Data Motifs: A Benchmark Proposal for Big Data and AI

Wanling Gao
ICT, CAS &UCAS

Bench18, Seattle, WA, USA

INSTITUTE OF COMPUTING TECHNOLOGY
中科院计算所

Hardware and Software Co-design

A New Golden Age for Computer Architecture—Domain-specific co-design
>>John Hennessy and David Patterson

- Database system
  - Relation algebra
  - Five primitive and fundamental operators
- Big data and AI systems
  - Demand-driven
  - Case-by-case

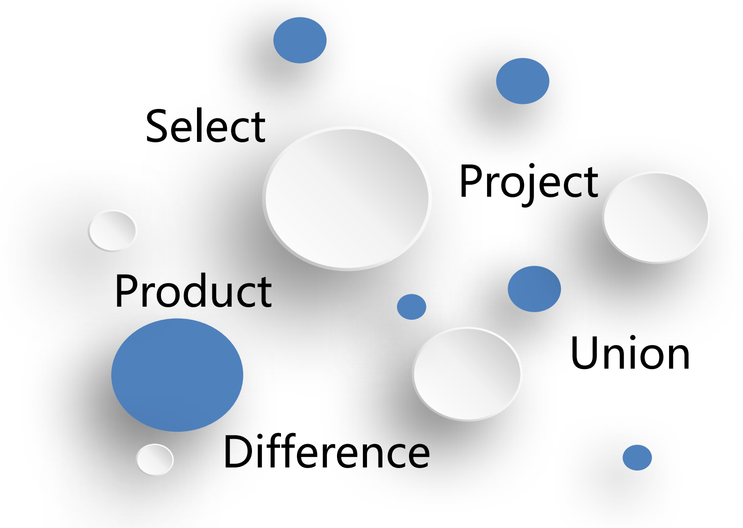

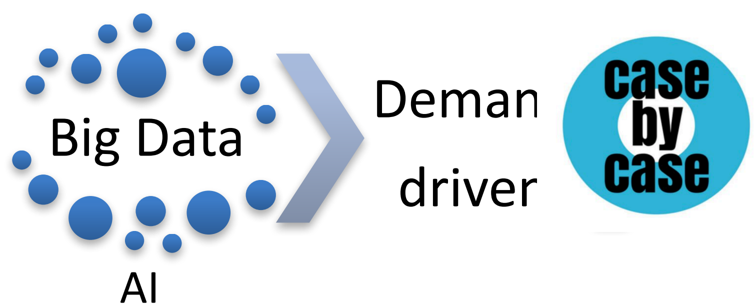

Costly and time-consuming !

Why Big Data and AI Benchmarking?

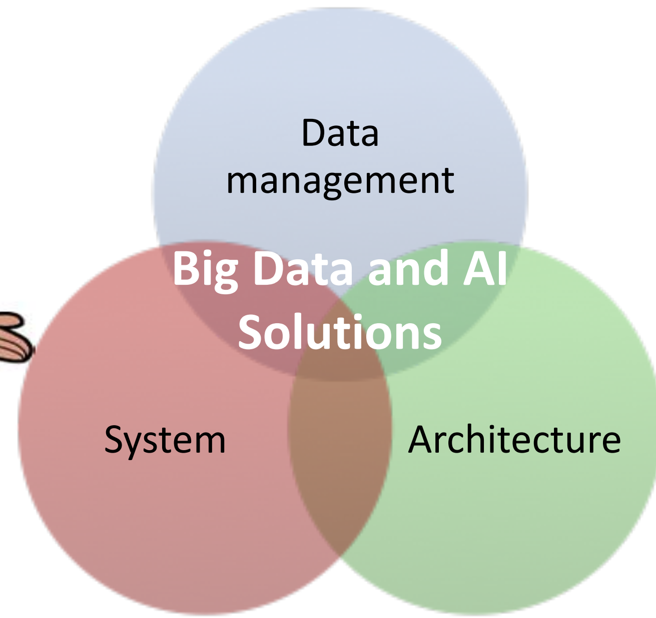

Measuring systems and architectures quantitatively

Benchmark Construction

- Top-down:
  - representative program selection
  - can yield accurate representations of the program space of interest
  - usually impossible to make any form of hard statements about the representativeness
- Bottom-up:
  - diverse range of characteristics
  - program characteristics are quantities that can be measured and compared
  - not all portions of the characteristics space are equally important

-- *C. Bienia. Benchmarking modern multiprocessors. PhD thesis, Princeton University, 2011.*

Current Benchmark Methodology

- Popularity
  - CloudSuite
  - Parsec
  - BigDataBench 3.0
- Based on one application domain
  - TPC-C、TPC-DS

Benchmarking Cost

- Relevant
- Portable
- Scaleble
- Simple

What makes a good benchmark?
---Jim Gray

- Complexity and diversity of big data and AI workloads
- Impact of data on workload behaviors

- How about benchmarking scalability
  - Benchmark constructing cost
  - Benchmark evaluation cost

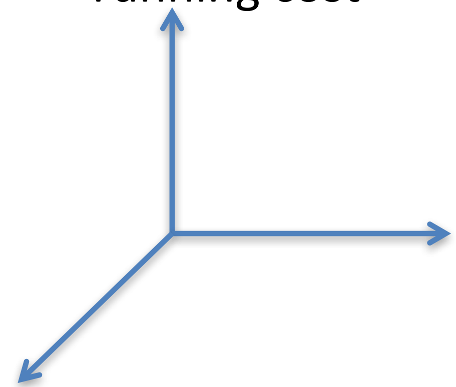

Benchmark Requirements for Big Data and AI

- Large-scale system and architecture evaluation
  - Portability for earlier stage
  - Comprehensiveness and reality for later stage
- Consistency across different communities
  - System： performance、cluster scalability
  - Architecture： micro-architecture behaviors
  - Machine learning： execution time, model prediction precision
- Reproducibility and interpretability

