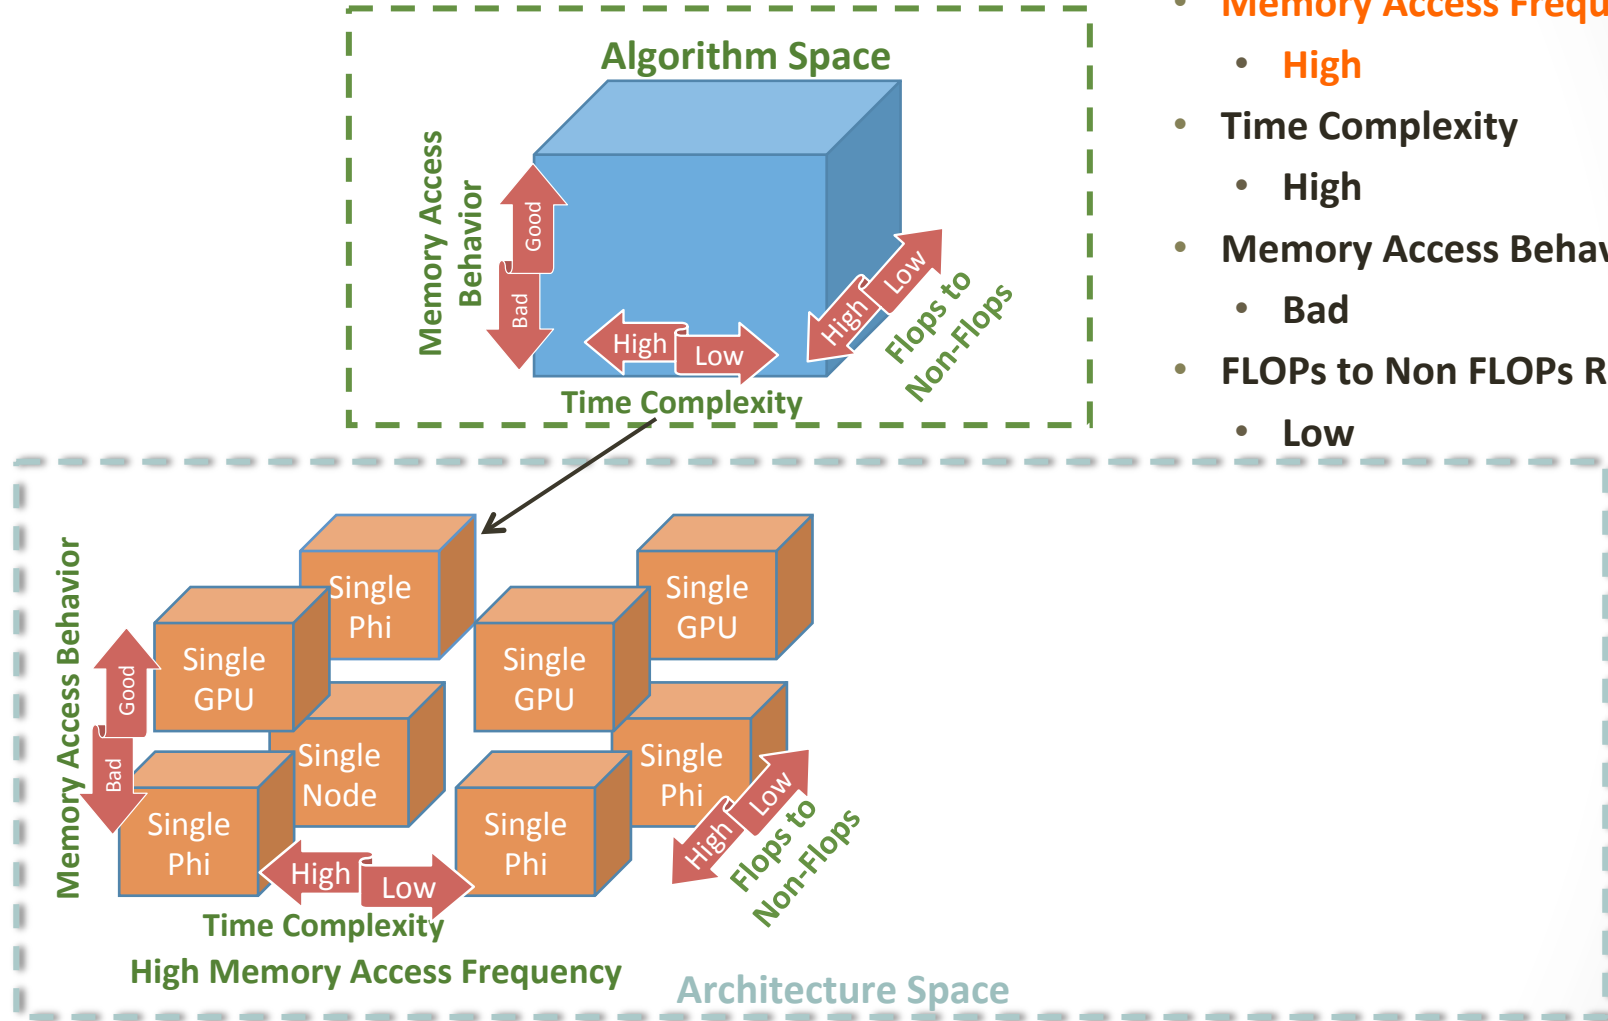


## SWAT

- Memory Access Frequency
  - High
- Time Complexity
  - High
- Memory Access Behavior
  - Bad
- FLOPs to Non FLOPs Ratio
  - Low

