Scalable I/O, File System and Run-Time Software to Enable Large-Scale Data Analytics

Alok Choudhary, Professor and Chair
Dept. of Electrical Engineering and Computer Science
And Professpr Kellogg School of Management
Director: Center for Ultra-Scale Computing and Security
Northwestern University
choudhar@eecs.northwestern.edu

Acknowledgements:
NSF and DOE
NU Students and Collaborators: Ramanathan, Berkin, Dan Honbo, Sanchit Misra, Prabhat, Joe Zambreno (now faculty at Iowa State), Wei-Keng Liao (Research prof at NWU), Kui Gao (postdoc, NU)

External Collaborators: Rajeev Thakur and Rob Ross (ANL)

Performance and Knowledge Discovery

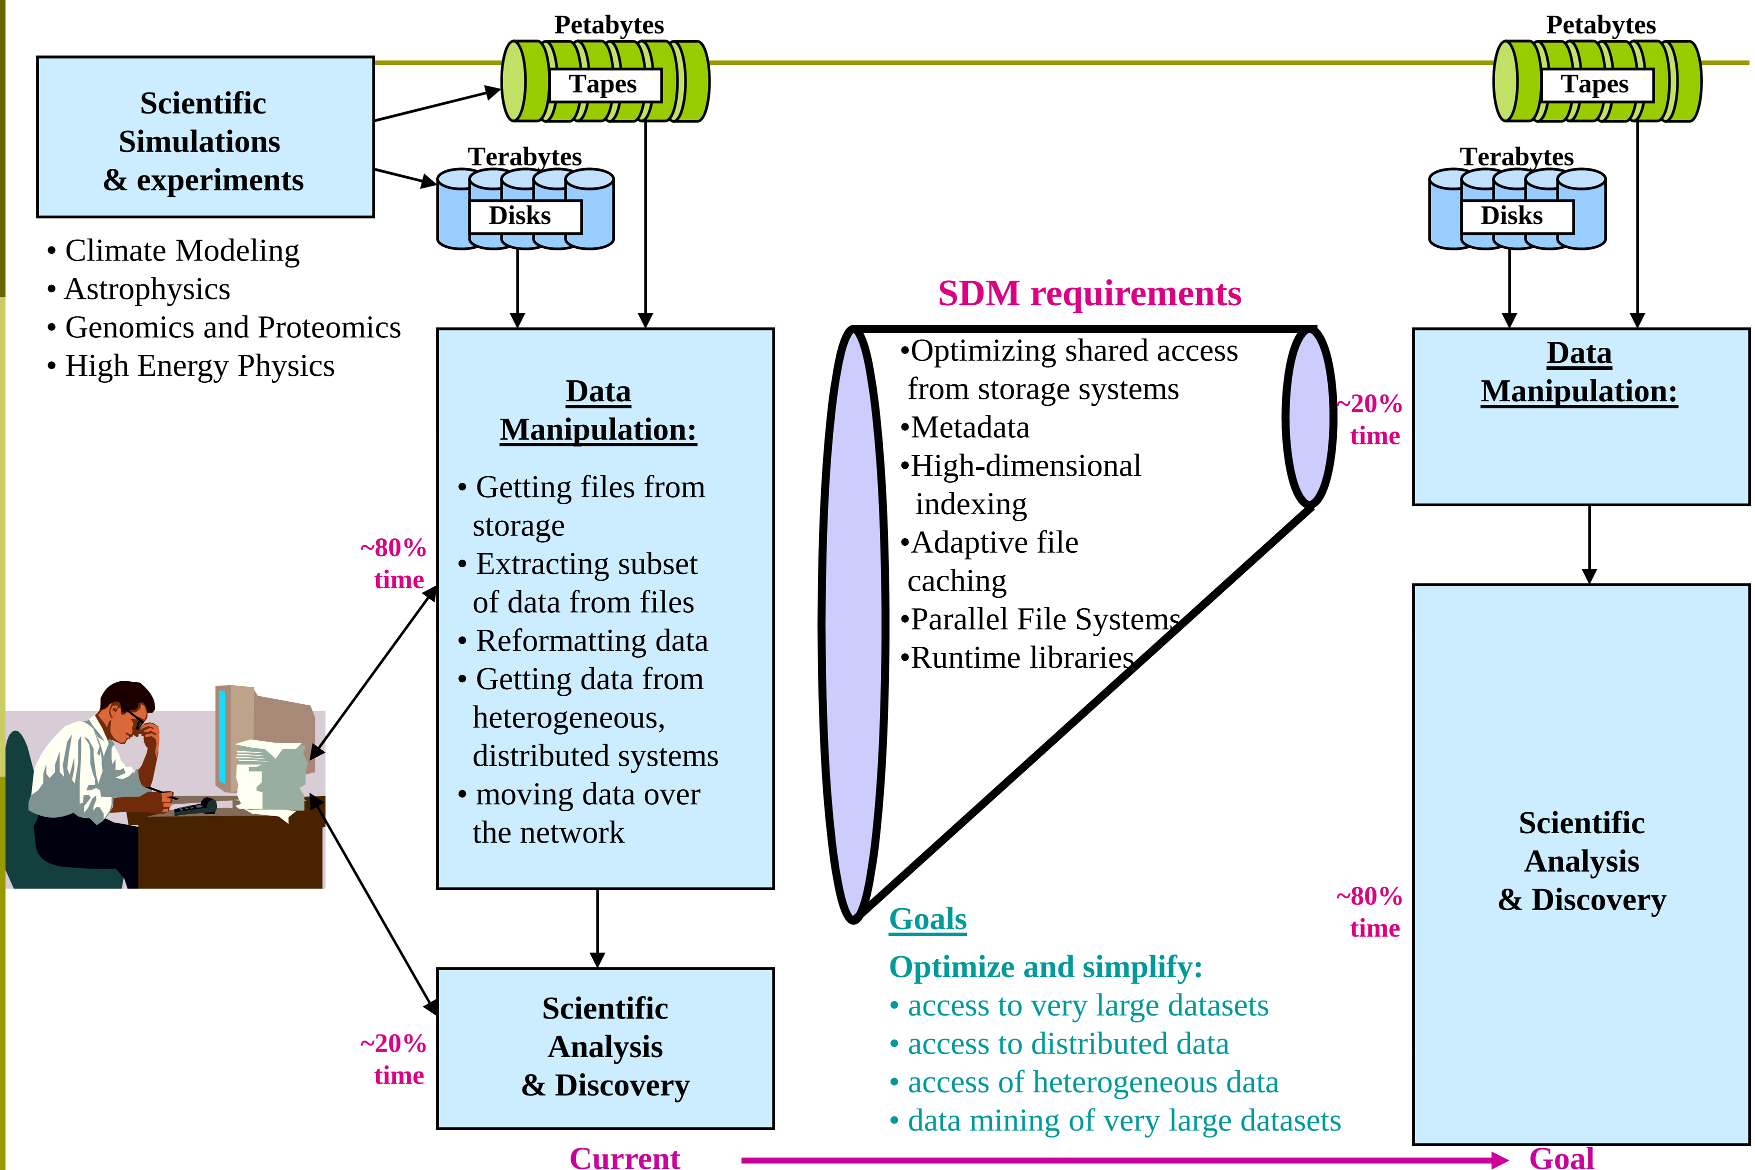

Challenges in Scientific Knowledge Discovery (Performance and Productivity)

