

Advanced Topics in Distributed Systems

Philippe Cudré-Mauroux

September 2013

Erasmus Mundus Program, KTH--Sweden

Lecture 1 – Introduction to Big Data

One Slide About the Lecturer

- M.Sc. from EPFL + Eurécom Institute
 - M.Sc. from INRIA / U. Sophia-Antipolis
 - Ph.D. from EPFL (EPFL best thesis award, EPFL Press Mention)
 - Postdoc M.I.T. Database Systems group
 - Also worked for Hewlett-Packard (CH), IBM Watson (USA), U.C. Berkeley (USA), Microsoft Research Asia (China)
 - Taught courses at EPFL, U.N. (Geneva), M.I.T., Royal Inst. Tech.
 - Associate NSF Prof @ U. Fribourg (CH) since 09/2010
- Research interests
 - Exascale Information Management
 - Scientific data management
 - Linked Data, Semantic Web, Social Web



Instant Quiz

- **Foreign-Key?**
- **Normal Forms?**
- **Two-Phase Commit?**
- **ACID? BASE?**
- **CAP?**
- **Hadoop?**
- **RDF?**

Exascale Data Deluge

- Science
 - Biology
 - Astronomy
 - Remote Sensing
- Web companies
 - Ebay
 - Yahoo
- Financial services, retail companies, governments, etc.



- ➔ New data formats
- ➔ Peta & exa-scale data sets



Big Data Central Theorem

Data+Technology → Actionable Insight → \$\$

Big Data Buzz

Between now and 2015, the firm expects big data to create some **4.4 million IT jobs** globally; of those, 1.9 million will be in the U.S. Applying an economic multiplier to that estimate, Gartner expects each new big-data-related IT job to create work for three more people outside the tech industry, for a total of almost 6 million more U.S. jobs.

Office of Science and Technology
Executive Office of the President
New Executive Office Building
Washington, DC 20502

FOR IMMEDIATE RELEASE
March 29, 2012

Contact: Rick Weiss 202 456-6037 rweiss@ostp.eop.gov
Lisa-Joy Zgorski 703 292-8311 lisajoy@ostp.eop.gov

Growth in the Asia Pacific Big Data market is expected to accelerate rapidly in two to three years time, from a mere US\$258.5 million last year to in excess of **\$1.76 billion in 2016**, with highest growth in the storage segment.

OBAMA ADMINISTRATION UNVEILS “BIG DATA” INITIATIVE: ANNOUNCES \$200 MILLION IN NEW R&D INVESTMENTS

In order to make the most of the fast-growing volume of digital data, the Obama Administration today announced a “Big Data Research and Development Initiative.” Improving our ability to extract knowledge and insights from large and complex collections of digital data, the initiative promises to help solve some the Nation’s most pressing challenges.

To launch the initiative, six Federal departments and agencies today announced more than \$200 million in new commitments that, together, promise to greatly improve the tools and techniques needed to access, organize, and glean discoveries from huge

Big Data Everywhere!

- The Age of Big Data (NYTimes Feb. 11, 2012)
<http://www.nytimes.com/2012/02/12/sunday-review/big-datas-impact-in-the-world.html>

“Welcome to the Age of Big Data. The new megarich of Silicon Valley, first at Google and now Facebook, are masters at harnessing the data of the Web — online searches, posts and messages — with Internet advertising. At the World Economic Forum last month in Davos, Switzerland, Big Data was a marquee topic. A report by the forum, “Big Data, Big Impact,” declared **data a new class of economic asset, like currency or gold.**”

Big data can generate significant financial value across sectors



US health care

- \$300 billion value per year
- ~0.7 percent annual productivity growth




Europe public sector administration

- €250 billion value per year
- ~0.5 percent annual productivity growth



Global personal location data

- \$100 billion+ revenue for service providers
- Up to \$700 billion value to end users



US retail

- 60+% increase in net margin possible
- 0.5–1.0 percent annual productivity growth

