

Active Storage with Analytics Capabilities and I/O Runtime System for Petascale Systems

**PI: Alok Choudhary, Professor and Chair
Electrical Engineering and Computer Science
Director: Center for Ultra-Scale Computing and
Security**

Northwestern University

choudhar@eecs.northwestern.edu

**Co-PIs: Rajeev Thakur and Rob Ross
Argonne National Labs**

Challenges in Scientific Knowledge Discovery (Performance and Productivity)

Scientific Data Management

- Data management
- Query of Scientific DB
- Performance optimizations

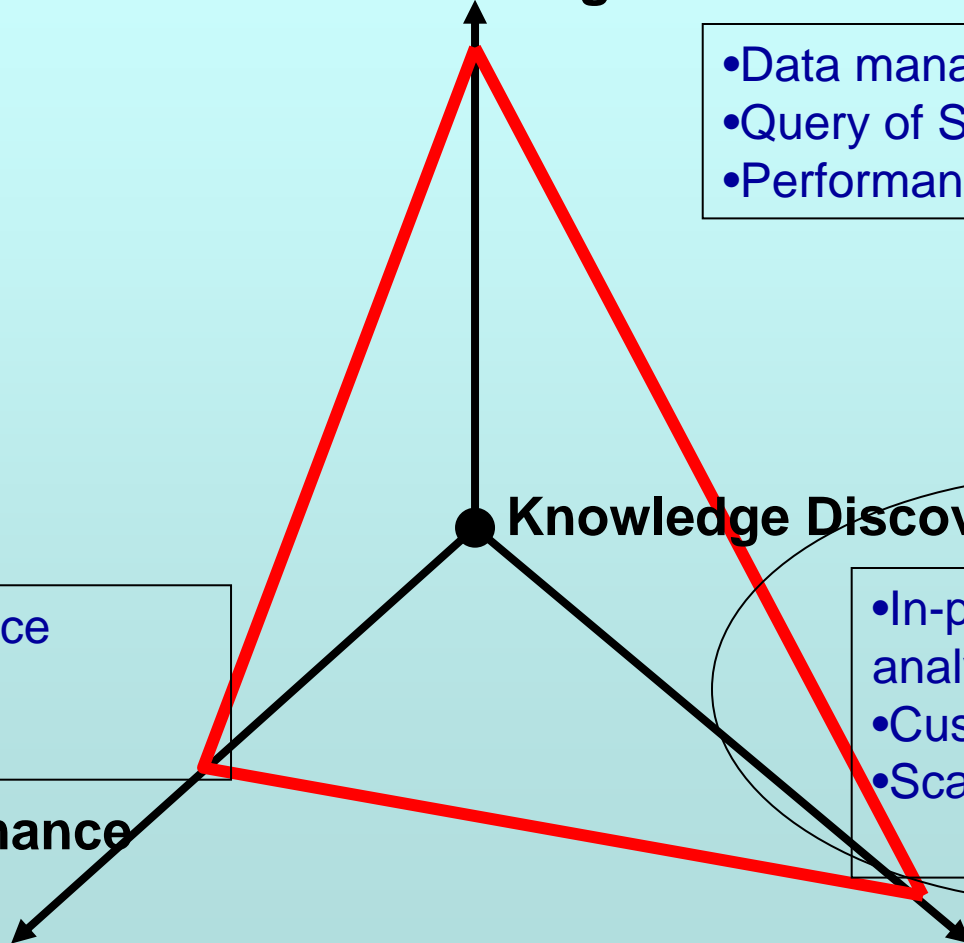
Knowledge Discovery

- High-level interface
- proactive
- What not How?