Data Motif-based Co-design & Benchmarking

- Data Motif Definition
  - Captures the common requirements of each class of unit of computation
    - Being reasonably divorced from individual implementations
  - ***A minimum set*** to represent ***maximum patterns***

Benchmarking scalability

Portability cost

Better interpretation of performance data

We need to understand What's the abstractions of frequently-appearing units of computation among big data and AI workloads (big data and AI motif)?

Overview

- **Looking back at history**
- What is Data Motif
- BigDataBench 4.0
  - Unified Big Data and AI Benchmark Suite
  - Proxy Benchmarks for Simulation
- Conclusion

Successful abstractions in other domains…

Abstraction - Relational Algebra

- Relational Algebra
  - Five primitive and fundamental operators
    - Theoretical foundation of database
    - Strong expression power
      - Compose complex queries

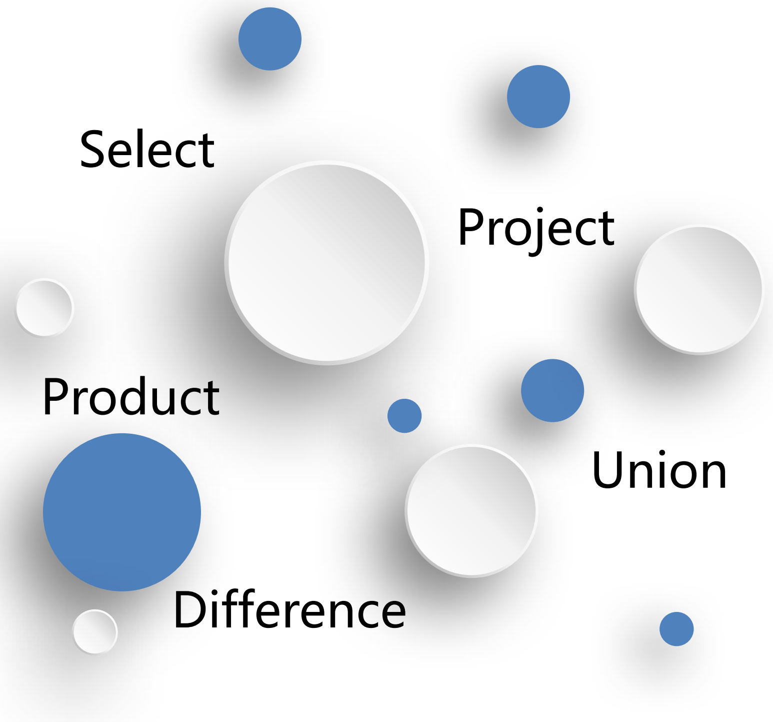

From E. F. Codd, A relational Model of Data for Large shared data banks. Communication of ACM, vol 13. no.6, 1970

Abstraction - Numerical Method

- Seven motifs would be important for the next decade

7 “Motifs”

Structu-red Grids

Unstru-ctured Grids

Dense linear algebra

Sparse linear algebra

- Phillip Colella proposed
- Simulation in the physical sciences is done out using various combinations of the following core algorithms
  - distinctive combination of computation and data access

FFT

Particles

Monte Carlo

From P. Colella, “Defining software requirements for scientific computing,” 2004.

Abstraction – Parallel Computing

■ Landscape of Parallel Computing Research

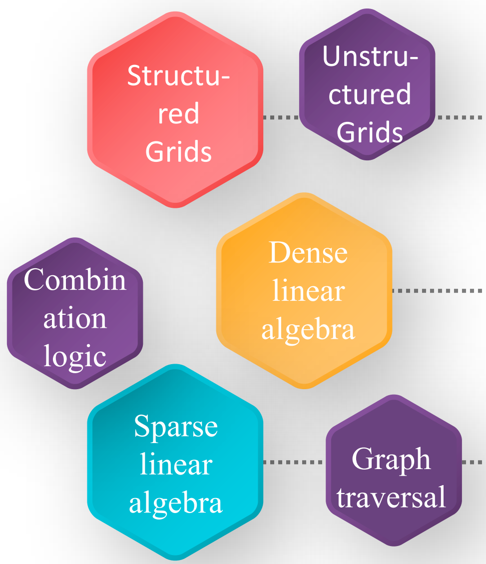

13 dwarfs

- Berkeley research group
- Define building blocks for creating libraries & frameworks
- A pattern of computation and communication

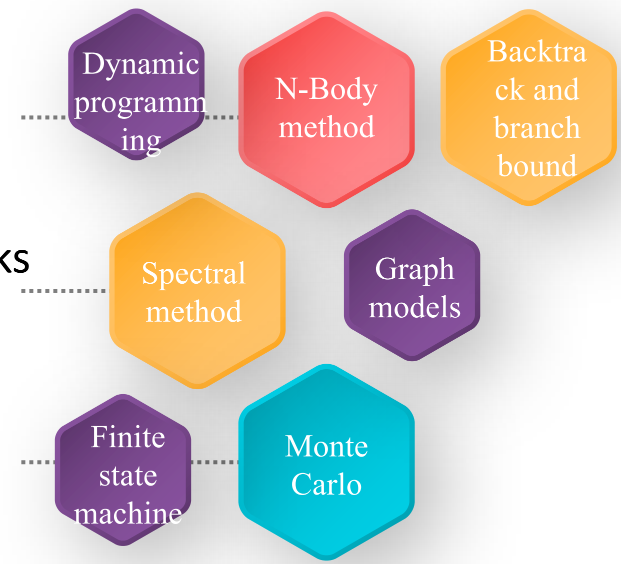

From K. Asanovic, R. Bodik, B. C. Catanzaro, J. J. Gebis, P. Husbands, K. Keutzer, D. A. Patterson, et al, "The landscape of parallel computing research: A view from berkeley," tech. rep., Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley, 2006.