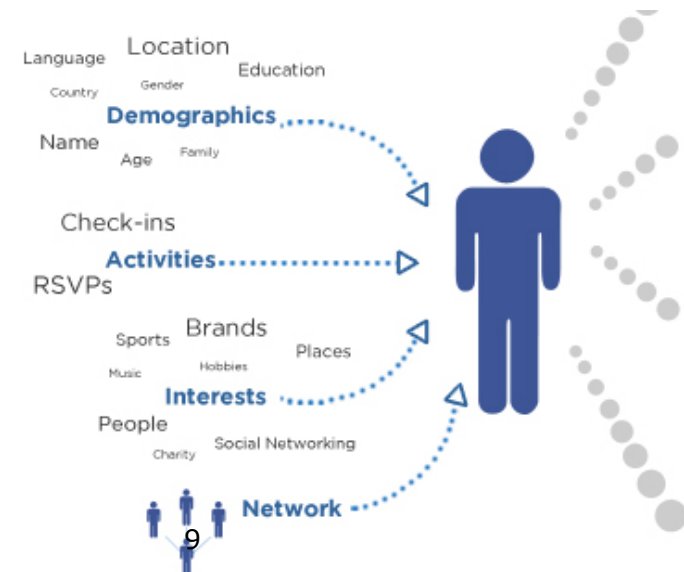


Manufacturing

- Up to 50 percent decrease in product development, assembly costs
- Up to 7 percent reduction in working capital

Typical Big Data Success Story

- **Modeling users** through Big Data
 - Online ads sale / placement [e.g., Facebook]
 - Personalized Coupons [e.g., Target]
 - Product Placement [Walmart]
 - Content Generation [e.g., Netflix]
 - Personalized learning [e.g., Duolingo]
 - HR Recruiting [e.g., Gild]



10 ways big data changes everything

- Some concrete examples
 - <http://gigaom.com/2012/03/11/10-ways-big-data-is-changing-everything/2/>

1. Can gigabytes predict the next Lady Gaga?
2. How big data can curb the world's energy consumption
3. Big data is now your company's virtual assistant
4. The future of Foursquare is data-fueled recommendations
5. How Twitter data-tracked cholera in Haiti
6. Revolutionizing Web publishing with big data
7. Can cell phone data cure society's ills?
8. How data can help predict and create video hits
9. The new face of data visualization
10. One hospital's embrace of big data



The 3-**V**s of Big Data

- **V**olume
 - Amount of data
- **V**elocity
 - speed of data in and out
- **V**ariety
 - range of data types and sources
- [Gartner 2012] *"Big Data are high-volume, high-velocity, and/or high-variety information assets that require new forms of processing to enable enhanced decision making, insight discovery and process optimization"*

What can you do with the data

© Mike Franklin

- Reporting
 - Post Hoc
 - Real time
- Monitoring (fine-grained)
- Exploration
- Finding Patterns
- Root Cause Analysis
- Closed-loop Control
- Model construction
- Prediction
- ...

More Data => Better Answers?

© Mike Jordan

- Not that easy...
- More Rows: Algorithmic complexity kicks in
- More Columns: Exponentially more hypotheses
- Another formulation of the problem:
 - Given an inferential goal and a fixed computational budget, provide a guarantee that the quality of inference will increase monotonically as data accrue (without bound)
- In other words:
 - => **Data should be a resource, not a load**

