Crafting Benchmarks for Big Data: The Big Data Top100 List

Chaitan Baru
Director, Center for Large-Scale Data Systems Research
San Diego Supercomputer Center
UC San Diego

Outline

- Welcome
- Center for Large-Scale Data Systems, CLDS
- Workshops on Big Data Benchmarking, WBDB
- Application-level benchmarks
  - TPC-like
  - Deep Analytics Pipeline
- Collaboration Opportunities

Welcome

Center for Large-Scale Data Systems Research, CLDS

- Center focused on technology and technology management research related to big data
- Current focus areas
  - Big Data Benchmarking
  - Data Value
- Center sponsors
  - Seagate, Pivotal, EMC, NetApp, Intel, GraphSQL
  - Workshop sponsors: Mellanox, Cisco

Big Data Benchmarking

- Genesis of the Big Data Benchmarking effort
  - Grant from NSF under the Cluster Exploratory (CluE) program
  - "Performance Evaluation of On-Demand Provisioning of Data Intensive Applications" (2009-2012)
    - Study of parallel relational database versus Hadoop systems
    - Premise: store all data in Hadoop, and "hot" data in RDBMS
- Identified need for objective benchmarks
- Launched Workshops on Big Data Benchmarking
  - First WBDB: May 2012, San Jose. Hosted by Brocade

1st WBDB: 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!

2nd WBDB: http://clds.sdsc.edu/wbdb2012.in

3rd WBDB: http://clds.sdsc.edu/wbdb2013.cn

4th WBDB: http://clds.sdsc.edu/wbdb2013.us

5th WBDB: http://clds.sdsc.edu/wbdb2014.de

WBDB Outcomes

- Big Data Benchmarking Community (BDBC) mailing list (~200 members from ~80 organizations)
  - Organized webinars every other Thursday
  - http://clds.sdsc.edu/bdbc/community
- Paper from First WBDB
  - *Setting the Direction for Big Data Benchmark Standards* C. Baru, M. Bhandarkar, R. Nambiar, M. Poess, and T. Rabl, published in *Selected Topics in Performance Evaluation and Benchmarking*, Springer-Verlag

WBDB Outcomes…

- Selected papers in Springer Verlag, *Lecture Notes in Computer Science*, Springer Verlag
  - Papers from 1st and 2nd WBDB published in Specifying Big Data Benchmarks, ISBN 978-3-642-53973-2, Editors: Rabl, Poess, Baru, Jacobsen
  - Papers from 3rd and 4th WBDB workshops in Vol II
  - Papers from 5th WBDB will be in Vol III
- Formation of TPC Subcommittee on Big Data Benchmarking
  - Working on TPCx-HS: TPC Express benchmark for Hadoop Systems, based on Terasort
  - http://www.tpc.org/tpcbd/
- Currently considering initiating a SPEC Research Group on Big Data Benchmarking

Types of Benchmarks

- Micro-benchmarks. To evaluate specific lower-level, system operations
  - E.g., A Micro-benchmark Suite for Evaluating HDFS Operations on Modern Clusters, Panda et al, OSU
- Functional benchmarks. Specific high-level function.
  - E.g. Sorting: Terasort
  - E.g. Basic SQL: Individual SQL operations, e.g. Select, Project, Join, Order-By, …
- Genre-specific benchmarks. Benchmarks related to type of data
  - E.g. Graph500. Breadth-first graph traversals
- Application-level benchmarks.
  - Measure system performance (hardware and software) for a given application scenario—with given data and workload

TPC Big Data Subcommittee

- TPCx-HS
  - TPC Express for Hadoop Systems
- Based on Terasort
  - Teragen, Terasort, Teravalidate
- Database size / Scale Factors
  - Should they be system-wide or per node?
  - i.e. 1TB run on 10 nodes, 100 nodes, 1000 nodes?
  - Or, 1TB/node on 10 nodes, 100 nodes, ...
  - Nodes can be physical or virtual
- Performance, Price/Performance
  - Response time
  - System acquisition cost; TCO, energy cost; etc

Application-Level Benchmark Design Issues from WBDB

- Audience: Who is the audience for the benchmark?
  - Marketing (Customers / End users)
  - Internal Use (Engineering)
  - Academic Use (Research and Development)
- Is the benchmark for innovation or competition?
  - If a competitive benchmark is successful, it will be used for innovation
- Application: What type of application should be modeled?
  - TPC: schema + transaction/query workload
  - BigData: Abstractions of a data processing pipeline, e.g. Internet-scale businesses

Design 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?
  - The benchmark should consist of individual components that ultimately make up an end-to-end benchmark
- Single benchmark specification: Is it possible to specify a single benchmark that captures characteristics of multiple applications?
  - Maybe: Create a single, multi-step benchmark, with plausible end-to-end scenario

Design Issues - 3

- Paper & Pencil vs Implementation-based. Should the implementation be specification-driven or implementation-driven?
  - Start with an implementation and develop specification at the same time
- 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, non-structured data
- Verifiability. Should there be a process for verification of results?
  - YES!

