Introduction of BigDataBench 4.0

Wanling Gao
http://prof.ict.ac.cn

ICT, Chinese Academy of Sciences

ASPLOS 2018, Williamsburg, VA, USA

中科院计算所
INSTITUTE OF COMPUTING TECHNOLOGY

BigDataBench Tutorial Program

- 8:30-9:30 Wanling Gao
  - Introduction of BigDataBench 4.0
- 9:30-10:00 Chen Zheng
  - How to use BigDataBench 4.0
- 10:00-10:30 Coffee break
- 10:30-11:15 Chen Zheng
  - Big data and AI proxy benchmarks for simulation

First part

- *Introduction of BigDataBench 4.0*
- BigDataBench Benchmarking Methodology
- Simulation Benchmarks
- Characterization

Why Big Data and AI Benchmarking?

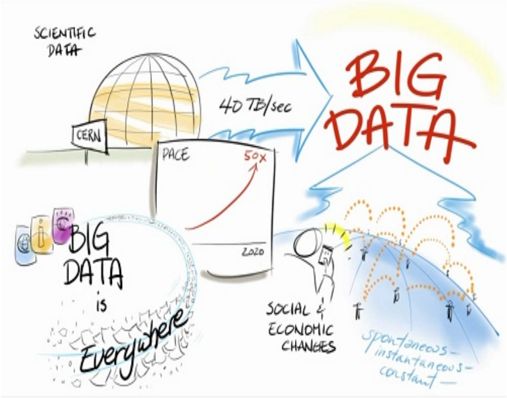

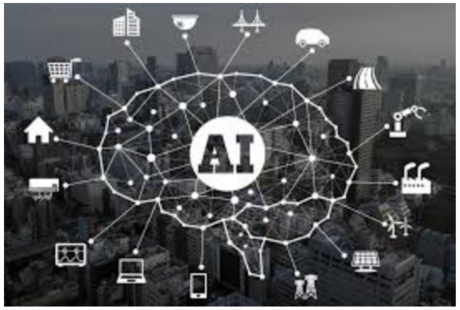

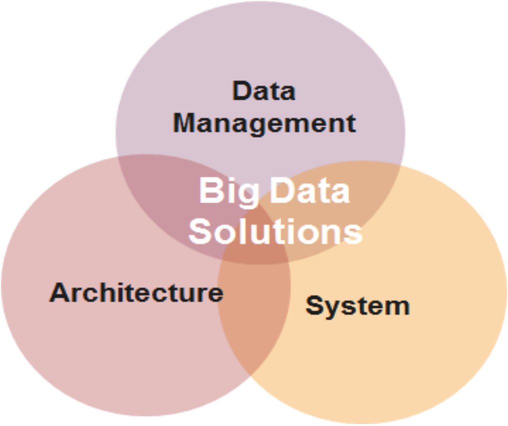

Measuring big data and AI systems, architectures quantitatively

Challenge #1 Complexity

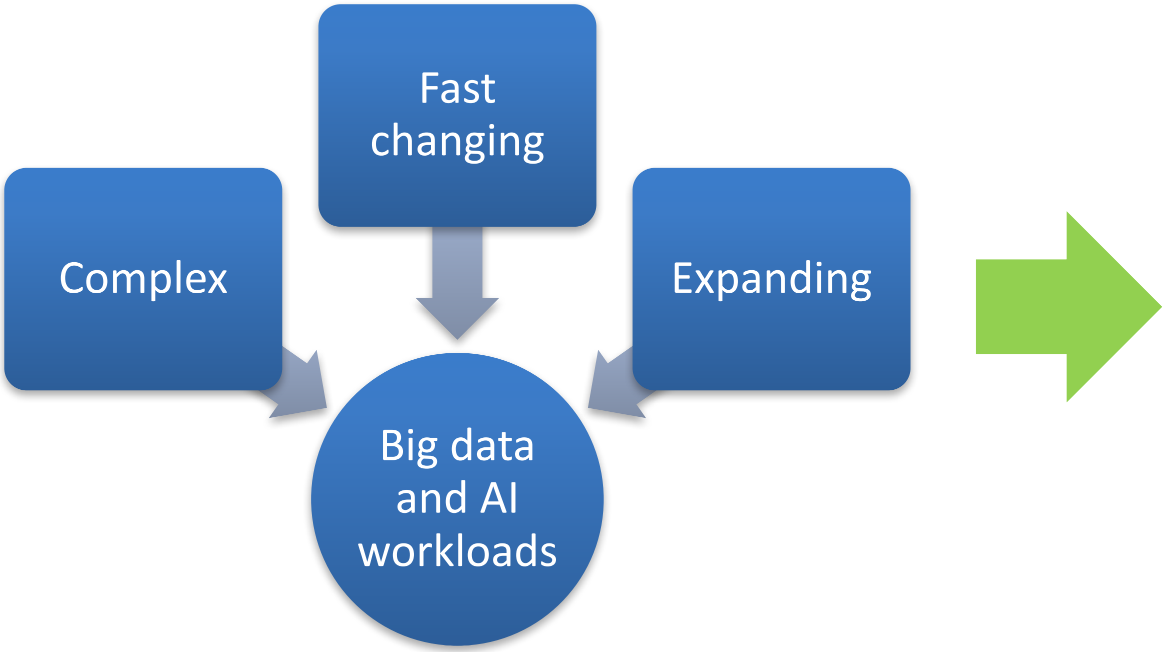

1) ***Portability cost***
2) ***Reproducibility and interpretability of performance data***

Traditional Benchmarking Methodology

- Creating a new benchmark or proxy for every possible workload
  - Case-by-case solution

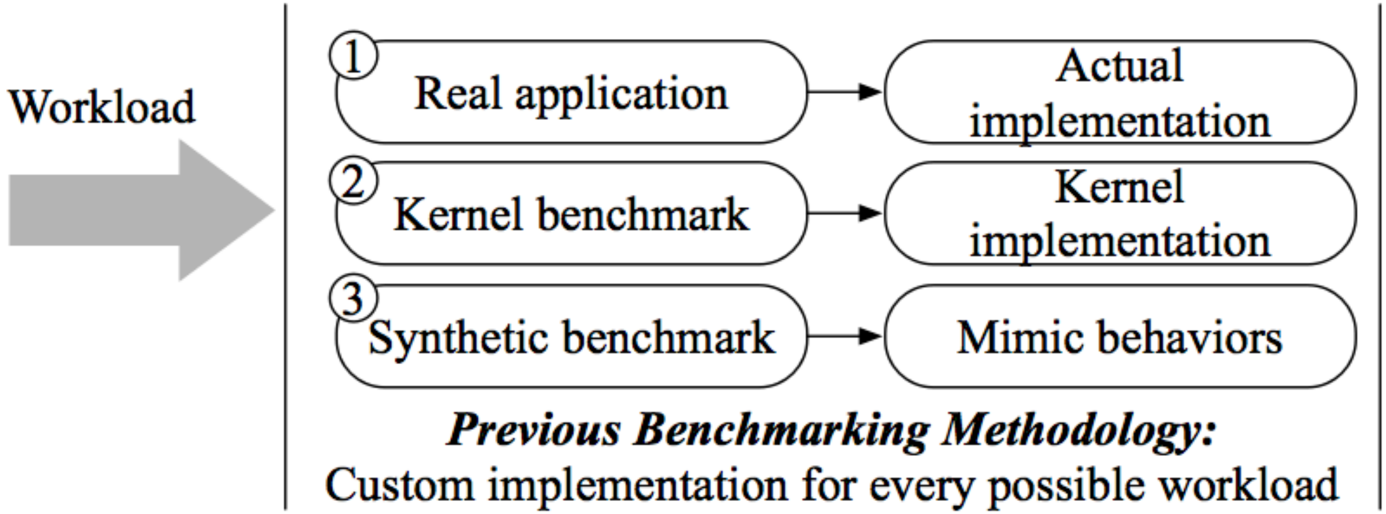

What's the Units of Computation?

- So how to define a representative big data and AI benchmark suite ?

Big data and AI dwarf*: *frequently-appearing units of computation in big data and AI workloads

-- a minimum set to represent ***maximum patterns***

Challenge #2 Fairness

- No one-size-fits-all solution
  - Impact of data set
  - Impact of workloads
  - Impact of software stacks
- Many classes of big data and AI applications without comprehensive characterization

Challenge #3 Consistency

- Requirement difference between community
  - System: performance evaluation on large-scale system deployments
  - Architecture: heavily relies upon simulator-based research, needing shorter (simulation) runtime
  - AI researcher: runtime and model's prediction precision

The benchmarks should be consistent across different communities for the co-design of software and hardware

Simulation for Big Data and AI

- Challenges
  - Simulators have limited supports on complex software stacks
    - For example: Hadoop modes
      - Standalone mode
      - Pseudo-distributed mode
      - Fully-distributed mode
  - Different modes have large behavior differences
- Long running time is unbearable
  - 1000+ times execution time than physical machine

What is *BigDataBench*?

- An open source big data benchmarking project
  - http://prof.ict.ac.cn
  - Search Google using "**BigDataBench**"

BigDataBench 4.0 Overview

BDGS(Big Data Generator Suite) for scalable data

Wikipedia Entries	Amazon Movie Reviews	Google Web Graph
Facebook Social Network	E-commerce Transaction	ProfSearch Resumes
ImageNet	CIFAR-10	LSUN
TED Talks	SoGou Data	MovieLens Dataset
	MNIST	

13 Real-world Data Sets

AI workloads
Offline analytics
Online service
Streaming
Graph analytics
Data warehouse
NoSQL workloads

- Micro benchmarks
- Component benchmarks
- Application benchmarks

47 Workloads with 7 types

JStorm (Alibaba), Flink, GraphLab (Carnegie Mellon), *NoSql*, *Impala*, Caffe, GraphX, ORACLE, MySQL, HIVE, *Shark*, TensorFlow™, *Hadoop RDMA*, MVAPICH, *MPI*, Spark (Lightning-Fast Cluster Computing), hadoop, *DataMPI*

16 Software Stacks

What's New in BigDataBench 4.0

Dwarf-based benchmarking methodology

- Micro, Component and Application Benchmarks

Seven workload types

- AI, Online service, Offline analytics, Graph analytics, Streaming, Data warehouse, NoSQL

