CPS 216: Data-intensive Computing Systems

Shivnath Babu

A Brief History

Relational database management systems

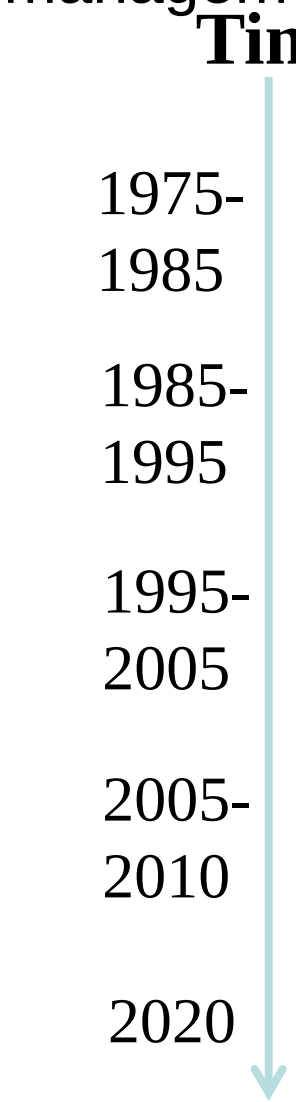

Let us first see what a relational database system is

Data Management

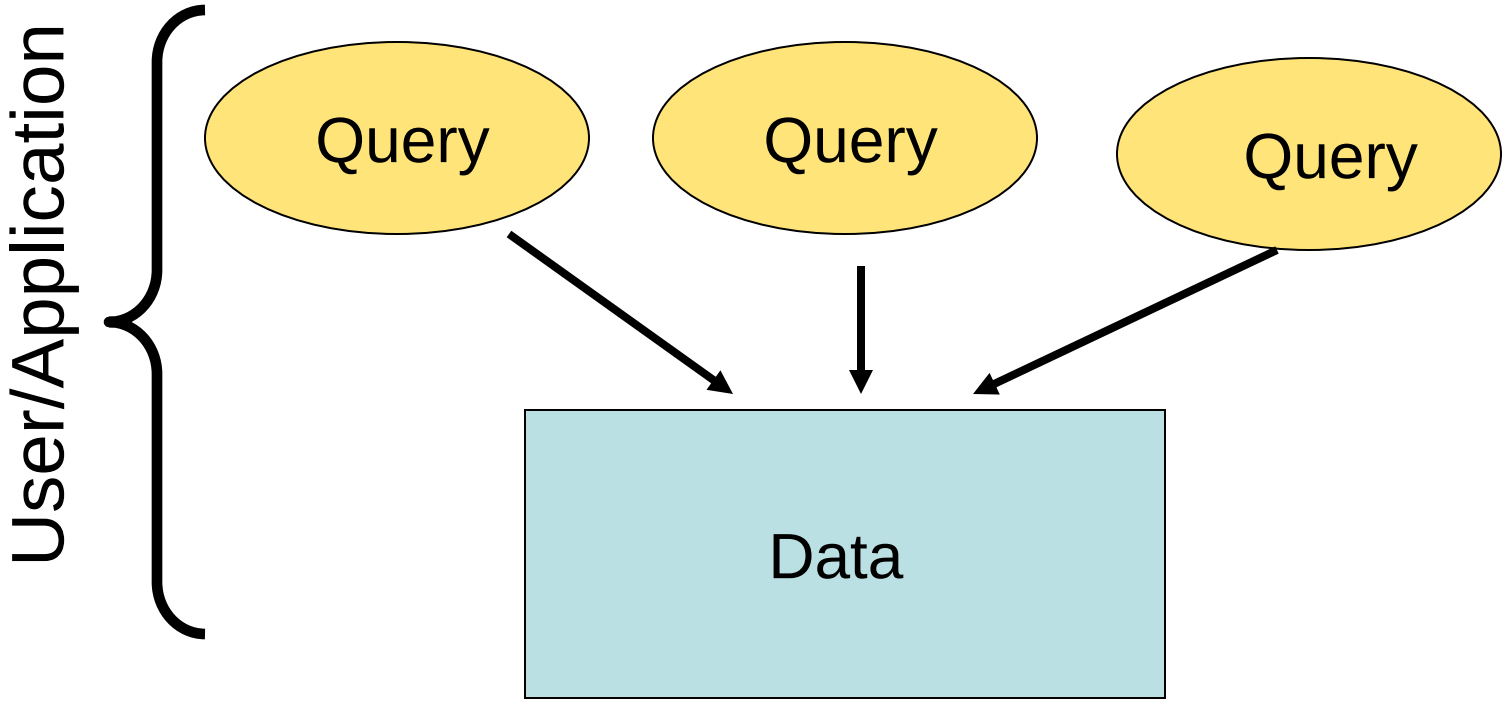

DataBase Management System (DBMS)

