Subsetting Big Data Workloads from BigDataBench

Zhen Jia

Institute of Computing Technology, Chinese Academy of Sciences

BigDataBench Tutorial
MICRO 2014 Cambridge, UK

Challenges in Understanding Big Data Apps

- A huge number of representative workloads
  - Hard to thoroughly understand behaviors
  - Prohibitively expensive for simulation-based research
- Having many software stacks aggravates the challenge

Revisit BigDataBench 3.0

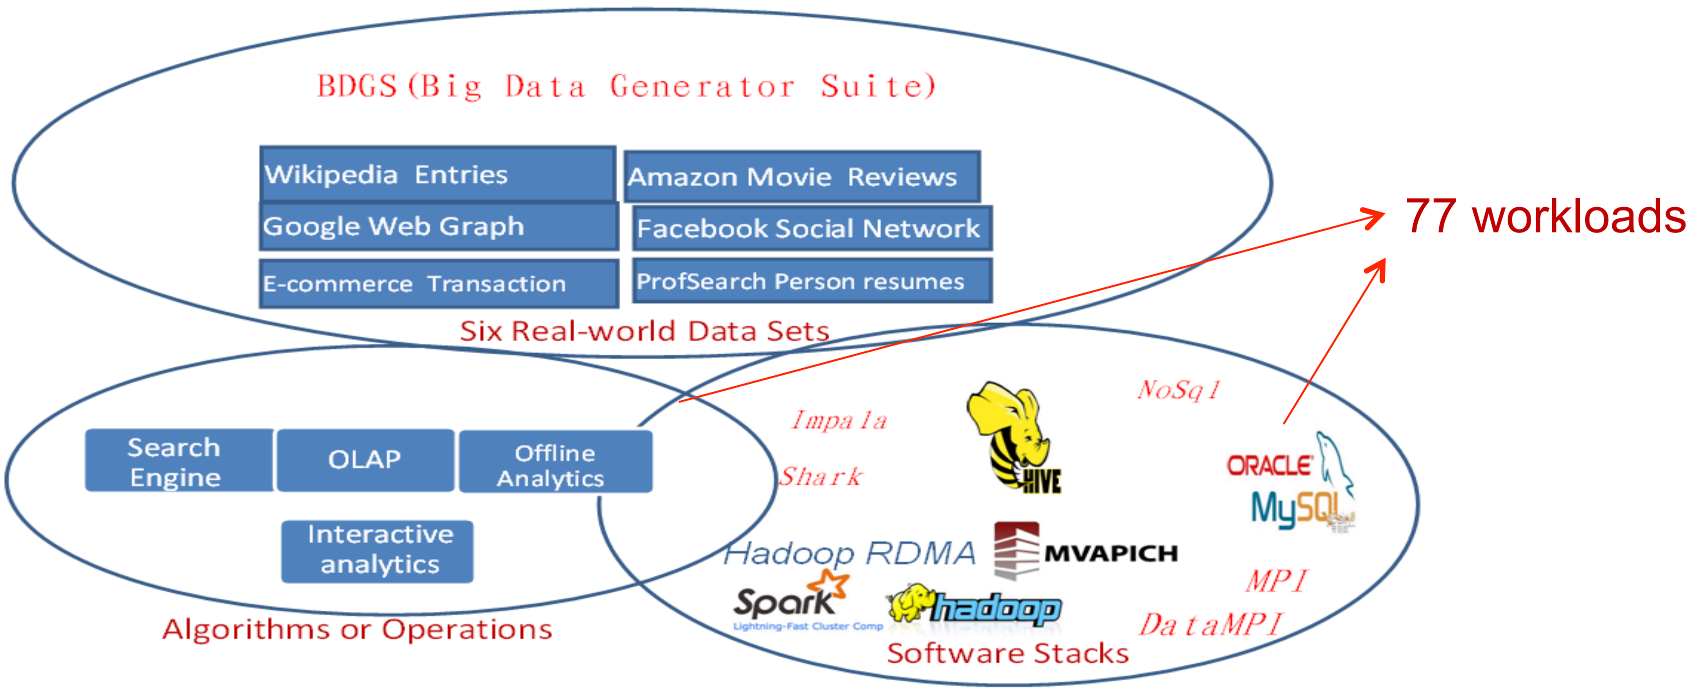

- Include multiple software stacks
- Too time-consuming to run and analyze them all

Why do we consider different software stacks?