Successful benchmarks based on abstractions …

TPC Functional Workload Model

- Application domain → encapsulate user cases
- Functions of abstraction
  - abstraction of the implementations of use cases in different application domains.
- Systems View and Physical View
  - Different systems and hardware

-- Yanpei Chen, Francois Raab, Randy Katz:
From TPC-C to Big Data Benchmarks: A Functional Workload Model, WBDB, 2012

TPC-C Methodology

- Functions of Abstraction
  - a mid-weight read-write transaction (i.e., New-Order)
  - a light-weight read-write transaction (i.e., Payment)
  - a mid-weight read-only transaction (i.e., Order-Status)
  - a batch of mid-weight read-write transactions (i.e., Delivery)
  - a heavy-weight read-only transaction (i.e., Stock-Level)
- Functional Workload Model
  - captures in an implementation-independent manner the load that the system needs to service

HPCC: Components

1. HPL (High Performance LINPACK) $Ax=b$

2. STREAM

name	kernel	bytes/iter	FLOPS/iter
COPY:	a(i) = b(i)	16	0
SCALE:	a(i) = q*b(i)	16	1
SUM:	a(i) = b(i) + c(i)	24	1
TRIAD:	a(i) = b(i) + q*c(i)	24	2

3. PTRANS $A \leftarrow A^T + B$

4. RandomAccess

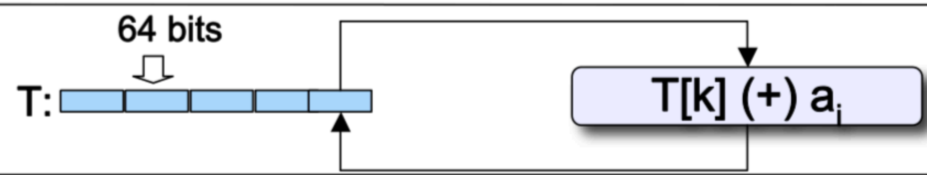

5. FFT

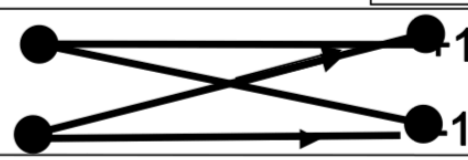

$$z_k = \Sigma x_j \exp(-2\pi\sqrt{-1}\, jk/n)$$

6. Matrix-matrix multiply $C \leftarrow s*C + t*A*B$

7. b_eff (effective bandwidth/latency)