Architecture Space

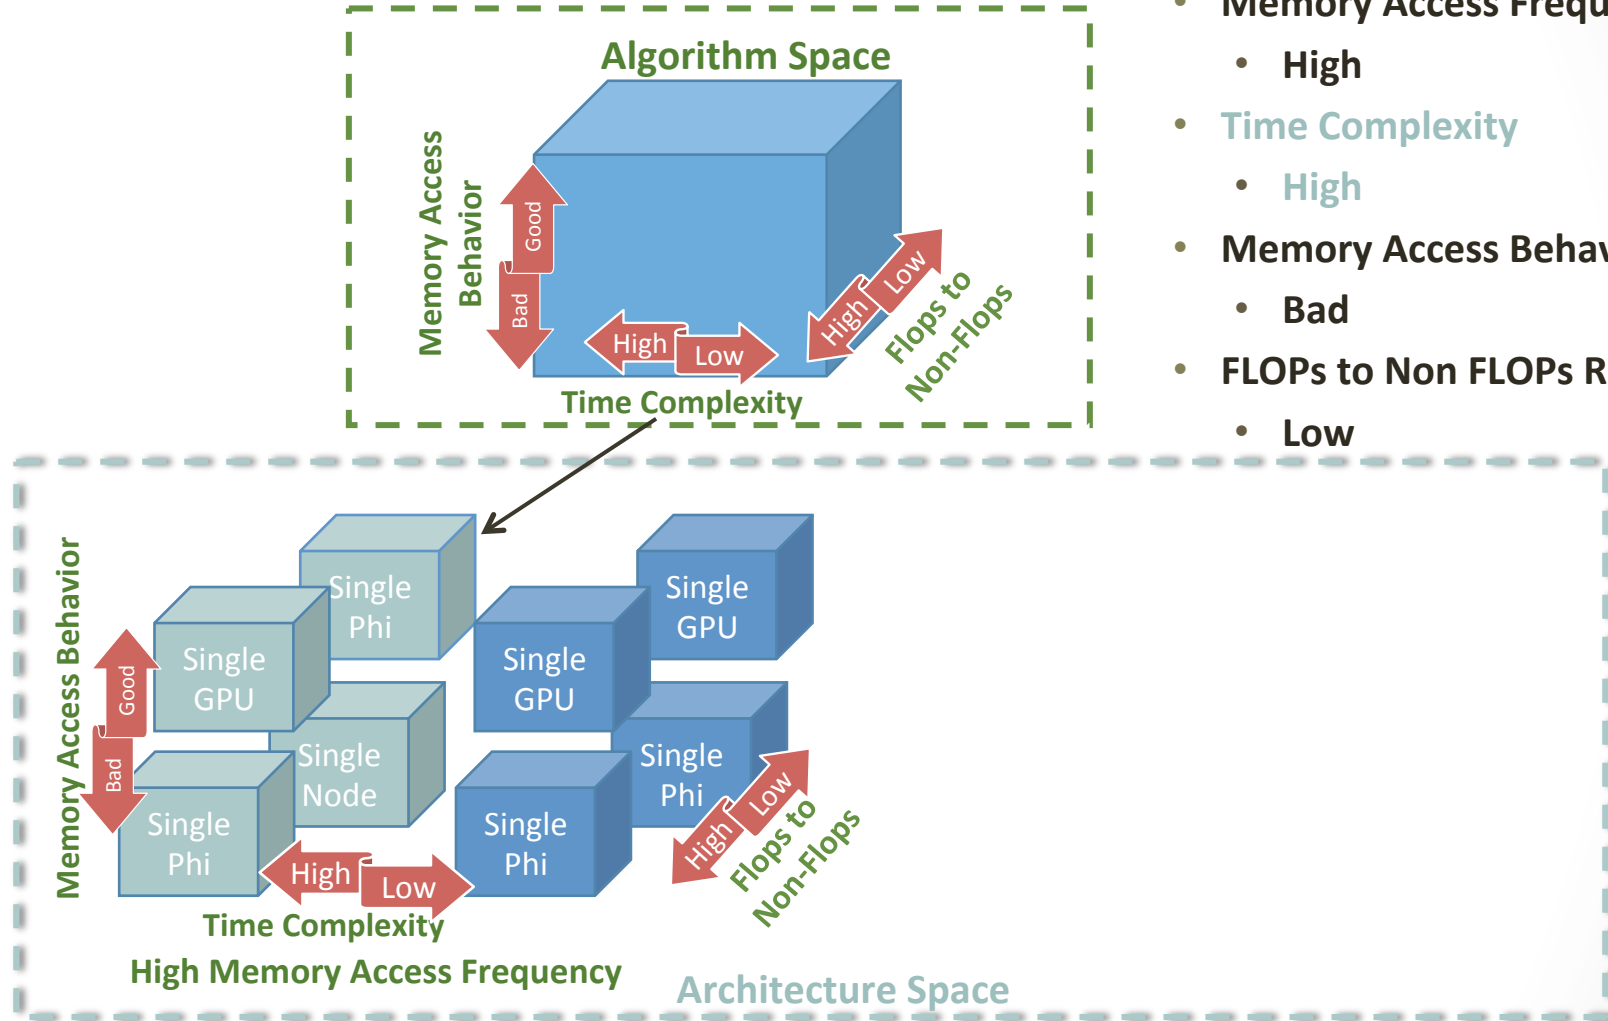
# TESSERACT



## SWAT

- **Memory Access Frequency**
  - High
- Time Complexity
  - High
- Memory Access Behavior
  - Bad
- FLOPs to Non FLOPs Ratio
  - Low