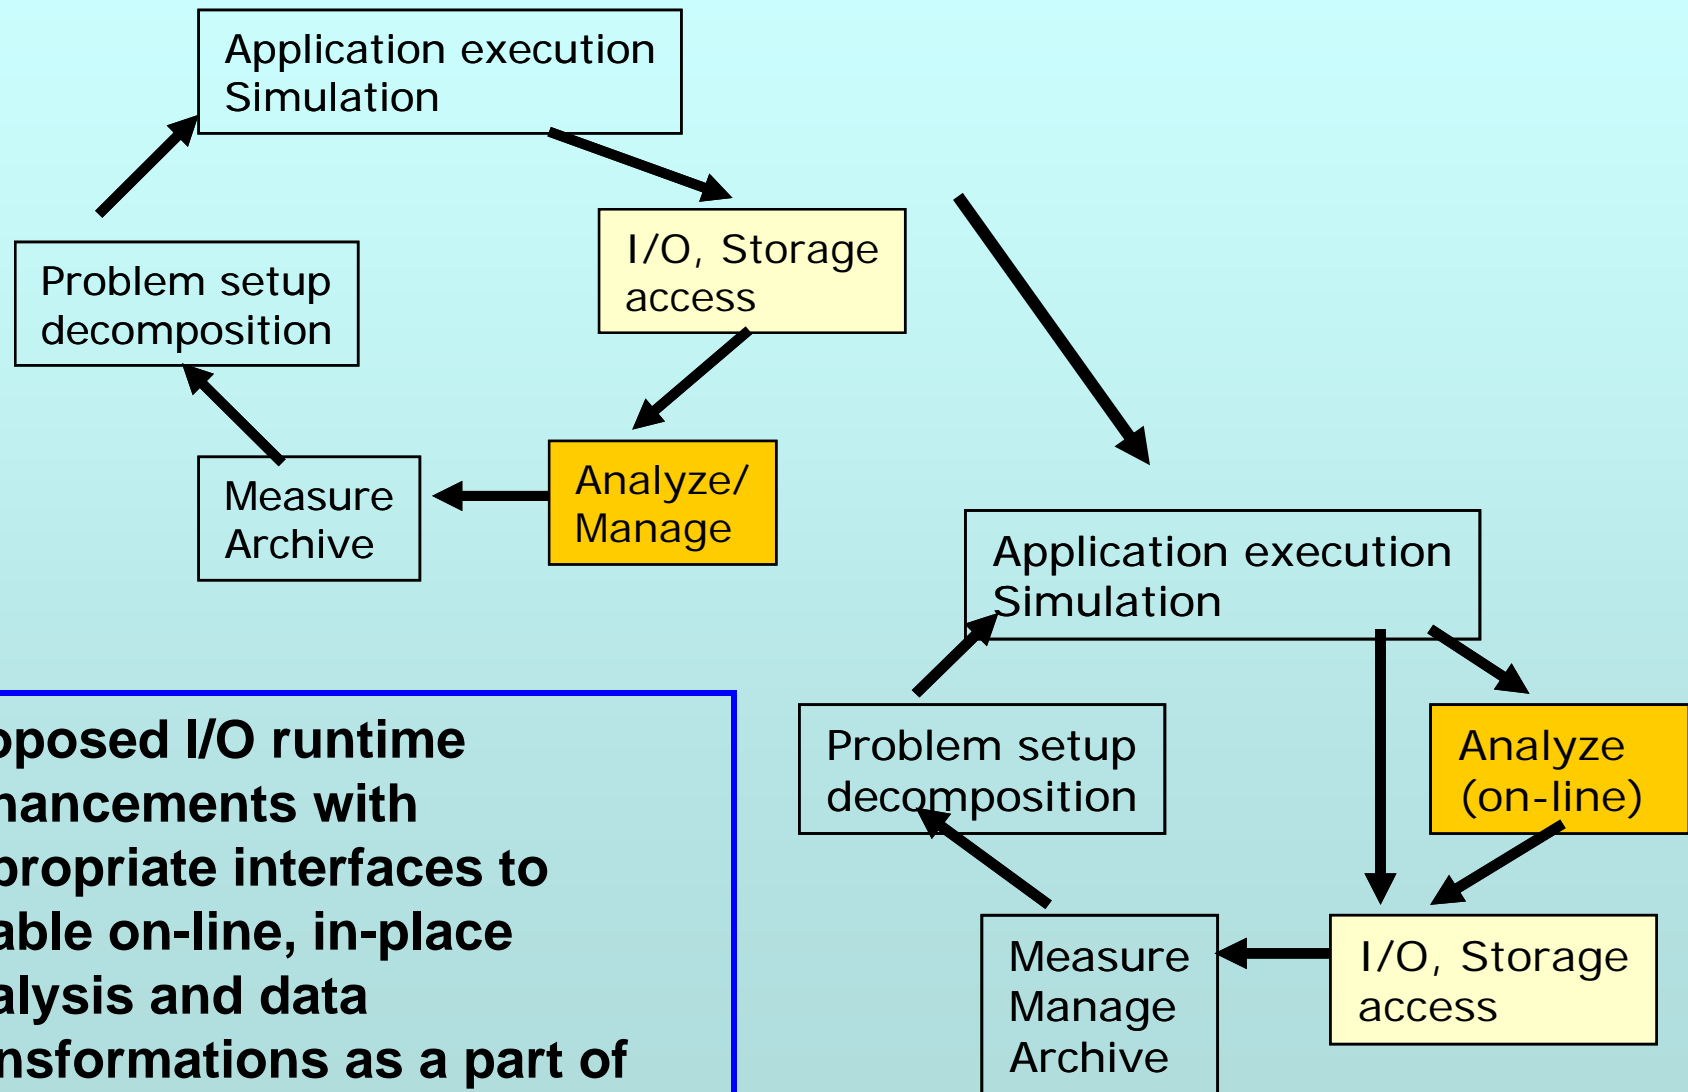
- In-place and on-line analytics
- Customized acceleration
- Scalable Mining

High-Performance
I/O

Analytics and
Mining



On-line Analysis and Its acceleration




Proposed I/O runtime enhancements with appropriate interfaces to enable on-line, in-place analysis and data transformations as a part of the I/O runtime.