Big Data Top100 List

- Modeled after Top500 and Graph500 in HPC community
- Proposal presented at Strata Conference, February 2013 (Baru & Bhandarkar)
- Based on application-level benchmarking
- Article in inaugural issue of the 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.

Abstracting the Big Data World

1. Enterprise Data Warehouse [illegible]her types of data
   - Structured enterprise d[illegible]e
   - Extend to incorp[illegible]ctured data, e.g. from weblogs, machin[illegible]ckstream, customer reviews
   - "Design time" sch[illegible]mas
2. Collection of heterogeneous data + Bi[illegible] processing
   - Enterprise data processin[illegible]rom data ingestion to transformation, [illegible]etting, machine learning, predictive an[illegible]
   - Data fr[illegible]ructured and non-structured sources
   - "R[illegible]emas – late binding, application-driven sch[illegible]

BigBench

Deep Analytics Pipeline (DAP)

The BigBench Proposal

- Ghazal, Hu, Rabl, Crolotte, Raab, Poess, Jacobsen, (Teradata, U.of Toronto, InfoSizing, Oracle)
  - Presentation at WBDB 1 and 2; SIGMOD 2013; WBDB Proceedings, Springer LNCS
- Derived from TPC-Decision Support (TPC-DS)
  - Multiple snowflake schemas with shared dimensions
  - 24 tables with an average of 18 columns
  - 99 distinct SQL 99 queries with random substitutions
  - More representative skewed database content
  - Sub-linear scaling of non-fact tables
  - Ad-hoc, reporting, iterative and extraction queries
  - ETL-like data maintenance