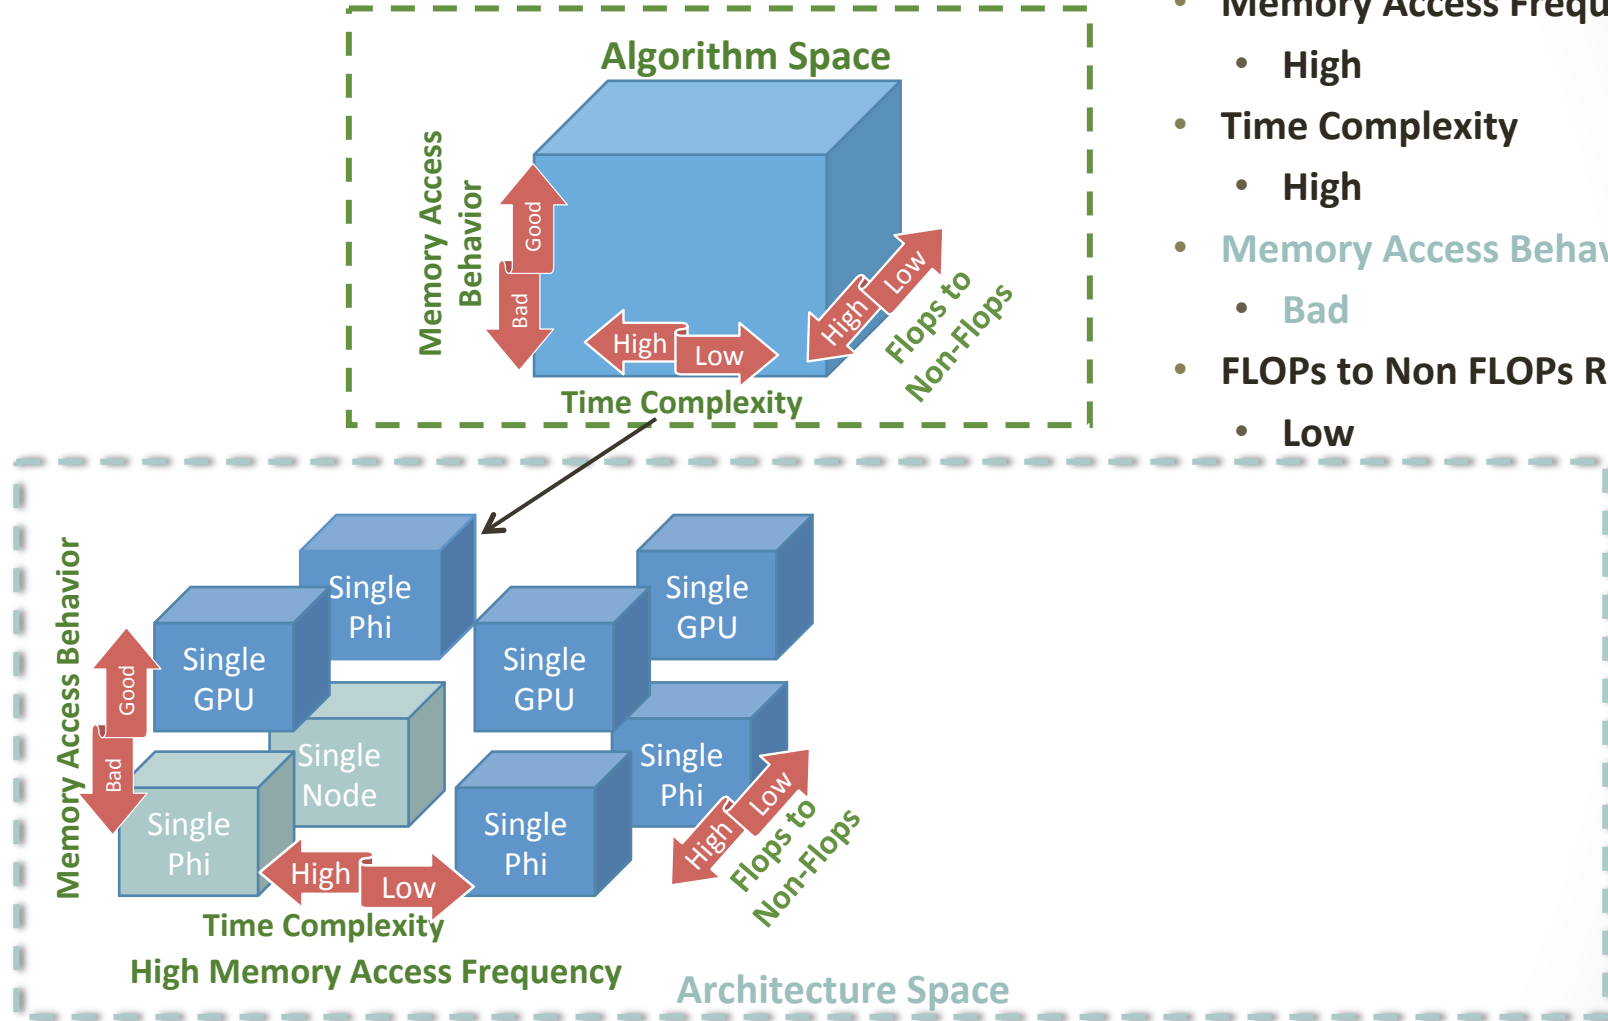
# TESSERACT



## SWAT

- Memory Access Frequency
  - High
- Time Complexity
  - High
- Memory Access Behavior
  - Bad
- FLOPs to Non FLOPs Ratio
  - Low