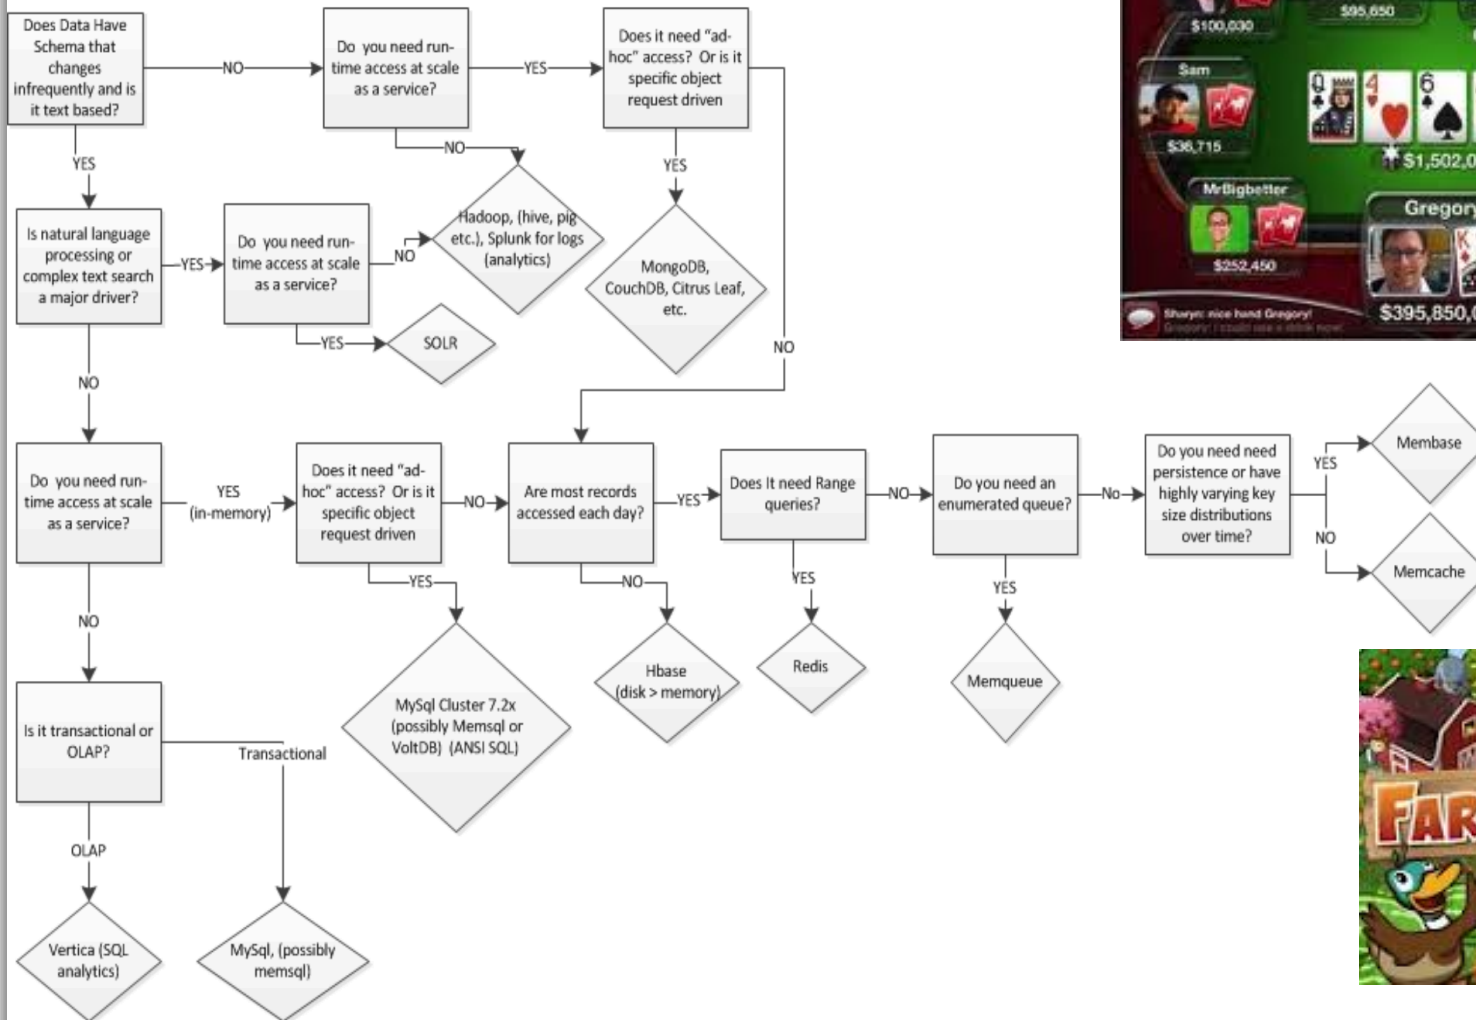
Big Data Today: A Mess

Big Data Landscape (Version 2.0)