Example: At a Company

Query 1: Is there an employee named “Nemo”?
Query 2: What is “Nemo’s” salary?
Query 3: How many departments are there in the company?
Query 4: What is the name of “Nemo’s” department?
Query 5: How many employees are there in the “Accounts” department?

Employee

ID	Name	DeptID	Salary	…
10	Nemo	12	**120K**	…
20	Dory	156	**79K**	…
40	Gill	89	**76K**	…
52	Ray	34	**85K**	…
…	…	…	…	…

Department

ID	Name	…
12	IT	…
34	Accounts	…
89	HR	…
156	Marketing	…
…	…	…

DataBase Management System (DBMS)

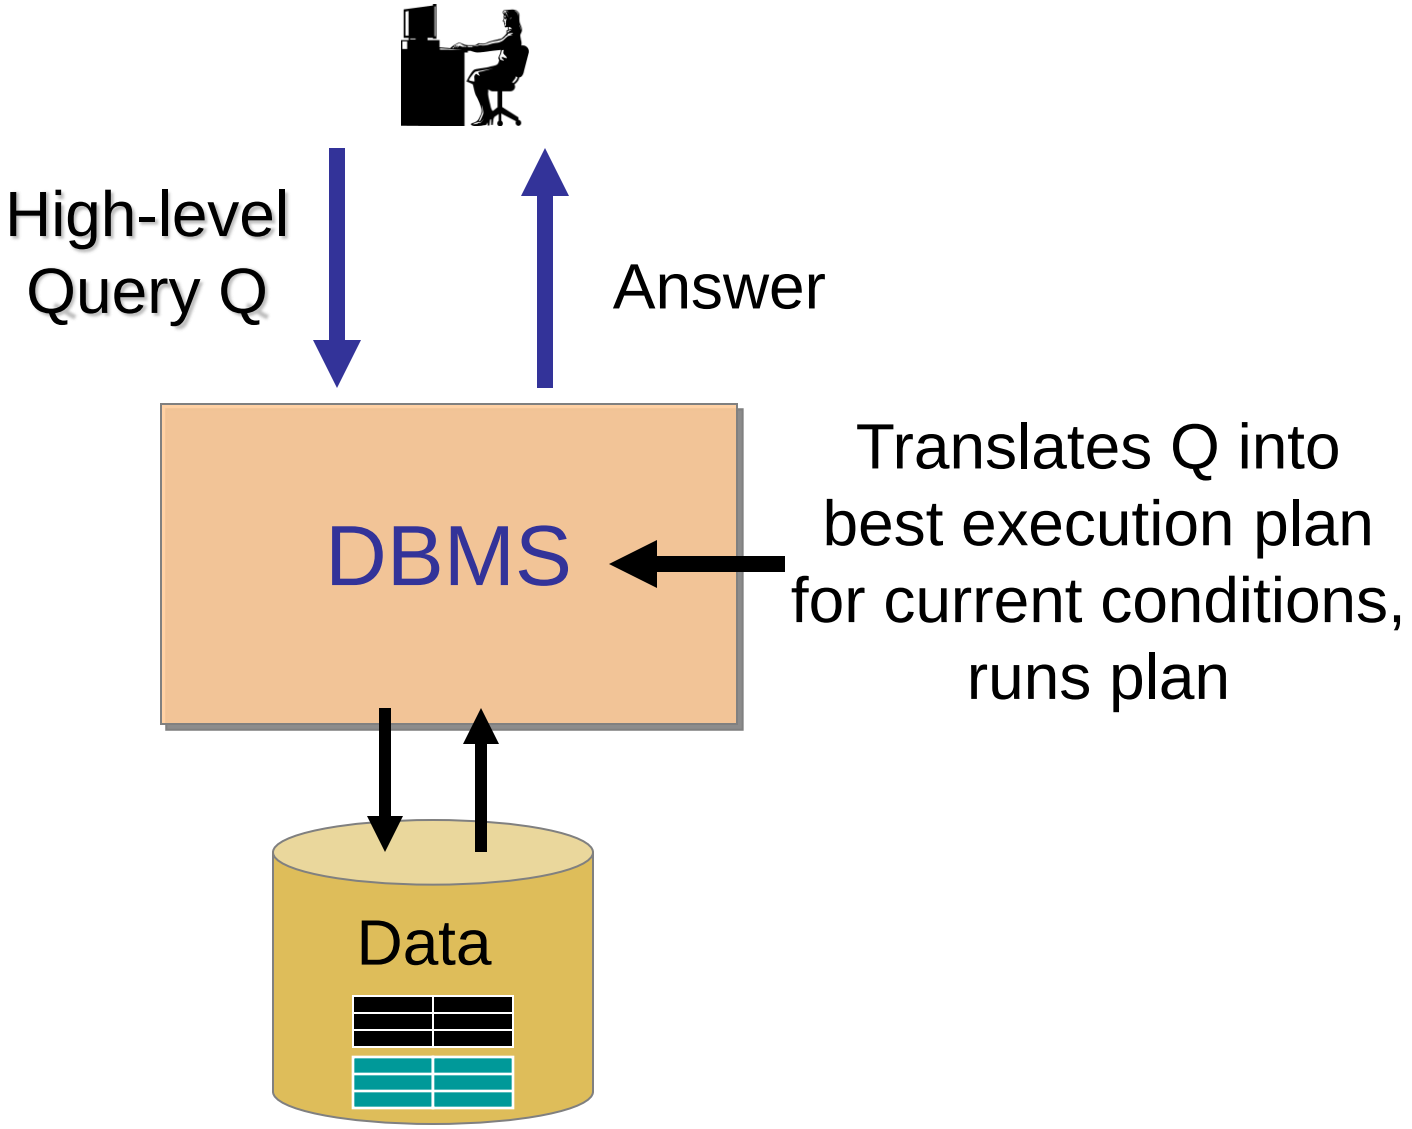

Example: Store that Sells Cars

Owners of Honda Accords who are <= 23 years old

Make	Model	OwnerID	ID	Name	Age
Honda	Accord	12	12	Nemo	**22**
Honda	Accord	156	156	Dory	**21**

Join (Cars.OwnerID = Owners.ID)

Filter (Make = Honda and Model = Accord)

Filter (Age <= 23)

Cars

Make	Model	OwnerID
Honda	Accord	12
Toyota	Camry	34
Mini	Cooper	89
Honda	Accord	156
…	…	…

Owners

ID	Name	Age
12	Nemo	**22**
34	Ray	**42**
89	Gill	**36**
156	Dory	**21**
…	…	**…**

DataBase Management System (DBMS)

High-level Query Q

Answer

DBMS

Translates Q into best execution plan for current conditions, runs plan

Keeps data safe and correct despite failures, concurrent updates, online processing, etc.

Data

A Brief History

Relational database management systems

Time

1975-1985

1985-1995

1995-2005

2005-2010

2020

Assumptions and requirements changed over time

Semi-structured and unstructured data (Web)

Hardware developments

Developments in system software

Changes in data sizes

Big Data: How much data?

- Google processes 20 PB a day (2008)
- Wayback Machine has 3 PB + 100 TB/month (3/2009)
- eBay has 6.5 PB of user data + 50 TB/day (5/2009)
- Facebook has 36 PB of user data + 80-90 TB/day (6/2010)
- CERN's LHC: 15 PB a year (any day now)
- LSST: 6-10 PB a year (~2015)