Dwarf-based simulation benchmarks

- 100X runtime speedup, 90+% average accuracy

BigDataBench Evolution

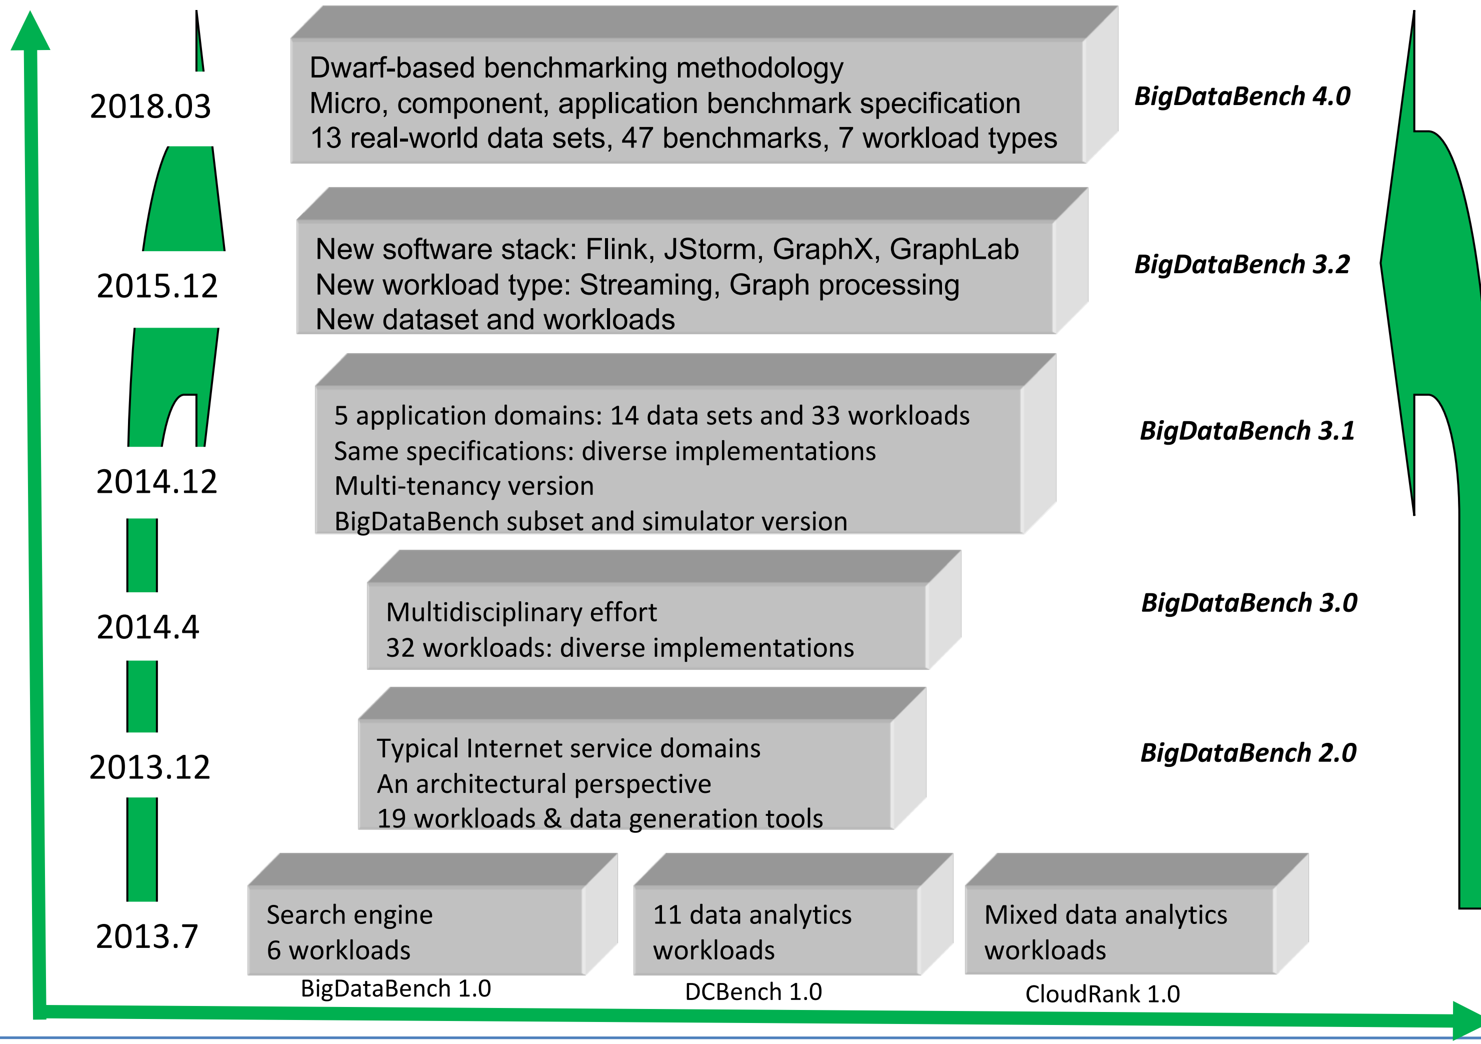

BigDataBench Users

- http://prof.ict.ac.cn/BigDataBench/users/
- Industry users
  - Accenture, BROADCOM, SAMSUMG, Huawei, IBM
- About 100 academia groups published papers using or citing BigDataBench
  - VLDB/SIGMOD, SC, FAST, ASPLOS, ISCA/Micro/ HPCA, ICPP and etc.

Why BigDataBench?

	Benchmarking Target	Methodology	Application domains	Workload types	Workloads	Scalable data sets abstracting from real data	Software Stacks
BigDataBench 4.0	Big data and AI systems and architecture	Dwarf-based	five	seven[1]	forty-seven	13 real data sets; 6 scalable data sets	sixteen
BigDataBench 2.0 [10]	Big data systems and architecture	Popularity	three	three	nineteen	6 real data sets; 6 scalable data sets	ten
BigBench 2.0 [11]	Big data systems	Application model	one	five	Proposal	Proposal	Proposal
BigBench 1.0 [8]	Big data analytics	Application model	one	one	ten	3 data generators	three
CloudSuite 3.0 [4]	Cloud services	Popularity	N/A	four	eight	3 data generators	three
HiBench 6.0 [12]	Big data systems	Popularity	N/A	six	nineteen	Random generate or with specific distribution	five
CALDA [13]	MapReduce system and parallel DBMSs	Popularity	N/A	one	five	N/A	three
YCSB [14]	Cloud serving systems	Performance model	N/A	one	six	N/A	four
LinkBench [15]	Database systems	Application model	N/A	one	ten	one data generator	two
AMP Benchmarks [16]	Data analytic systems	Popularity	N/A	one	four	N/A	five
Fathom [17]	AI systems	Popularity	N/A	one	eight	N/A	one

[1] The seven workload types are online service, offline analytics, graph analytics, artificial intelligence (AI), data warehouse, NoSQL, and streaming.

BigDataBench Publications

- BigDataBench: a Dwarf-based Big Data and AI Benchmark Suite. Technical Report. https://arxiv.org/pdf/1802.08254.pdf
- BOPS, Not FLOPS! A New Metric, Measuring Tool, and Roofline Performance Model For Datacenter Computing. Technical Report. https://arxiv.org/pdf/1801.09212.pdf
- Big Data Dwarfs: Towards Fully Understanding Big Data Analytics Workloads. Technical Report. https://arxiv.org/pdf/1802.00699.pdf
- BigDataBench: a Big Data Benchmark Suite from Internet Services. 20th IEEE International Symposium On High Performance Computer Architecture (**HPCA-2014**).
- Understanding Big Data Analytics Workloads on Modern Processors. **TPDS 2017**. https://arxiv.org/pdf/1504.04974.pdf
- Characterizing data analysis workloads in data centers. 2013 IEEE International Symposium on Workload Characterization (**IISWC 2013**)（Best paper award）

First part

- Introduction of BigDataBench 4.0
- *BigDataBench Benchmarking Methodology*
- Simulation Benchmarks
- Characterization

Dwarf-based Benchmarking Methodology

- A ***scalable*** dwarf-based benchmarking methodology
  - Combinations of dwarfs with different weights

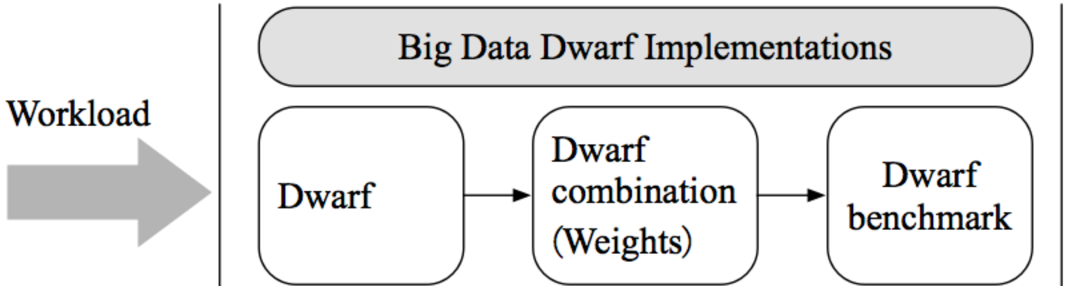

Dwarf-based Benchmarking Methodology:
DAG-like combination

Inspiration

Successful Compute Abstractions

- Relational algebra
  - 5 primitive operations
  - Select, Project, Product, Union, Difference
- Parallel computing
  - Computational & communication patterns
  - 13 dwarfs

Successful Benchmarks

- TPC-C
  - OLTP domain
  - Functions of abstraction
- HPCC
  - High performance computing
  - Seven basically tests

Dwarf Abstraction

- Big Data &AI Dwarf
  - Units of computation
- Dwarf Abstraction
  - Algorithmic analysis
  - Experimental analysis