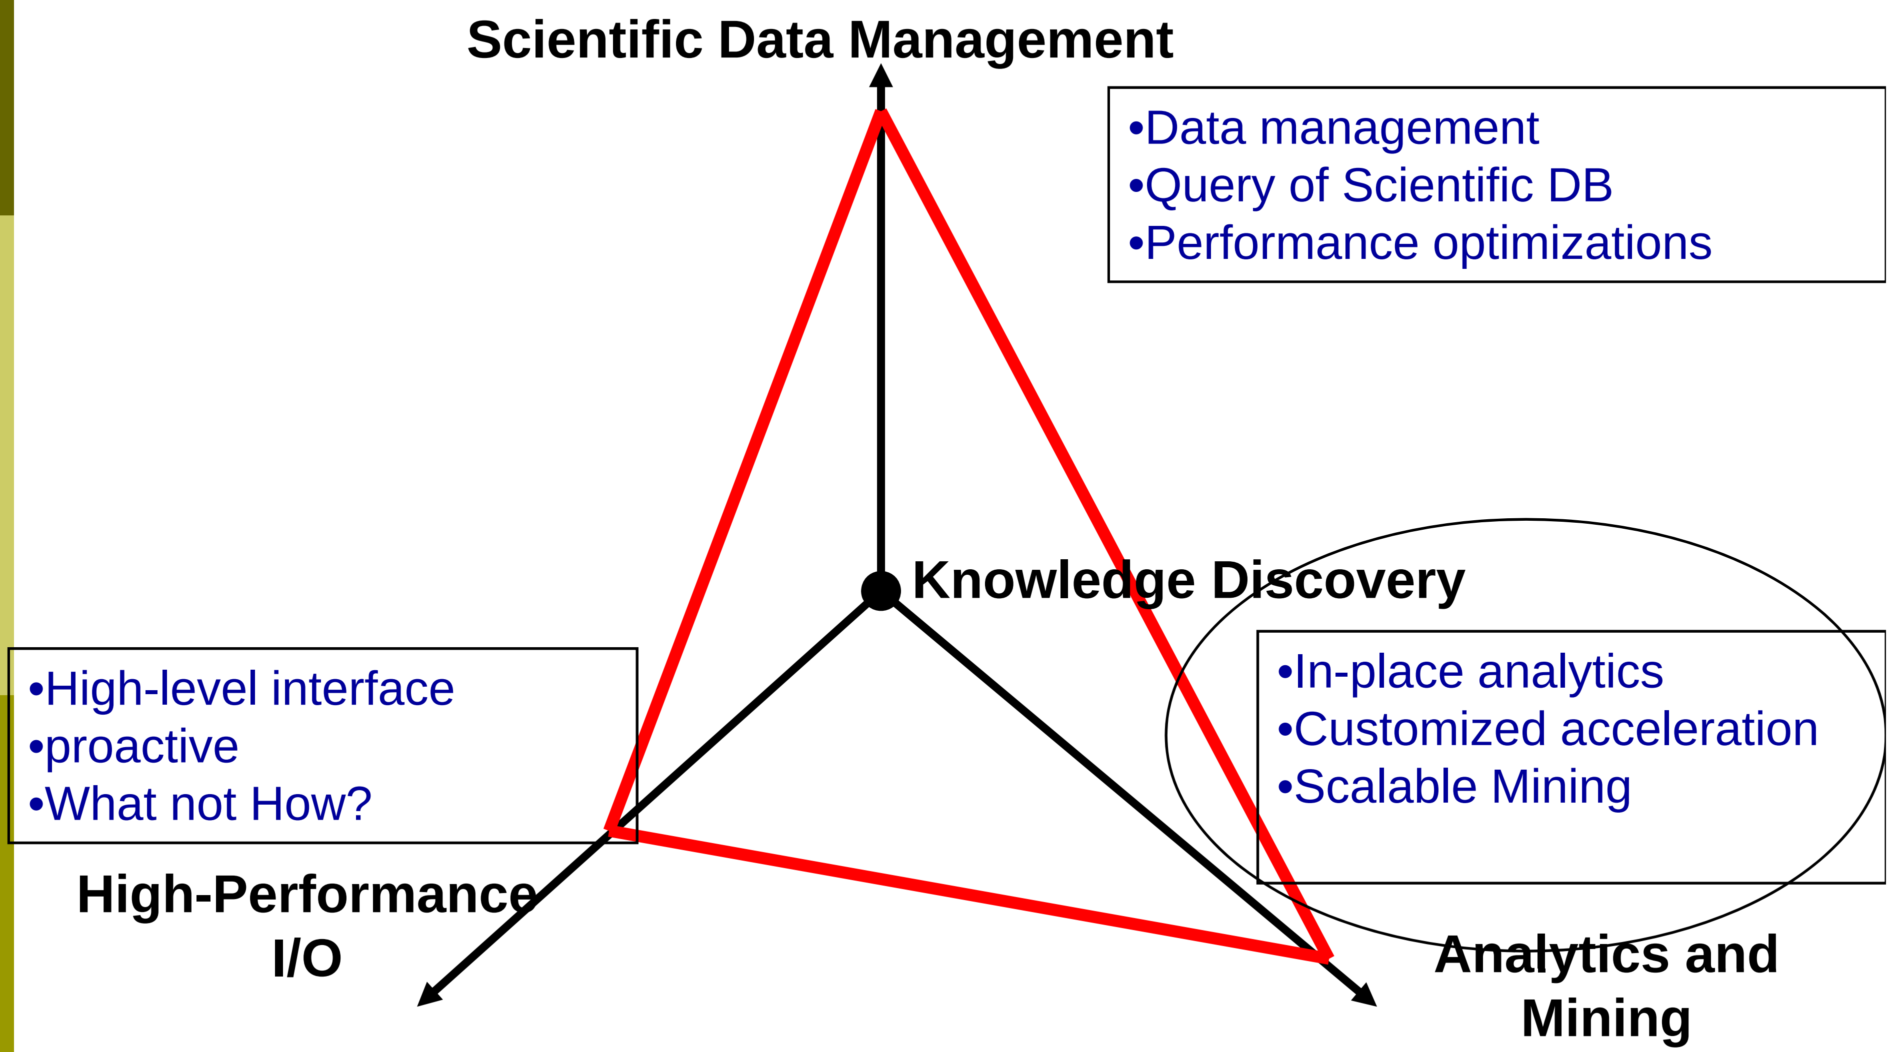

Analyzing and Mining Large Data Sets

- Today's digital society has seen enormous data growth in both commercial and scientific databases
- Data Mining is becoming a commonly used tool to extract information from large and complex datasets

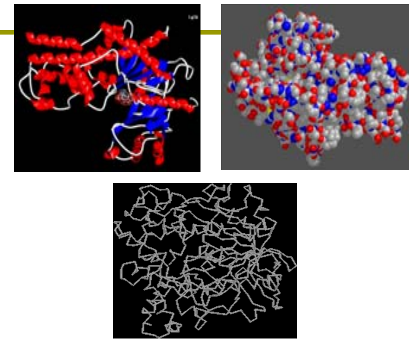

Biomedical Data

Homeland Security

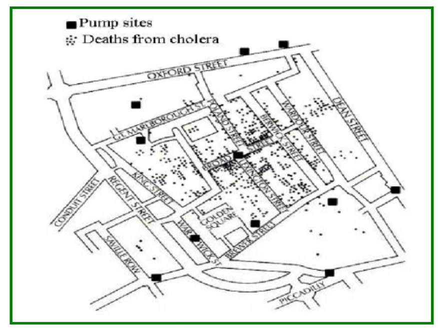

Geo-spatial intelligence

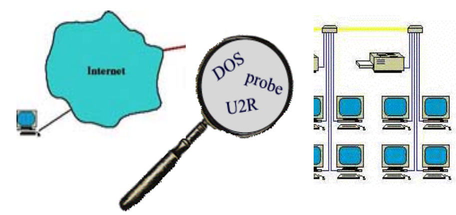

Information Assurance
Network Intrusion Detection

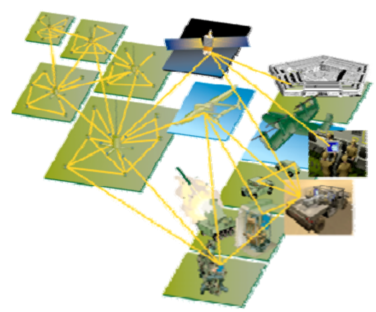

Sensor Networks

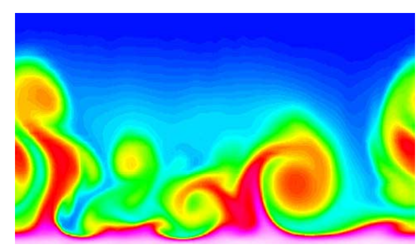

Computational Simulations

Discovery of Climate Patterns from Global Data Sets

Science Goal: Understand global scale patterns in biosphere processes

Earth Science Questions:
- When and where do ecological disasters occur?
- What is the scale and location of human-induced land cover change and its impact?
- How are ocean, atmosphere and land processes coupled?