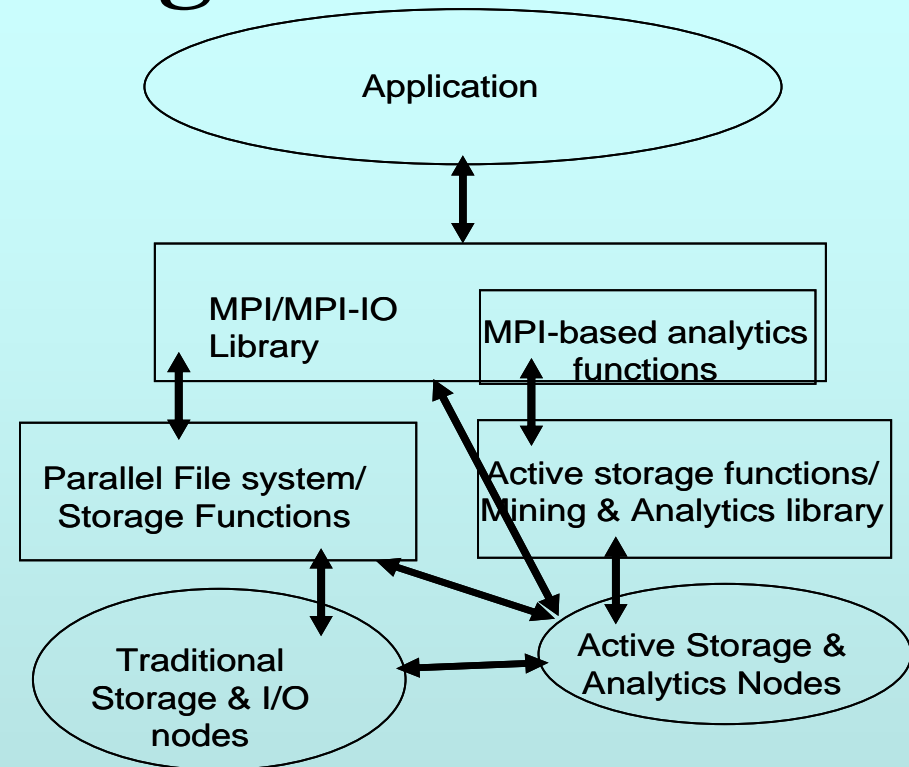
Project Objectives

- ❑ Parallel I/O runtime interface that will utilize customized active storage nodes to perform I/O operations
 - ❑ Parallel I/O runtime interface that will enable data analytics, mining, statistical operations within active storage
 - ❑ Build a prototype active storage cluster
 - ❑ Develop a library of functions, evaluate their performance and scalability
-

Concept and Software Architecture for using Active Storage

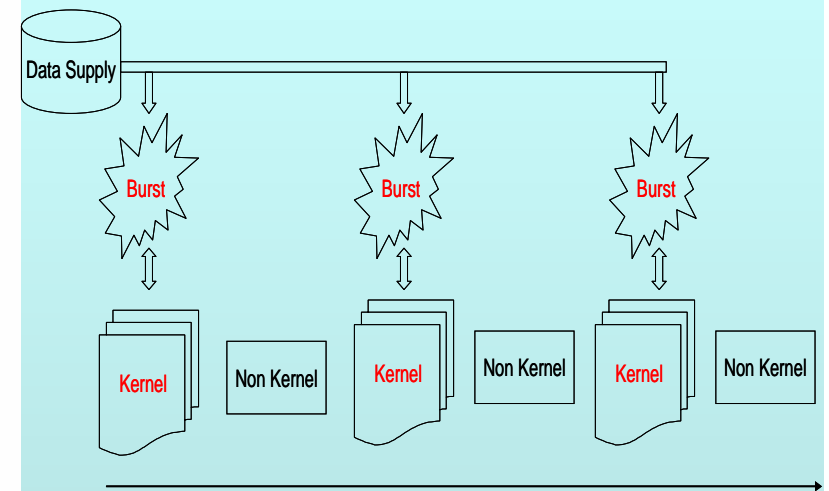
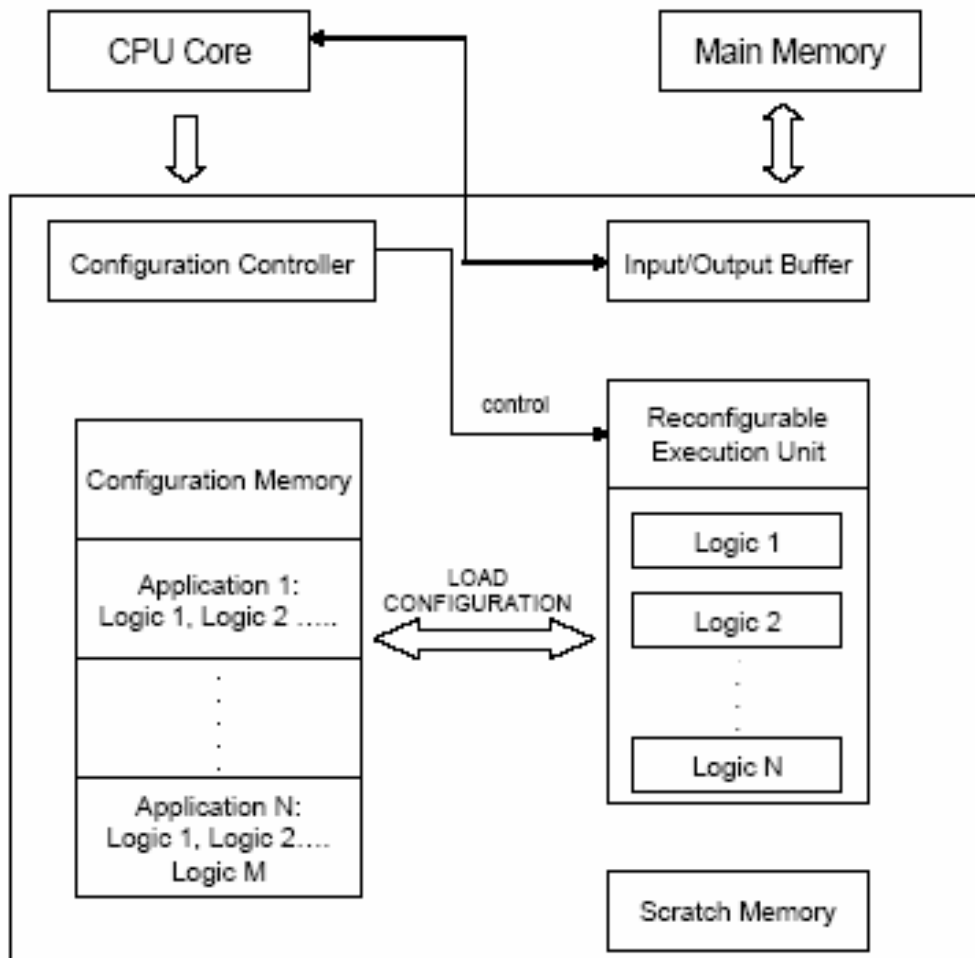


Customized nodes for I/O and active storage (conceptual)



Proposed runtime layers for compute nodes, I/O nodes and on active storage nodes.