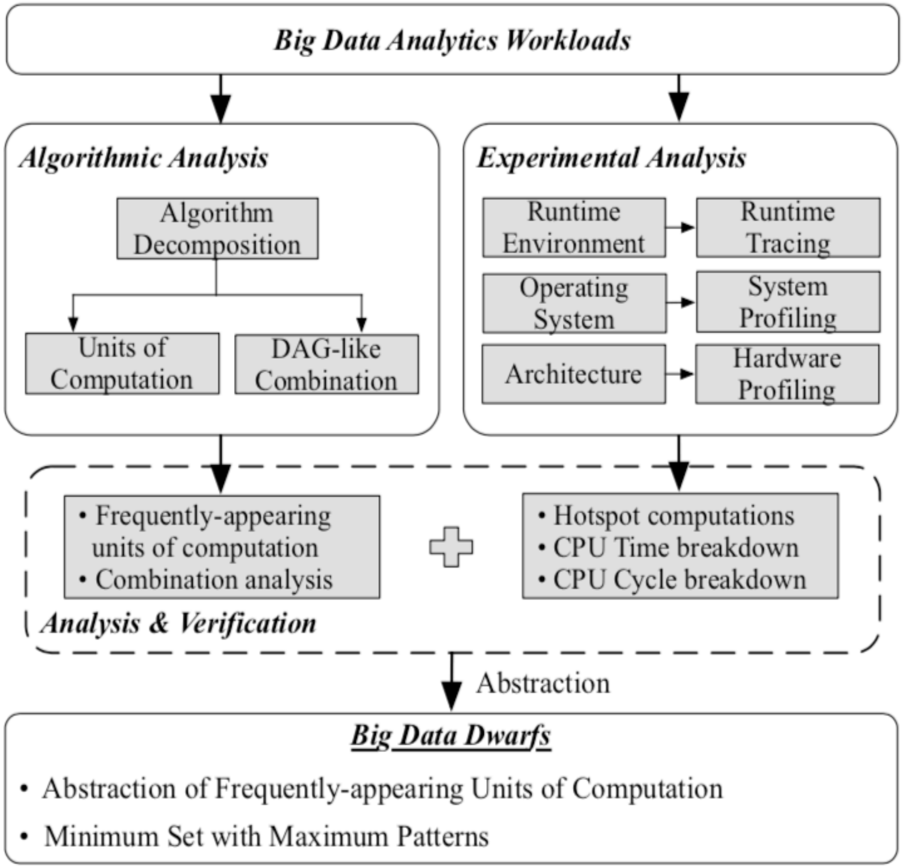

Units of Computation

- Importance of eight classes units of computation

Catergory	Application Domain	Workload	Unit of Computation
Graph Mining	Search Engine Community Detection	PageRank	Matrix, Graph, Sort
		BFS, Connected component(CC)	Graph
Deminsion Reduction	Image Processing Text Processing	Principal components analysis(PCA)	Matrix
		Latent dirichlet allocation(LDA)	Basic Statistic, Sampling
Deep Learning	Image Recognition Speech Recognition	Convolutional neural network(CNN)	Matrix, Sampling, Transform
		Deep belief network(DBN)	Matrix, Sampling
Recommendation	Association Rules Mining Electronic Commerce	Aporiori	Basic Statistic, Set
		FP-Growth	Graph, Set, Basic Statistic
		Collaborative filtering(CF)	Graph, Matrix
Classification	Image Recognition Speech Recognition Text Recognition	Support vector machine(SVM)	Matrix
		K-nearest neighbors(KNN)	Matrix, Sort, Basic Staticstic
		Naive bayes	Basic Statistic
		Random forest	Graph, Basic Statistic
		Decision tree(C4.5/CART/ID3)	Graph, Basic Statistic
Clustering	Data Mining	K-means	Matrix, Sort

Units of Computation (cont')

- Importance of eight classes units of computation

Feature Preprocess	Image Processing Signal Processing Text Processing	Image segmentation(GrabCut)	Matrix, Graph
		Scale-invariant feature transform(SIFT)	Matrix, Transform, Sampling, Sort, Basic Statistic
		Image Transform	Matrix, Transform
		Term Frequency-inverse document frequency (TF-IDF)	Basic Statistic
Sequence Tagging	Bioinformatics	Hidden Markov Model(HMM)	Matrix
	Language Processing	Conditional random fields(CRF)	Matrix, Sampling
Indexing	Search Engine	Inverted index, Forward index	Basic Statistic, Logic, Set, Sort
Encoding/Decoding	Multimedia Processing	MPEG-2	Matrix, Transform
	Security	Encryption	Matrix, Logic
	Cryptography	SimHash, MinHash	Set, Logic
	Digital Signature	Locality-sensitive hashing(LSH)	Set, Logic
Data Warehouse	Business intelligence	Project, Filter, OrderBy, Union	Set, Sort

Big Data and AI Dwarfs

- Matrix computations
- Sampling computations
- Transform computations
- Graph computations
- Logic computations
- Set computations
- Basic statistic computations
- Sort computations

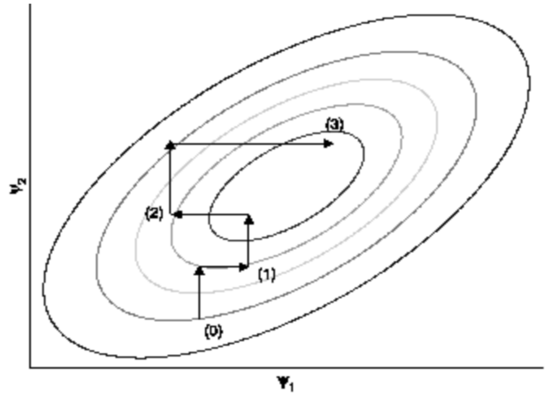

Figure 3.4: Gibbs sampling algorithm in two dimensions starting from an initial point and then completing three iterations

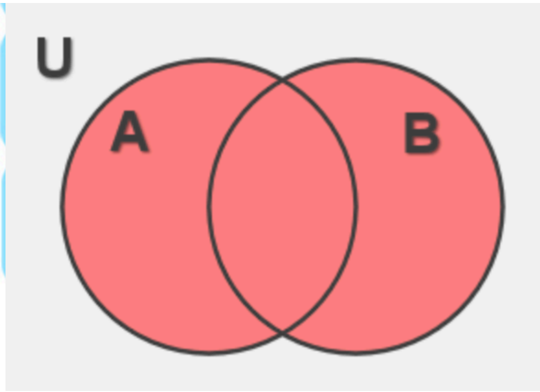

One Combination Example

- Feature extraction – SIFT Workload

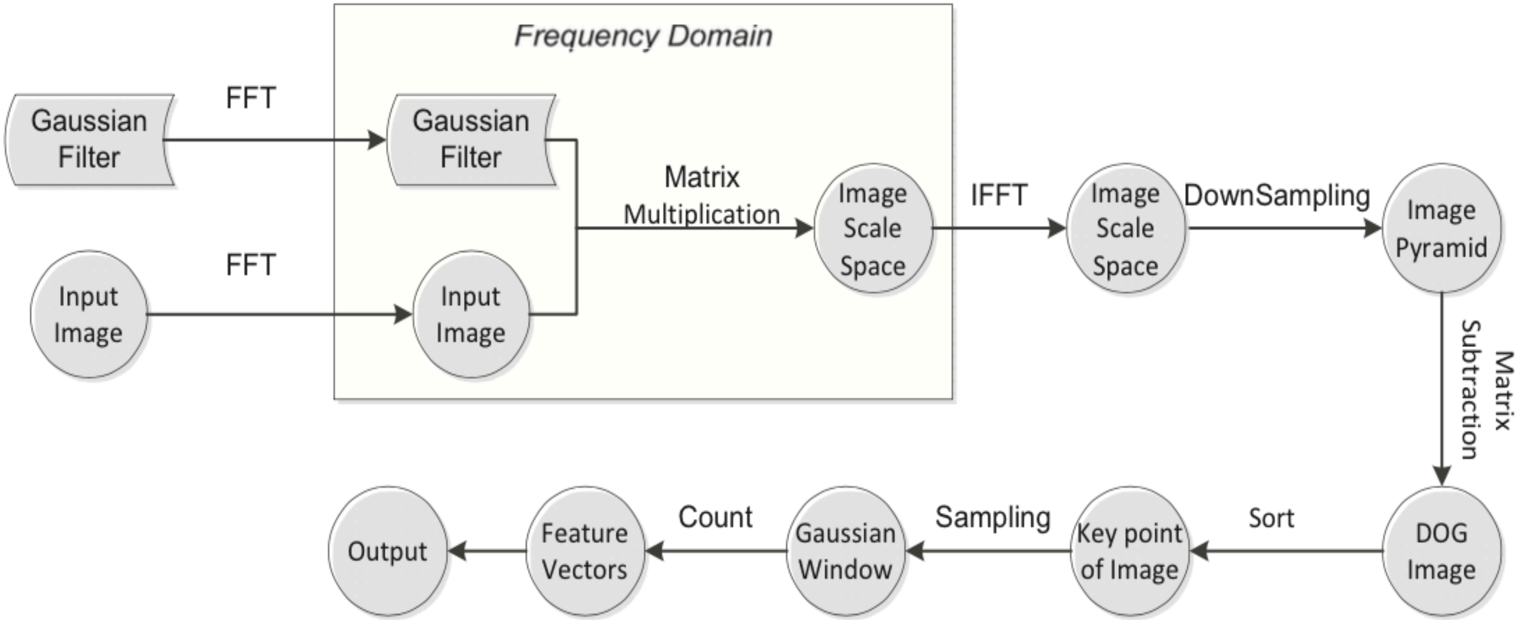

Several dwarfs: ***transform*** computations(FFT, IFFT), ***sampling*** computations(downsampling), ***matrix*** computations(matrix multiplication/subtraction), ***sort*** computations(sort), ***basic statistic*** computations(count)

Methodology Principle

Separating specification from implementation.