Data sources:
- Weather observation stations
- High-resolution EOS satellites
  1982-2000 AVHRR at 1° x 1° resolution (~115kmx115km)
  2000-present MODIS at 250m x 250m resolution
- Model-based data from forecast and other models
- Data sets created by data fusion

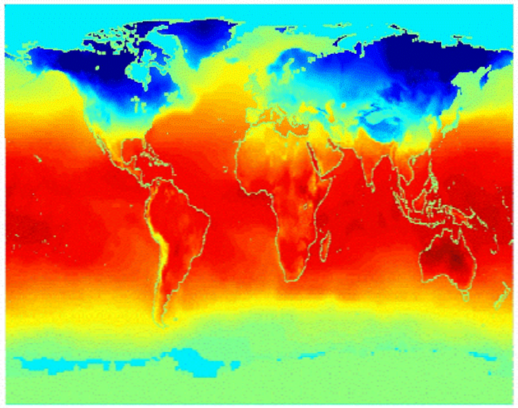

Monthly Average Temperature

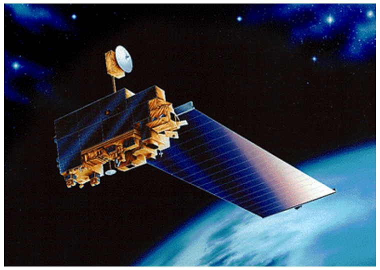

Earth Observing System

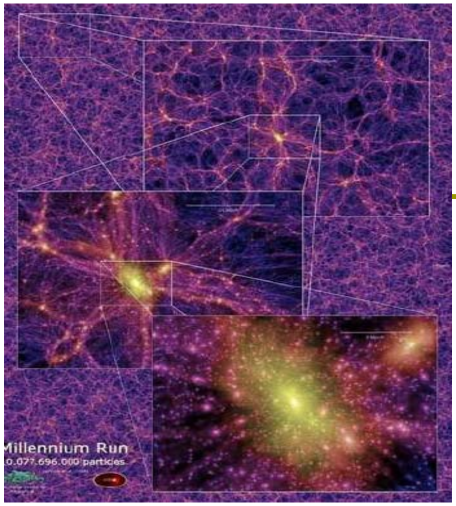

Astrophysics

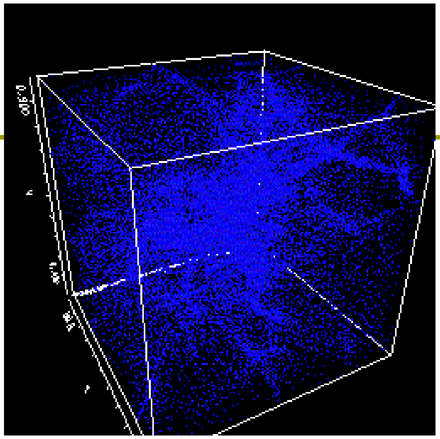

Cosmological Simulations

•Simulate formation and evolution of galaxies

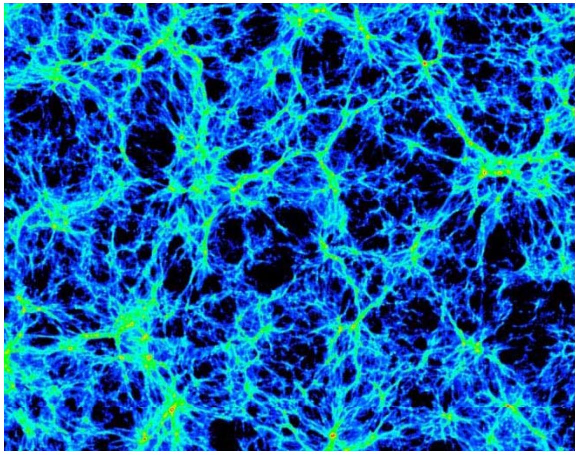

•What is dark matter?
•What is the nature of dark energy?
•How did galaxies, quasars, and supermassive black holes form from the initial conditions in the early universe.

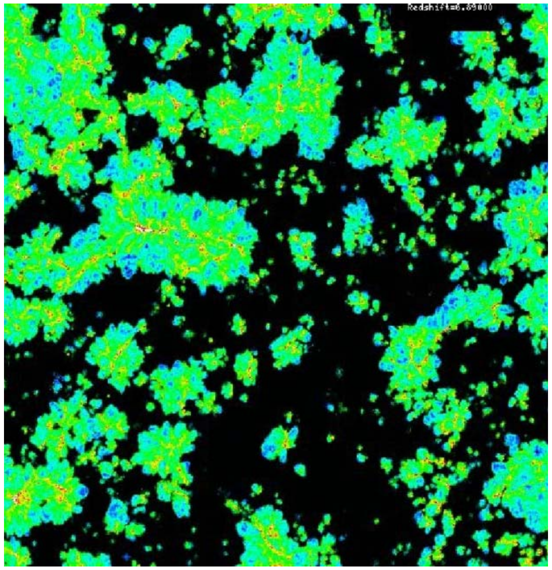

Snapshot from a pure N-body simulation showing the distribution of dark matter at the present time (light colors represent greater density of dark matter). 1B particles

Postprocessed to demonstrate the impact of ionizing radiation from galaxies.

Combustion understanding and modeling: Detection and tracking of autoignition features on-line

Direct simulation of a 3D turbulent flame with detailed chemistry
(200 million grids, 12 species, 5 TB raw data, 5 TB derived data

