

Second Workshop on Big Data Benchmarking: WBDB2012.in

WELCOME !



Acknowledgements

- Local organization:
 - Chandrashekhar Sahasrabudhe, Persistent Systems
 - Arun Kadekode, Soft-Corner
 - Infosys
 - Computer Society of India Div II, SIGDATA/COMAD
- Other WBDB industry sponsors
 - Seagate, Greenplum, NetApp, Brocade, Cisco, Mellanox
- US National Science Foundation, NSF

Workshop Objectives

- Further the objective of defining Big Data Benchmarks
- Build upon efforts so far, since early 2012
 - First WBDB workshop
 - Meetings of the CLDS center, San Diego Supercomputer Center
- Make progress in defining rules and parameters for the BigData100 List
- Set the parameters for the next WBDB in China in July

Meeting Structure

- Interactive meeting
 - Don't get distracted with email/chat/tweet
- Invited talks plus Submitted papers
- Parallel discussion sessions
- Reconvene in plenary with a report back
- Flexible structure...

Agenda Day 1: Monday 17th

0830 0900 Breakfast

- 0900 0930 Opening Remarks, Chaitan Baru, SDSC
- 0930 1015 Sponsor talk - Persistent Systems, Mukund Deshpande, VP, Persistent Systems, Lead, Big Data Group
- 1015 1030 Participant introductions
- 1030 1100 Lessons from Industry Standard Benchmarking, Raghunath Nambiar, Cisco

1100 1130 Coffee Break

- 1130 1200 Introduction to PDGF, Tilmann Rabl, U of Toronto
- 1200 1220 BigBench, Ahmad Ghazal, Teradata
- 1220 1240 Making Sense of System Performance at Scale, Vinayak Borkar, UC Irvine
- 1240 1300 Stratosphere, Kostas Tzoumas

1300 1430 Lunch

- 1430 1450 Presentation of the discussion topic: BigData100, Raghunath Nambiar
- 1450 1630 Discussion

1630 1700 Tea / Coffee Break

- 1700 1730 Discussion and Summarization

1730 1800 Dinner at Courtyard Marriott

Agenda Day 2: Tuesday 18th

0830 0900 Breakfast

0900 0910 Recap of Day 1, Chaitan Baru, SDSC

0910 0930 UIDAI, Rajendra Kumar, UIDAI

0930 0950 Enterprise Use Cases for Big Data Platforms, Susheel Kaushik, Greenplum

0950 1010 Benchmarking MapReduce in a High Performance Computing Network Environment, Lessons Learned and Results Found, Sreevathsa Doddabalapur, Mellanox

1010 1030 Brocade presentation, Edgar Dias, Brocade

1030 1050 Optimizing Hadoop Deployment on Gordon Data Intensive Supercomputer
Amit Majumdar, SDSC

1050 1110 A Micro-benchmark Suite for Evaluating HDFS Operations on Modern Clusters, Nusrat Islam, Ohio State University

1110 1140 Coffee Break

Agenda: Day 2, Tuesday 18th

1110 1140 Coffee Break

1140 1200 The implications of diverse data sets of different scales for big data benchmark, Jia, Zhen

1200 1220 Benchmarking Large Arrays in Databases, Heinrich Stamerjohanns

1220 1240 Big Data Provenance: Challenges and Potential Implications for Benchmarking, Boris Glavic

1240 1300 Unleashing Semantics of Research Data, Florian Stegmaier

1300 1400 Lunch

1400 1420 NetApp Presentations, TBN

1420 1435 Late Binding, Stephen Brobst

1435 1445 Big Data, T. V. Gopal

1445 1500 Presentation of the discussion topic: 1000-node Big Data Challenge

1500 1600 Discussion

1600 1630 Tea / Coffee Break

1630 1700 Discussion and Summarization

1700 1730 Workshop Conclusion and Next Steps

Context and Background

Chaitan Baru

Director, Center for Large-Scale Data Systems Research
(CLDS)

Associate Director Data Initiatives, San Diego
Supercomputer Center

University of California San Diego

First Workshop on Big Data Benchmarking, May 2012, San Jose

Invited Attendee Organizations

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

Topics discussed

- Audience: Who is the audience for such a benchmark?
- Application: What is the application that should be modeled?
- Single benchmark spec: Is it possible to develop a single benchmark to capture characteristics of multiple applications?
- 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)?
- Paper and Pencil vs Implementation-based. Should the implementation be specification-driven or implementation-driven?
- Reuse. Can we reuse existing benchmarks?
- Benchmark Data. Where do we get the data from?
- Innovation or competition? Should the benchmark be for innovation or competition?

Audience: Who is the primary audience for a big data benchmark?

- Customers
 - ➔ Workload should preferably be expressed in
 - English
 - Or, a declarative Language (unsophisticated user)
 - But, not a procedural language (sophisticated user)
 - Want to compare among different vendors
- Vendors
 - Would like to sell machines/systems based on benchmarks
- Computer science/hardware research is also an audience
 - Niche players and technologies will emerge out of academia
 - Will be useful to train students on specific benchmarking

Applications: What application should we model?

