The BigData Top100 List Initiative

Speakers:
Chaitan Baru, San Diego Supercomputer Center, UC San Diego
Milind Bhandarkar, Greenplum/EMC

Outline

- Background
- Benchmark Context and Technical Issues
- Next Steps and Moving Forward
- Specifying a Data Analytics Pipeline Workload

What's new

- Article in inaugural issue of Big Data Journal
  - *Big Data Benchmarking and the Big Data Top100 List* by Baru, Bhandarkar, Nambiar, Poess, Rabl, Big Data Journal, Vol.1, No.1, 60-64, Anne Liebert Publications.

What's New - 2

- BigDataTop100.org Website

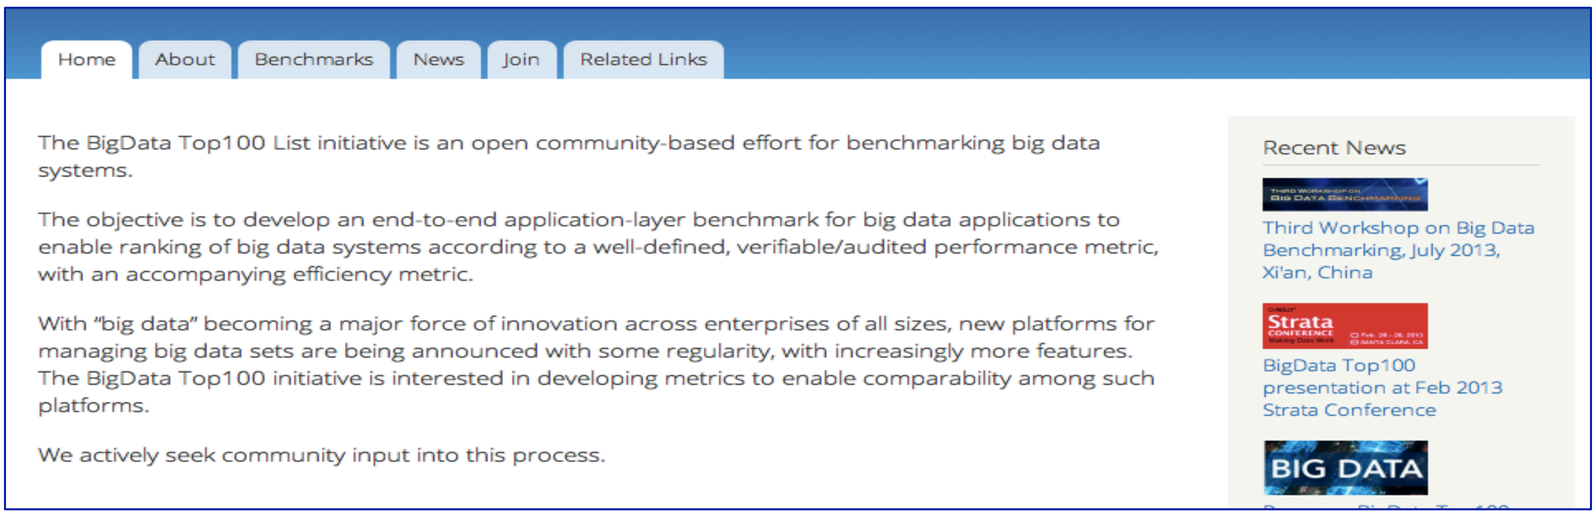

Home | About | Benchmarks | News | Join | Related Links

The BigData Top100 List initiative is an open community-based effort for benchmarking big data systems.

The objective is to develop an end-to-end application-layer benchmark for big data applications to enable ranking of big data systems according to a well-defined, verifiable/audited performance metric, with an accompanying efficiency metric.

With "big data" becoming a major force of innovation across enterprises of all sizes, new platforms for managing big data sets are being announced with some regularity, with increasingly more features. The BigData Top100 initiative is interested in developing metrics to enable comparability among such platforms.

We actively seek community input into this process.