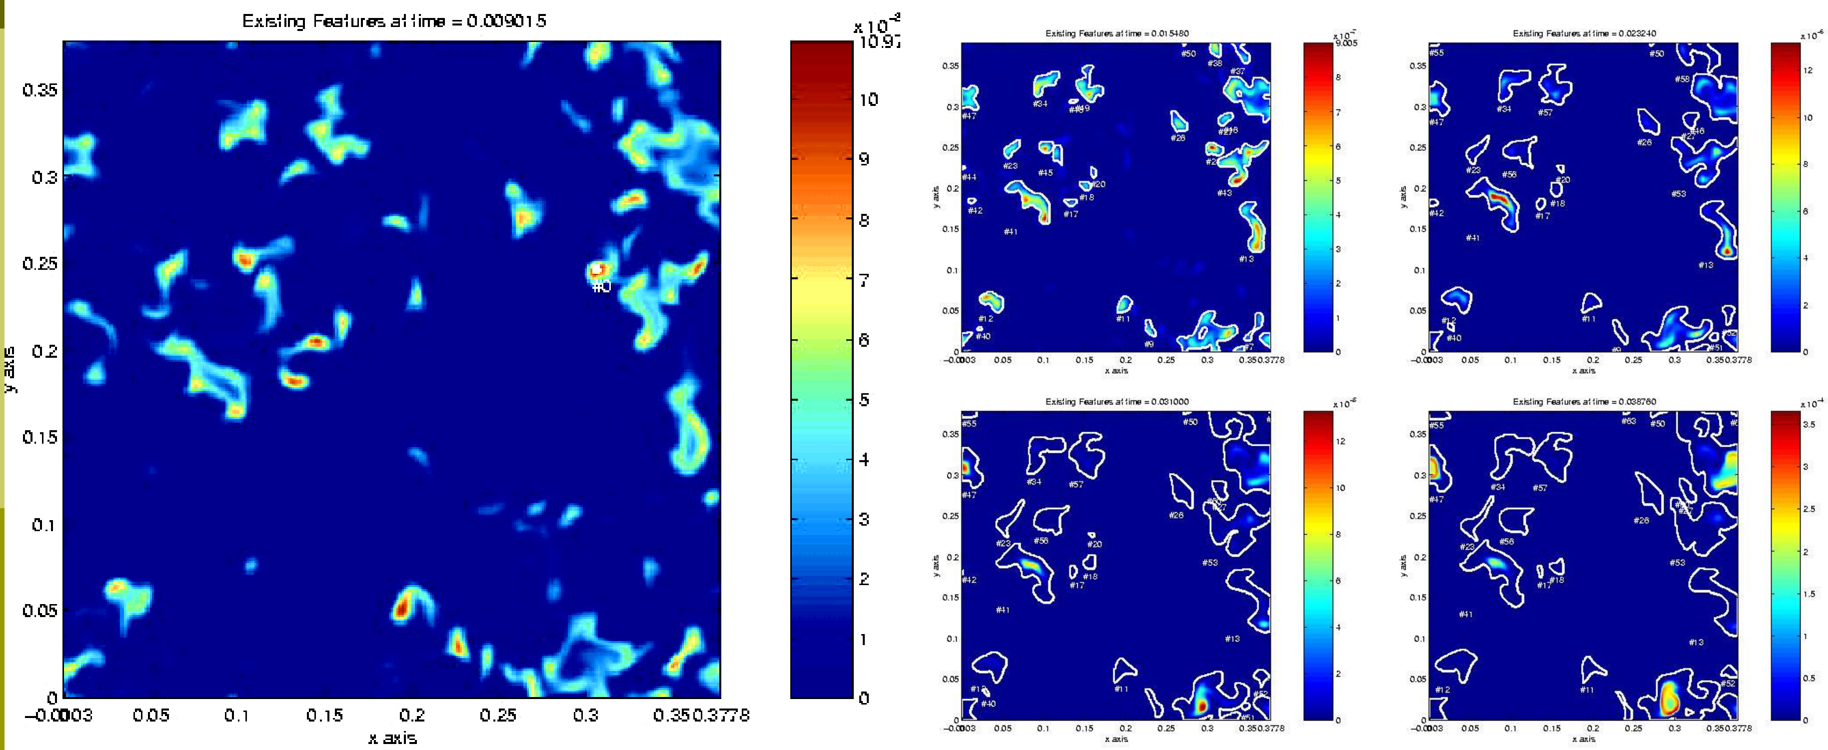

ACK: Jackeline Chen, SNL

Scientific Data Workflow

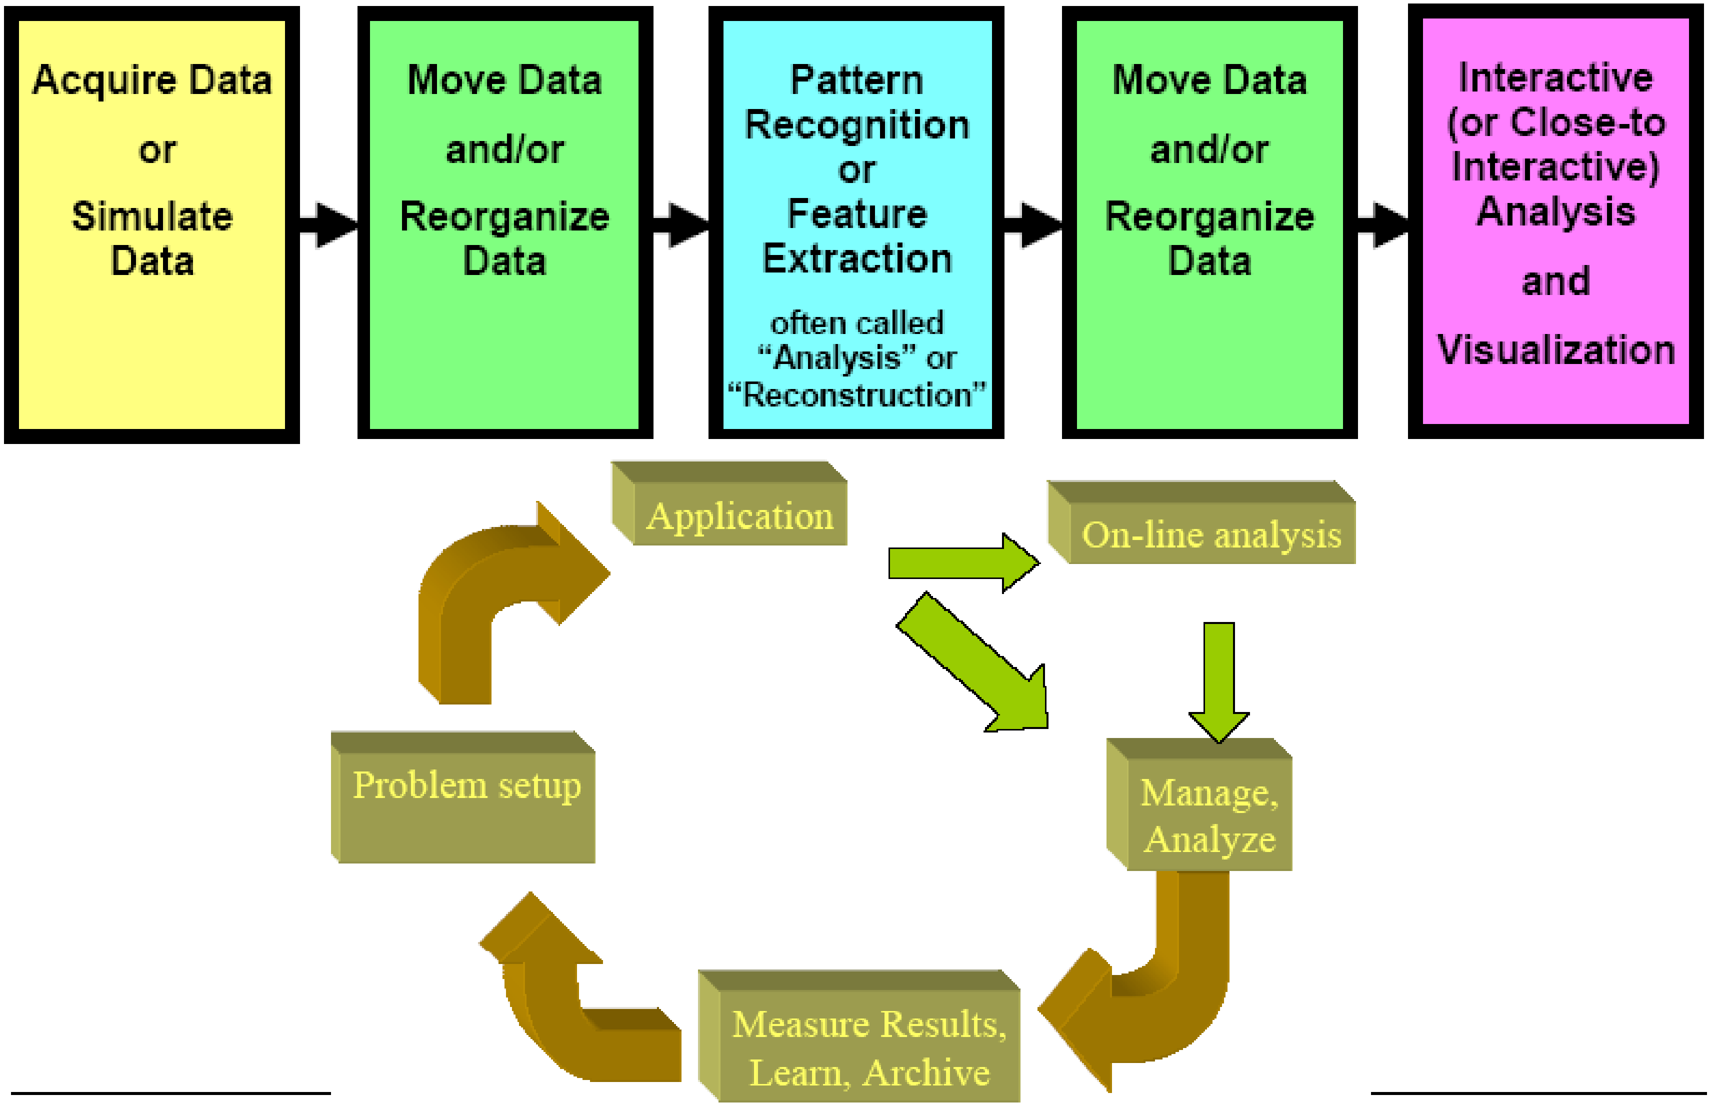

Data Mining and Analytics Characteristics

- Multi-phased operations
  - Cyclic data+compute (large) nature

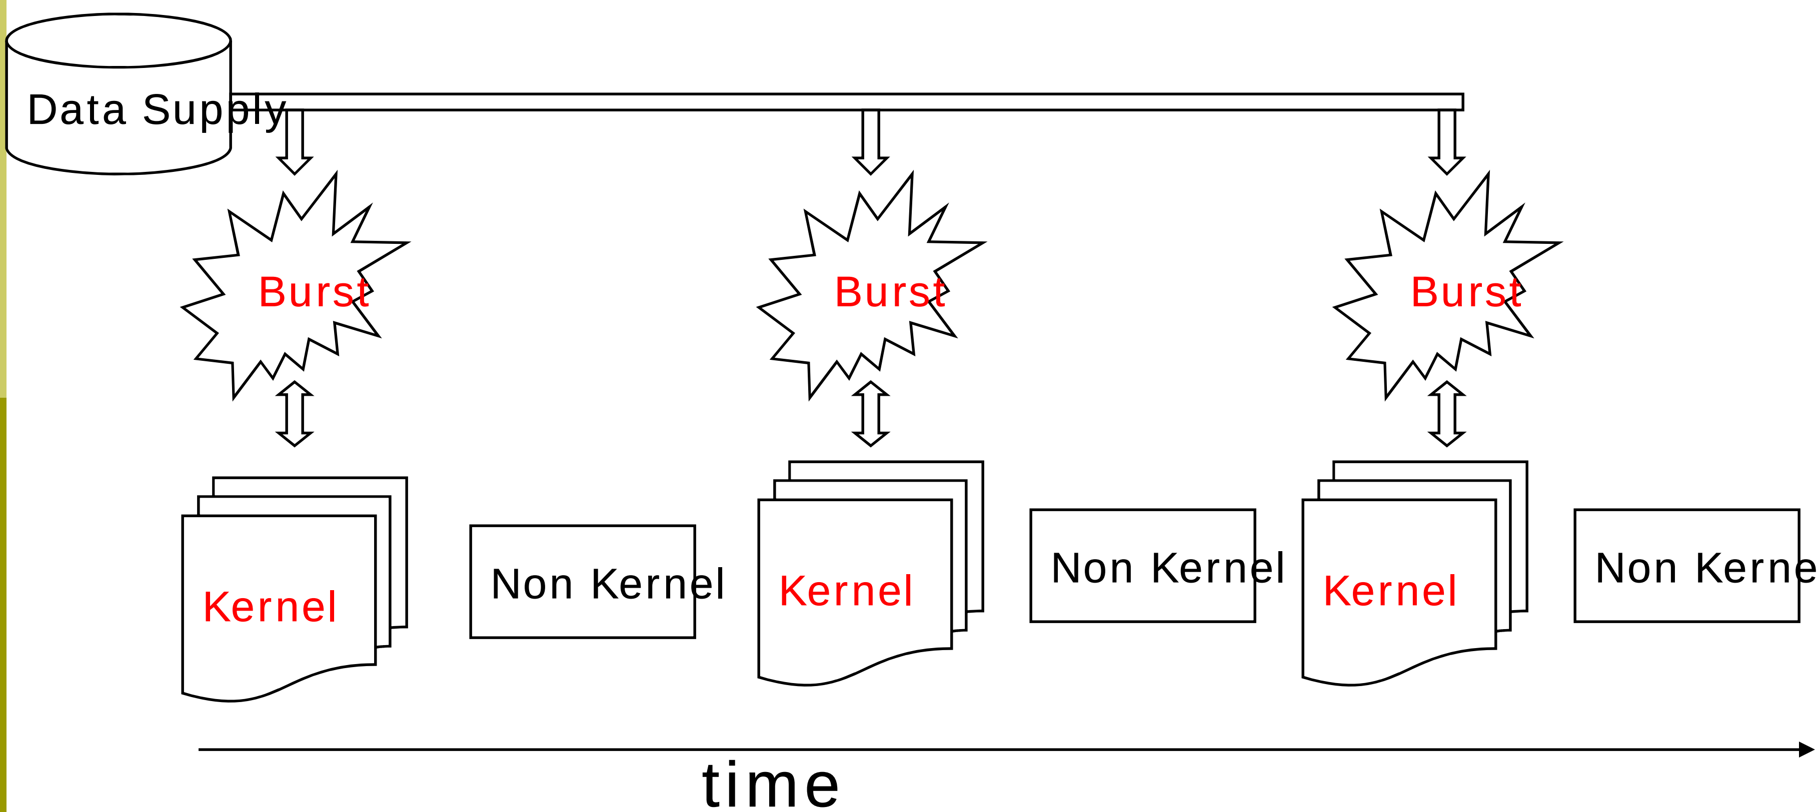

Kernels of the Applications

Kernel Distribution: % of the total execution time

Application	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

- Kernels could be prominent/spread across
- Common kernels across applications: distance, variance

Ramanathan Narayanan, Berkin Ozisikyilmaz, Joseph Zambreno, Jayaprakash Pisharath, Gokhan Memik, and Alok Choudhary. MineBench: A Benchmark Suite for Data Mining Workloads. In *Proceedings of the International Symposium on Workload Characterization (IISWC)*, October 2006.

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

In-Place On-Line Analytics (illustration)