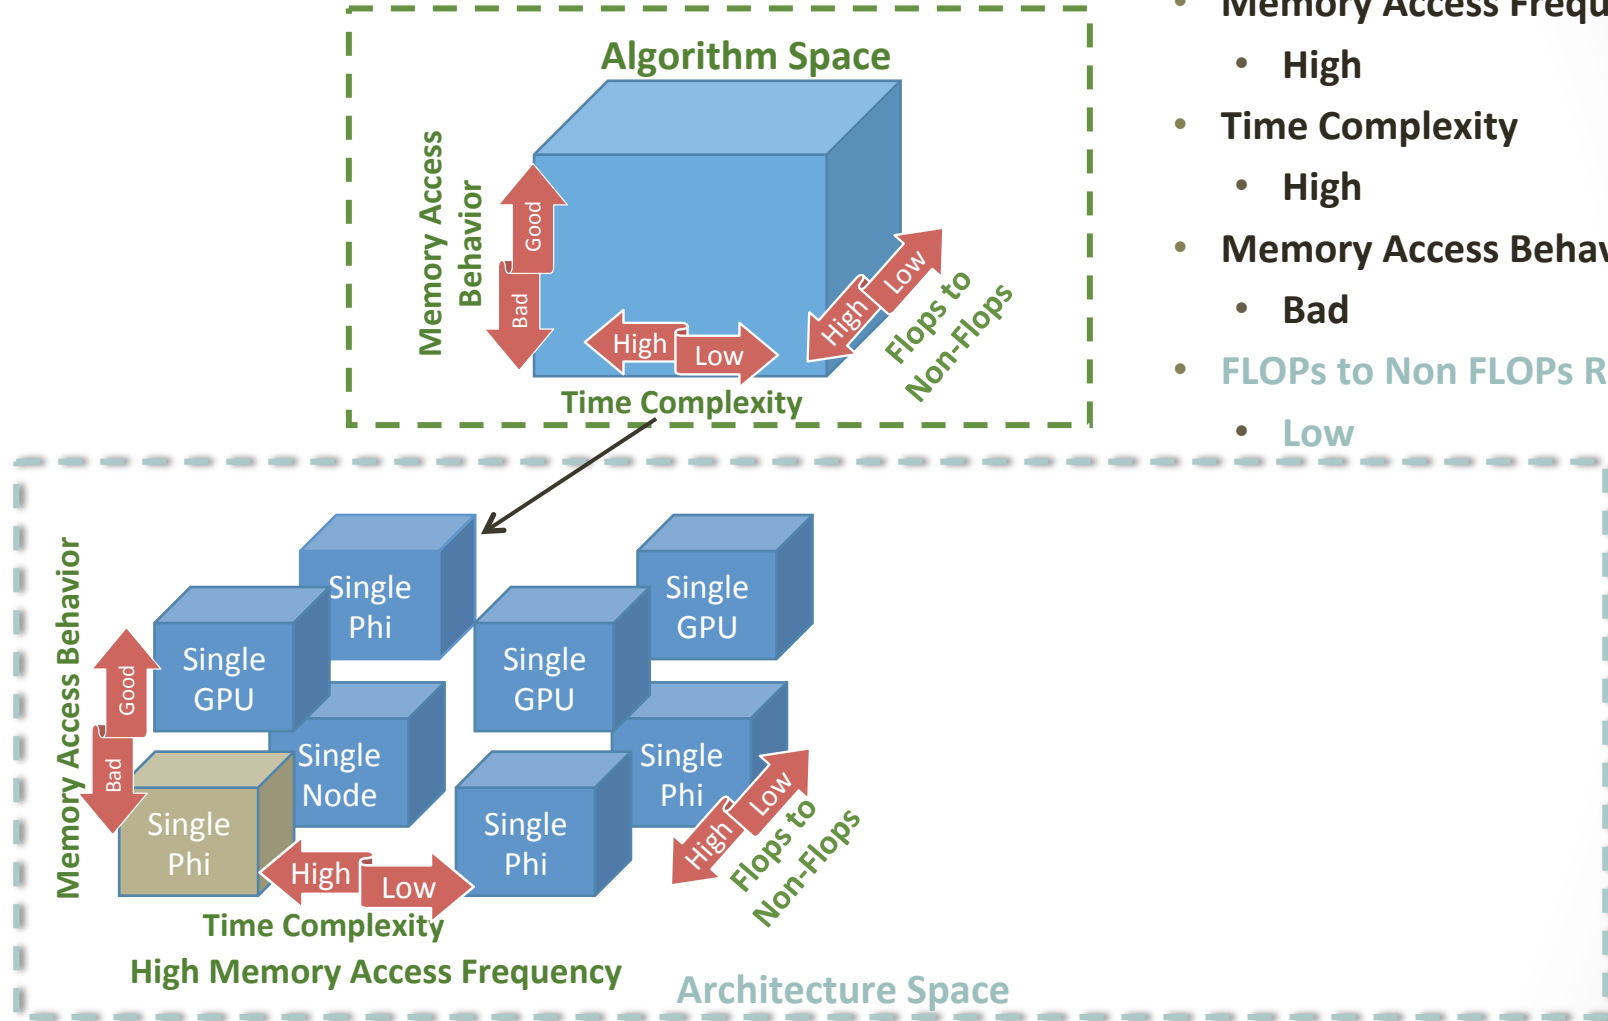
# TESSERACT



## SWAT

- Memory Access Frequency
  - High
- Time Complexity
  - High
- Memory Access Behavior
  - Bad
- FLOPs to Non FLOPs Ratio
  - Low