From: Ghazal et al, 2nd WBDB

Data Model

7/19/2013

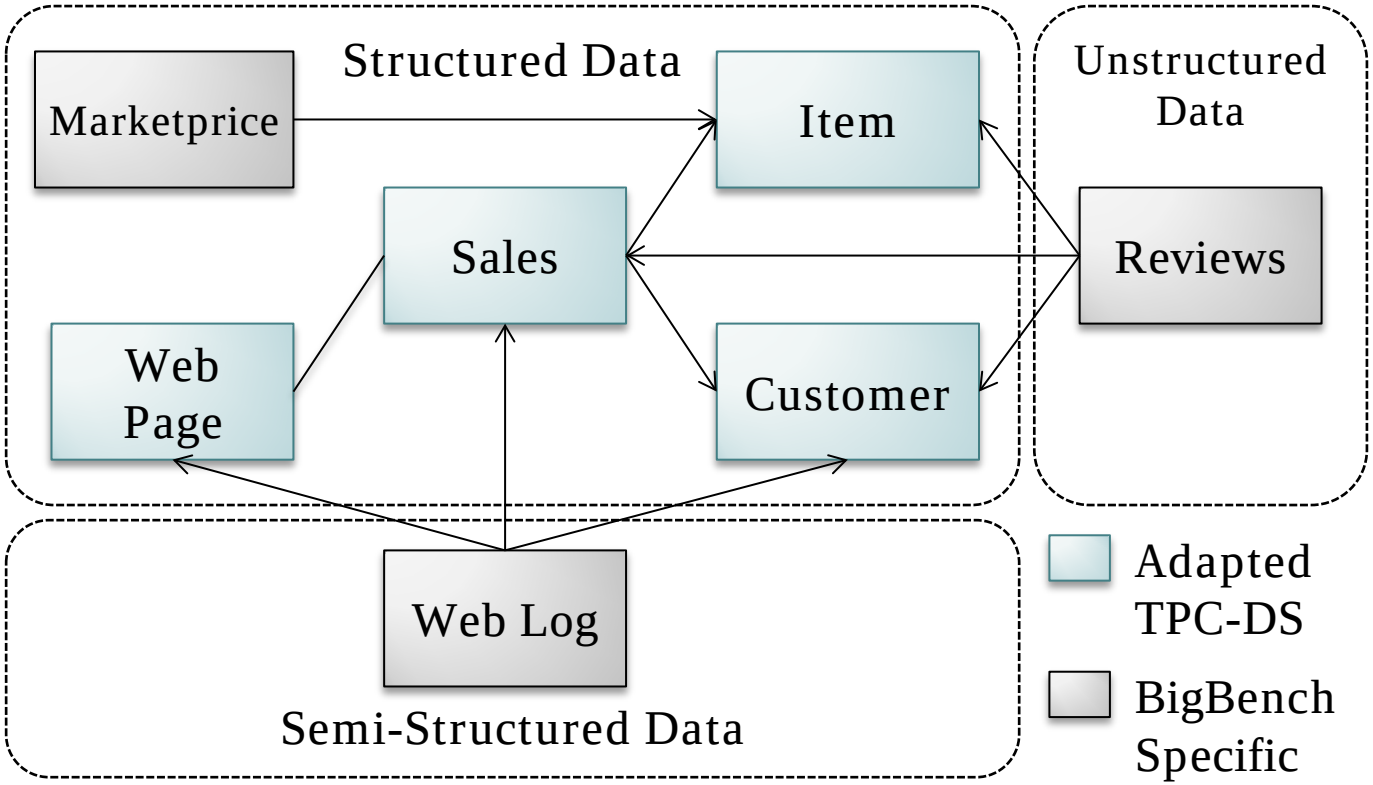

- Structured: TPC-DS + market prices
- Semi-structured: website click-stream
- Unstructured: customers' reviews

From: Ghazal et al, 2nd WBDB

Data Model – 3 Vs

- **Variety**
  - Different schema parts
- **Volume**
  - Based on scale factor
  - Similar to TPC-DS scaling
  - Weblogs & product reviews also scaled
- **Velocity**
  - Periodic refreshes for all data
  - Different velocity for different areas
    - $V_{structured} < V_{unstructured} < V_{semistructured}$
  - Queries run with refresh

From: Ghazal et al, 2nd WBDB

SQL-MR Query 1

```
SELECT category_cd1 AS category1_cd,
    category_cd2 AS category2_cd , COUNT (*) AS cnt
FROM basket_generator (
    ON
        ( SELECT i.i_category_id AS category_cd ,
            s.ws_bill_customer_sk AS customer_id
        FROM web_sales s INNER JOIN item i
        ON s.ws_item_sk = i.item_sk )
    PARTITION BY customer_id
    BASKET_ITEM ('category_cd')
    ITEM_SET_MAX (500)
)
GROUP BY 1,2
ORDER BY 1, 3, 2;
```

Other Benchmarks discussed at WBDB

- Big Decision, Jimmy Zhao, HP
- HiBench, Yan Li, Intel
- BigDataBench, Jianfeng Zhan, Chinese Academy of Sciences
- …

Big Decision, Jimmy Zhao, HP, 4th WBDB

- **Benchmark for A DSS/Data Mining solutions**
  - Everything running in the same system
  - Engine of Analytics
  - Reflecting the real business model
- **Huge data volume**
  - Data from Social
  - Data from Web log
  - Data from Comments
- **Broader Data support**
  - Semi-structured data
  - Un-structured data
- **Continuous Data Integration**
  - ETL just a normal job of the system
  - Data Integration whenever there's data
- **Big Data Analytics**

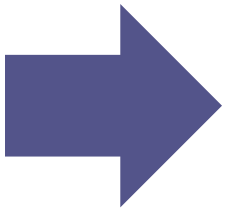

Big Decision – Big TPC-DS!

TPC-DS

- Mature and proved workload for BI
- Mix workloads
- Well defined scale factors

SoMoized TPC-DS

- Additional data and dimension from new data
- Semi-structured and unstructured data
- TB to PB or event Zeta Byte support

NEW TPC-DS generator – Agile ETL

- Continuously data generation and injection
- Consider as part of the workloads

New massive parallel processing technologies

- Convert queries to SQL liked queries
- Include interactive & regular Queries
- Include Machine Learning jobs