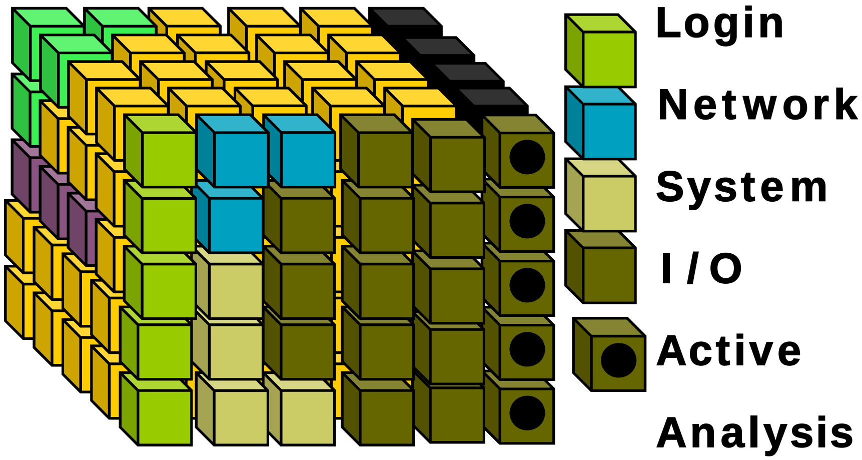

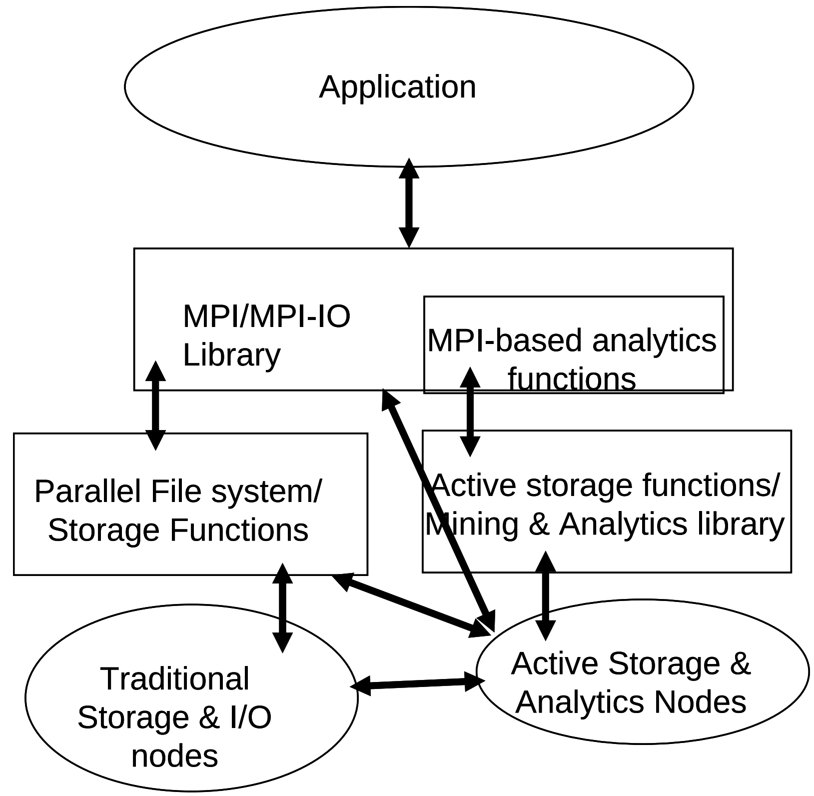

Example 1: Accelerating Mining using FPGAs

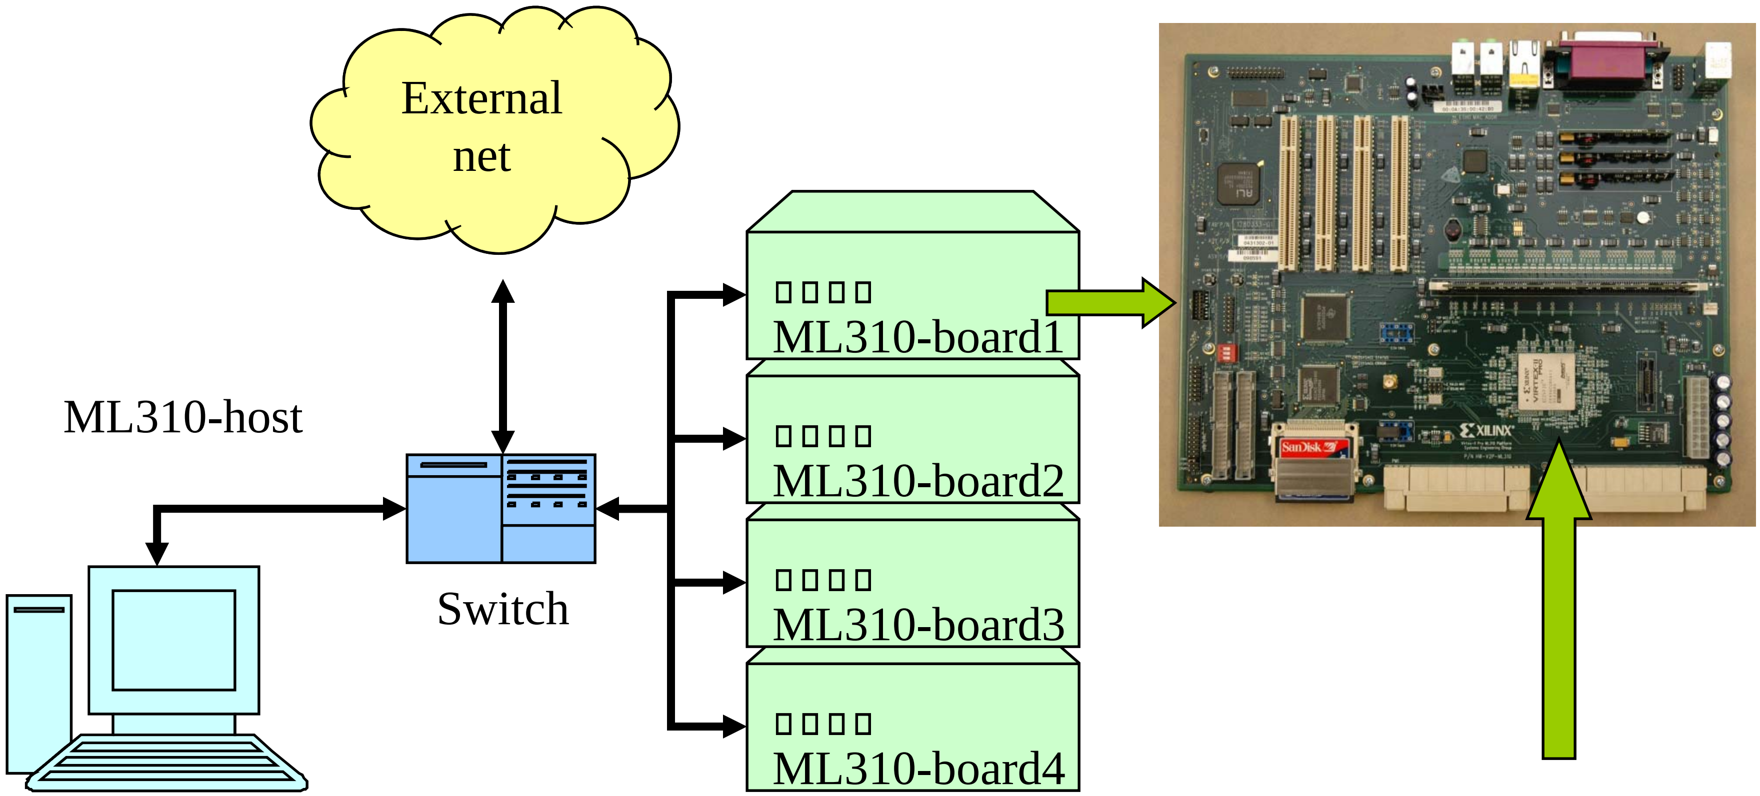

- Software:
  - Data Mining
  - Encryption
  - Functions and runtime libs
  - Linux micro-kernel
- Xilinx XC2VP30 Virtex-II Pro family
  - 30,816 logic cells (3424 CLBs)
  - 2 PPC405 embedded cores
  - 2,448 Kb (136 18 Kb blocks) BRAM
  - 136 dedicated 18x18 multiplier blocks