Recent News

Third Workshop on Big Data Benchmarking, July 2013, Xi'an, China

BigData Top100 presentation at Feb 2013 Strata Conference

Background

- Outcome from an NSF research project
  - *Performance Evaluation of On-demand Provisioning of Data-intensive Applications*, 2009-2012, NSF IIS-0844530
  - Blending use of DBMS and Hadoop
  - Evaluation of shared-nothing DBMS vs Hadoop for geospatial data
- Identified need for objective benchmarks for big data applications
  - Reached out to the TPC world

Context

- Launched the Workshop series on Big Data Benchmarking (WBDB)
  - First workshop, May 2012, San Jose. Hosted by Brocade.
  - Second workshop, December 2012, Pune, India. Hosted by Persistent Systems / Infosys.
  - Third workshop, July 2013, Xi'an, China. Hosted by Xi'an University.
- First paper:
  - Setting the Direction for Big Data Benchmark Standards by C. Baru, M. Bhandarkar, R. Nambiar, M. Poess, and T. Rabl, published in *Selected Topics in Performance Evaluation and Benchmarking*, Springer-Verlag
- Selected WBDB papers in Springer Verlag Lecture Notes in Computer Science

1st and 2nd Workshop Attendees

- Actian
- AMD
- Anna University
- BMMsoft
- Brocade
- CA Labs
- Cisco
- Cloudera
- Convey Computer
- CWI/Monet
- DBSync LLC
- Dell
- EPFL
- Facebook
- Google
- Greenplum
- Hewlett-Packard
- Hortonworks
- Indiana Univ / Hathitrust Research Foundation
- IIT Chicago
- InfoSizing
- Informatica
- Infosys
- Intel
- Jacobs University
- LinkedIn
- MapR/Mahout
- Mellanox
- Microsoft
- NSF
- NetApp
- NetApp/OpenSFS
- Oracle
- Persistent
- Red Hat
- San Diego Supercomputer Center
- SAS
- Scripps Research Institute
- Seagate
- Shell
- SNIA
- SoftCorner
- Teradata Corporation
- Twitter
- UC Irvine
- UC San Diego
- Univ. of Minnesota
- University of Passau
- Univ. of Toronto
- Univ. of Washington
- VMware
- WhamCloud
- Yahoo!

Technical Issues - 1

- Audience: Who is the audience for such a benchmark?
  - Marketing, Internal Use, Academic Use
- Application: What is the application that should be modeled?
  - Abstractions of a data pipeline, e.g. Internet-scale business
- Single benchmark spec: Is it possible to develop a single benchmark to capture characteristics of multiple applications?
  - Single, multi-step benchmark, with plausible end-to-end scenario

Technical Issues - 2

- Component vs. end-to-end benchmark. Is it possible to factor out a set of benchmark "components", which can be isolated and plugged into an end-to-end benchmark(s)?
  - The benchmark should consist of individual components that ultimately make up an end-to-end benchmark
- Paper and Pencil vs Implementation-based. Should the implementation be specification-driven or implementation-driven?
  - Start with an implementation and develop specification at the same time

Technical Issues - 3

- Reuse. Can we reuse existing benchmarks?
  - Leverage existing work and built-up knowledgebase
- Benchmark Data. Where do we get the data from?
  - Synthetic data generation: structured, semistructured, unstructured data
- Innovation or competition? Should the benchmark be for innovation or competition?
  - Successful competitive benchmarks will be used for innovation

Keep in mind principles for good benchmark design

- Self-scaling, e.g. TPC-C
- Comparability between scale factors
  - Results should be comparable at different scales
- Technology agnostic (if meaningful to the application)
- Simple to run

More considerations

- Extrapolating Results
  - To larger configurations
- Elasticity and durability
  - Big data systems should be intrinsically elastic and durable
  - TPC runs ACID outside the performance window
- Performance + Price/performance
  - Try to capture price in a simple, intuitive, meaningful way…
  - For **price/performance:** what is the most useful quantity for price?

Next Steps…

- Don't let the perfect be the enemy of the good…
- "A good plan executed … now, rather than a perfect plan next week"
- BigData Top100 Board of Directors
  - Chaitan Baru, San Diego Supercomputer Center
  - Milind Bhandarkar, EMC/Greenplum
  - Dhruba Borthakur, Facebook
  - Eyal Gutkind, Mellanox
  - Jian Li, IBM
  - Raghunath Nambiar, Cisco
  - Ken Osterberg, Seagate
  - Scott Pearson, Brocade
  - Meikel Poess, Oracle
  - Tilmann Rabl, University of Toronto
  - Richard Treadway, NetApp
  - Jerry Zhao, Google