- Model relevant domains

State-of-the-art algorithms and technologies

- Implementation keep in pace with the improvement

Data impact

- Representative data sets considering typical types and sources

Five Steps

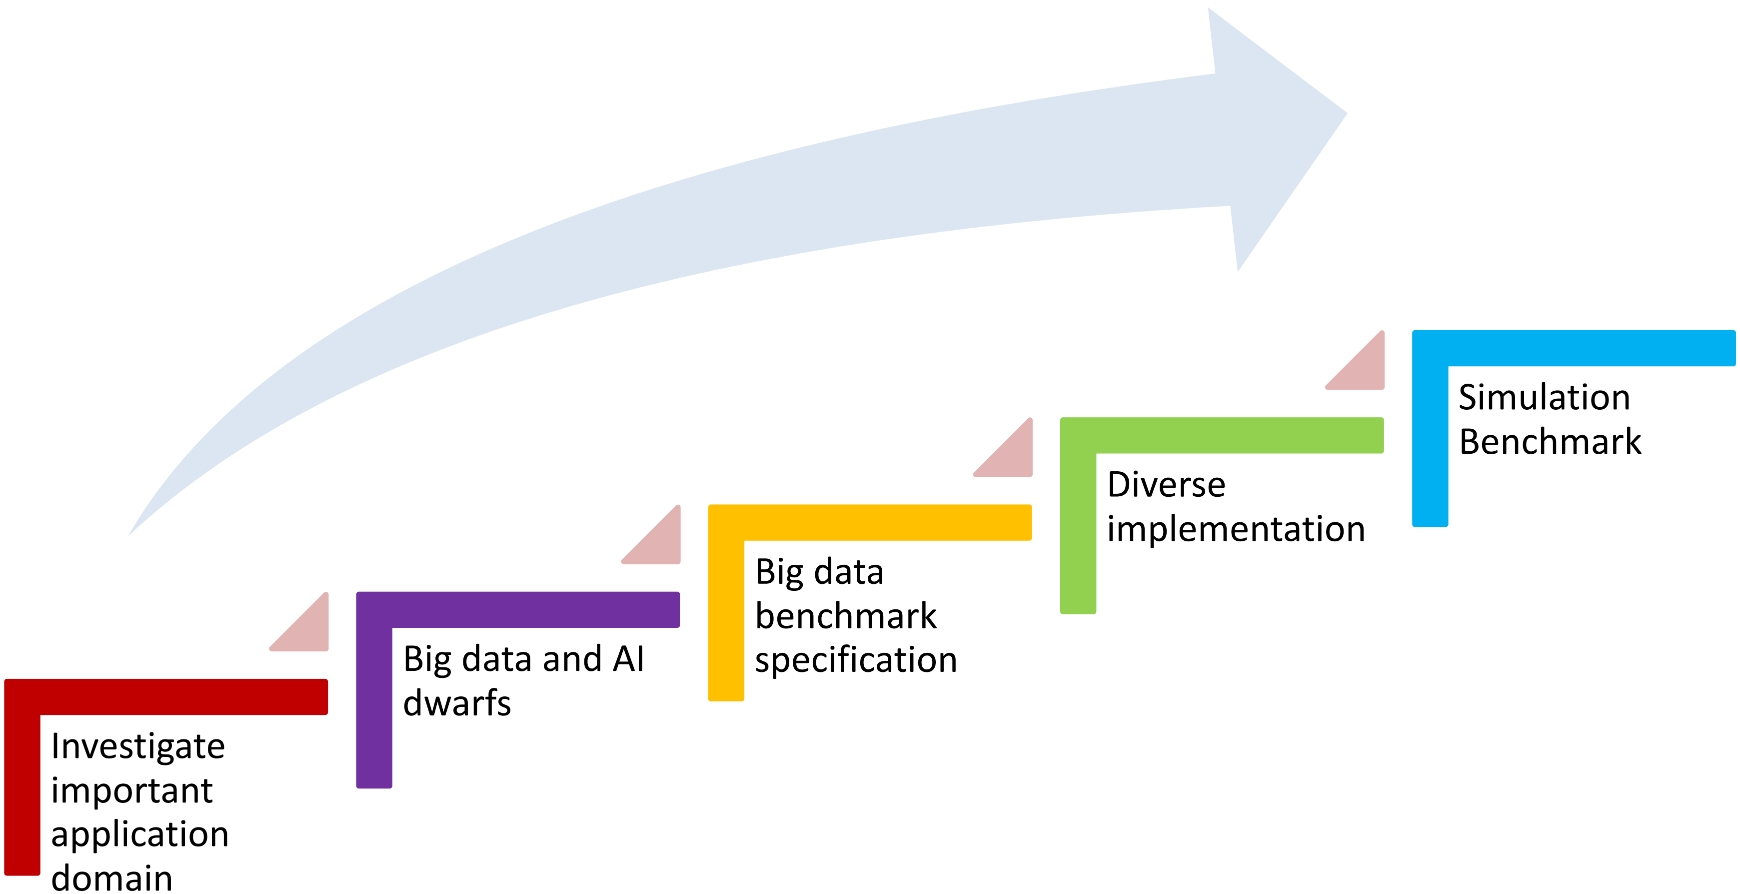

Investigate important application domain
Big data and AI dwarfs
Big data benchmark specification
Diverse implementation
Simulation Benchmark