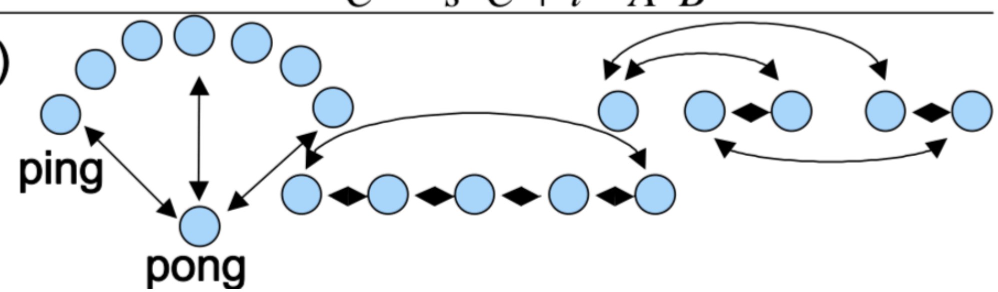

https://pdfs.semanticscholar.org/bf41/25645dd8a4c0b7f0b1c5502c3713d4ffe68c.pdf

HPCC Methodology

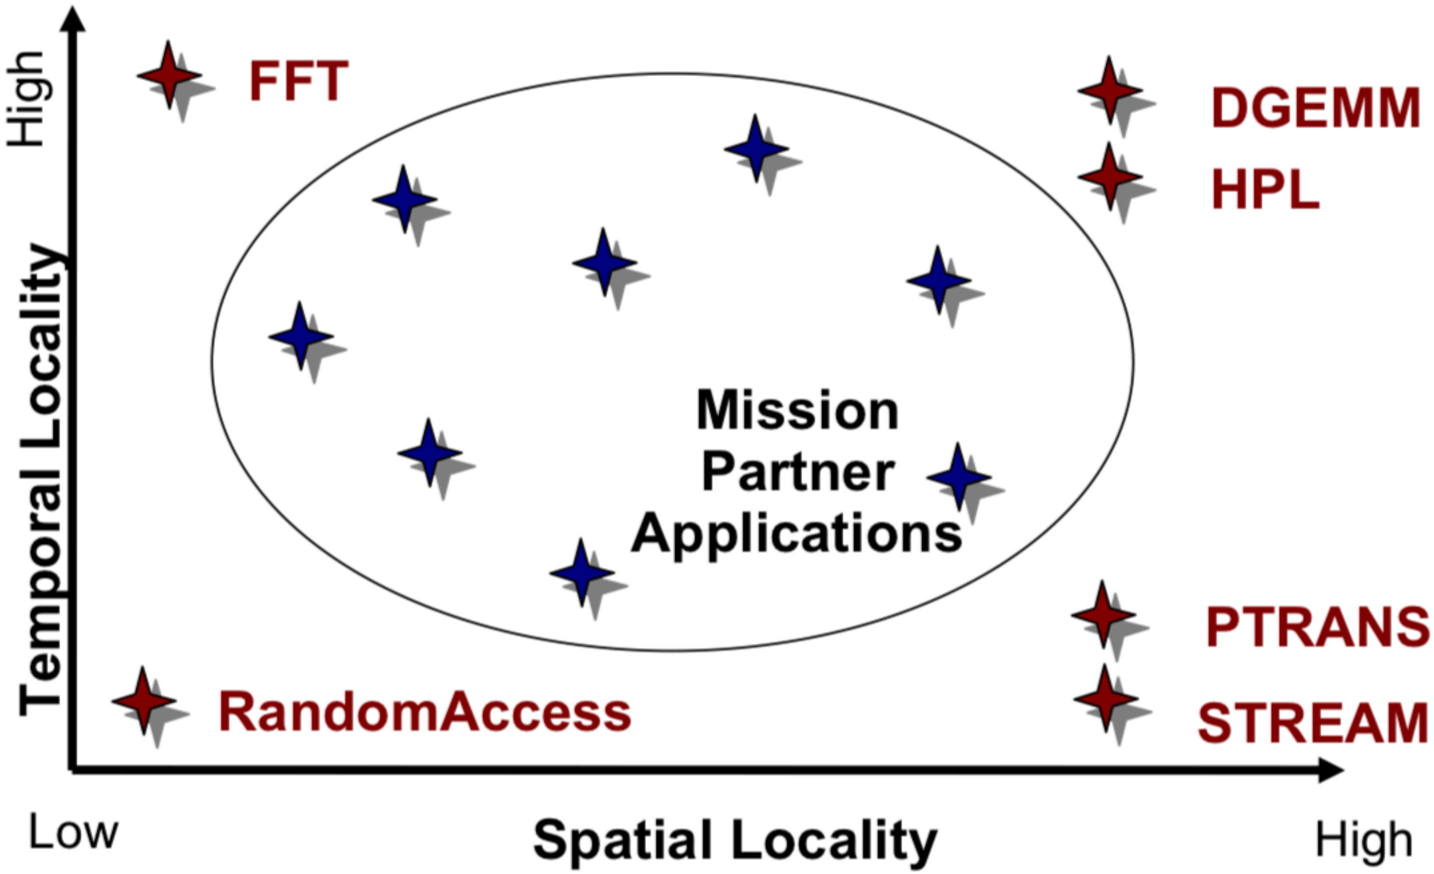

http://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.482.1783&rep=rep1&type=pdf

Overview

- Looking back at history
- **What is Data Motif**
- BigDataBench 4.0
  - Unified Big Data and AI Benchmark Suite
  - Proxy Benchmarks for Simulation
- Conclusion

AI Workload --- AlexNet

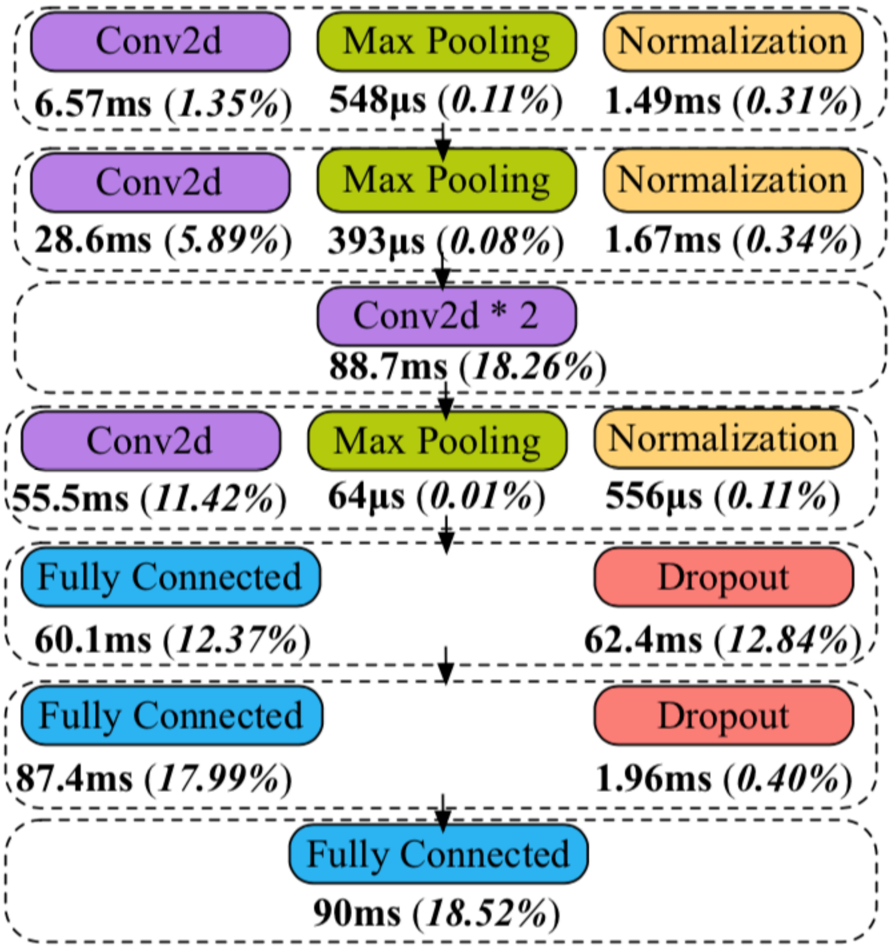

AlexNet
Units of Computation:

1) Convolution: 36.91%
----Conv2d

2) Sampling: 13.45%
----Max Pooling
----Dropout

3) Matrix Multiply: 48.87%
----Fully Connected

4) Basic Statics: 0.76%
----Normalization

Feature Extraction --- SIFT

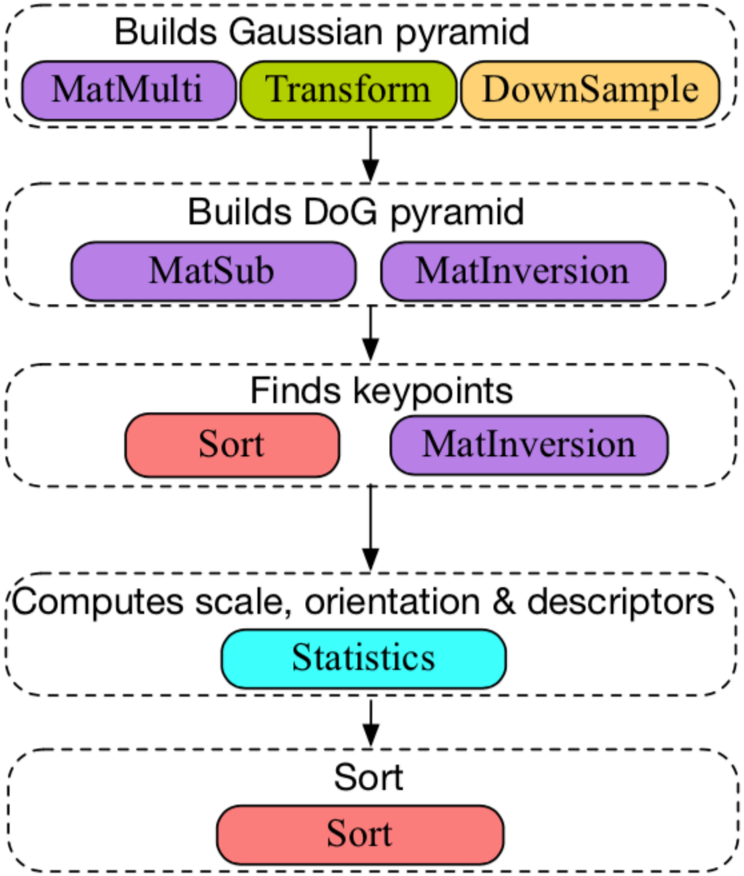

SIFT: Units of Computation

1) Builds Gaussian pyramid:13.16%

——Matrix Multiplication

——Transform

——DownSample

2) Builds DoG pyramid: 4.17%

——Matrix Subtraction

——Matrix Inversion

3) Finds keypoints: 26.01%

——Sort

——Matrix Inversion

4) Compute scale, orientation & descriptors: 53.11%

——Statistic

5) Sort: 0.53%

——Sort

What is Data motif?

- Data Motif---unified computation abstraction
  - abstractions of time-consuming units of computation

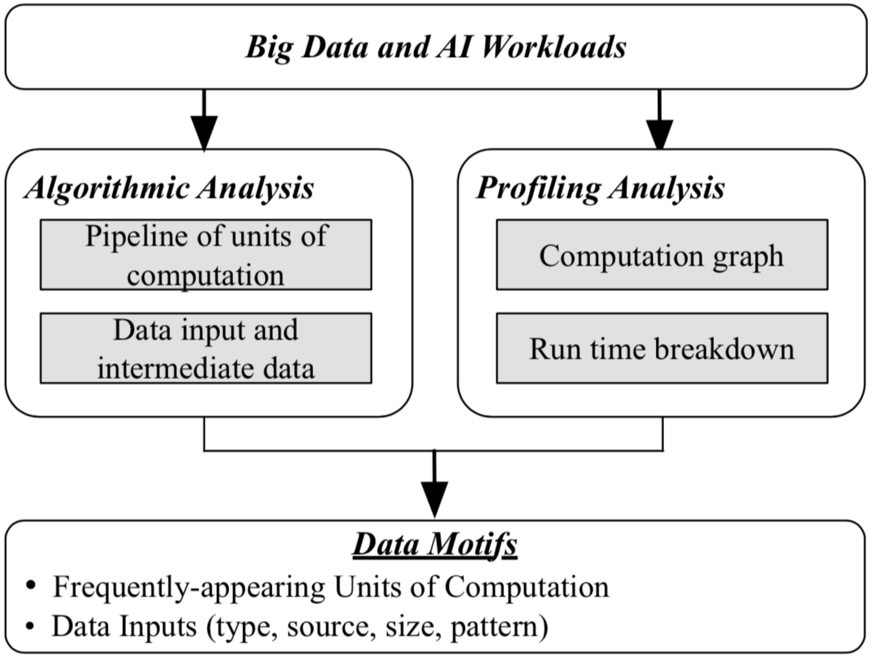

Wanling Gao, Jianfeng Zhan, Lei Wang, et al. Data Motif: A Lens towards Fully Understanding Big Data and AI Workloads. **PACT 2018.**