A Concrete Example: Zynga

Dan's Scalable Database Decision Matrix at Zynga, 2012



Let's get started

- What is a database?
 - A collection of files storing related data
- Give examples of databases
 - Accounts database; payroll database; students database; Amazon's products database; airline reservation database

Data Management

- Data is *extremely* valuable but hard and costly to manage
 - All aspects of any large-scale, modern project eventually are preserved / analyzed in a database
 - Data management is terribly *hard* today, much harder than 10 years ago
- What operations do we want to perform on this data?
- What functionality do we need to manage this data?

Required Functionality

1. Describe real-world entities in terms of stored data
2. Create & persistently store large datasets
3. Efficiently query & update
 1. Must handle complex questions about data (**OLAP**)
 2. Must handle sophisticated updates (**OLTP**)
 3. Performance matters
4. Change structure (e.g., add attributes)
5. Concurrency control: enable simultaneous updates
6. Crash recovery
7. Access control, security, integrity
8. Distribution, *cloudiness*, etc.

Difficult and costly to implement all these features

Managing Data in the 21 Century

- Creating data management infrastructures is perhaps the most challenging part in today's computing ecology
 - Why? Because distributed data management is, strictly speaking, impossible (CAP theorem)
- A few recent stories:
 - Facebook VS [Friendster](#)
 - Goldman Sachs VS UBS
 - Amazon.com VS the rest of the world
 - Google & Yahoo VS the rest of the world