Programming Model for ASA

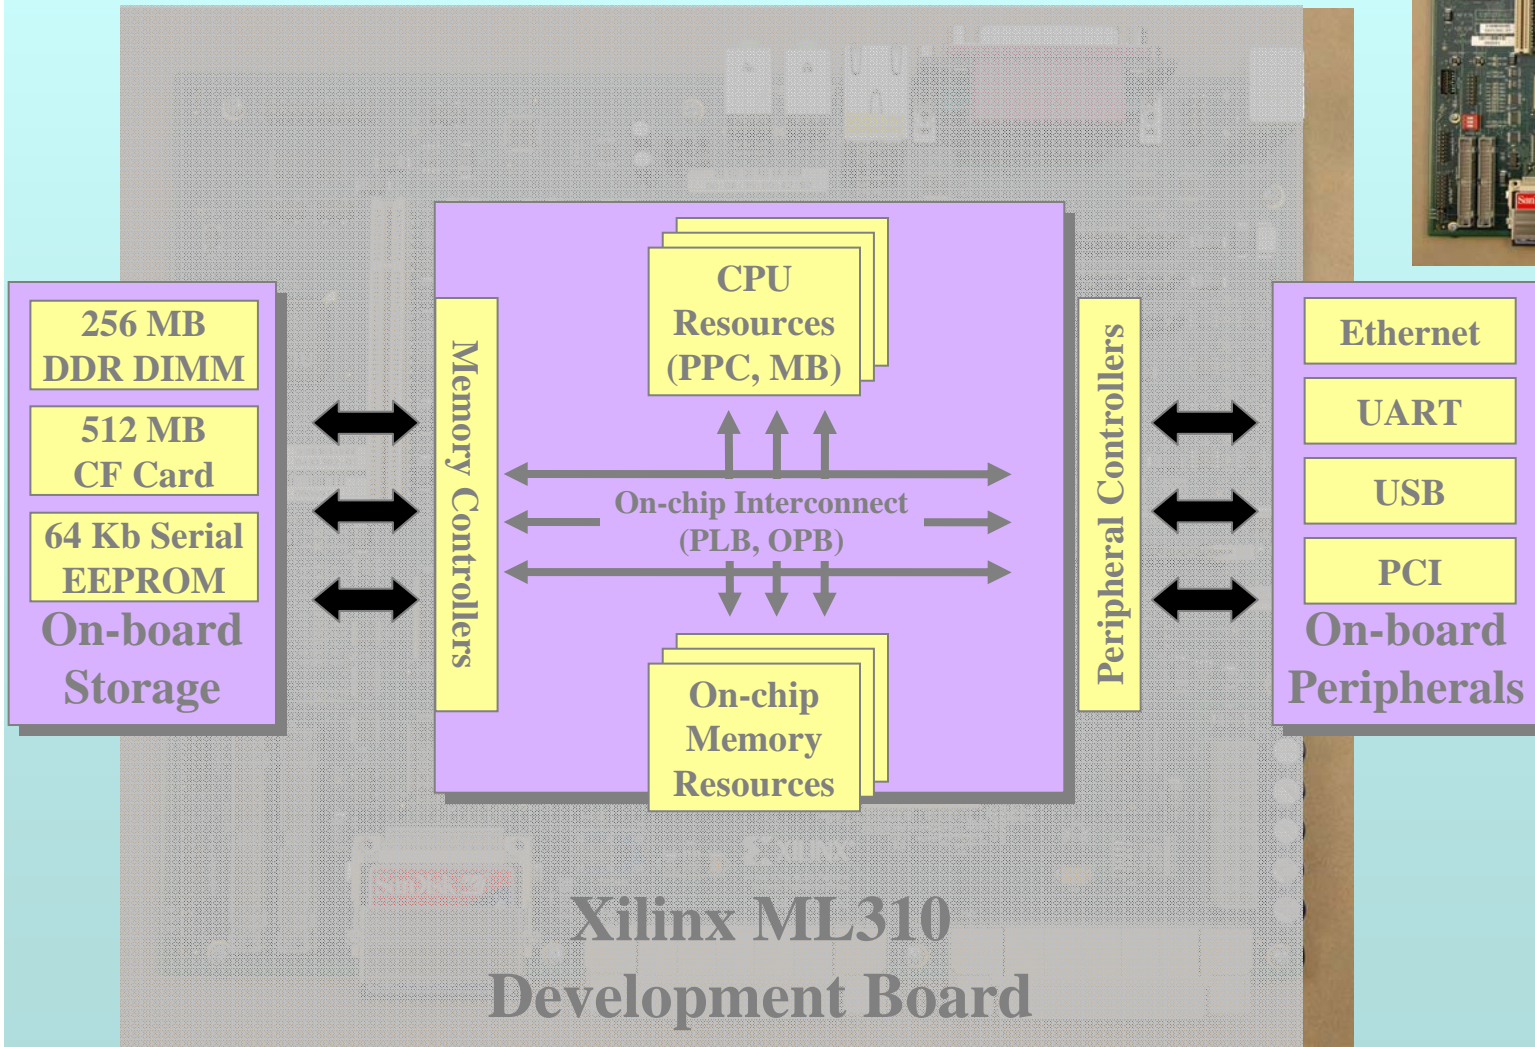


Kernels and Functions (examples)