- Software stacks have significant impact on workload behaviors--even greater than benchmark algorithms [1]
  - Deep software stacks
  - Integrated mechanisms

Easy to write a big data app

App code << software stack

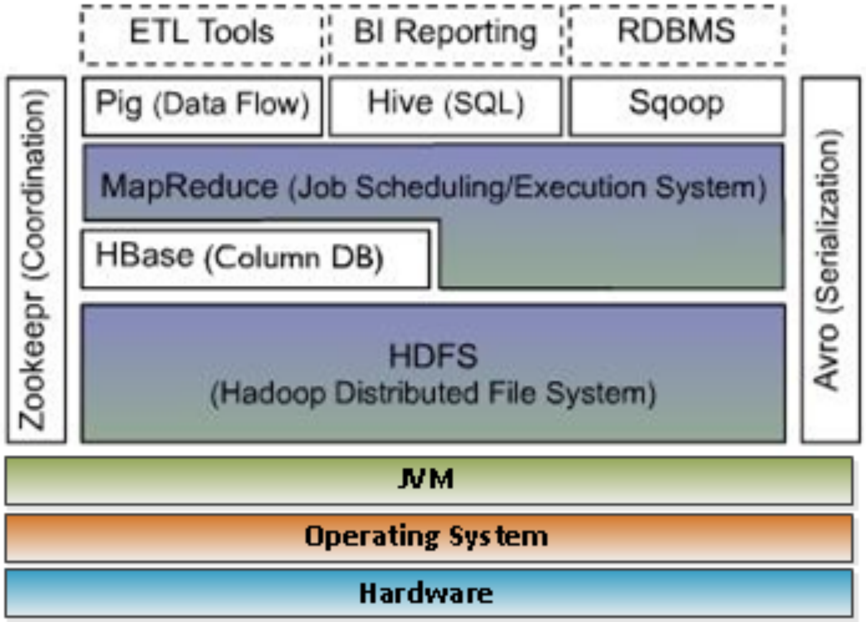

[1]Jia et al. Characterizing and Subsetting Big Data Workloads. IISWC 2014

Goals

- Find a way to downsize BigDataBench 3.0
  - BigDataBench suite contains 77 workloads
  - Should be shrunk to a manageable number
- Reduce the evaluation time of the Big Data research
  - Especially for architecture research using simulator

Subsetting Methodology

--From a view of microarchitecture

Microarchitectural Metric Selection

↓

Correlation Removing

↓

Clustering

↓

Choosing Representative Ones

Metric Selection

- 45 total metrics, including:
  - Instruction Mix
  - Cache Behavior
  - TLB Behavior
  - Branch Execution
  - Pipeline Behavior
  - Offcore Requests and Snoop Responses
  - Parallelism
  - Operation Intensity
- PMCs accessed via *perf*
- Hard to analyze 77 workloads with 45 metrics

Correlation

- Many program characteristics (metrics) are correlated
  - e.g. long latency cache misses => pipeline stalls
- Correlated data can skew analysis
  - May overemphasize a particular property's importance

PCA

- Use PCA (Principal Components Analysis) to eliminate correlated data.
  - PCA: A statistical procedure that uses an orthogonal transformation to convert a set of observations of possibly correlated variables into a set of values of linearly uncorrelated

- Principal Components Analysis (PCA) computes Principal Component (Y_i):
  - Yis that are linear combinations of original metrics xis called PCs (Principal Components):
    - $Y_i = a_{i1}x_1 + a_{i2}x_2 + \ldots a_{ip}x_p$; $i=1..p$

PC properties

- PCs are derived in decreasing order of importance, and they are orthogonal
  - First principal component is the direction of greatest variability (covariance) in the data
  - Second is the next orthogonal (uncorrelated) direction of greatest variability
- Keep PCs with eigenvalues >=1
  - Data is ensured to be uncorrelated while capturing most of the original information **(Kaiser criterion)**

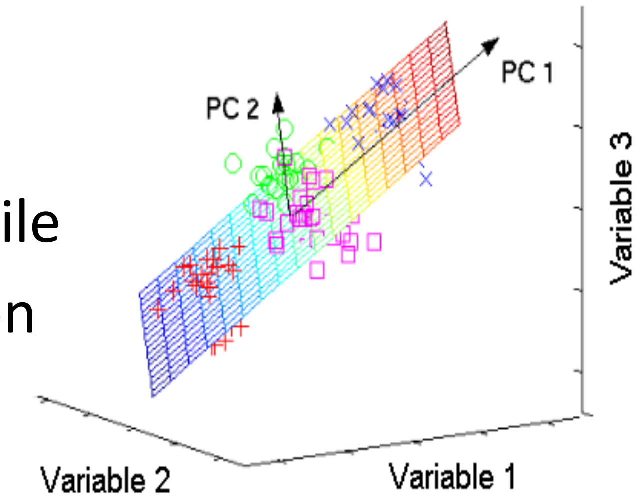

Clustering

- Use K-means algorithm to partition workloads into K clusters.