- Possibilities
 - An application that somebody could donate
 - An application based on empirical data
 - Examples from scientific applications
 - Multi-channel retailer-based application, like the amended TPC-DS for Big Data?
 - Mature schema, large scale data generator, execution rules, audit process exists.
 - “Abstraction” of an Internet-scale application, e.g. data management at the Facebook site, with synthetic data

Single Benchmark vs Multiple

- Is it possible to develop a single benchmark to represent multiple applications?
- Yes, but not desired if there is no synergy between the benchmarks, e.g. say, at the data model level
 - Synthetic Facebook application might provide context for a single benchmark
 - Click streams, data sorting/indexing, weblog processing, graph traversals, image/video data, ...

Component benchmark vs. end-to-end benchmark

- Are there components that can be isolated and plugged into an end-to-end benchmark?
- The benchmark should consist of individual components that ultimately make up an end-to-end benchmark
- The benchmark should include a component that extracts large data
 - Many data science applications extract large data and then visualize the output
 - Opportunity for “pushing down” viz into the data management system

Paper and Pencil / Specification driven versus Implementation driven

- Start with an implementation and develop specification at the same time
- Some post-workshop activity has begun in this area
 - Data generation; sorting; some processing

Where Do we Get the Data From?

- Downloading data is not an option
- Data needs to be generated (quickly)
- Examples of actual datasets from scientific applications
 - Observational data (e.g. LSST), simulation outputs
- Using existing data generators (TPC-DS, TPC-H)
- Data that is generic enough with good characteristics is better than specific data

Should the benchmark be for innovation or competition?

- Innovation and competition are not mutually exclusive
 - Should be used for both
 - The benchmark should be designed for competition, such a benchmark will then also be used internally for innovation
- TPC-H is a prime example of a benchmark model that could drive competition and innovation (if combined correctly)

Can we reuse existing benchmarks?

- Yes, we could but we need to discuss:
 - How much augmentation is necessary?
 - Can the benchmark data be scaled
 - If the benchmark uses SQL, we should not require it
- Examples: but none of the following could be used unmodified

- Statistical Workload Injector for Map Reduce (SWIM)
- GridMix3 (lots of shortcomings)
 - Open source
- TPC-DS
- YCSB++ (lots of shortcomings)
- Terasort – strong sentiment for using this, perhaps as part of an “end-to-end” scenario

Big Data Benchmark Data Model

- TPC-DS
 - Decision support benchmark from Transaction Processing Performance Council
 - http://www.tpc.org/tpcds/spec/tpcds_1.1.0.pdf
- Why build on top of TPC-DS?
- Volume :
 - No theoretical limit
 - Tested up to 100 TB
- Velocity : rolling updates
- Variety
 - Rich relational model and data, 26 tables, 7 fact tables
 - Easy to add the other two sources

4

TERADATA | aster data

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

TPC

- + Longevity: TPC-C has carried the load for 20 years
- + Comparability
 - Audit requirements and strict detailed run rules mean one can compare results published by two different entities
- + Scaling
 - Results just as meaningful at the high-end of the market as at the low-end; as relevant on clusters as on single servers
- Hard and expensive to run
- No kit
- DeWitt clauses

Reza Taheri, VMWare

Other considerations

- **Extrapolating Results**
 - TPC benchmarks typically run on “over-specified” systems
 - i.e. Customer installations may have less hardware than benchmark installation (SUT)
 - **Big Data Benchmarking may be opposite**
 - May need to run benchmark on systems that are smaller than customer installations
- **Elasticity and durability**
 - TPC runs ACID outside the performance window
 - Big data systems may need to be intrinsically elastic and able to cope with failures

Price / cost

- For a **price/performance** metric, what is the most useful quantity for price?
- How can we capture price in a simple, intuitive, meaningful way...?

Outcomes from first workshop

- Formation of a Big Data Benchmarking Community (BDBC)
 - Biweekly phone conferences
 - Contact baru@sdsc.edu if you are interested in being on the mailing list for this
- Paper in TPCTC
 - “Setting the direction for big data benchmark standards,” in TPCTC 2012, VLDB2012, Aug 27-31, Istanbul, Turkey

Next Steps

- Defining the BigData100 list
 - Session on *Big Data Benchmarking and the BigData100 List* at Strata Conference, February 2013, Santa Clara
- Third WBDB in Xi'an, China, July 9-10, 2013
 - Local organization: Shanxi Supercomputing Center and IBM China

Connecting on social media

- LinkedIn: www.linkedin.com/groups/CLDS_BigData
- Twitter – @CLDS_BigData

Other Next Steps

- Short/Medium-term
 - Generate 100TB TPC-DS + 10x (1PB) semistructured and/or structured data?
 - To understand the data generation issues
 - Obtain genomic, spatial(?), graph data and run simple operations?
 - To assess what it takes, and do it in a BDB context
- Medium-term (with short-term deadline)
 - Proposal to NSF on alpha, beta phases of the big data benchmarking effort
 - Establish benchmarking resources
 - Hardware testbeds at SDSC; large cluster at Greenplum; Google shared cycles