Algorithms

40+ algorithms with a broad spectrum

- Data mining/Machine learning
- Natural language processing
- Computer vision
- Bioinformatics

Data Mining & Machine Learning

C4.5/CART/ID3, Logistic regression, SVM, KNN, HMM, Maximum-entropy markov model, Conditional random field, PageRank, HITS, Aporiori, FP-growth, Principal component analysis, Back Propagation, Adaboost, MCMC, Connected component, Random forest, CF, CNN, DBN

Natural Language Processing

Latent semantic indexing, pLSI, Latent dirichlet allocation, Index, Porter Stemming, Sphinx speech recognition

Investigated Algorithms

Computer Vision

MPEG-2, Scale-invariant feature transform, Image segmentation, Ray Tracing

Bioinformatics

Needleman-Wunsch, Smith-Waterman, Basic local alignment search tool (BLAST)

Catergory	Application Domain	Workload	Unit of Computation
Deep Learning	Image Recognition	Convolutional neural network(CNN)	Matrix, Sampling, Transform
	Speech Recognition	Deep belief network(DBN)	Matrix, Sampling
Graph Mining	Search Engine	PageRank	Matrix, Graph, Sort
	Community Detection	BFS, Connected component(CC)	Graph
Dimension Reduction	Image Processing	Principal components analysis(PCA)	Matrix
	Text Processing	Latent dirichlet allocation(LDA)	Statistics, Sampling
Recommendation	Association Rules Mining Electronic Commerce	Aporiori	Statistics, Set
		FP-Growth	Graph, Set, Statistics
		Collaborative filtering(CF)	Graph, Matrix
Classification	Image Recognition Speech Recognition Text Recognition	Support vector machine(SVM)	Matrix
		K-nearest neighbors(KNN)	Matrix, Sort, Statistics
		Naive bayes	Statistic
		Random forest	Graph, Statistics
		Decision tree(C4.5/CART/ID3)	Graph, Statistics
Clustering	Data Mining	K-means	Matrix, Sort
Feature Preprocess	Image Processing Signal Processing Text Processing	Image segmentation(GrabCut)	Matrix, Graph
		Scale-invariant feature transform(SIFT)	Matrix, Transform, Sampling, Sort, Statistics
		Image Transform	Matrix, Transform
		Term Frequency-inverse document frequency (TF-IDF)	Statistics
Sequence Tagging	Bioinformatics	Hidden Markov Model(HMM)	Matrix
	Language Processing	Conditional random fields(CRF)	Matrix, Sampling
Indexing	Search Engine	Inverted index, Forward index	Statistics, 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

Eight Data Motifs

- Matrix
- Sampling
- Transform
- Graph
- Logic
- Set
- Statistics
- Sort

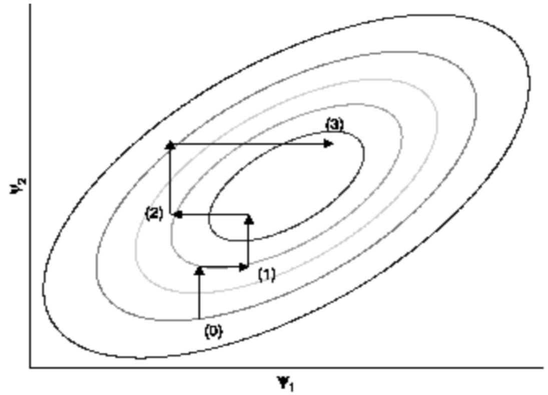

Figure 3.4: Gibbs sampling algorithm in two dimensions starting from an initial point and then completing three iterations

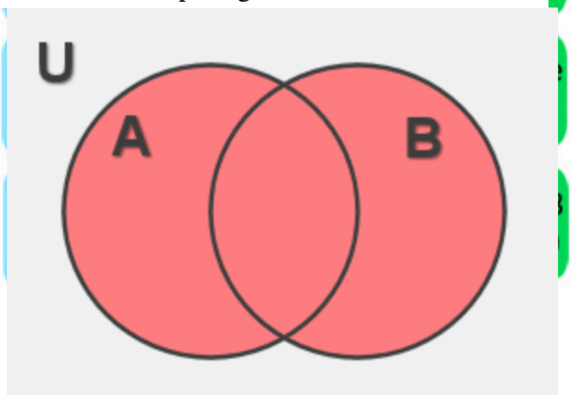

Difference with Kernels

- Data motifs
  - behaviors are affected by the sizes, patterns, types, and sources of different data inputs
  - reflect not only **computation** patterns, **memory access** patterns, but also **disk and network I/O** patterns

Domain-specific Hardware and Software Co-design

- Tailoring the system and architecture to characteri stics of data motifs

  - New architecture/accelerator design
  - Data motif-based libraries
  - Bottleneck identification and optimization

Overview

- Looking back at history
- What is Data Motif
- **BigDataBench 4.0**
  - **Unified Big Data and AI Benchmark Suite**
  - Proxy Benchmarks for Simulation
- Conclusion

Scalable Benchmark Methodology

- Data motif-based Scalable Methodology
  - Micro Benchmark---**Single** data motif
  - Component Benchmark---**Data motif combination with different weights**
  - Application Benchmark---**End-to-end** application model