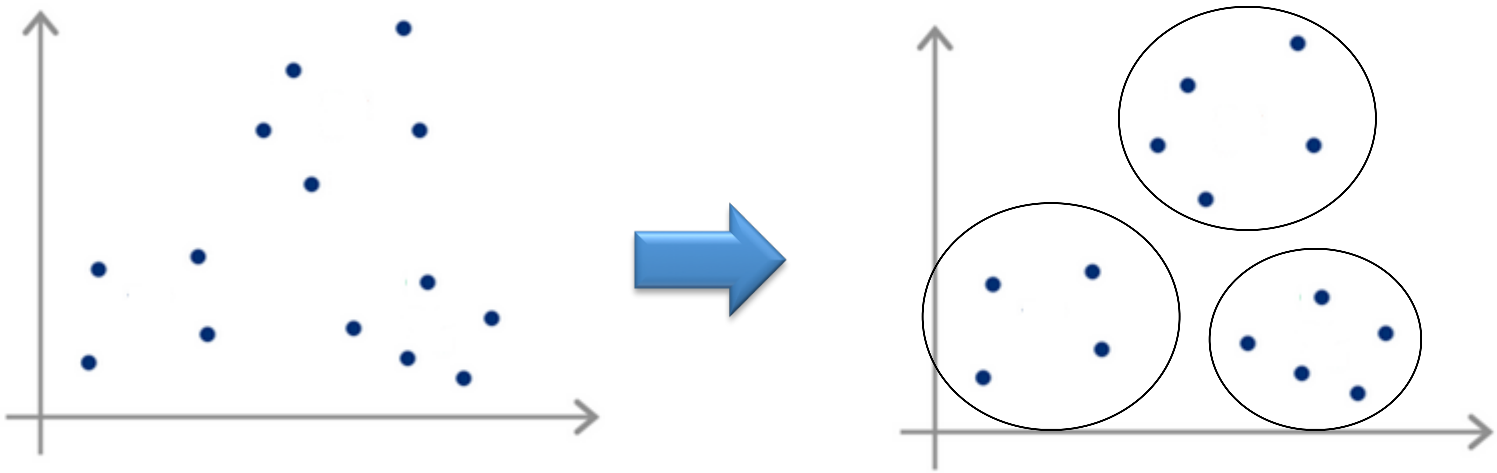

- The problem: how to chose K?

BIC

- Use Bayesian Information Criterion (BIC) to choose proper K value
  - Measures how well the clustering fits the data set
  - Larger BIC scores are better
  - We choose the K with highest BIC scores

$$BIC(D,K) = l(D|K) - \frac{p_j}{2} log(R)$$

$$l(D|K) = \sum_{i=1}^{K} \left(-\frac{R_i}{2} log(2\pi) - \frac{R_i \cdot d}{2} log(\sigma^2) - \frac{R_i - K}{2} + R_i log R_i - R_i log R\right)$$

$$\sigma^2 = \frac{1}{R-K} \sum_i (x_i - \mu(i))^2$$

Selecting Representative Workloads

- Select workloads near the cluster center
- ✓ Select workloads near the cluster boundary

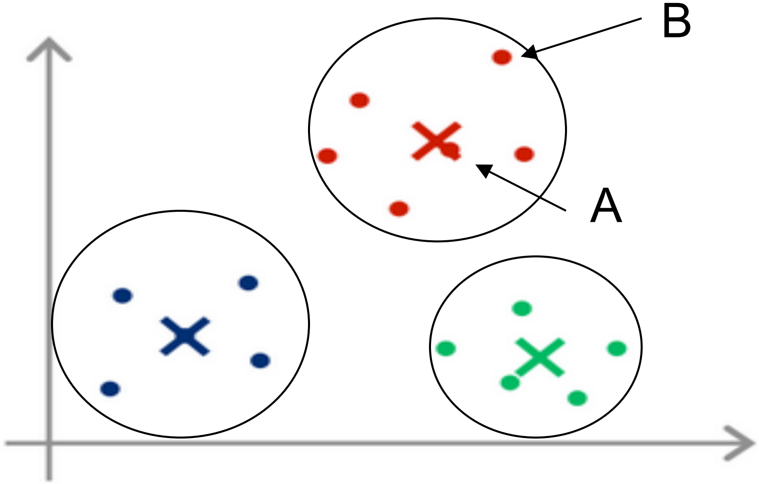

Architecture Subset Workloads

No. of cluster	Workload name	Number of workloads in cluster
1	Cloud-OLTP-Read	10
2	Hive-Difference	9
3	Impala-SelectQuery	9
4	Hive-TPC-DS-query3	9
5	Spark-WordCount	8
6	Impala-Orderby	7
7	Hadoop-Grep	7
8	Shark-TPC-DS-query10	4
9	Shark-Project	4
10	Shark-Orderby	3
11	Spark-Kmeans	1
12	Shark-TPC-DS-query8	1
13	Spark-Pagerank	1
14	Spark-Grep	1
15	Hadoop-WordCount	1
16	Hadoop-NaiveBayes	1
17	Spark-Sort	1

What do those 17 workloads do?

- Offline analytics:
  - Sort, Grep, Word Count, Page Rank, K-means, Bayes
- No-SQL operation: Read
- TPC-DS queries:
  - Query 3, 8, 10
- Basic relational algebra operations:
  - Difference
  - Select query to filter data
  - Sorting
  - Project

Further Work

- Workloads in Big Data change frequently
  - New workloads may be introduced
  - Out-of-date workloads will be removed
- Subsetting is a continuing process
  - The subset may change over time.

Thank You!


QUESTIONS
And
Answers