Moving Forward…

- BigData Top100 will select a benchmark specification for ranking systems Two alternatives for workload specification
  - **BigBench**: based on TPC-DS
    - Extended with semistructured and unstructured data and operations on those data. See 1st WBDB and 2nd WBDB program online for slides.
  - **Data Analytics Pipeline (DAP)**
    - Proposed end-to-end pipeline, from data ingestion to predictive modeling
  - **Track bigdatatop100.org/benchmarks for info**
- Use Kaggle for competitions
  - Propose data generation programs
  - Propose / evaluate operations for steps in the Data Analytics Pipeline

Deadlines…

- March 15, 2013
  - BigBench turned into first draft of a benchmark spec
  - First draft of data and operations for the Data Analytics Pipeline
- April 15, 2013
  - Receive comments from the community
- May 1, 2013:
  - Short paper submission deadline for 3rd WBDB, July 16-17, Xi'an, China.
  - By 3rd WBDB, July 2013: Proposals for metrics, execution rules, audit rules and reporting rules.
- July 16, 2013
  - Reference implementations of BigBench and Data Analytics Pipeline
- August 31, 2013: Release of a benchmark specification

Data Analytics Pipeline

Defining A Typical Big Data Workload

Quest for Typical Workload

- Tune systems for broadly applicable workloads
- Benchmarks most relevant if representative
- Designing optimized systems: Make common tasks fast, other tasks possible

Is There a Typical Big Data Workload?

- Big Data systems are characterized by flexibility
- Multiple Interfaces : SQL, MapReduce, Streaming, GraphLab,…
- Workloads evolving rapidly

Encouraging Early Results

- Analyzed characteristics of 1M+ real Hadoop jobs on production clusters at Yahoo, 100+ features
- Identified 8 Job types
- Verified with GridMix 3
- Characterization of Hadoop Jobs Using Unsupervised Learning, Sonali Aggarwal, Shashank Phadke & **Milind Bhandarkar**, in 2010 IEEE Second International Conference on Cloud Computing Technology and Science, Indianapolis, Indiana, December 2010, http://doi.ieeecomputersociety.org/10.1109/CloudCom.2010.20

Big Data Sources

- Events
  - Direct - Human Initiated
  - Indirect - Machine Initiated
- Software Sensors (Clickstreams, Locations)
- Public Content (blogs, tweets, Status updates, images, videos)

User Modeling

- Objective: Determine User-Interests by mining user-activities
- Large dimensionality of possible user activities
- Typical user has sparse activity vector
- Event attributes change over time

User Modeling Pipeline

- Data Acquisition
- Sessionization
- Feature and Target Generation
- Model Training
- Offline Scoring & Evaluation
- Batch Scoring & Upload to serving

Data Acquisition

- Batched and collected at the edge
- Loaded incrementally
- Simple ETL/ELT
- Append / Replace

Denormalization / Cleansing

- Augment raw events data with attributes
- Look up dictionaries
- Multi-way Joins with dimension tables

Find Sessions with Target Events

- User-Actions of Interest
  - Clicks on Ads & Content
  - Site & Page visits
  - Conversion Events
    - Purchases, Quote requests
    - Sign-Up for membership etc

Feature Selection

- Summary of user activities over a time-window
- Aggregates, moving averages, rates over various time-windows
- Incrementally updated

Join Targets and Features

- Target rates very low: 0.01% ~ 1%
- First, construct targets
- Filter user activity without targets
- Join feature vector with targets

Model Training

- Regressions
- Boosted Decision Trees
- Naive Bayes
- Support Vector Machines
- Maximum Entropy modeling

Offline Scoring & Evaluations

- Apply model weights to features
- Pleasantly parallel
- Sort by scores and compute metrics
- Evaluate metrics

Batch Scoring

- Apply models to features from all user activity
- Upload scores to serving systems

5 Different Classes

- Tiny (10K Entities, 100GB)
- Small (100K Entities, 1TB)
- Medium (1M Entities, 10 TB)
- Large (10M Entities, 100 TB)
- Huge (1B Entities, 1PB)

Deadlines…

- March 15, 2013
  - BigBench turned into first draft of a benchmark spec
  - First draft of data and operations for the Data Analytics Pipeline
- April 15, 2013
  - Receive comments from the community
- May 1, 2013:
  - Short paper submission deadline for 3rd WBDB, July 16-17, Xi'an, China.
  - By 3rd WBDB, July 2013: Proposals for metrics, execution rules, audit rules and reporting rules.
- July 16, 2013
  - Reference implementations of BigBench and Data Analytics Pipeline
- August 31, 2013: Release of a benchmark specification