Benchmarking Methodology

- Specification
  - Micro, component and application benchmark

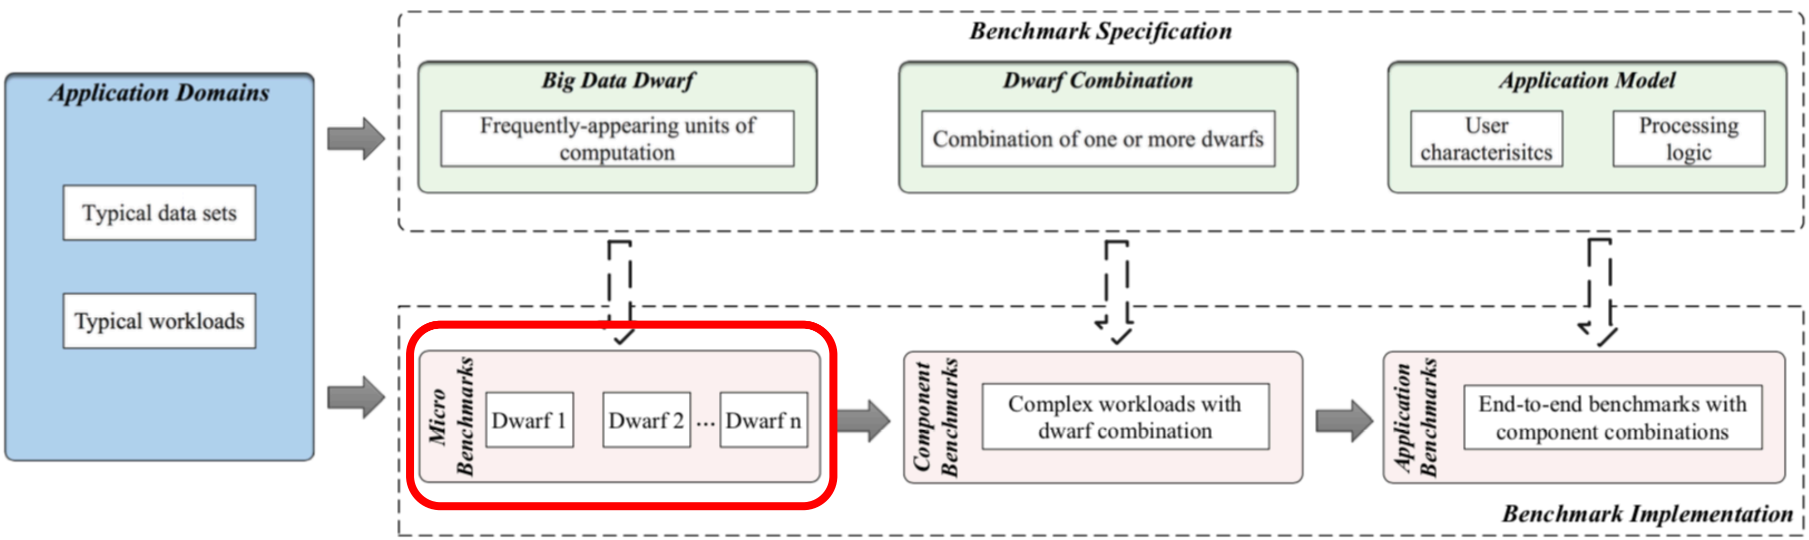

Five Application Domains

DDBJ/EMBL/GenBank database Growth

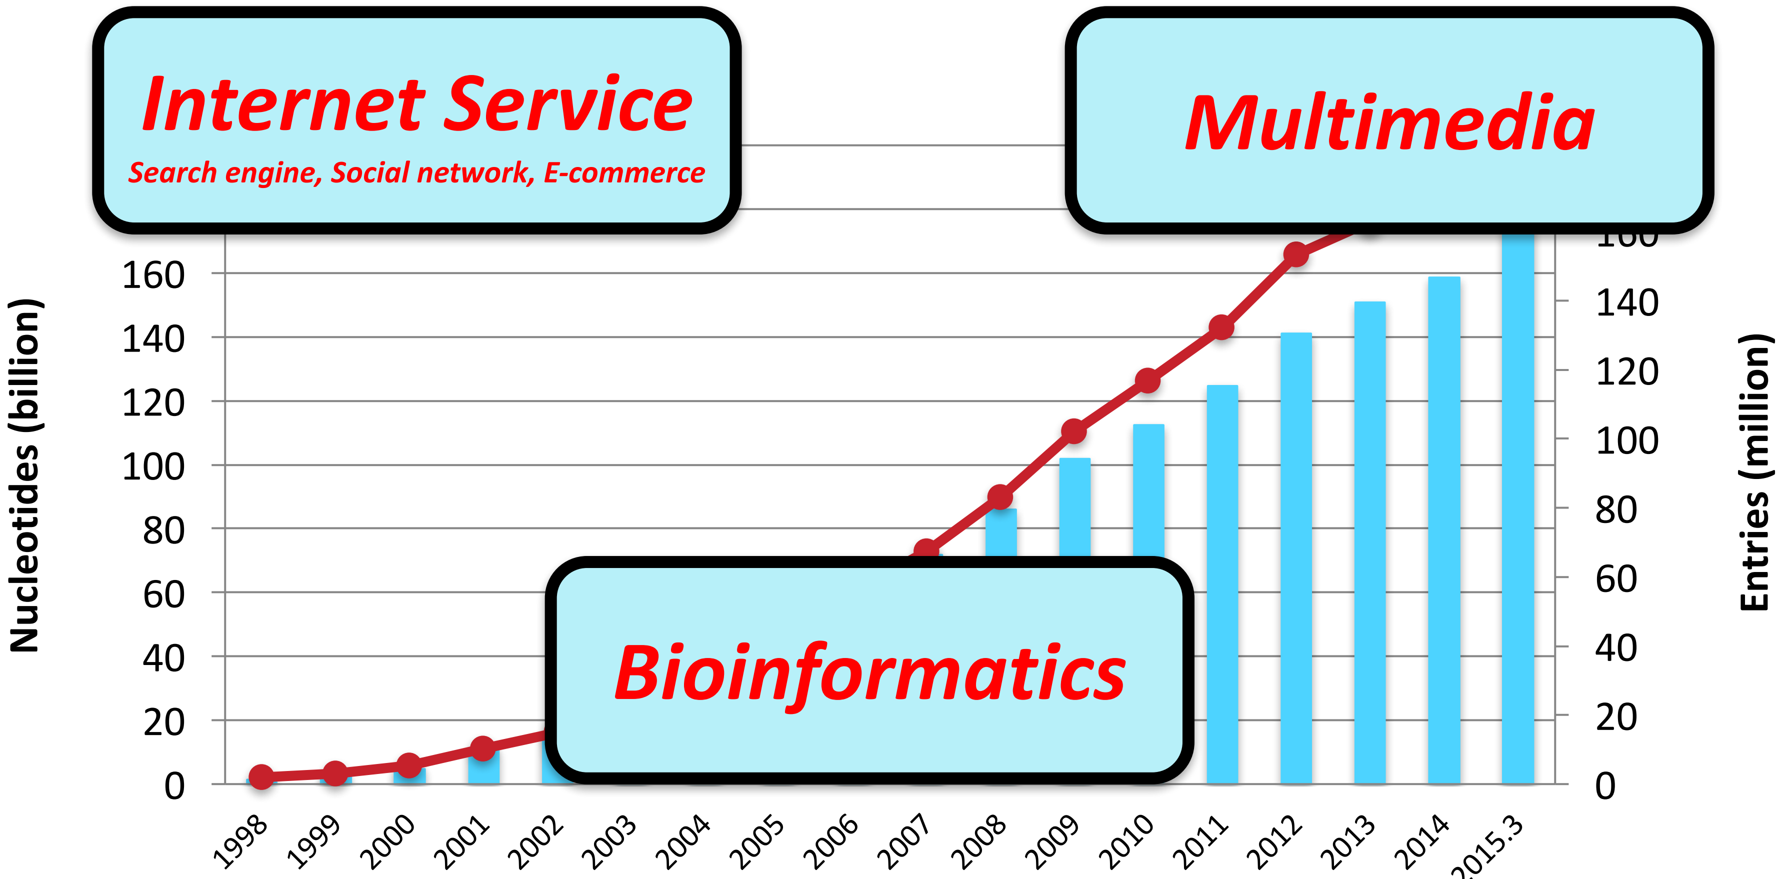

http://www.ddbj.nig.ac.jp/breakdown_stats/dbgrowth-e.html#dbgrowth-graph

BigDataBench 4.0 - Dataset

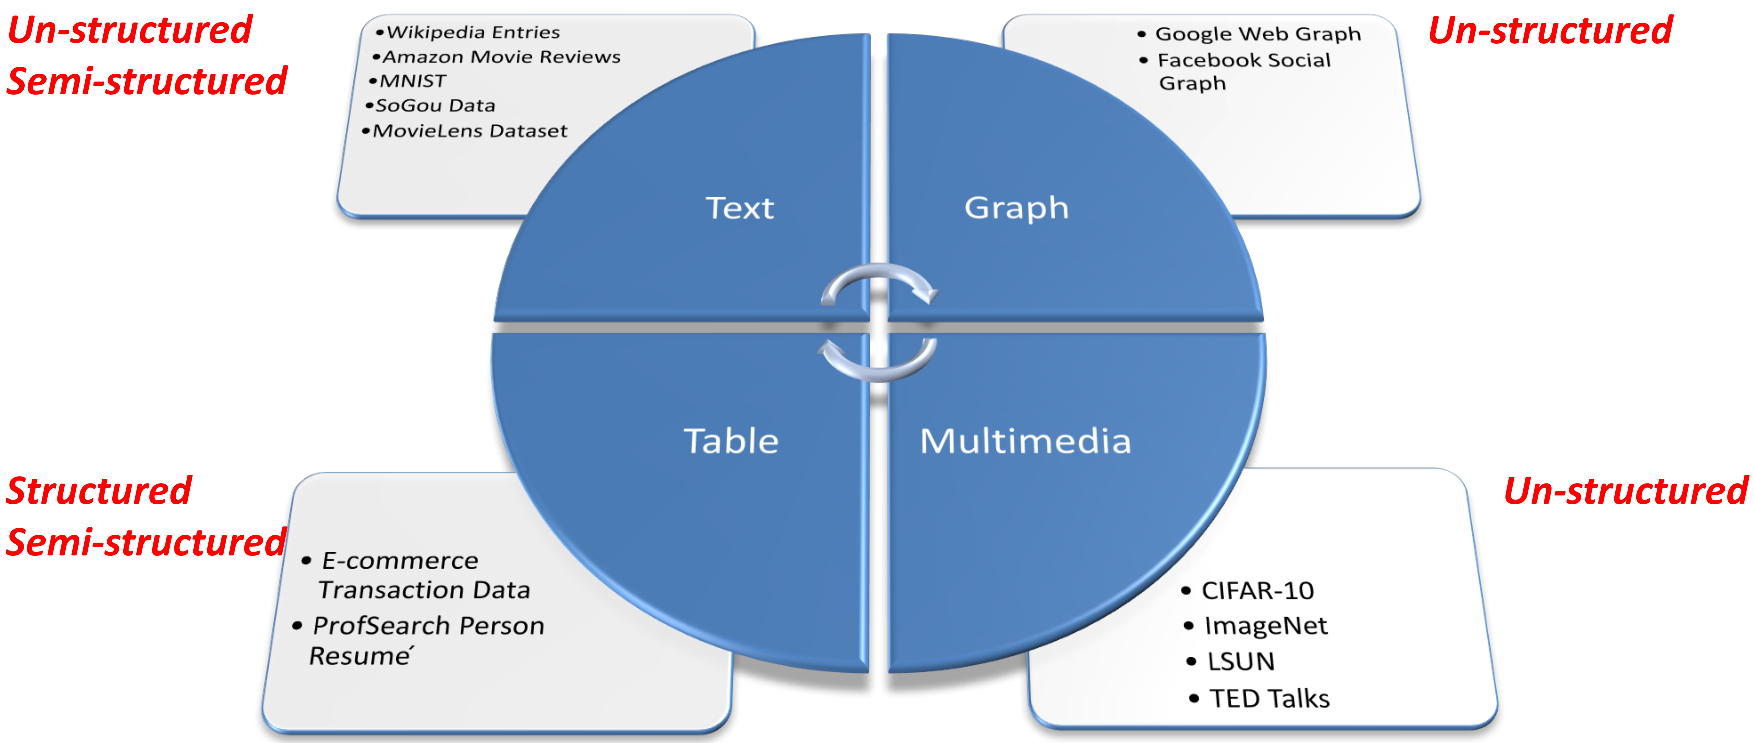

Big Data Generator Suite

- BDGS Architecture

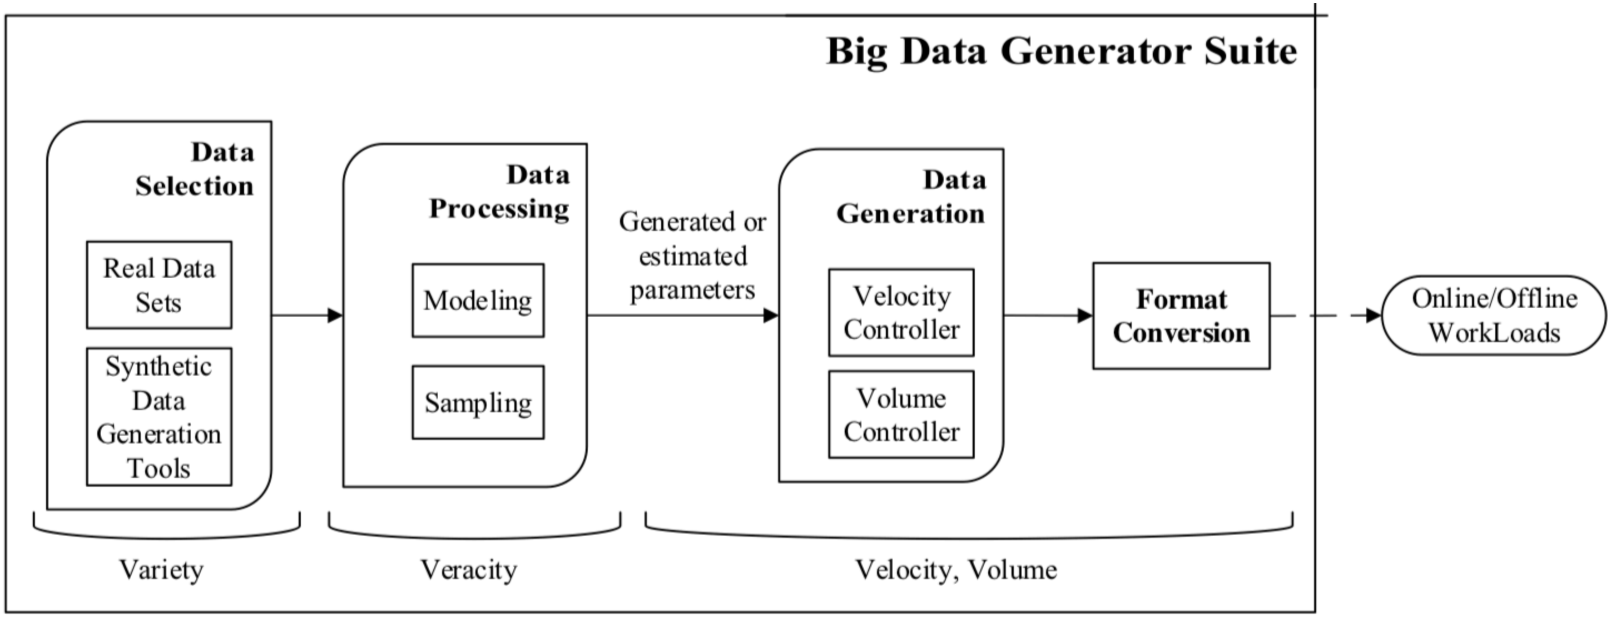

BDGS: Text Data Generator

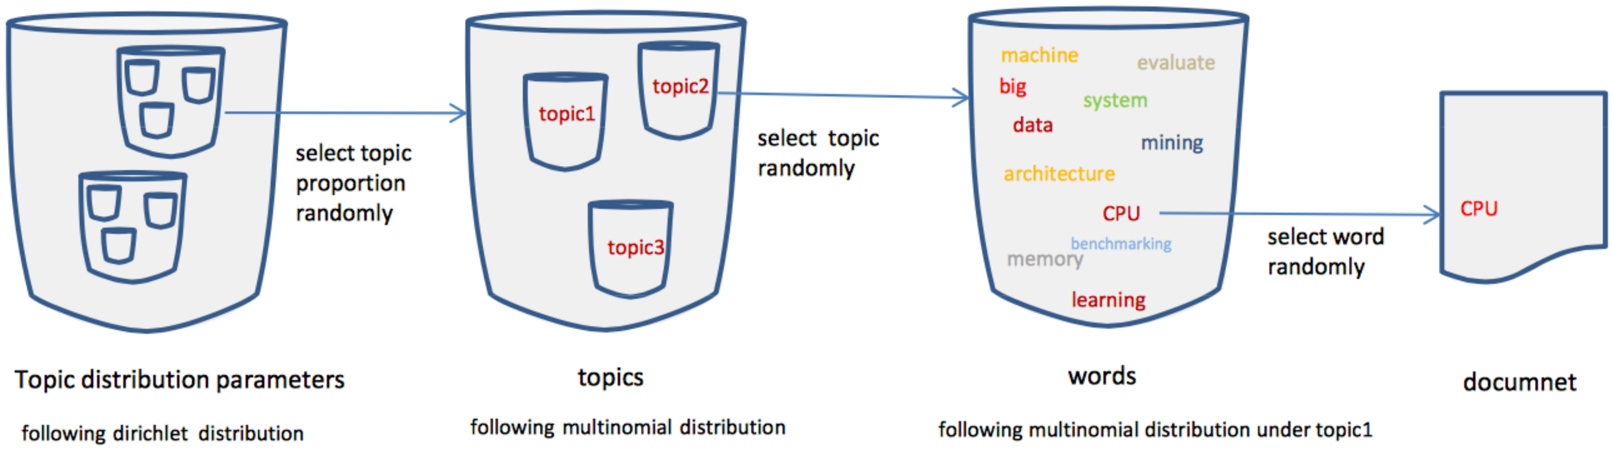

- Modeling on topic and word level
- Words are drew from distribution under particular topic
- Topics are selected from different distribution with parameters following a dirichlet distribution