Database Management System

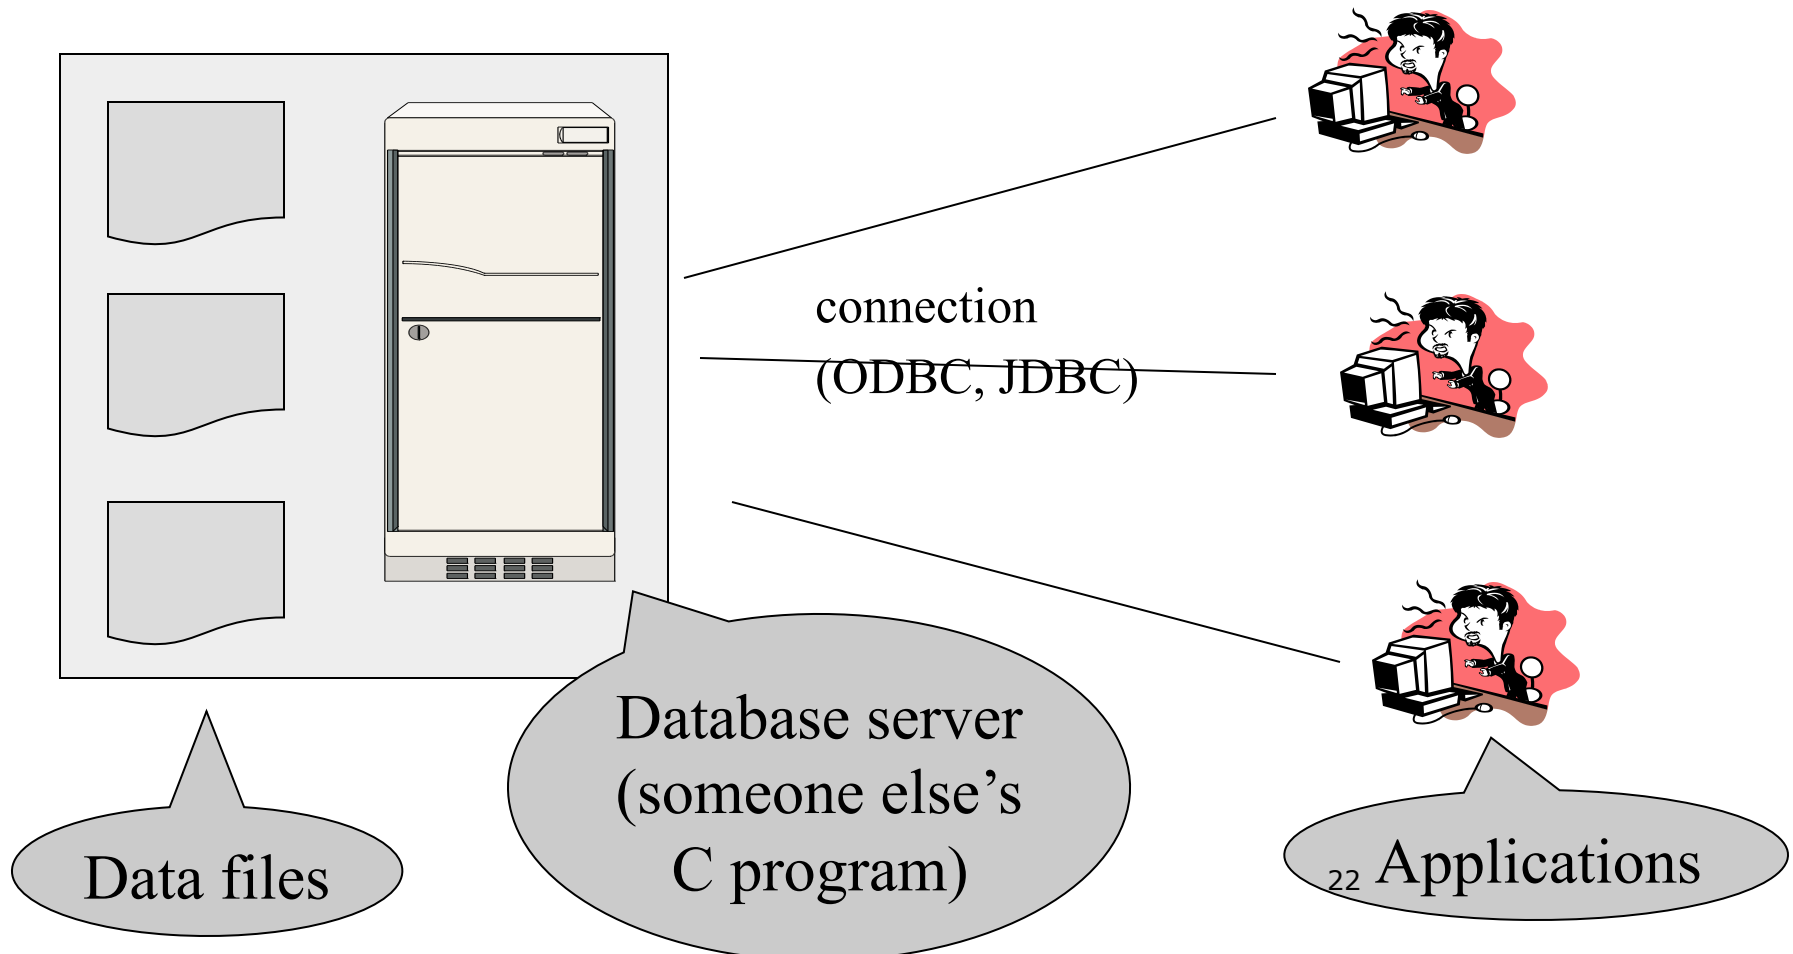
- A DBMS is a software system designed to provide data management services
- Examples of DBMS
 - Oracle, DB2 (IBM), SQL Server (Microsoft),
 - PostgreSQL, MySQL, SciDB
 - Teradata, Vertica
 - Google BigTable, Yahoo! PNUTS, (Facebook) Cassandra

Market Shares

- In 2004 (from www.computerworld.com)
 - IBM, 35% market with \$2.5 billion in sales
 - Oracle, 33% market with \$2.3 billion in sales
 - Microsoft, 19% market with \$1.3 billion in sales
- Quite different today with new players and start-ups
 - no-SQL systems
 - Vertical systems
 - High-end analytics
 - Stream databases
 - Cloud databases
 - Main-memory transactional systems
 - Etc.