Big Decision Block Diagram

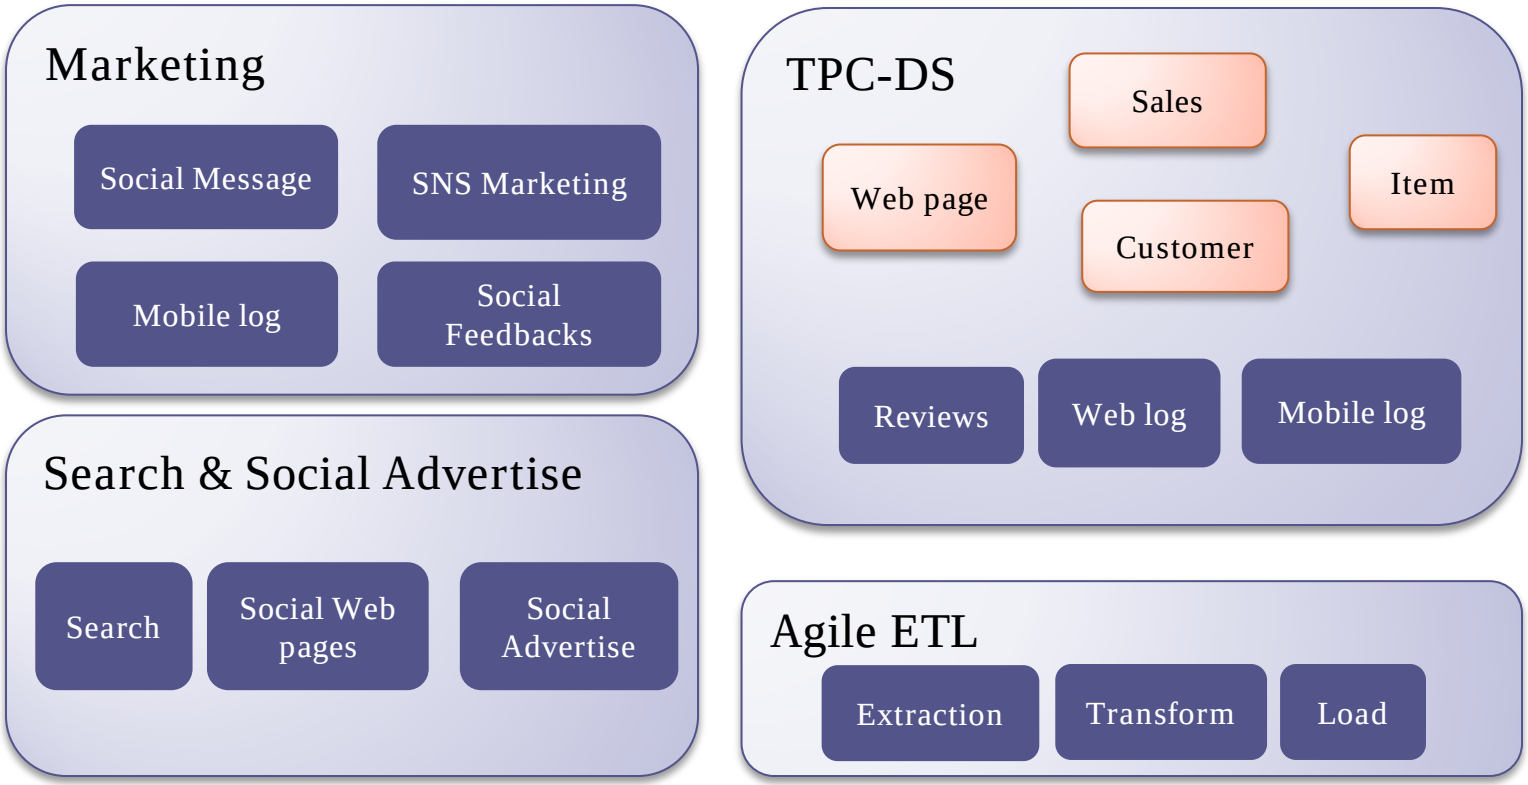

HiBench, Lan Yi, Intel, 4th WBDB

1. Different from GridMix, SWIM?
2. Micro Benchmark?
3. Isolated components?
4. End-2-end Benchmark?
5. We need ETL-Recommendation Pipeline

See our paper "The HiBench Suite: Characterization of the MapReduce-Based Data Analysis" in ICDE'10 workshops (WISS'10)

ETL-Recommendation (hammer)