Micro Benchmarks

Offline analytics

Graph analytics

Streaming

NoSQL

AI

Micro Benchmark	Involved Dwarf	Application Domain	Workload Type	Data Set	Software Stack
Sort	Sort	SE, SN, EC, MP, BI	Offline analytics	Wikipedia entries	Hadoop, Spark, Flink, MPI
Grep	Set		Offline analytics	Wikipedia entries	Hadoop, Spark, Flink, MPI
			Streaming	Random Generate	Spark streaming
WordCount	Basic statistics		Offline analytics	Wikipedia entries	Hadoop, Spark, Flink, MPI
MD5	Logic		Offline analytics	Wikipedia entries	Hadoop, Spark, MPI
Connected Component	Graph	SN	Graph analytics	Facebook social network	Hadoop, Spark, Flink, GraphLab, MPI
RandSample	Sampling	SE, MP, BI	Offline analytics	Wikipedia entries	Hadoop, Spark, MPI
FFT	Transform	MP	Offline analytics	Two-dimensional matrix	Hadoop, Spark, MPI
Matrix Multiply	Matrix	SE, SN, EC, MP, BI	Offline analytics	Two-dimensional matrix	Hadoop, Spark, MPI
Read	Set	SE, SN, EC	NoSQL	ProfSearch resumes	HBase, MongoDB
Write	Set	SE, SN, EC	NoSQL	ProfSearch resumes	HBase, MongoDB
Scan	Set	SE, SN, EC	NoSQL	ProfSearch resumes	HBase, MongoDB
Convolution	Transform	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
Fully Connected	Matrix	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
Relu	Logic	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
Sigmoid	Matrix	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
Tanh	Matrix	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
MaxPooling	Sampling	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
AvgPooling	Sampling	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
CosineNorm [36]	Basic Statistics	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
BatchNorm [37]	Basic Statistics	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
Dropout [38]	Sampling	SN, EC, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch

Component Benchmarks

Online service

Streaming

Offline analytics

Graph analytics

Data warehouse

AI

Component Benchmark	Involved Dwarf	Application Domain	Workload Type	Data Set	Software Stack
Xapian Server	Get, Put, Post	SE	Online service	Wikipedia entries	Xapian
PageRank	Matrix, Sort, Basic statistics, Graph	SE	Graph analytics	Google web graph	Hadoop, Spark, Flink, GraphLab, MPI
Index	Logic, Sort, Basic statistics, Set	SE	Offline analytics	Wikipedia entries	Hadoop, Spark
Rolling top words	Sort, Basic statistics	SN	Streaming	Random generate	Spark streaming, JStorm
Kmeans	Matrix, Sort, Basic statistics	SE, SN, EC, MP, BI	Offline analytics	Facebook social network	Hadoop, Spark, Flink, MPI
			Streaming	Random generate	Spark streaming
Collaborative Filtering	Graph, Matrix	EC	Offline analytics	Amazon movie review	Hadoop, Spark
		EC	Streaming	MovieLens dataset	JStorm
Naive Bayes	Basic statistics, Sort	SE, SN, EC	Offline analytics	Amazon movie review	Hadoop, Spark, Flink, MPI
SIFT	Matrix, Sampling, Transform, Sort	MP	Offline analytics	ImageNet	Hadoop, Spark, MPI
LDA	Matrix, Graph, Sampling	SE	Offline analytics	Wikipedia entries	Hadoop, Spark, MPI
OrderBy	Set, Sort	EC	Data warehouse	E-commerce transaction	Hive, Spark-SQL, Impala
Aggregation	Set, Basic statistics	EC		E-commerce transaction	Hive, Spark-SQL, Impala
Project, Filter	Set	EC	Data warehouse	E-commerce transaction	Hive, Spark-SQL, Impala
Select, Union	Set	EC		E-commerce transaction	Hive, Spark-SQL, Impala
Alexnet	Matrix, Transform, Sampling, Logic, Basic statistics	SN, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
Googlenet		SN, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
Resnet		SN, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
Inception Resnet V2		SN, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
VGG16		SN, MP, BI	AI	Cifar, ImageNet	TensorFlow, Caffe, pyTorch
DCGAN		SN, MP, BI	AI	LSUN	TensorFlow, Caffe, pyTorch
WGAN		SN, MP, BI	AI	LSUN	TensorFlow, Caffe, pyTorch
GAN	Matrix, Sampling, Logic, Basic statistics	SN, MP, BI	AI	LSUN	TensorFlow, Caffe, pyTorch
Seq2Seq		SE, EC, BI	AI	TED Talks	TensorFlow, Caffe, pyTorch
Word2vec	Matrix, Basic statistics, Logic	SE, SN, EC	AI	Wikipedia entries, Sogou data	TensorFlow, Caffe, pyTorch

Software Stacks

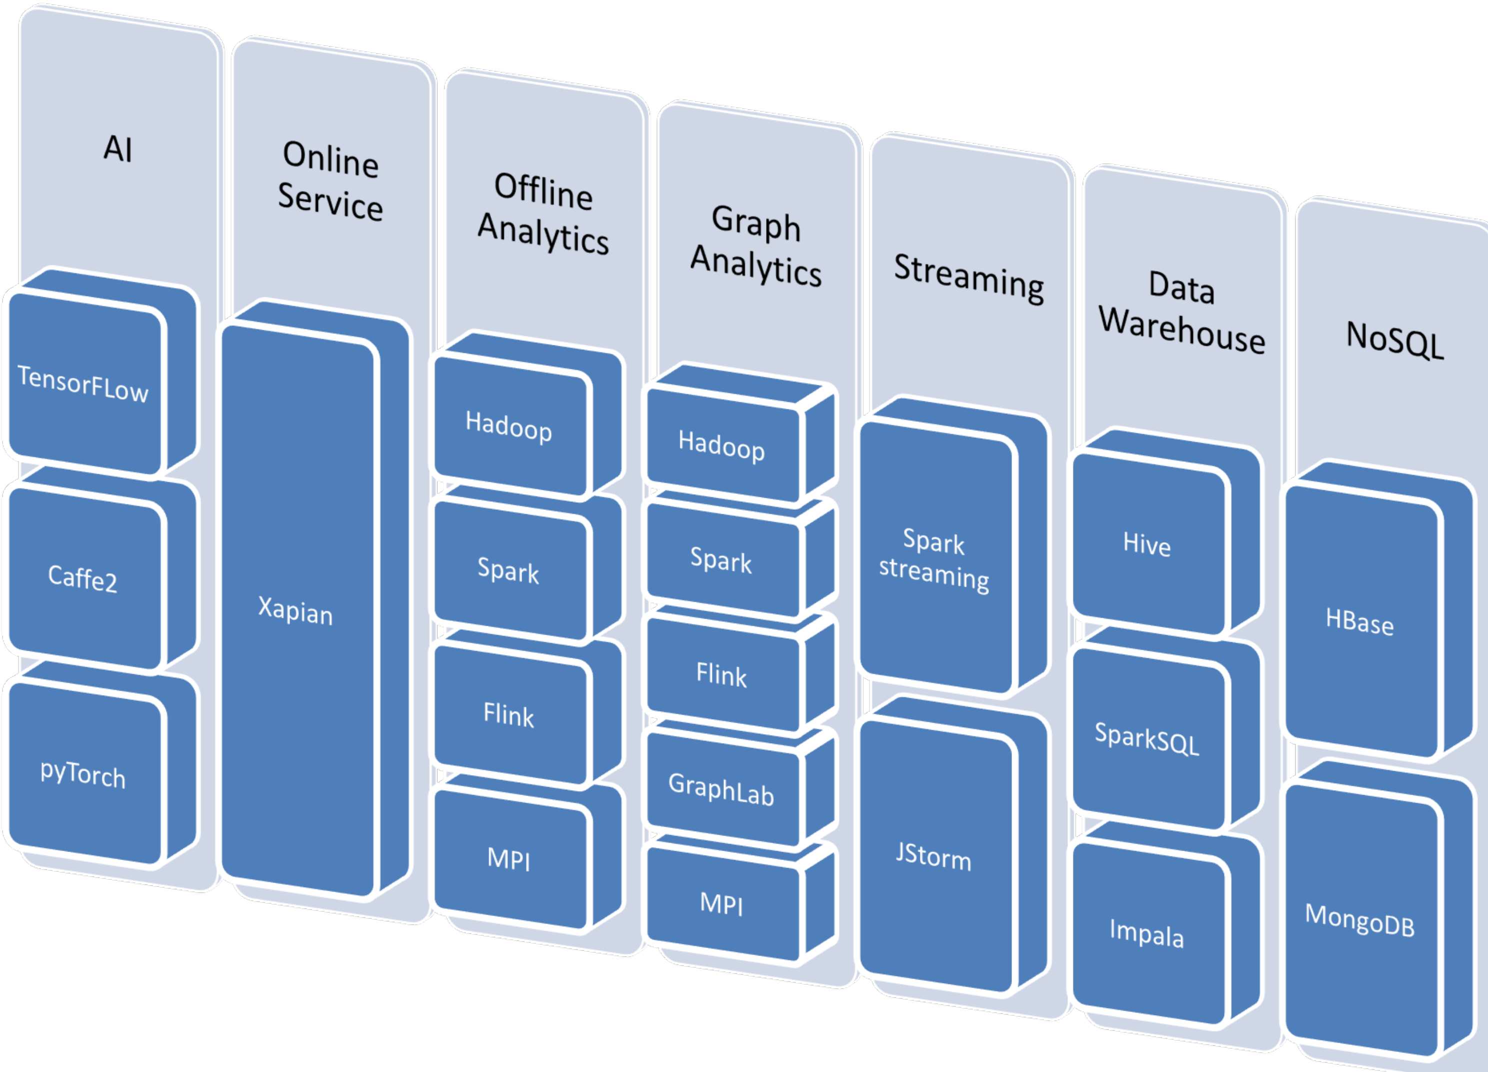

First part

- Introduction of BigDataBench 4.0
- BigDataBench Benchmarking Methodology
- *Simulation Benchmarks*
- Characterization