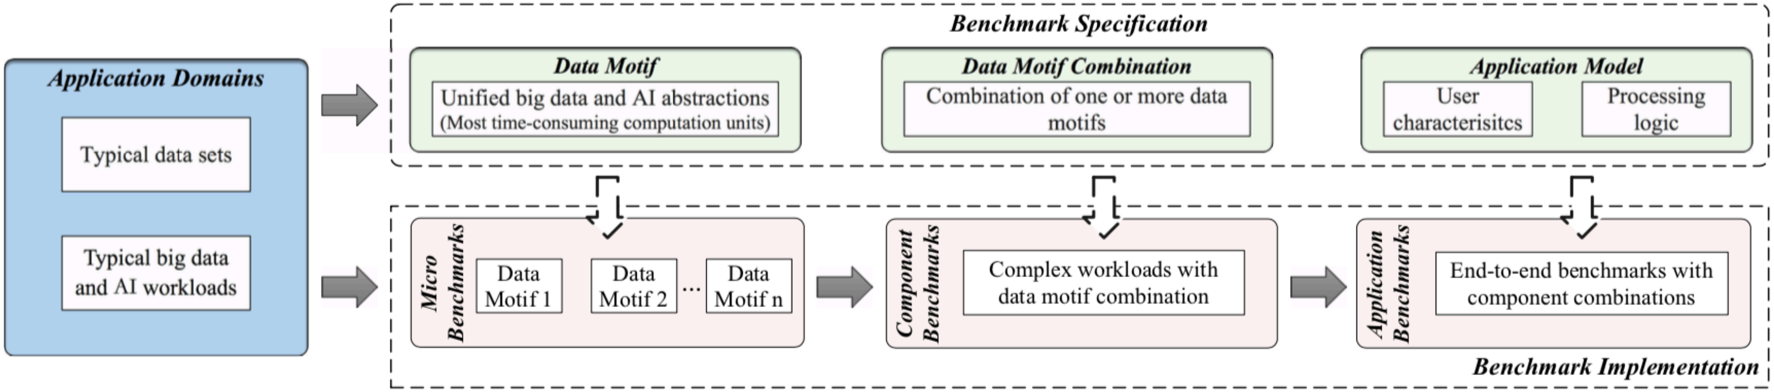

Why Data Motif-based Benchmark

- Using the combination to represent a wide variety of big data and AI workloads
  - No need to create a new benchmark or proxy for every possible workload

Big Data and AI Motif

- Matrix
- Sampling
- Transform
- Graph
- Logic
- Set
- Statistics
- Sort

Combinations with different weights →

Diverse big data and AI workloads

- Coverage of fundamental units of computation
- Provide the methodology of choosing typical workloads
- Reduce workload redundancy

Representative Application Domain

- Search Engine
- Social Network
- Electronic Commerce

Taking up 80% of internet services according to page views and daily

Internet Service

Search engine, Social network, E-commerce

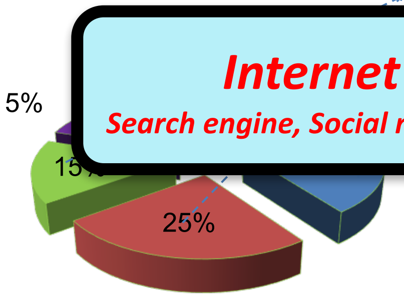

Top 20 websites

http://www.alexa.com/topsites/global;0

Multimedia

http://www.oldcolony.us/wp-content/uploads/2014/11/whatisbigdata-DKB-v2.pdf

DDBJ/EMBL/GenBank database Growth

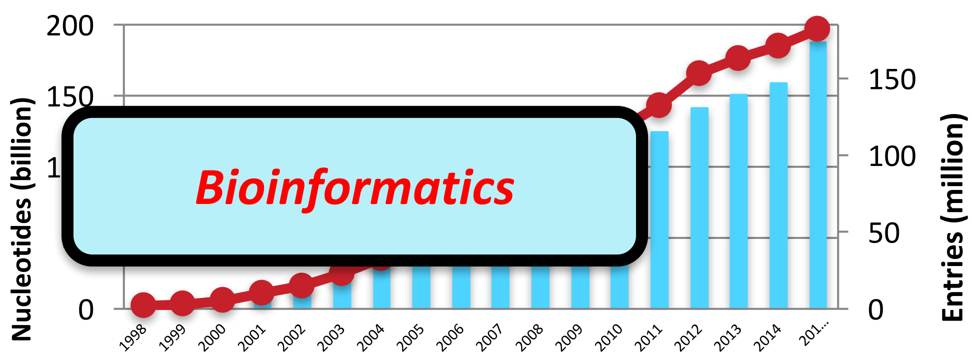

Bioinformatics

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

BigDataBench 4.0

Unified Big Data and AI Benchmark Suite --- http://prof.ict.ac.cn/BigDataBench

- Micro Benchmark
- Component Benchmark
- Application Benchmark

Large scale system-level benchmarks

- 100X Runtime Speedup
- 90%+ Average Accuracy
- Data Adaptability
- Configuration Adaptability
- Cross Architecture

Proxy benchmarks for simulation

- Table
- Text
- Graph
- Matrix
- Image
- Audio
- Structured
- Semi-structured
- Un-structured

Real-world dataset and data generation tools

- AI
- Graph
- Offline Analytics
- Online Service
- Streaming
- Data Warehouse
- NoSQL

47 Workloads covering 7 types

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

16 Software stack

Micro Benchmarks

Micro Benchmark	Involved Motif	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

Component Benchmark	Involved Motif	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

Overview

- Looking back at history
- What is Data Motif
- **BigDataBench 4.0**
  - Unified Big Data and AI Benchmark Suite
  - **Proxy Benchmarks for Simulation**
- Conclusion

Simulation

Problem

- Requirements of Huawei HiSilicon:
- Simulation Time & Accuracy

Technology

- Data Motif-based proxy benchmarks, average system and micro-architectural accuracy >= 90%
- 100X runtime speedup (1000s vs 10s)
- Supporting different architectures (X86 vs ARM)

Wanling Gao, Jianfeng Zhan, Lei Wang, et al. Data Motif-based Proxy Benchmarks for Big Data and AI Workloads. Workload Characterization (IISWC 2018).