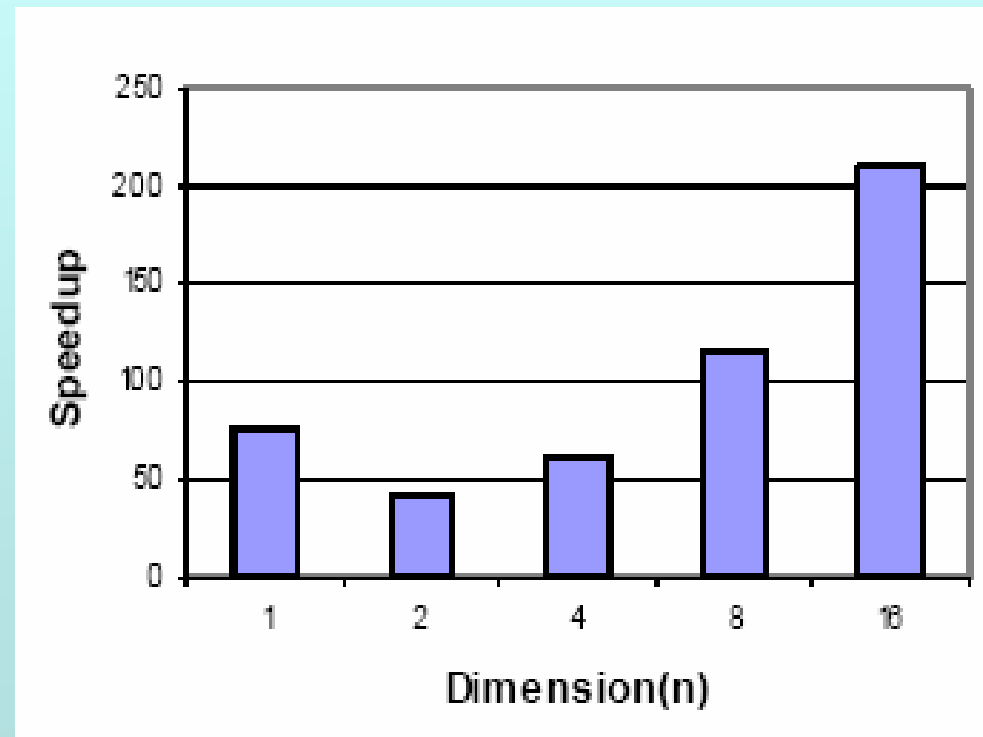
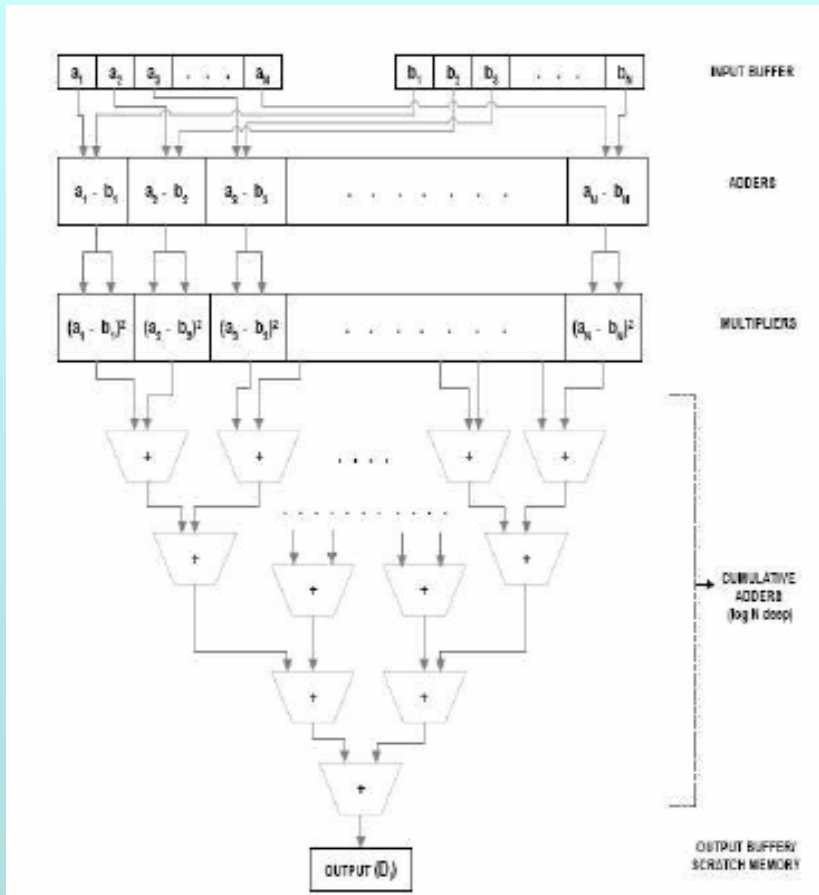
<i>Application</i>	Top 3 Kernels (%)			Sum %
	Kernel 1 (%)	Kernel 2 (%)	Kernel 3 (%)	
k-Means	distance (68%)	clustering (21%)	minDist (10%)	99
Fuzzy k-Means	clustering (58%)	distance (39%)	fuzzySum (1%)	98
BIRCH	distance (54%)	variance (22%)	redistribution (10%)	86
HOP	density (39%)	search (30%)	gather (23%)	92
Naïve Bayesian	probCal (49%)	variance (38%)	dataRead (10%)	97
ScalParC	classify (37%)	giniCalc (36%)	compare (24%)	97
Apriori	subset (58%)	dataRead (14%)	increment (8%)	80
Eclat	intersect (39%)	addClass (23%)	invertClass (10%)	72