Typical System Architecture

“Two tier system” or “client-server”



Main DBMS Features

- Data independence
 - Data model
 - Data definition language
 - Data manipulation language
- Efficient data access
- Data integrity and security
- Data administration
- Concurrency control
- Crash recovery
- Reduced application development time

How to decide what features should go into the DBMS?

When not to use a DBMS?

- DBMS is optimized for a certain workload
- Some applications may need
 - A completely different data model
 - Completely different operations
 - A few time-critical operations
- Examples
 - Text processing
 - Scientific analysis

Relation Definition

- Database is collection of relations
- Relation R is subset of $S_1 \times S_2 \times \dots \times S_n$
 - Where S_i is the domain of attribute i
 - n is number of attributes of the relation
- Relation is basically a table with rows & columns
 - SQL uses word *table* to refer to relations

Properties of a Relation

- Each row represents an n-tuple of R
- Ordering of rows is immaterial
- All rows are distinct
- Ordering of columns is significant
 - Because two columns can have same domain
 - But columns are labeled so
 - Applications need not worry about order
 - They can simply use the names
- Domain of each column is a primitive type
- Relation consists of a **relation schema** and **instance**

More Definitions

- **Relation schema**: describes column heads
 - Relation name
 - Name of each field (or column, or attribute)
 - Domain of each field
- **Degree (or arity) of relation**: nb attributes
- **Database schema**: set of all relation schemas

Even More Definitions

- **Relation instance**: concrete table content
 - Set of tuples (also called records) matching the schema
- **Cardinality of relation instance**: nb tuples
- **Database instance**: set of all relation instances

Example

- Relation schema

Supplier(sno: integer, sname: string, scity: string, scanton: string)

- Relation instance

sno	sname	scity	scanton
1	s1	city 1	FR
2	s2	city 1	FR
3	s3	city 2	FR
4	s4	city 2	FR

Integrity Constraints

- **Integrity constraint**
 - Condition specified on a database schema
 - Restricts data that can be stored in db instance
- DBMS enforces integrity constraints
 - Ensures only legal database instances exist
- Simplest form of constraint is domain constraint
 - Attribute values must come from attribute domain

Key Constraints

- **Key constraint:** “certain minimal subset of fields is a **unique identifier** for a tuple”
- **Candidate key**
 - Minimal set of fields
 - That uniquely identify each tuple in a relation
- **Primary key**
 - One candidate key can be selected as primary key

Foreign Key Constraints

- A relation can refer to a tuple in another relation
- **Foreign key**
 - Field that refers to tuples in another relation
 - Typically, this field refers to the primary key of other relation
 - Can pick another field as well

Key Constraint SQL Examples

```
CREATE TABLE Supply(  
    sno integer,  
    pno integer,  
    qty integer,  
    price integer,  
    PRIMARY KEY (sno,pno) ,  
    FOREIGN KEY (sno) REFERENCES Supplier,  
    FOREIGN KEY (pno) REFERENCES Part  
);
```

ACID Properties

- One of the main benefits of using a DBMS
 - Don't worry about concurrent accesses, hardware failures, etc.
- **Atomicity**: Either all changes performed by transaction occur or none occurs
- **Consistency**: A transaction as a whole does not violate integrity constraints (only valid tuples are written)
- **Isolation**: Transactions appear to execute one after the other in sequence
- **Durability**: If a transaction commits, its changes will survive failures