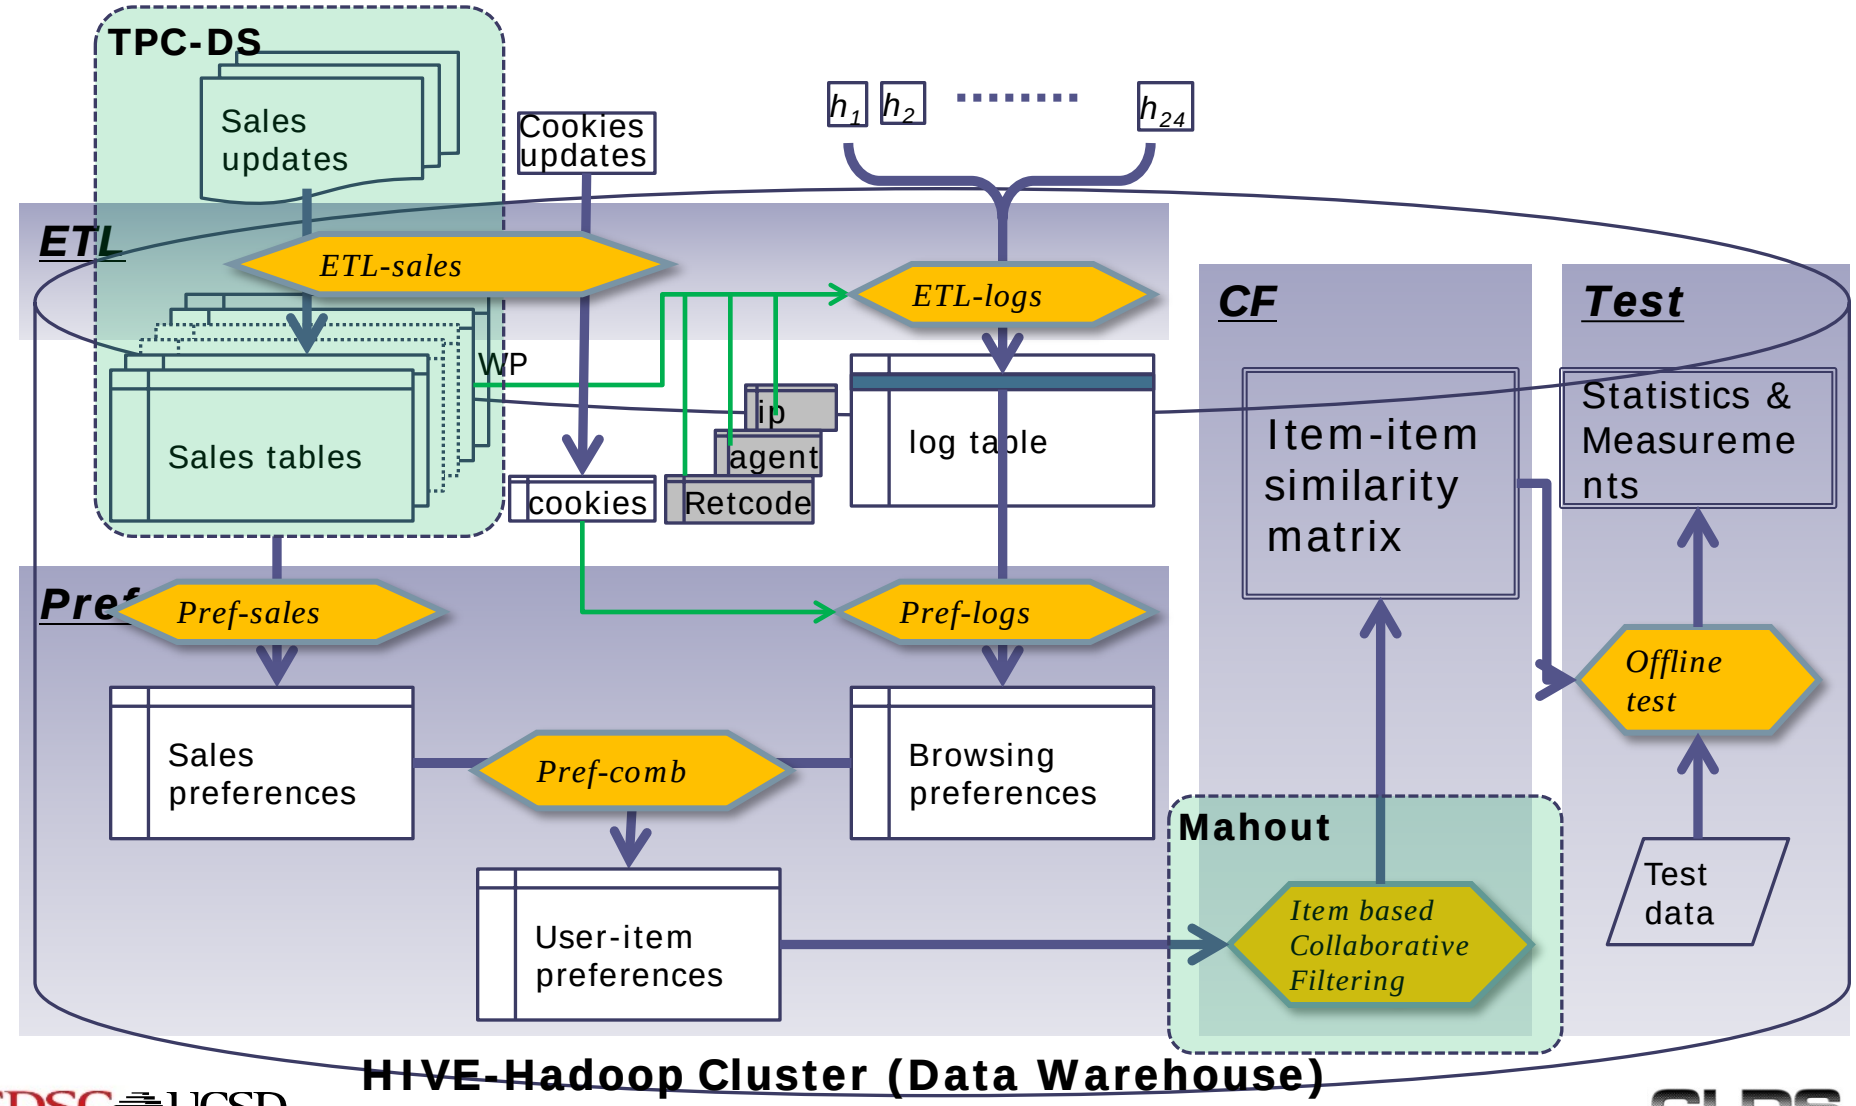

ETL-Recommendation (hammer)