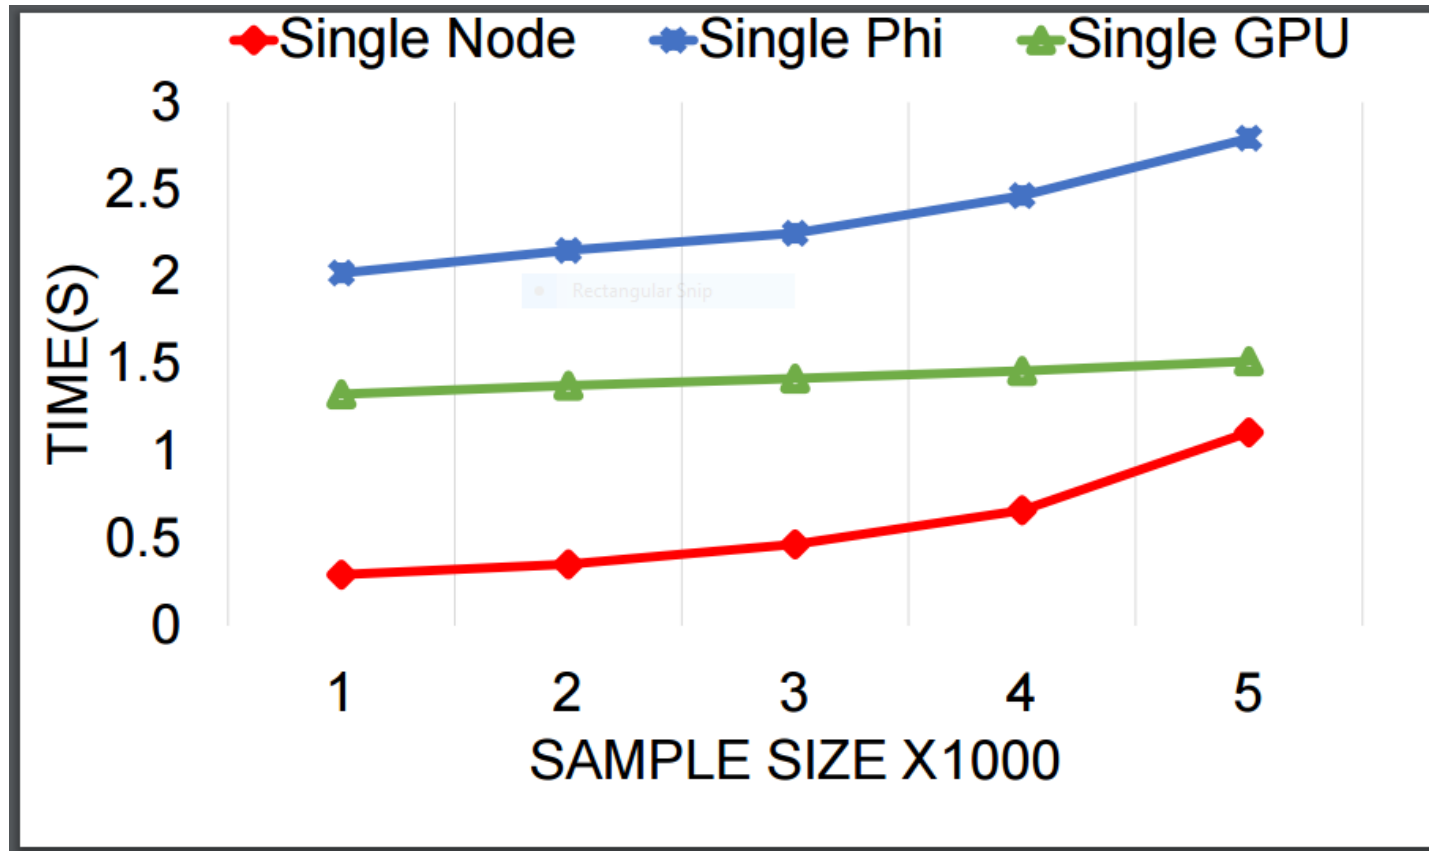
# TESSERACT




## SWAT

- Memory Access Frequency
  - High
- Time Complexity
  - High
- Memory Access Behavior
  - Bad
- FLOPs to Non FLOPs Ratio
  - Low