Illustration of Acceleration (1) Classification (2) PCA

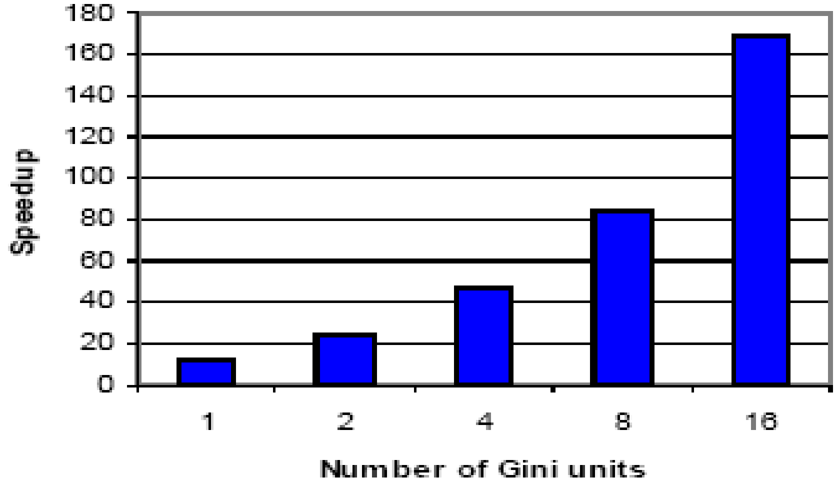

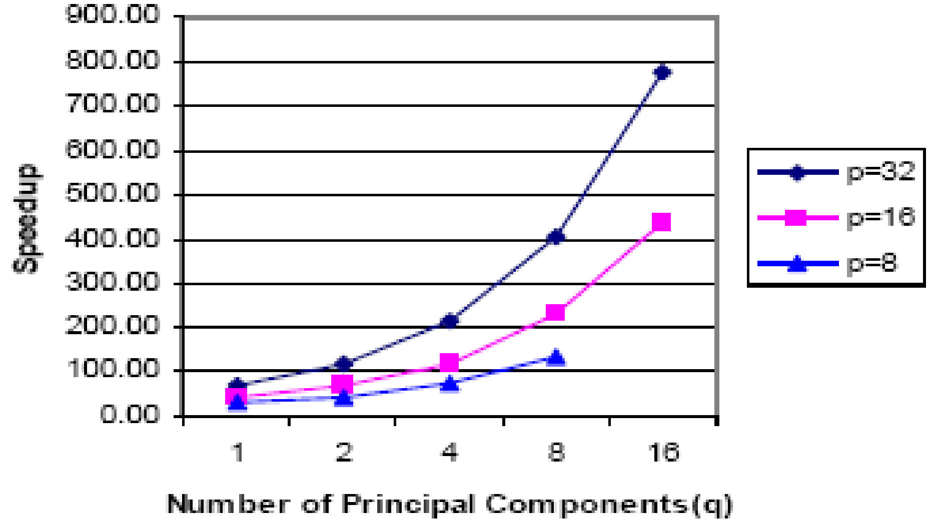

Example 2: Accelerating Mining using GPUs

- Compared to CPUs, GPUs offer 10x higher computational capability and 10x greater memory bandwidth.
  - Lower operating speed, but higher transistor count.
  - More transistors devoted to computation.
- In the past, general purpose computation on GPUs was difficult.
  - Hardware was specialized.
  - Programming required knowledge of the rendering pipeline.
- Now, however, GPUs look much more like SIMD machines.
  - More of the GPU's resources can be applied toward general-purpose computation.
  - Coding for the GPU no longer requires background knowledge in graphics rendering.
- Performance gains of 1-2 orders of magnitude are possible for data-parallel applications.

k-Means

1. Randomly choose initial centers

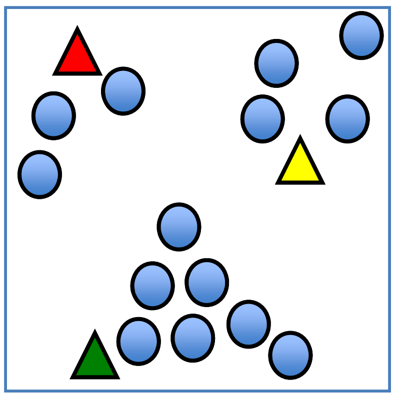

2. Assign each point to the nearest center

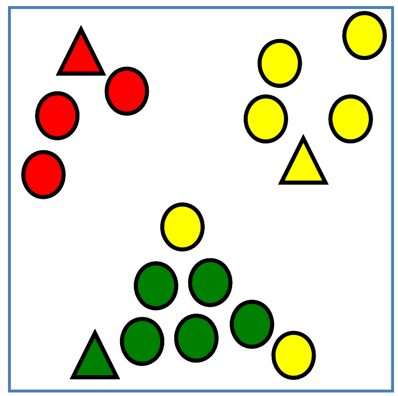

3. Update centers (mean of members)

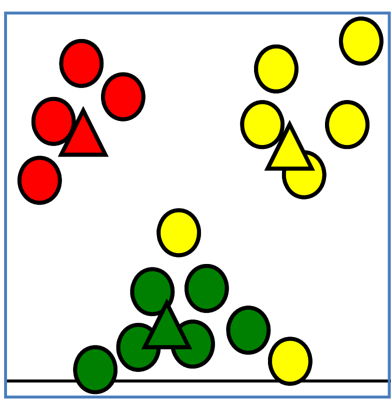

4. Repeat until convergence

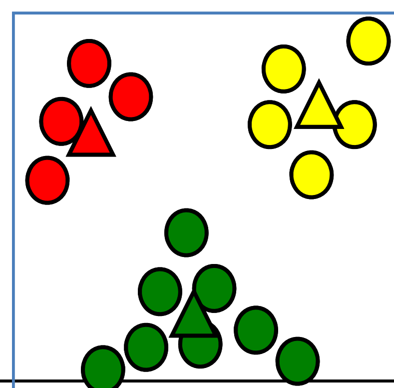

- Minimum distance calculation (Step 2) is data parallel.
- Center update (Step 3) is a reduce operation, which features very predictable memory access patterns.
- Fuzzy k-Means is similar in structure, but features many more computations per memory access.

k-Means Performance (compared with host processor)