- Task Dependences

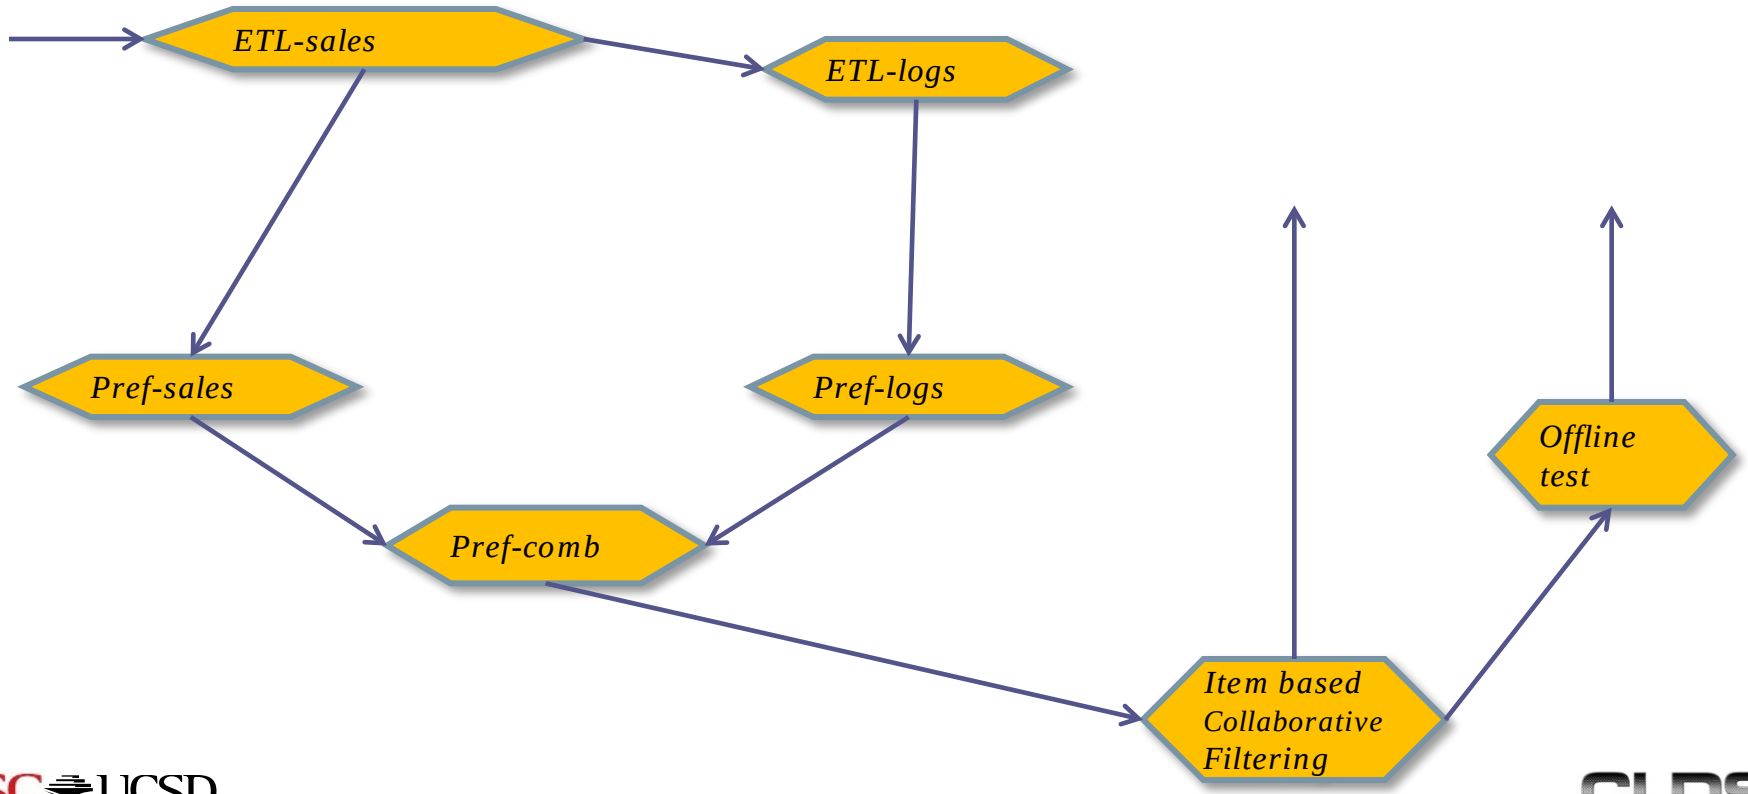

Courtesy: Milind Bhandarkar, Pivotal

The Deep Analytics Pipeline Proposal

- "User Modeling" pipelines
- Generic use case: Determine user interests or user categories by mining user activities
  - Large dimensionality of possible user activities
  - Typical user represents a sparse activity vector
  - Event attributes change over time

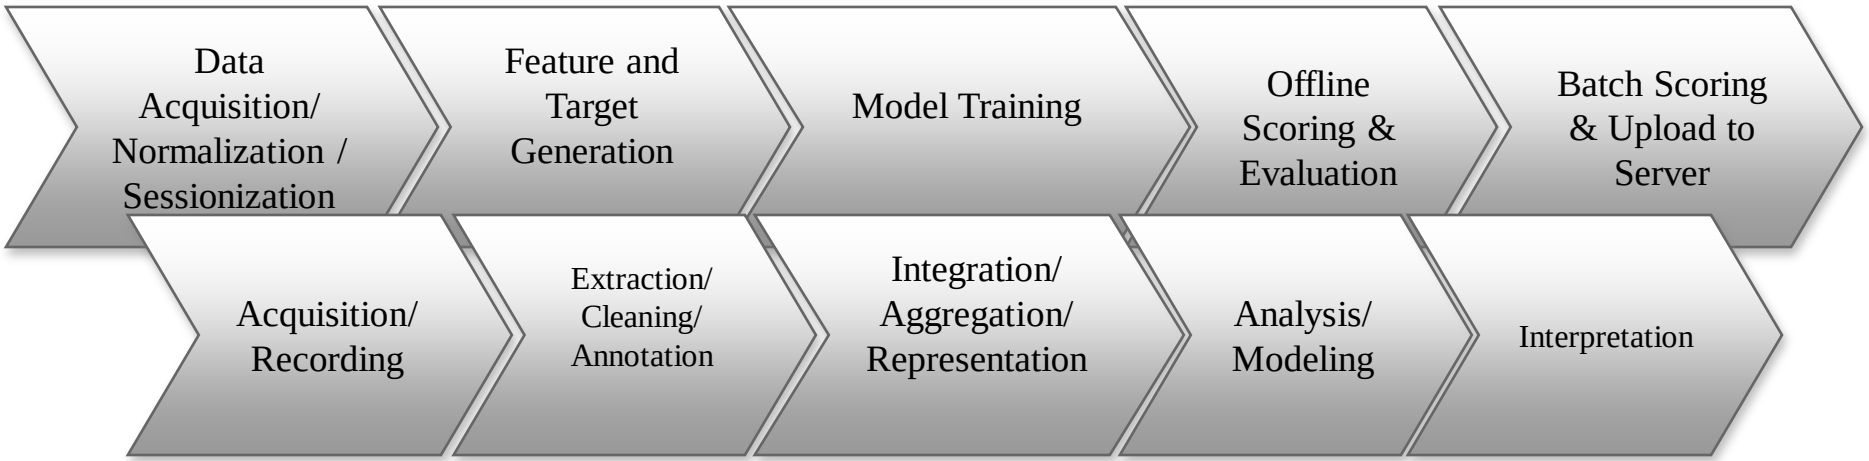

Courtesy: Milind Bhandarkar, Pivotal

Example Application Domains

- Retail
  - Events: clicks on purchases, ad clicks, FB likes, …
  - Goal: Personalized product recommendations
- Datacenters
  - Events: log messages, traffic, communications events, …
  - Goal: Predict imminent failures
- Healthcare
  - Events: Doctor visits, medical history, medicine refills, …
  - Goal: Prevent hospital readmissions
- Telecom
  - Events: Calls made, duration, calls dropped, location, social graph, …
  - Goal: Reduce customer churn
- Web Ads
  - Events: Clicks on content, likes, reposts, search queries, comments, …
  - Goal: Increase engagement, increase clicks on revenue-generation content

Steps in the Pipeline

- Acquisition and normalization of data
  - Collate, consolidate data
- Join targets and features
  - Construct targets; filter out user activity without targets; join feature vector with targets
- Model Training
  - Multi-model: regressions, Naïve Bayes, decision trees, Support Vector Machines, …
- Offline scoring
  - Score features, evaluate metrics
- Batch scoring
  - Apply models to all user activity; upload scores