From http://www.umiacs.umd.edu/~jimmylin/

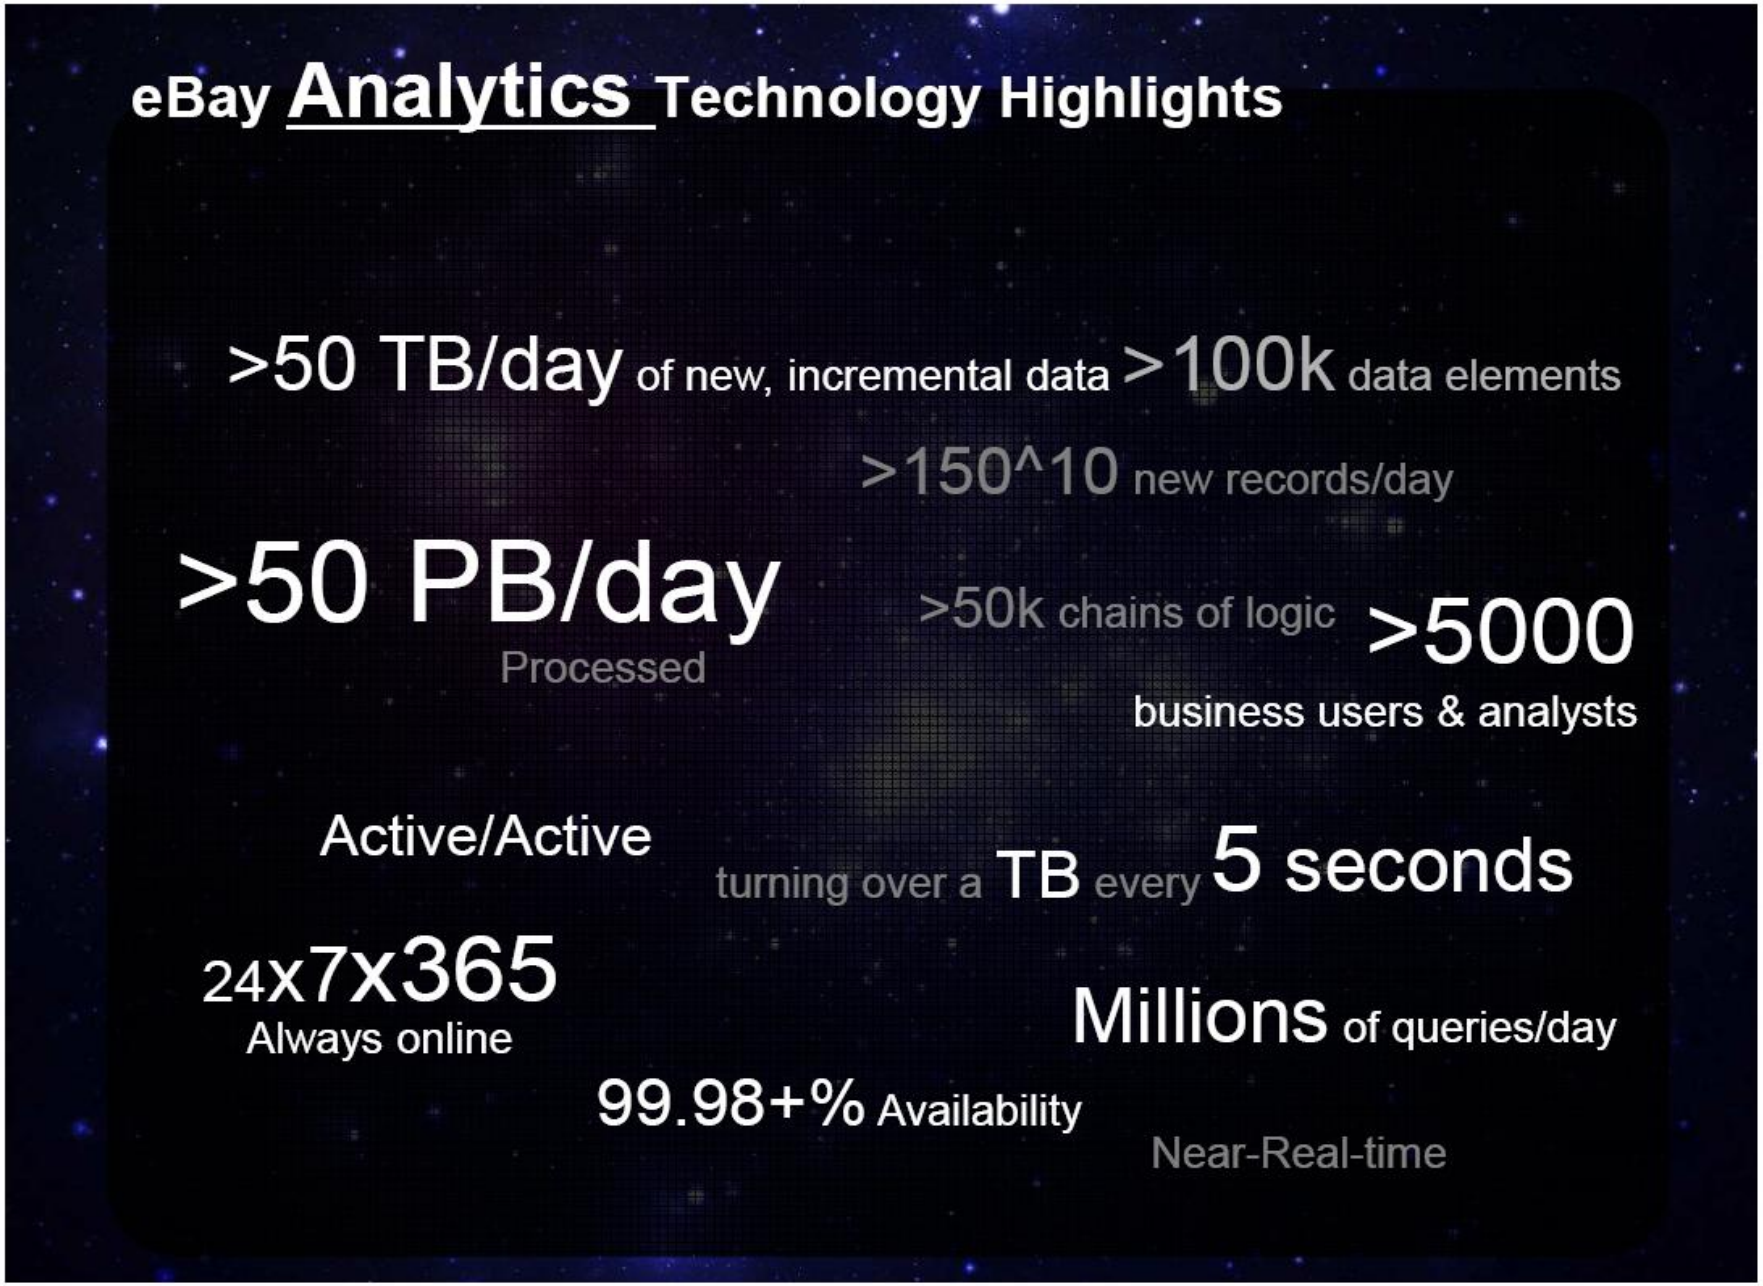

From: http://www.cs.duke.edu/smdb10/

NEW REALITIES

The quest for knowledge used to begin with grand theories.

Now it begins with massive amounts of data.

Welcome to the Petabyte Age.

From: http://db.cs.berkeley.edu/jmh/

FOX AUDIENCE NETWORK

- Greenplum parallel DB
  - 42 Sun X4500s ("Thumper") *each* with:
    - 48 500GB drives
    - 16GB RAM
    - 2 dual-core Opterons
- Big and growing
  - 200 TB data (mirrored)
  - Fact table of 1.5 trillion rows
  - Growing 5TB per day
    - 4-7 Billion rows per day

- Also extensive use of R and Hadoop

Yahoo! runs a 4000 node Hadoop cluster (probably the largest). Overall, there are 38,000 nodes running Hadoop at Yahoo!

From: http://db.cs.berkeley.edu/jmh/

As reported by FAN, Feb, 2009

A SCENARIO FROM FAN

How many female WWF fans under the age of 30 visited the Toyota community over the last 4 days and saw a Class A ad?

How are these people similar to those that visited Nissan?