# SWAT

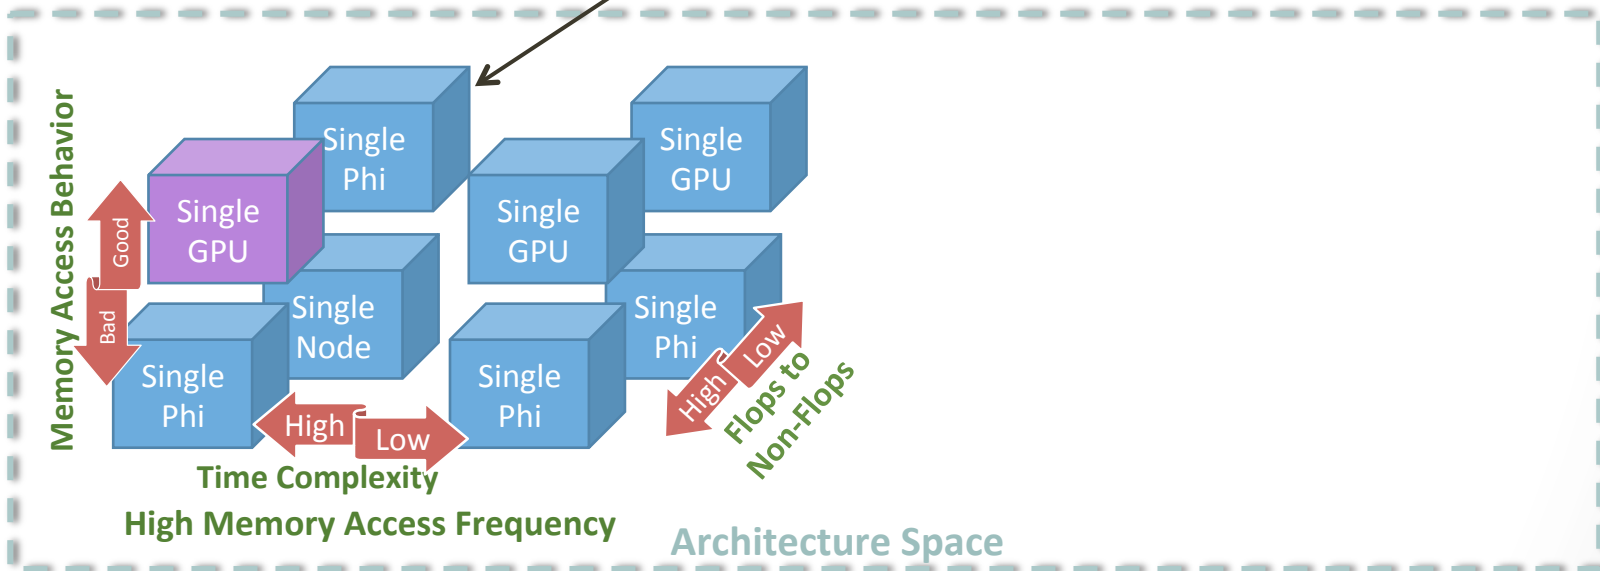


# TESSERACT

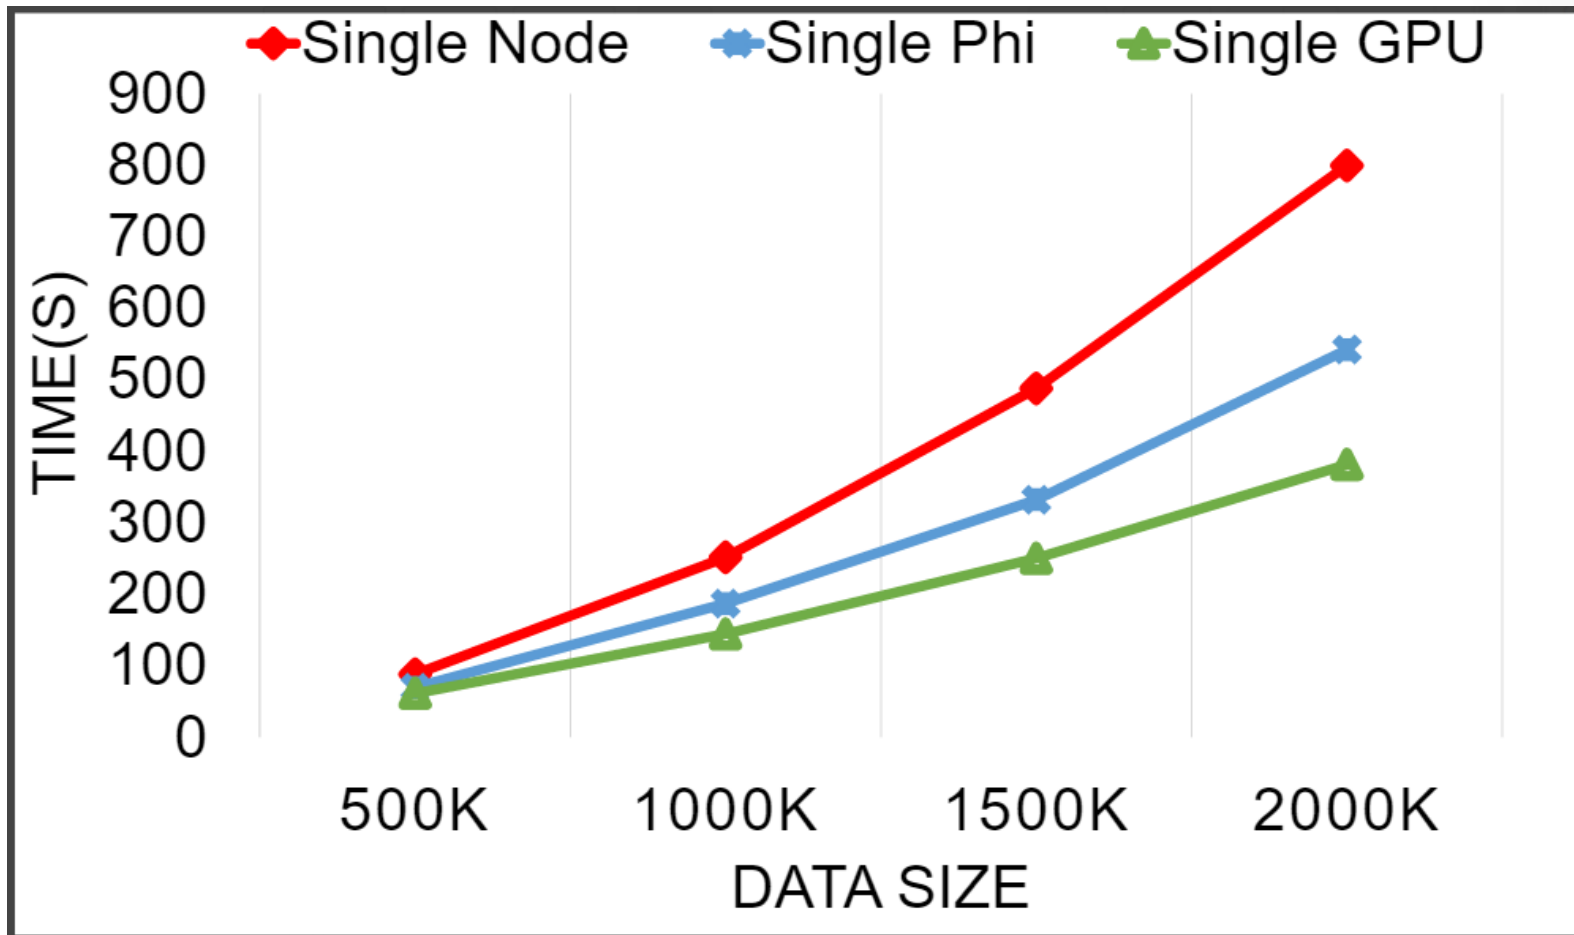


## K-Means Varying Data

- Memory Access Frequency
  - High
- Time Complexity
  - High
- Memory Access Behavior
  - Good
- FLOPs to Non FLOPs Ratio
  - High