Illustrative top kernels in terms of computation requirements for several illustrative analytics and data mining algorithms

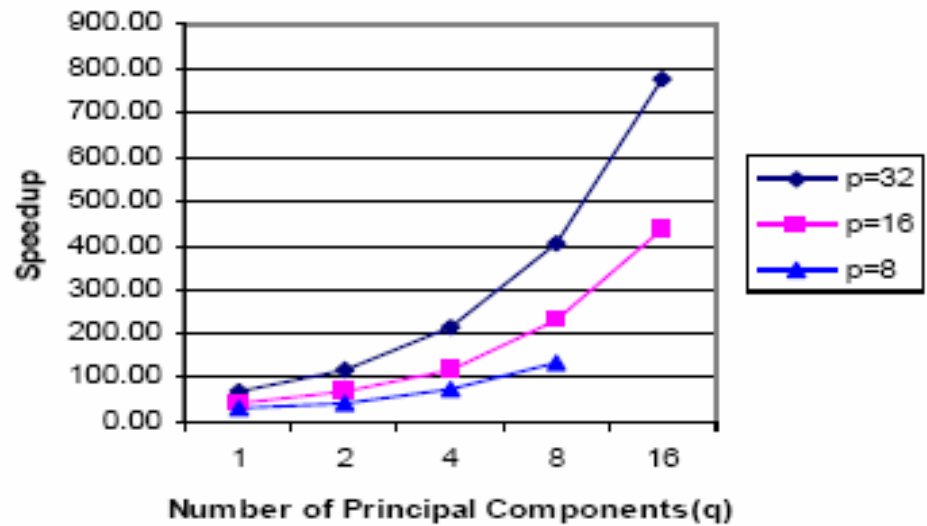
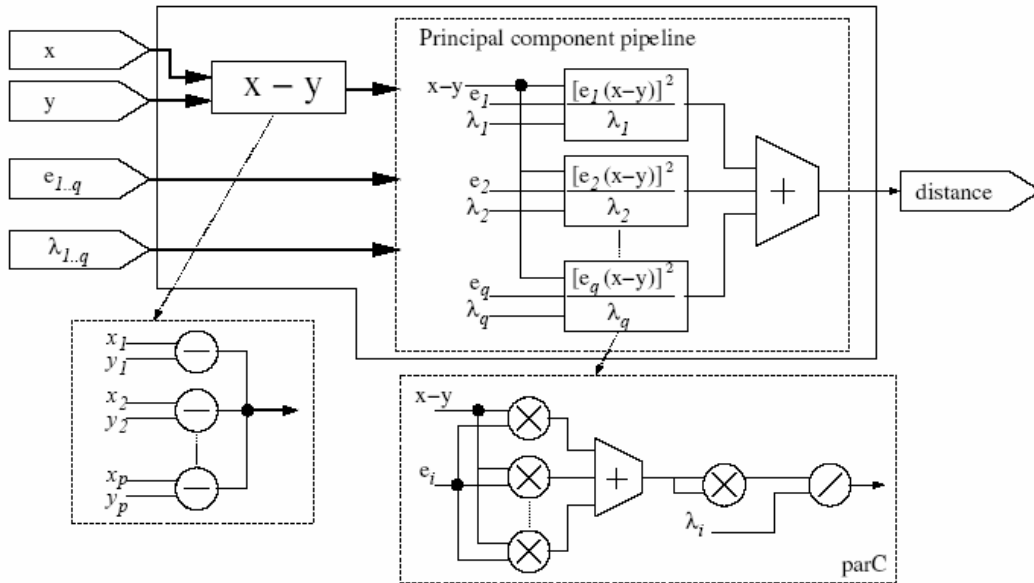
Example: Experimental Acceleration Hardware