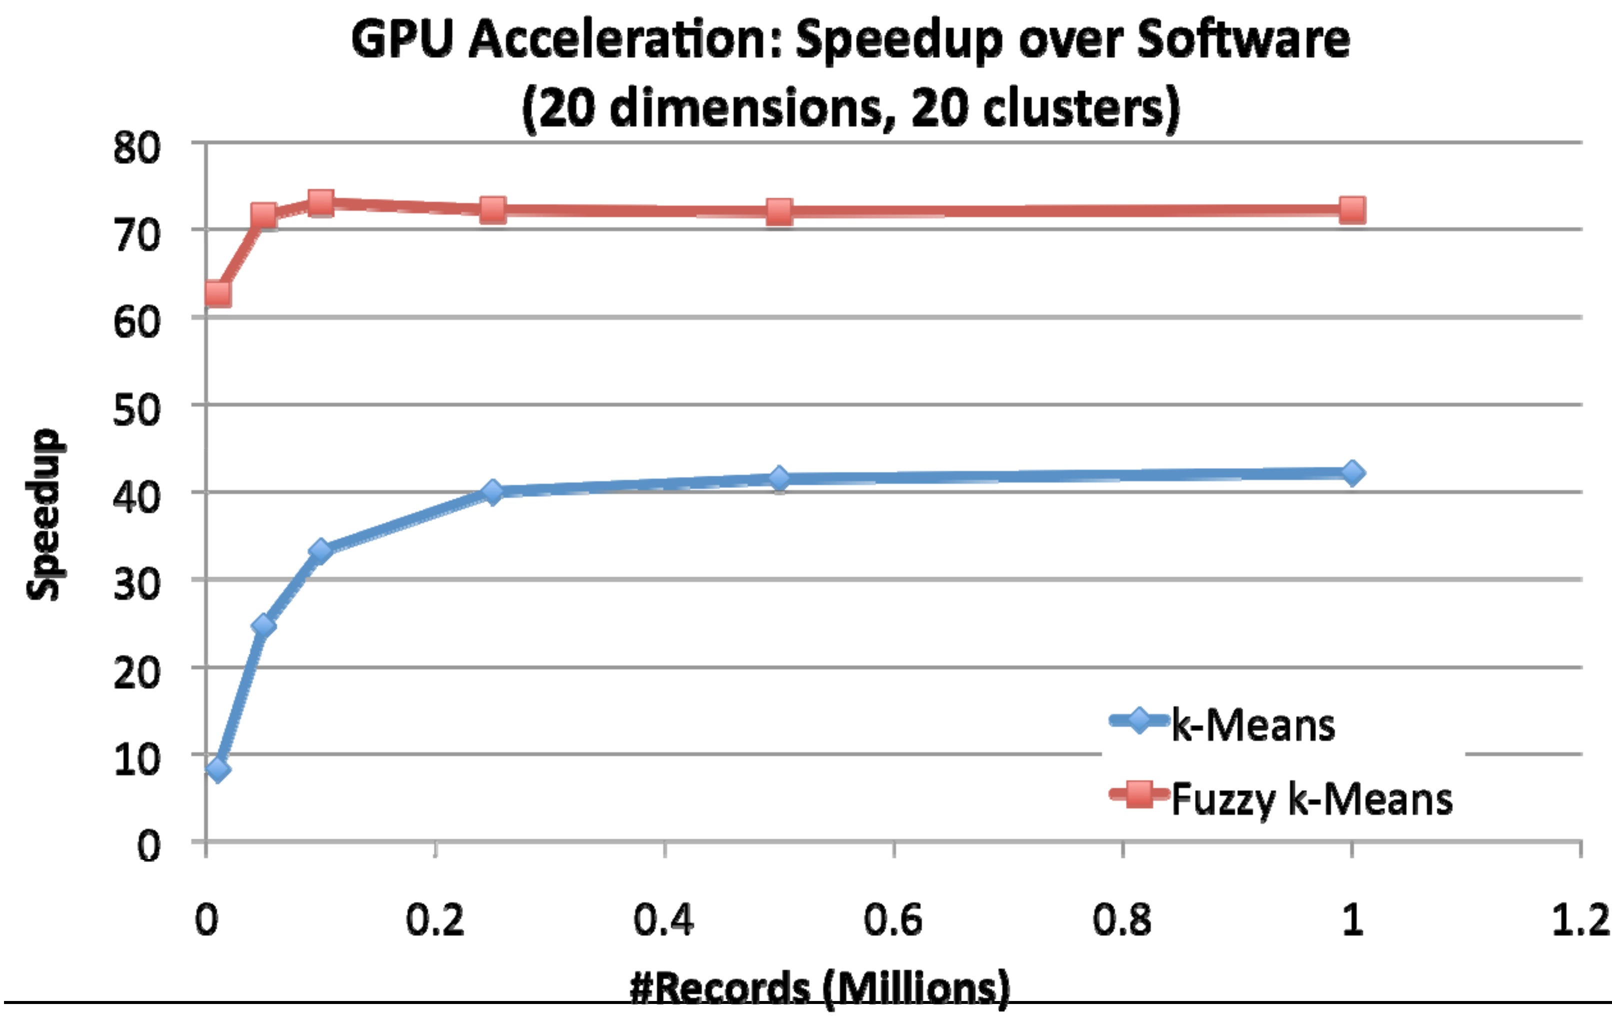

PCA

- Uses:
  - Identify patterns in data to highlight similarities and differences.
  - Data compression by using smaller number of dimensions without much loss of information
- Algorithm overview:
  - Subtract the mean
  - Calculate the covariance matrix
  - Calculate eigenvalues of covariance
  - Compute eigenvectors in the order of significance

PCA in CUDA

- Let the data set be represented by the matrix A, each data of size N,on a column
- Step 1: Subtract Mean
  - N threads to compute average of column
  - N threads to subtract the average from each element of the column
- Step 2: Calculate the covariance matrix
  - Let A' be the matrix obtained in Step 1
  - Covariance Matrix = A'*trans(A')
  - Each Thread computes a block of the covariance matrix

PCA in CUDA (2)

- Step 3: Calculate Eigenvalues and Eigenvectors
  - Tridiagonalization using Householder transformation
    - Converts the covariance matrix into a tridiagonal matrix, (main diagonal and one sub and one super diagonal)
    - N Threads updates each row/column
  - Eigenvalue computation using Bisection Method
    - Use Gerschgorin's theorem to obtain bounds for the eigenvalues
    - Use bisection method to estimate the eigenvalues
    - Threads work on intervals in parallel

Example Result

- Matrix size : 2048

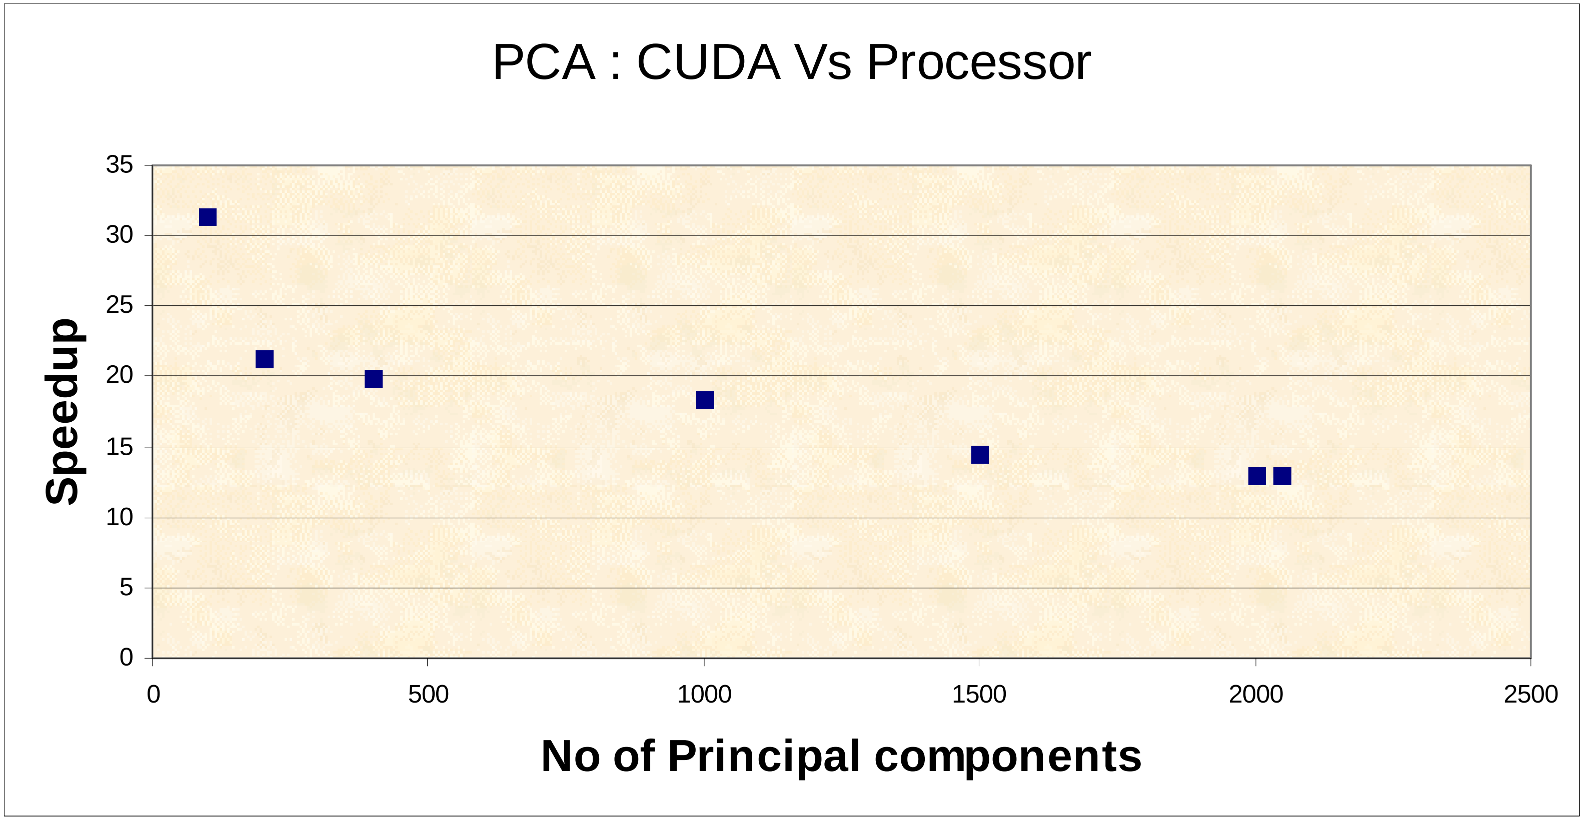

MineBench Project Homepage

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

Summary: Some Challenges in Data Analytics

- Size,
- Scaling,
- Complexity,
- Types of data,
- Processing requirements,
- discover knowledge from the data in a timely fashion,
- Sharing and collaboration patterns among scientists,
- Impact of discovery and results in driving experiments, simulations etc