Dwarf-based Simulation Methodology

- DAG-like combinations of dwarfs
  - Different weights
  - Computation logic

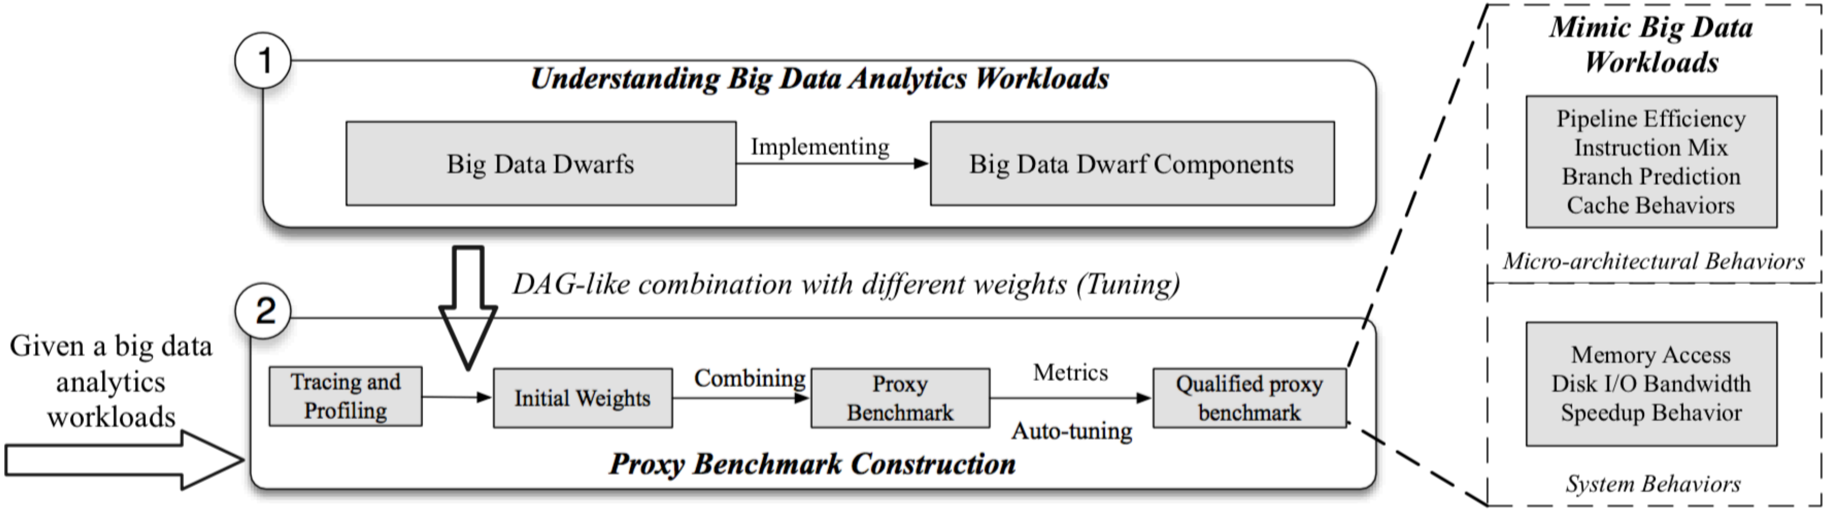

Simulation Benchmark for Big Data

- Simulation benchmarks for Hadoop workloads
  - 100X runtime speedup
  - 90+% data accuracy

  - OpenMP & Pthread Implementations
    - Provide a unified memory management module
      - Mimic JVM garbage collection (GC) process

Memory Management Module

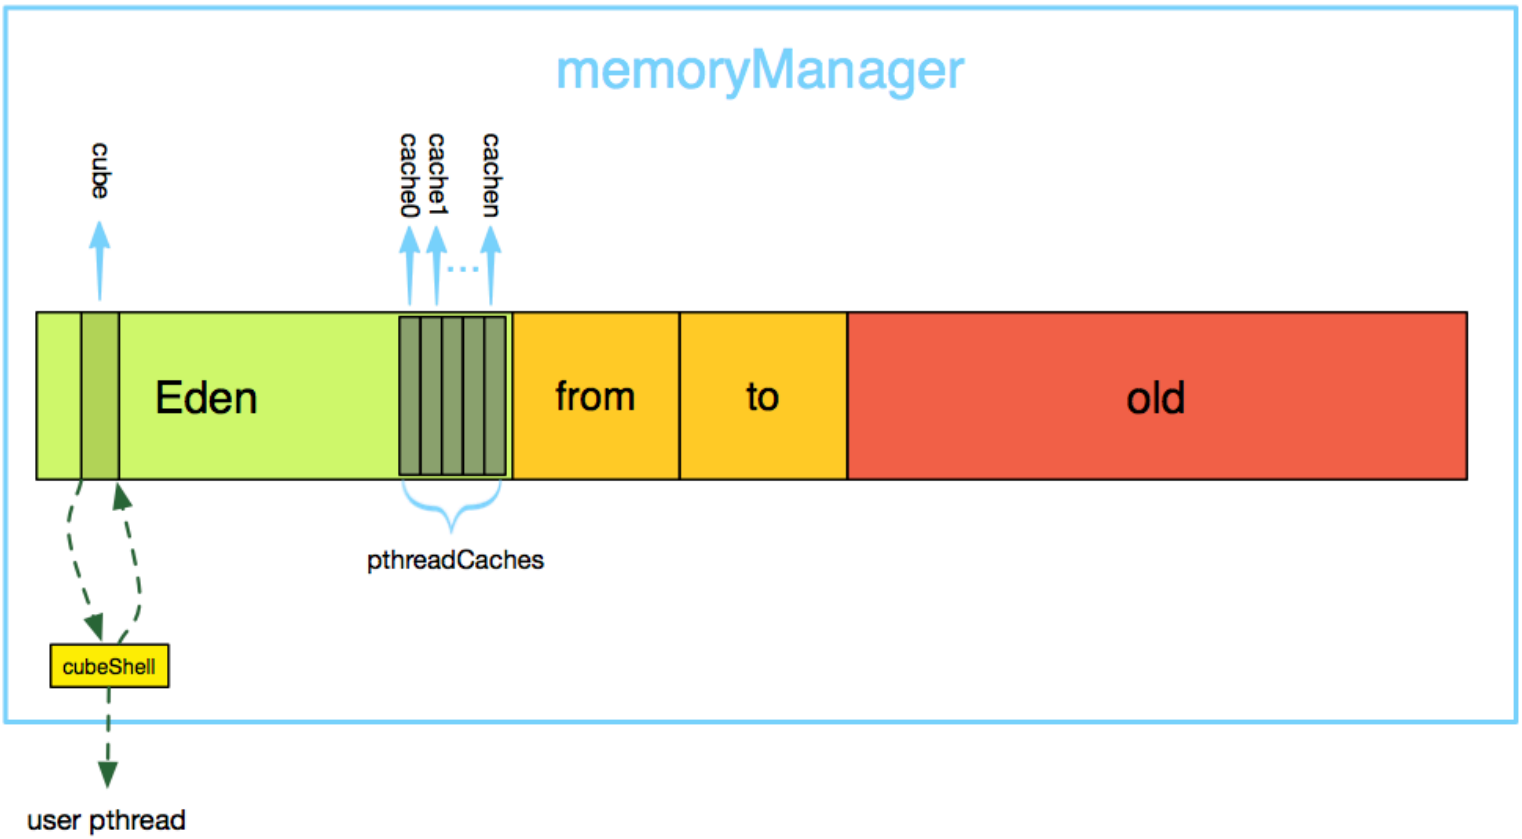

Dwarf Components for Big Data

- Data generation tools
- Dwarf implementations (OpenMP & Pthreads)

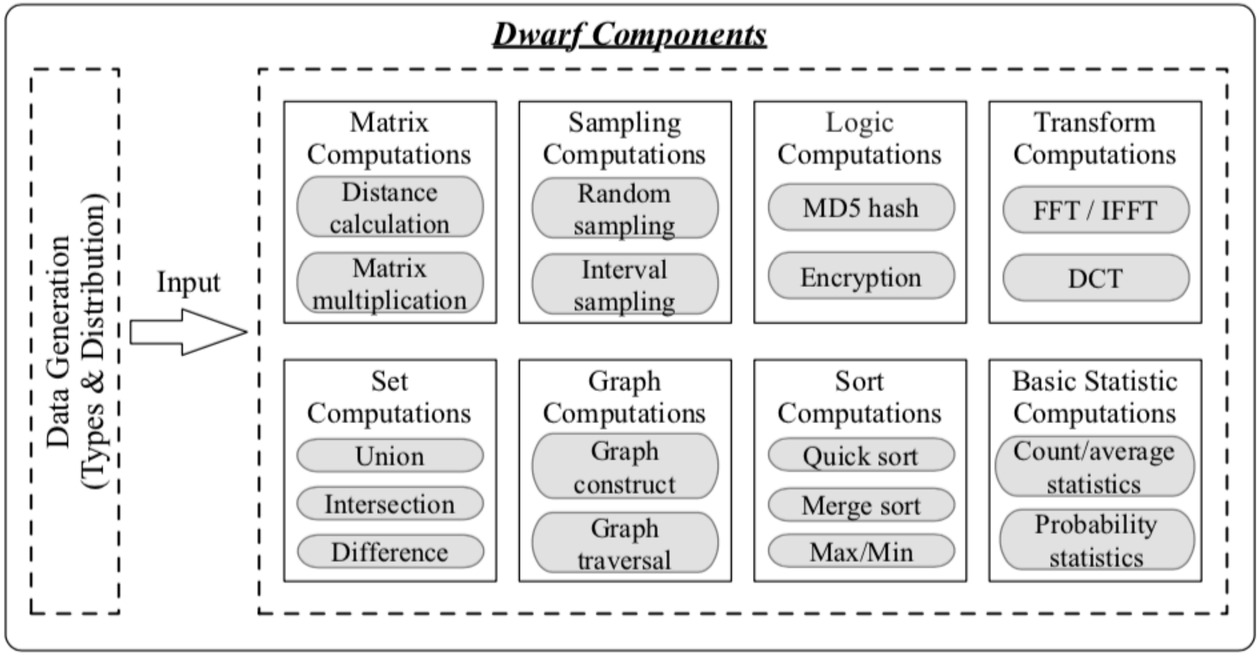

Simulation Benchmarks for AI

- Dwarf implementations (OpenMP & Pthreads)