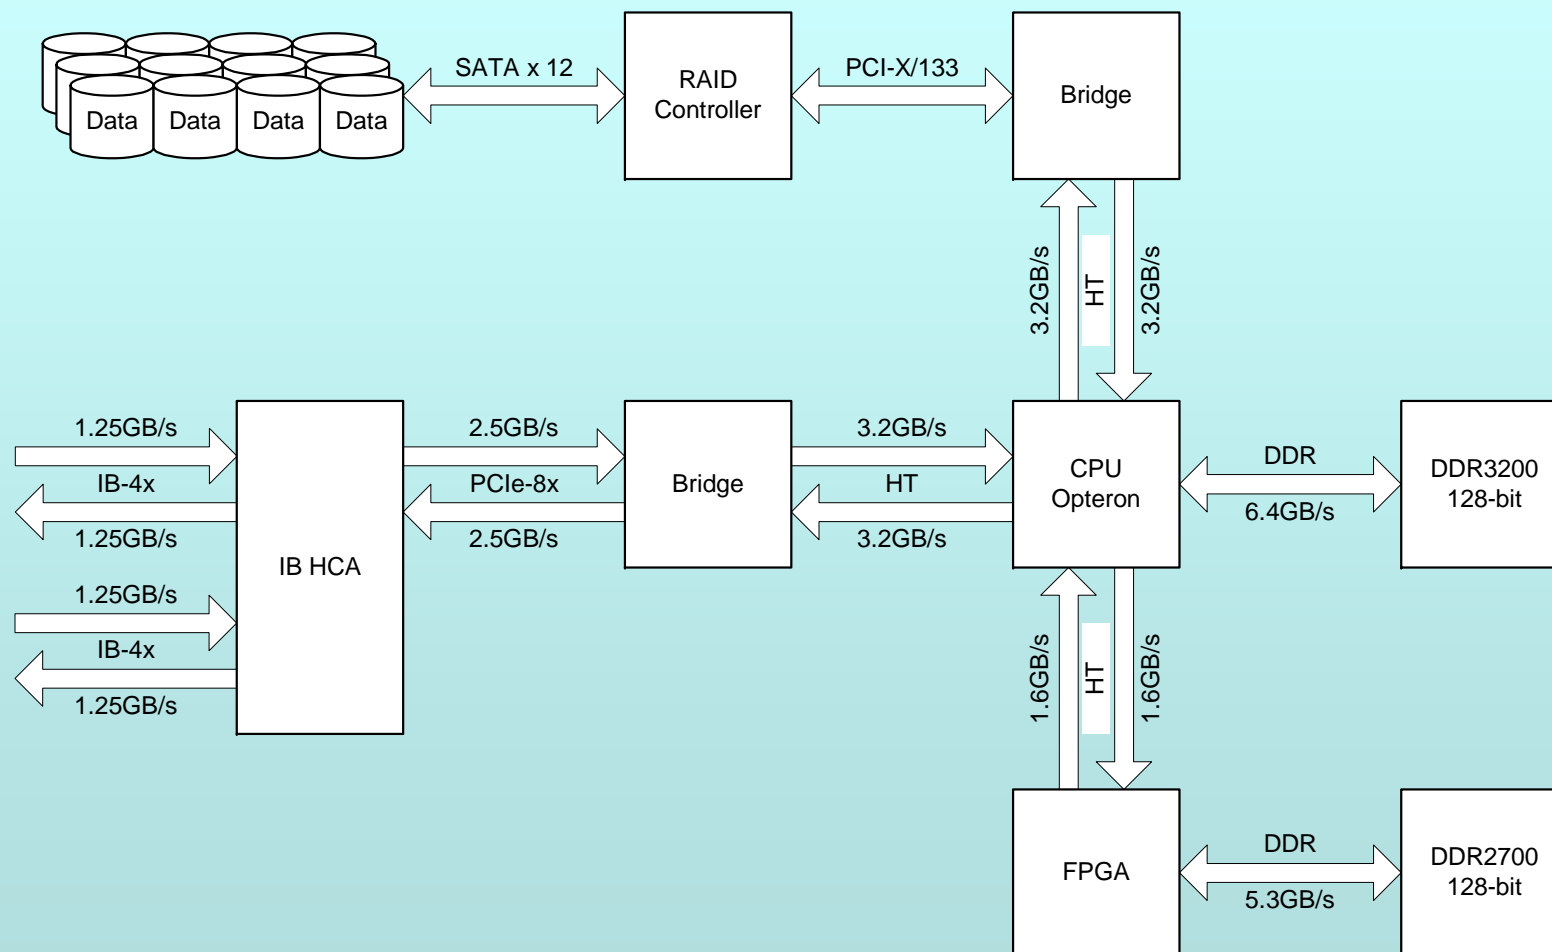
Distance kernel in Clustering data mining: Speedup over a 2.4GHz AMD Opteron