Application Classes

- Widely varying number of events per entity
- Multiple classes of applications, based on size, e.g.:
  - Tiny (100K entities, 10 events per entity)
  - Small (1M entities, 10 events per entity)
  - Medium (10M entities, 100 events per entity)
  - Large (100M entities, 1000 events per entity)
  - Huge (1B entities, 1000 events per entity)

Proposal for Pipeline Benchmark Results

- Publish results for every stage in the pipeline
- Data pipelines for different application domains may be constructed by mix and match of various pipeline stages
- Different modeling techniques per class
- So, need to publish performance numbers for every stage

A Different Type of Processing Pipeline: For A Scientific Application

- Whole genome assembly

- By: Kristopher Standish*^, Tristan M. Carland*, Glenn K. Lockwood+^, Mahidhar Tatineni+^, Wayne Pfeiffer+^, Nicholas J. Schork*^
- Scripps Translational Science Institute, +San Diego Supercomputer Center, ^UC San Diego

- Project funding provided by Janssen R&D

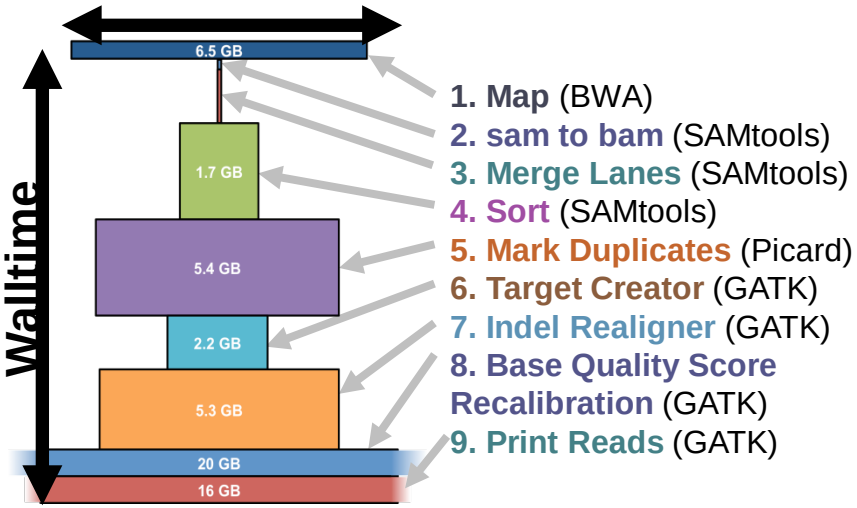

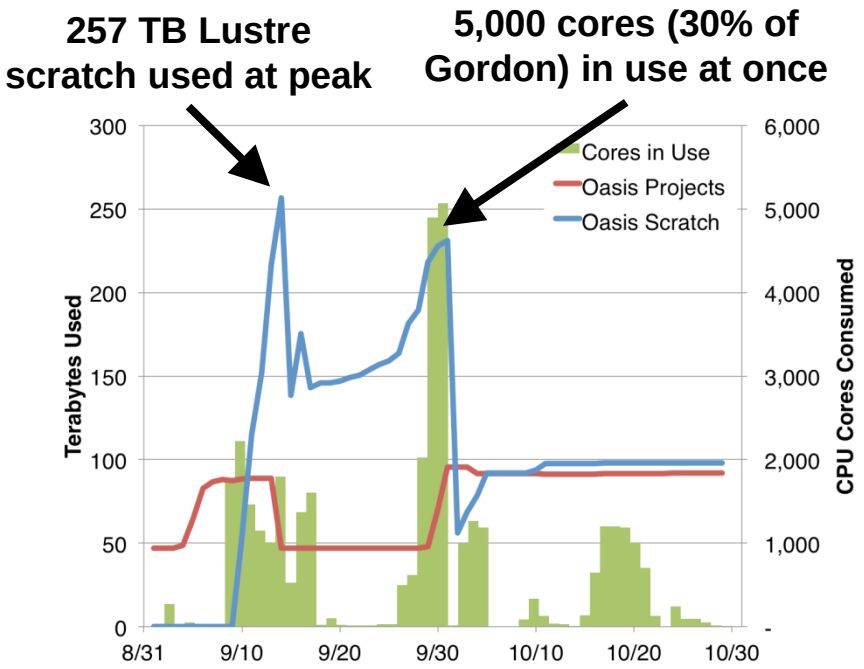

Collaboration Opportunities

- Need your workflows / pipelines!
  - Your data and workloads
- Join our benchmarking project
- Help establish a Benchmarking Laboratory
  - Graduate Projects Program: MS students work on benchmarking with your hardware/software
  - Run benchmark on different system configurations
  - Tune system performance
- Join the Top100 Effort

Thank You !

- Chaitan Baru, baru@sdsc.edu