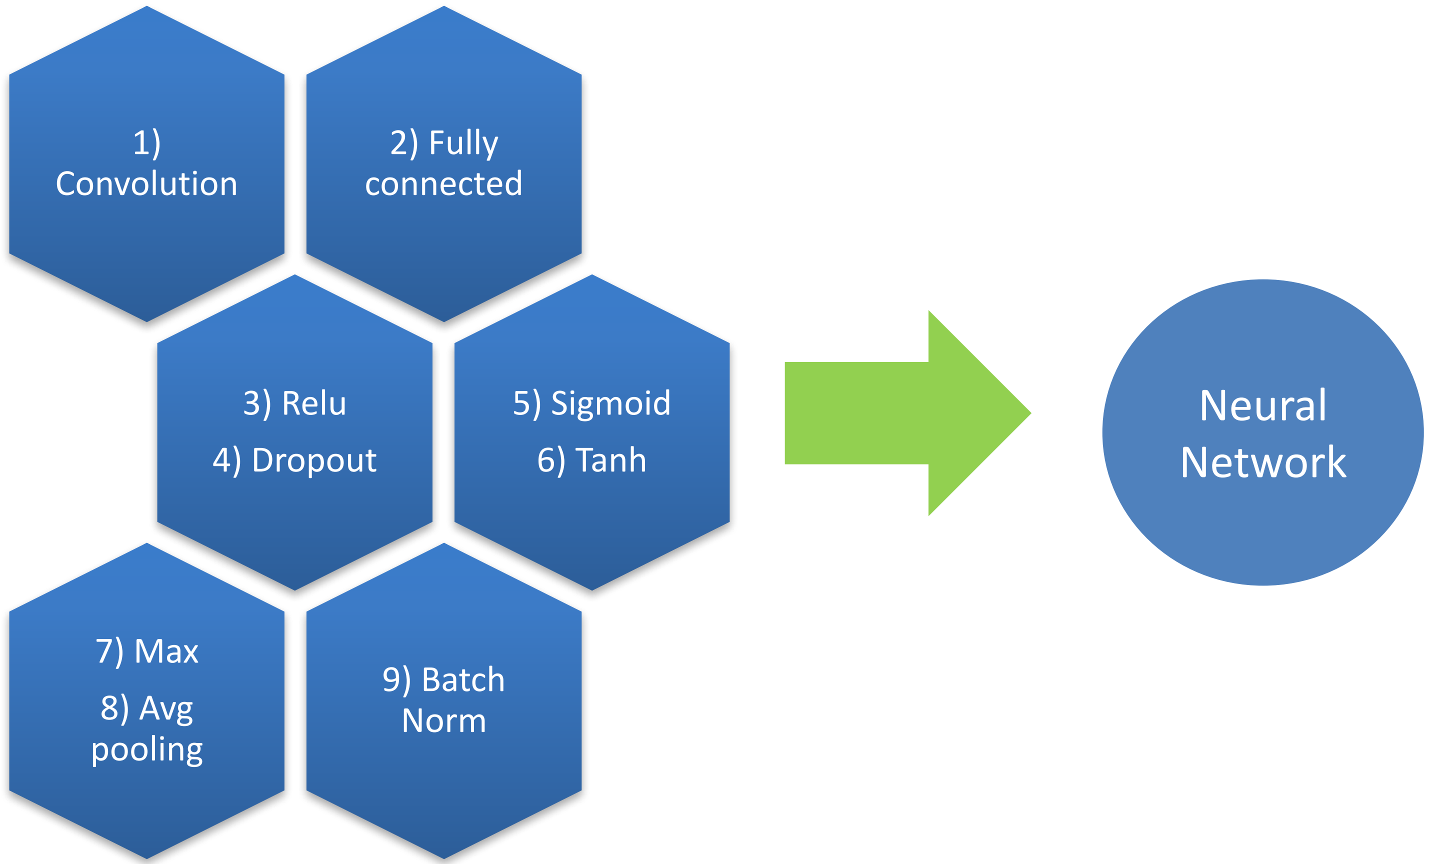

First part

- Introduction of BigDataBench 4.0
- BigDataBench Benchmarking Methodology
- Simulation Benchmarks
- *Characterization*

Workload Characterization

Hardware

- 3-node Hadoop cluster
  - Network: 1 Gb Ethernet network
  - Processor: Intel Xeon E5-2620 v3 (Haswell)

Hardware Configurations			
CPU Type		Intel CPU Core	
Intel ®Xeon E5-2620 V3		12 cores@2.40G	
L1 DCache	L1 ICache	L2 Cache	L3 Cache
12 × 32 KB	12 × 32 KB	12 × 256 KB	15MB
Memory		64GB,DDR4	
Disk		SATA@7200RPM	
Ethernet		1Gb	
Hyper-Threading		Disabled	

Software Configurations	
Operating System	CentOS 7.2
Linux Kernel	4.1.13
JDK Version	1.8.0_65
Hadoop Version	2.7.1
Hive Version	0.9.0
HBase Version	1.0.1
Spark Version	1.5.2
Tensorflow Version	1.0

Software

- Software version
  - CentOS 7.2, Kernel 4.1.13
  - JDK version: 1.8.0_65
  - Hadoop version: 2.7.1
- Compared benchmarks
  - SPEC CPU2006
  - HPCC 1.4.0
  - PARSEC 2.0

Benchmark

- Seven workload types

Execution Performance

- ILP and MLP
  - AI: ILP slightly lower than SPECCPU, MLP similar with HPCC
  - Big data has lower ILP and MLP than AI for almost all types, except Hive based data warehouse type

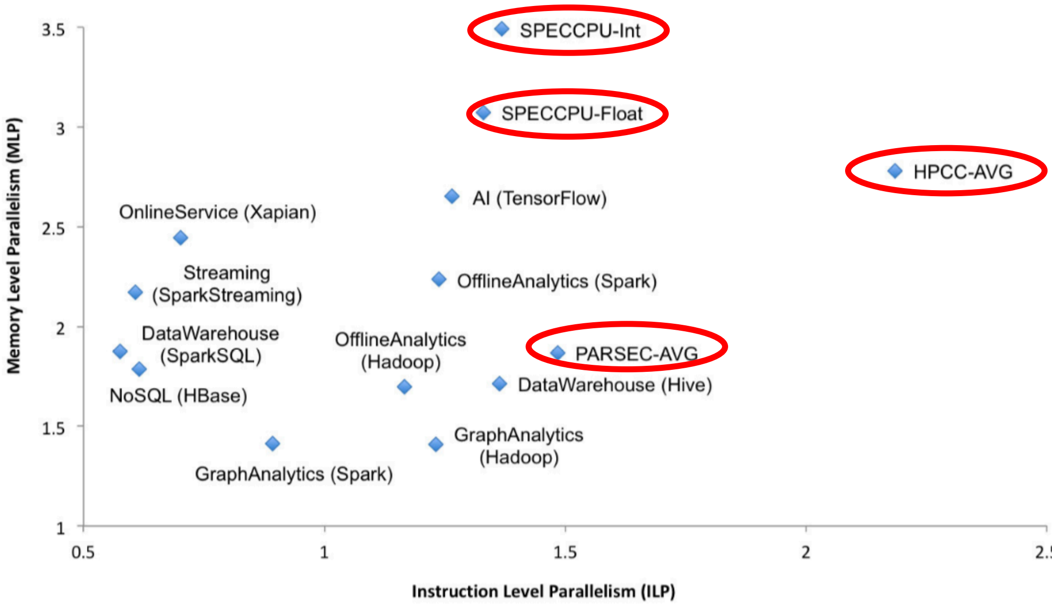

Pipeline Efficiency (Level 1)

- AI reflect similar pipeline behaviors with the traditional benchmarks
  - retiring (35% v.s. 39.8%), bad speculation (6.3% v.s. 6.1%), frontend bound (both about 9%), and backend bound (49.7% v.s. 45.1%)
- Big data and AI have a small fraction of bad speculation

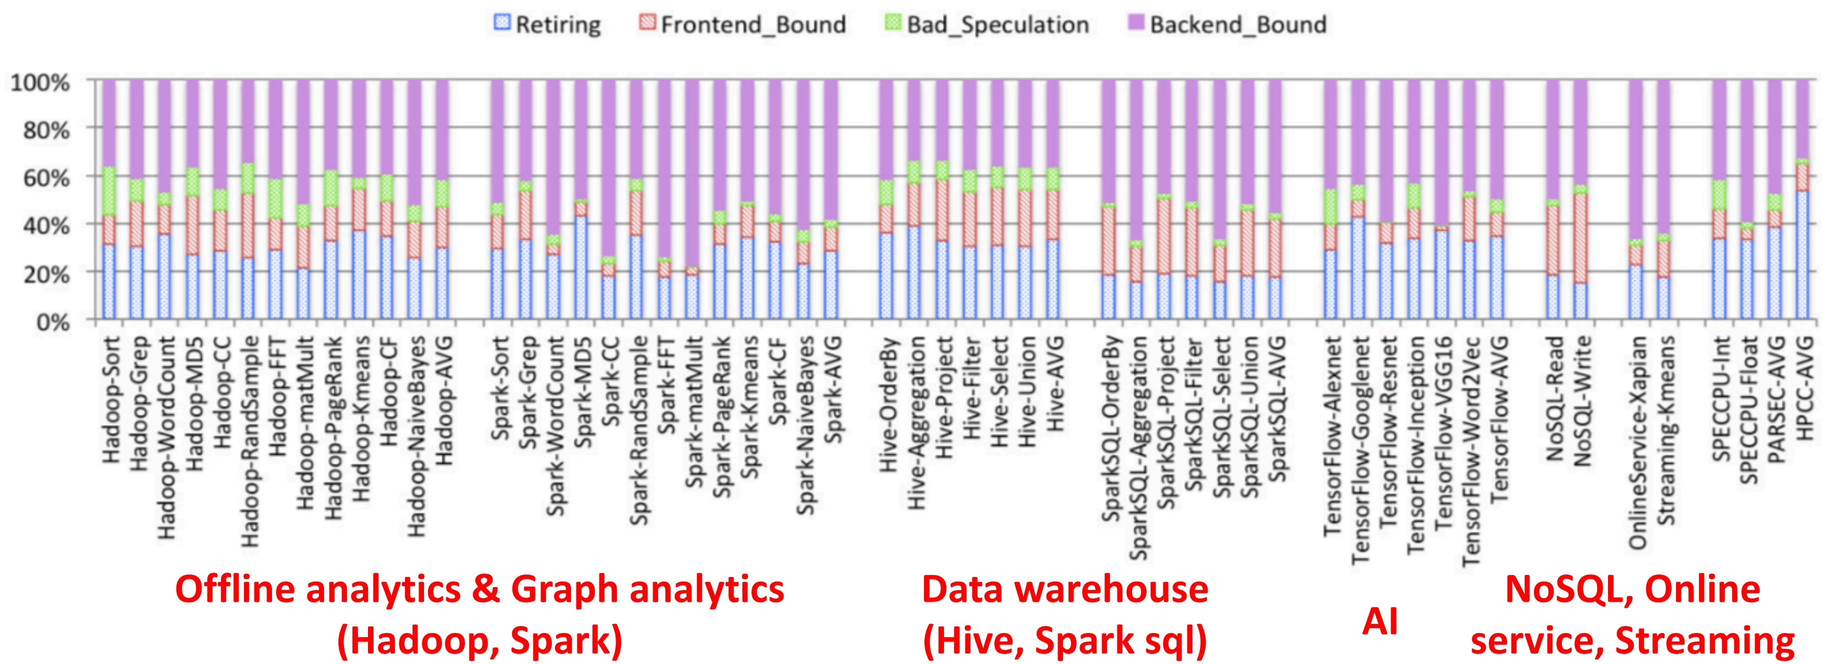

About Detailed Characterization

- BPOE workshop tomorrow will give detailed characterization results

- Look forward to your participation！

Download

- http://prof.ict.ac.cn/download

- ***Packing & Testing now !***
- ***Release soon (April 1st, 2018)***

Conclusion

- BigDataBench 4.0
  - An open source dwarf-based big data and AI benchmark suite

- Website: http://prof.ict.ac.cn
- Technical Reports:
  - https://arxiv.org/pdf/1802.08254.pdf
  - https://arxiv.org/pdf/1801.09212.pdf


QUESTIONS
And
Answers