PCA Illustration



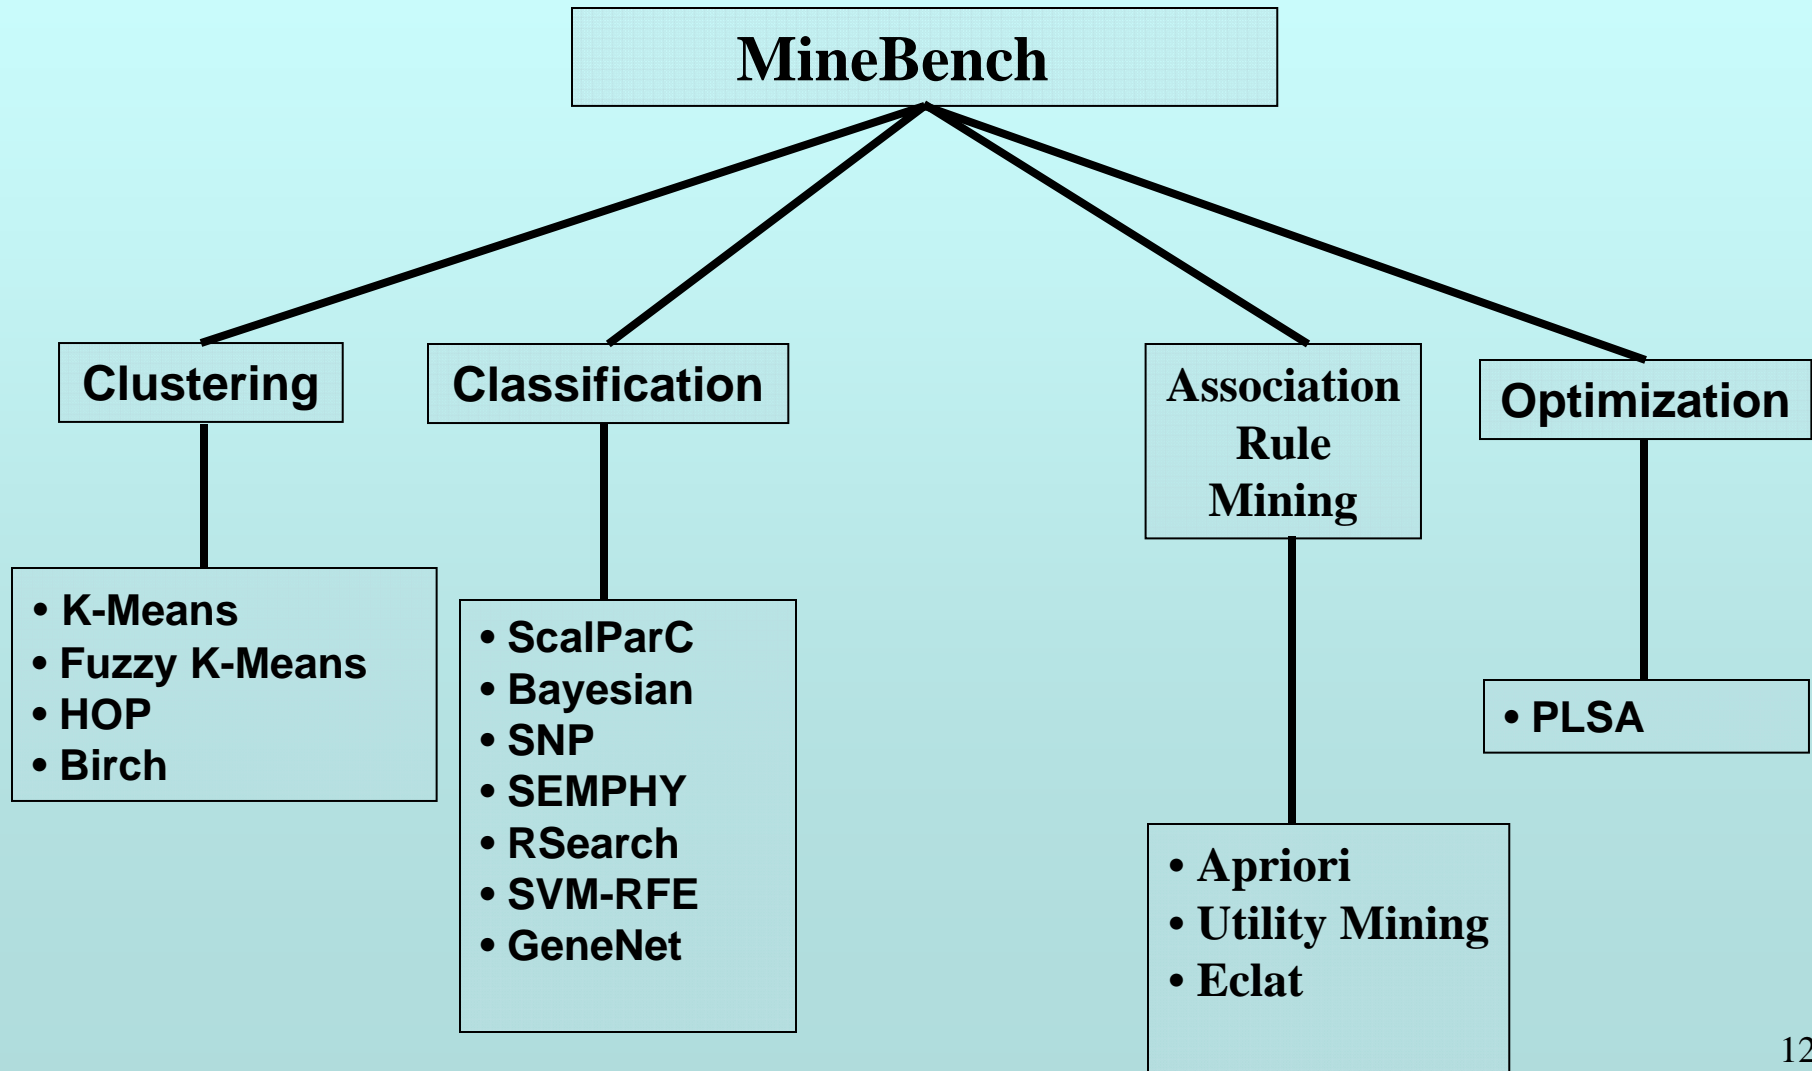
Node of an active storage cluster



Design still under revision

NU-Minebench Benchmark Suite

<http://cucis.ece.northwestern.edu/projects/DMS>



MineBench Overview

Non-bioinformatics workload, includes applications from:

- a) Decision trees
- b) Clustering/
Hierarchical Clus.
- c) Utility Mining
- d) Predictive Modeling
- e) Market Basket Analysis

Data taken from:

- a) Image processing
- b) Astrophysics
- c) Grocery chain
- d) Pharmaceutical

Bioinformatics workload (algorithms used in other fields as well)

Application	Instruction Count (billions)				Size (kB)
	1 Processor	2 Processors	4 Processors	8 Processors	
ScalParC	23.664	24.817	25.550	27.283	154
Naïve Bayesian	23.981	N/A	N/A	N/A	207
K-means	53.776	54.269	59.243	77.026	154
Fuzzy K-means	447.039	450.930	477.659	564.280	154
HOP	30.297	26.920	26.007	26.902	211
BIRCH	15.180	N/A	N/A	N/A	609
Apriori	42.328	42.608	43.720	47.182	847
Eclat	15.643	N/A	N/A	N/A	2169
Utility	13.640	19.902	20.757	22.473	853
SNP	429.703	299.960	267.596	241.680	14016
GeneNet	2,244.470	2,263.410	2,307.663	2,415.428	13636
SEMPHY	2,344.533	2,396.901	1,966.273	2,049.658	7991
Rsearch	1,825.317	1,811.043	1,789.055	1,772.200	676
SVM-RFE	51.370	55.249	63.053	82.385	1336
PLSA	4,460.823	4,526.160	2,080.610	4,001.675	836