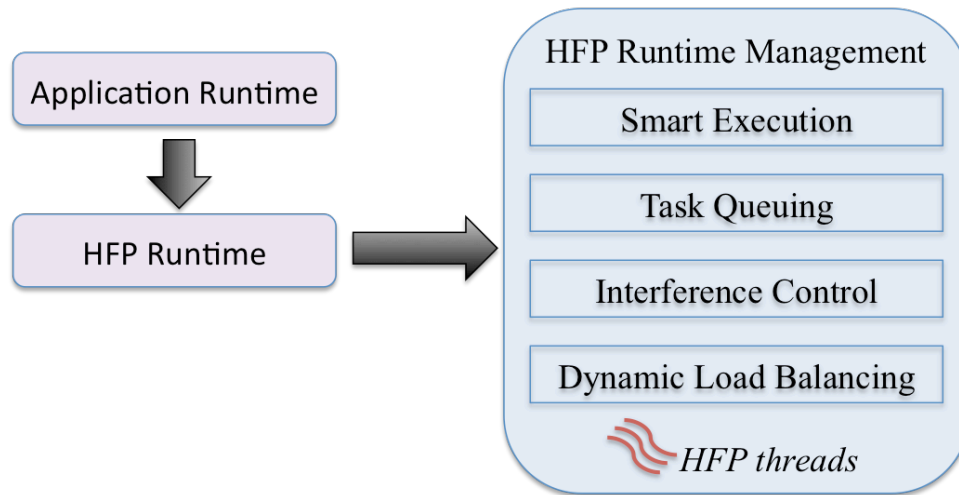


# KMEANS



# Functional Partitioning (FP) Framework

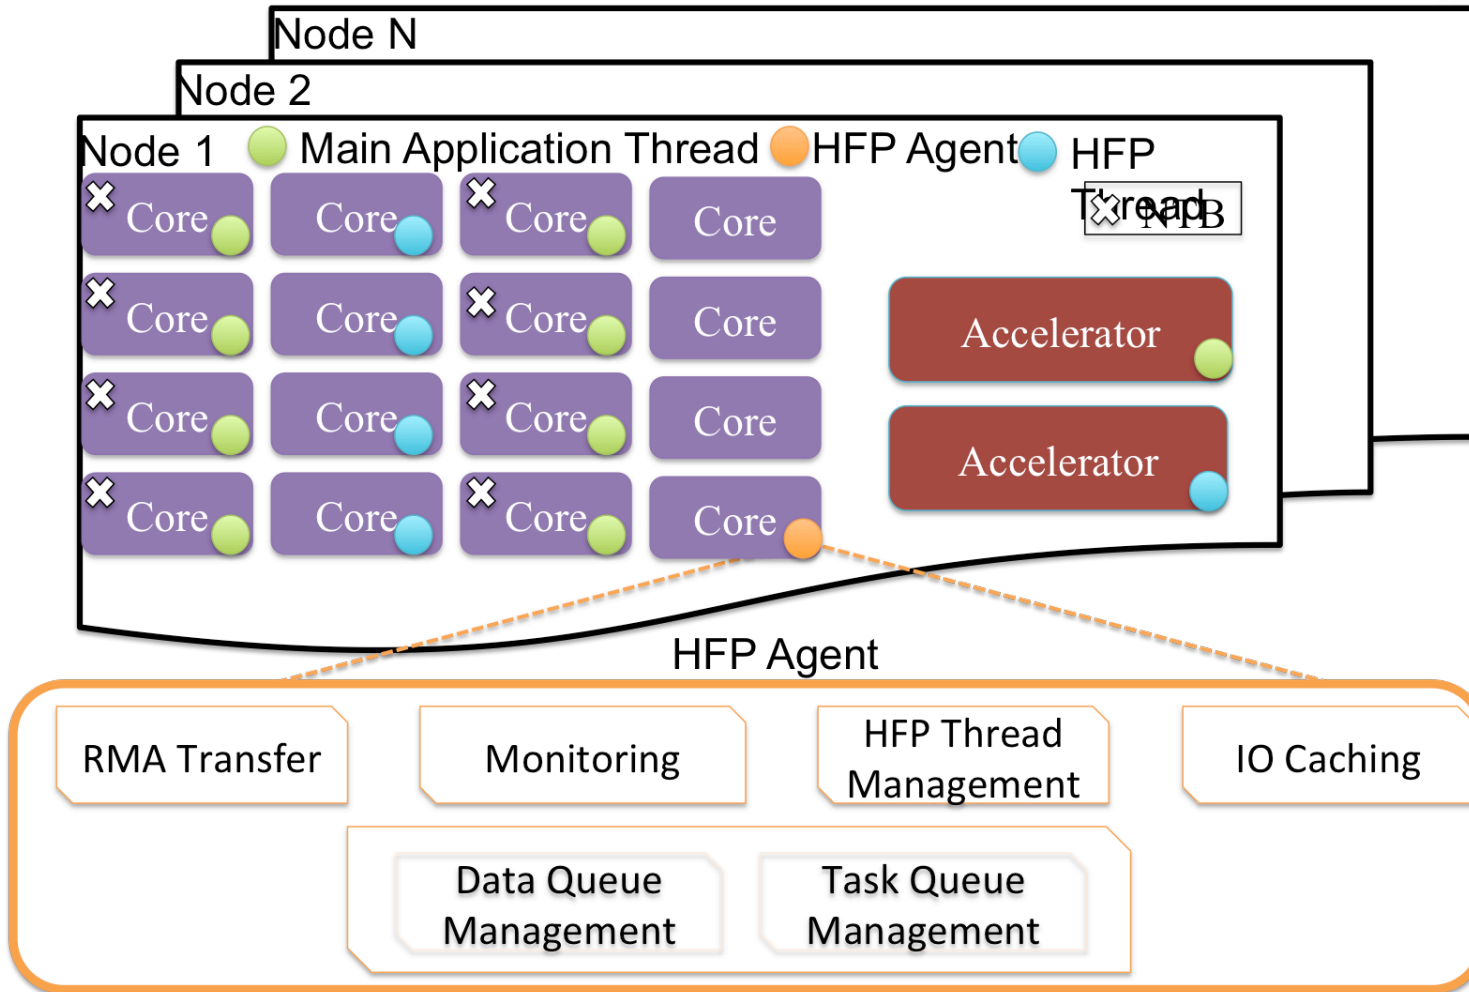
- C based framework designed to utilize heterogeneous architectures efficiently