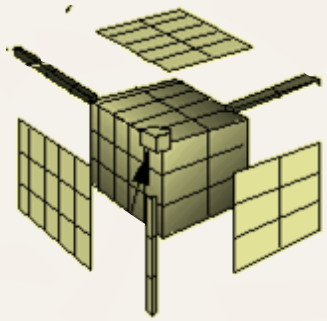

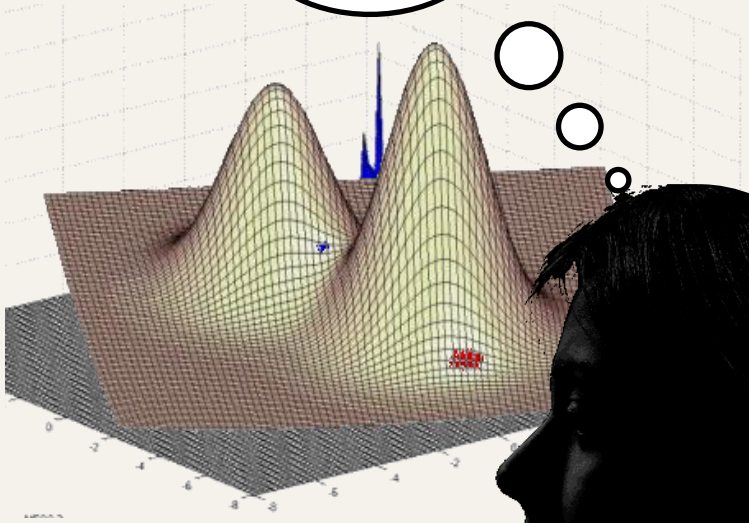

Open-ended question about statistical *densities (distributions)*

From: http://db.cs.berkeley.edu/jmh/

MULTILINGUAL DEVELOPMENT

- SQL or MapReduce
- Sequential code in a variety of languages
  - Perl
  - Python
  - Java
  - R
- Mix and Match!

From: http://db.cs.berkeley.edu/jmh/

The Next Gen = Cloud Computing

From: http://outsideinnovation.blogs.com/pseybold/2009/03/-sun-will-shine-in-blue-cloud.html

What we will cover

- Principles of query processing ***(35%)***
  - Indexes
  - Query execution plans and operators
  - Query optimization
- Data storage ***(15%)***
  - Databases Vs. Filesystems (Google/Hadoop Distributed FileSystem)
  - Data layouts (row-stores, column-stores, partitioning, compression)
- Scalable data processing ***(40%)***
  - Parallel query plans and operators
  - Systems based on MapReduce
  - Scalable key-value stores
  - Processing rapid, high-speed data streams
- Concurrency control and recovery ***(10%)***
  - Consistency models for data (ACID, BASE, Serializability)
  - Write-ahead logging

Course Logistics

- Web: http://www.cs.duke.edu/courses/fall11/cps216
- TA: Rozemary Scarlat
- Books:
  - ***(Recommended)*** *Hadoop: The Definitive Guide*, by Tom White
  - *Cassandra: The Definitive Guide*, by Eben Hewitt
  - *Database Systems: The Complete Book*, by H. Garcia-Molina, J. D. Ullman, and J. Widom
- Grading:
  - Project 25% (Hopefully, on Amazon Cloud!)
  - Homeworks 25%
  - Midterm 25%
  - Final 25%

Projects + Homeworks (50%)

- Project 1 (Sept to late Nov):
  1. Processing collections of records: Systems like Pig, Hive, Jaql, Cascading, Cascalog, HadoopDB
  2. Matrix and graph computations: Systems like Rhipe, Ricardo, SystemML, Mahout, Pregel, Hama
  3. Data stream processing: Systems like Flume, FlumeJava, S4, STREAM, Scribe, STORM
  4. Data serving systems: Systems like BigTable/HBase, Dynamo/Cassandra, CouchDB, MongoDB, Riak, VoltDB
- Project 1 will have regular milestones. The final report will include:
  1. What are properties of the data encountered?
  2. What are concrete examples of workloads that are run? Develop a benchmark workload that you will implement and use in Step 5.
  3. What are typical goals and requirements?
  4. What are typical systems used, and how do they compare with each other?
  5. Install some of these systems and do an experimental evaluation of 1, 2, 3, & 4
- Project 2 (Late Nov to end of class). Of your own choosing. Could be a significant new feature added to Project 1
- Programming assignment 1 (Due third week of class ~Sept 16)
- Programming assignment 2 (Due fifth week of class ~Sept 30)
- Written assignments for major topics