Proxy benchmarks

- Data Motif-based proxy benchmark generating methodology
  - A DAG-like combination of data motifs
  - An auto-tuning tool using machine learning model
    - Mimic system and micro-architectural behaviors

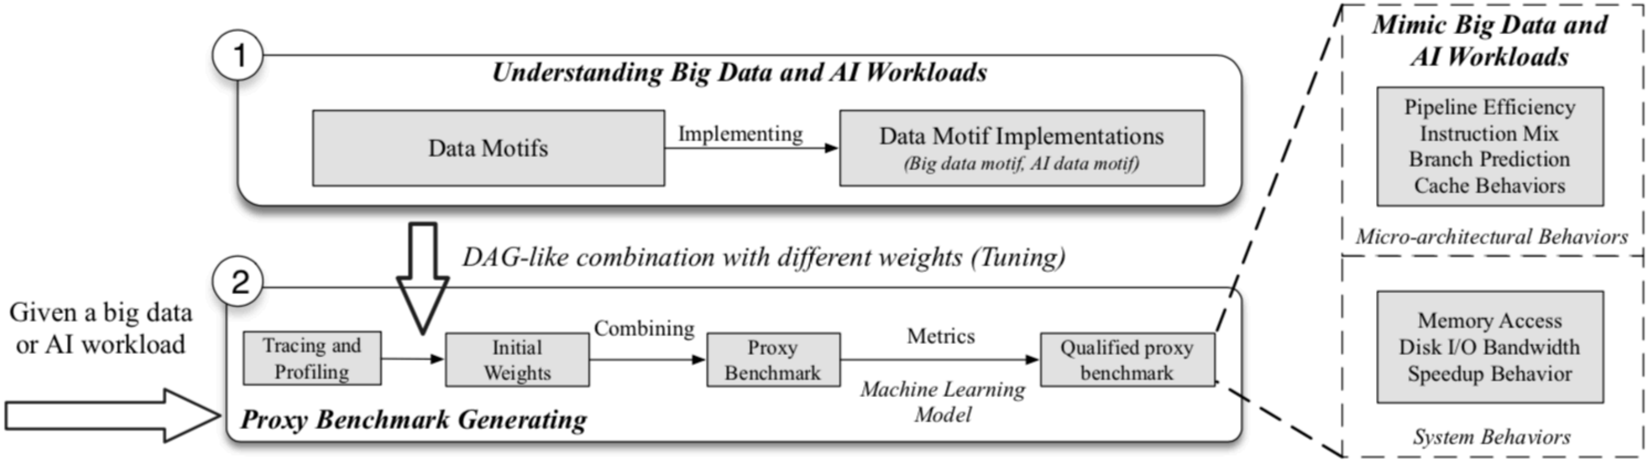

Overview

- Looking back at history
- What is Data Motif
- BigDataBench 4.0
  - Unified Big Data and AI Benchmark Suite
  - Proxy Benchmarks for Simulation
- **Conclusion**

Executive summary

BigDataBench 4.0 --- Unified Big Data and AI Benchmark Suite

Workloads:
- AI
- Graph
- Offline Analytics
- Online Service
- Streaming
- Data Warehouse
- NoSQL

Proxy Benchmarks:
- 100X Runtime Speedup
- 90%+ Average Accuracy

A Data Motif-based Scalable Benchmarking Methodology

Diverse Implementations and Data Motif Combinations

Data Motifs： A Lens Towards Fully Understanding Big Data and AI Workloads

Most Time-consuming Units of Computation:
- Matrix
- Sampling
- Transform
- Logic
- Graph
- Set
- Statistics
- Sort

Data Characteristics:
- Type
- Source
- Pattern
- Scale

Big Data and AI Benchmark

- Collaborate between industry and academia
  - Industry: workload & dataset
  - Academia: benchmark specification
- Various application domains !
  - Internet service
  - Scientific data
  - Medical domain
  - ……

Publications

- Data Motifs: A Lens Towards Fully Understanding Big Data and AI Workloads. **PACT'18.**
- BigDataBench: a Motif-based Big Data and Artificial Intelligence Benchmark Suite. **Technical Report**.
- Understanding Big Data Analytics Workloads on Modern Processors. **TPDS'16**
- Auto-tuning Spark Big Data Workloads on POWER8: Prediction-Based Dynamic SMT . **PACT'16**
- BigDataBench: a Big Data Benchmark Suite from Internet Services. **HPCA'14**
- CVR: Efficient Vectorization of SpMV on X86 Processors. **CGO'18.**
- BOPS, Not FLOPS! A New Metric, Measuring Tool, and Roofline Performance Model For Datacenter Computing. **Technical report**.
- Data Motif-based Proxy Benchmarks for Big Data and AI Workloads. **IISWC 2018.**


QUESTIONS
And
Answers