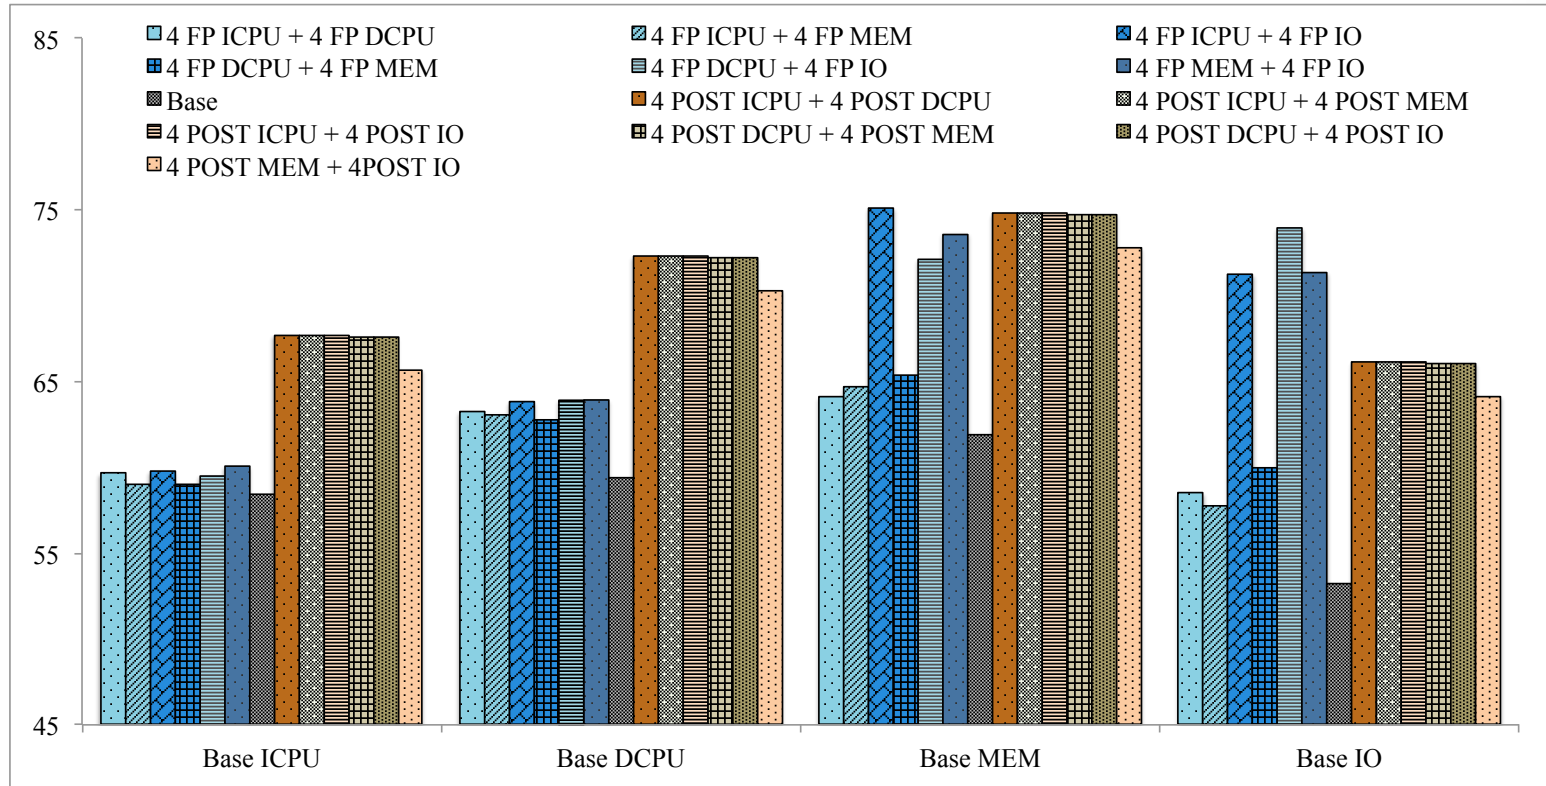
HFP Agent

- Specify multiple post-processing tasks performed concurrently based on resources
- FP services allow a complex work flow


# Framework Overview



# Micro-Benchmark Analysis



# Benchmark Performance Analysis



# Questions ?