
1 NEW DEFINITION IS ADDED ON urban
1,600+ READS ON Scribd.
13,000+ HOURS MUSIC STREAMING ON PANDORA
12,000+ NEW ADS POSTED ON craigslist
370,000+ MINUTES VOICE CALLS ON skype
98,000+ TWEETS
New Craigslist Ads
THE LARGEST
20,000+ NEW POSTS ON tumblr.
320+ NEW twitter ACCOUNTS
13,000+ iPhone APPLICATIONS DOWNLOADED
100+ NEW Linked in ACCOUNTS
IN 60 SECONDS
1 associatedcontent NEW ARTICLE IS PUBLISHED
THE WORLD'S LARGEST COMMUNITY CREATED CONTENT
QUESTIONS ASKED ON THE INTERNET...
100+ Answers.com
40+ YAHOO! ANSWERS
6,600+ NEW PICTURES ARE UPLOADED ON flickr
600+ NEW VIDEOS
You Tube
50+ WORDPRESS DOWNLOADS
25+ HOURS TOTAL DURATION
70+ DOMAINS REGISTERED
60+ NEW BLOGS
168 MILLION EMAILS ARE SENT
694,445 SEARCH QUERIES
1,700+ Firefox DOWNLOADS
695,000+ facebook STATUS UPDATES
125+ PLUGIN DOWNLOADS
1,500+ BLOG POSTS
79,364 WALL POSTS

Big data all around

Genoveva Vargas-Solar
CR1, CNRS, LIG-LAFMIA
Genoveva.Vargas@imag.fr

http://vargas-solar.com, Evry, 23rd June, 2015

What everybody says about Big Data

Digital information scale

Megabyte(10^6) Gigabyte (10^9) Terabytes (10^{12}), petabytes (10^{15}), exabytes (10^{18}) and zettabytes (10^{21})

Unit	Size	Meaning
Bit (b)	1 or 0	Short for binary digit, after the binary code (1 or 0) computers use to store and process data
Byte (B)	8 bits	Enough information to create an English letter or number in computer code. It is the basic unit of computing
Kilobyte (KB)	2^{10} bytes	From "thousand' in Greek. One page of typed text is 2KB
Megabyte (MB)	2^{20} bytes	From "large" in Greek. The complete works of Shakespeare total 5 MB. A typical pop song is 4 MB.
Gigabyte (GB)	2^{30} bytes	From "giant" in Greek. A two-hour film ca be compressed into 1-2 GB.
Terabyte (TB)	2^{40} bytes	From "monster" in Greek. All the catalogued books in America's Library of Congress total 15TB
Petabyte (PB)	2^{50} bytes	All letters delivered in America's postal service this year will amount ca. 5PB. Google processes 1PB per hour.
Exabyte (EB)	2^{60} bytes	Equivalent to 10 billion copies of The Economist
Zettabyte (ZB)	2^{70} bytes	The total amount of information in existence this year is forecast to be around 1,27ZB
Yottabyte (YB)	2^{80} bytes	Currently too big to imagine

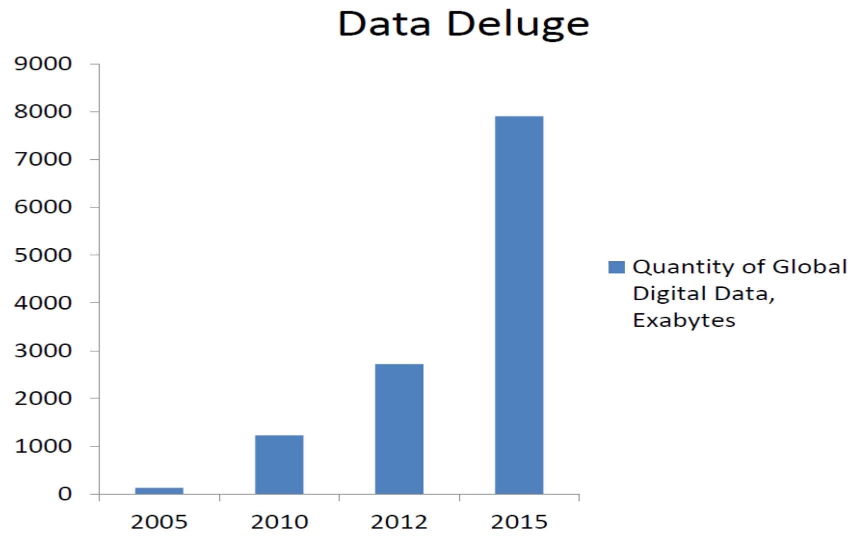

Source: EMC/IDC Digital Universe Study, 2011

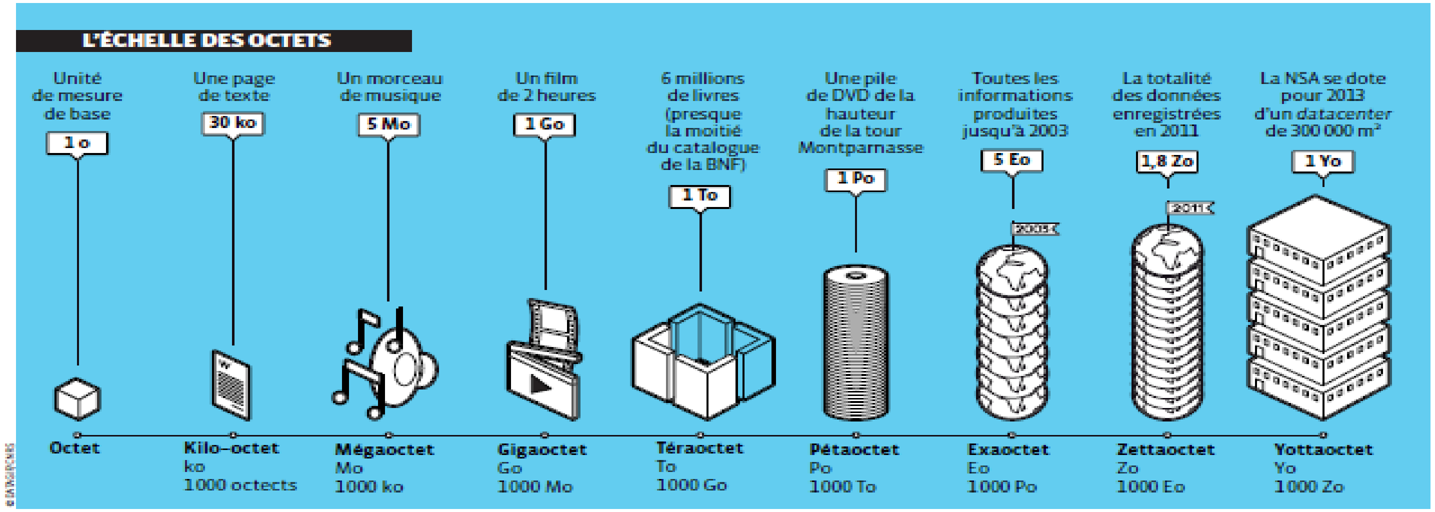

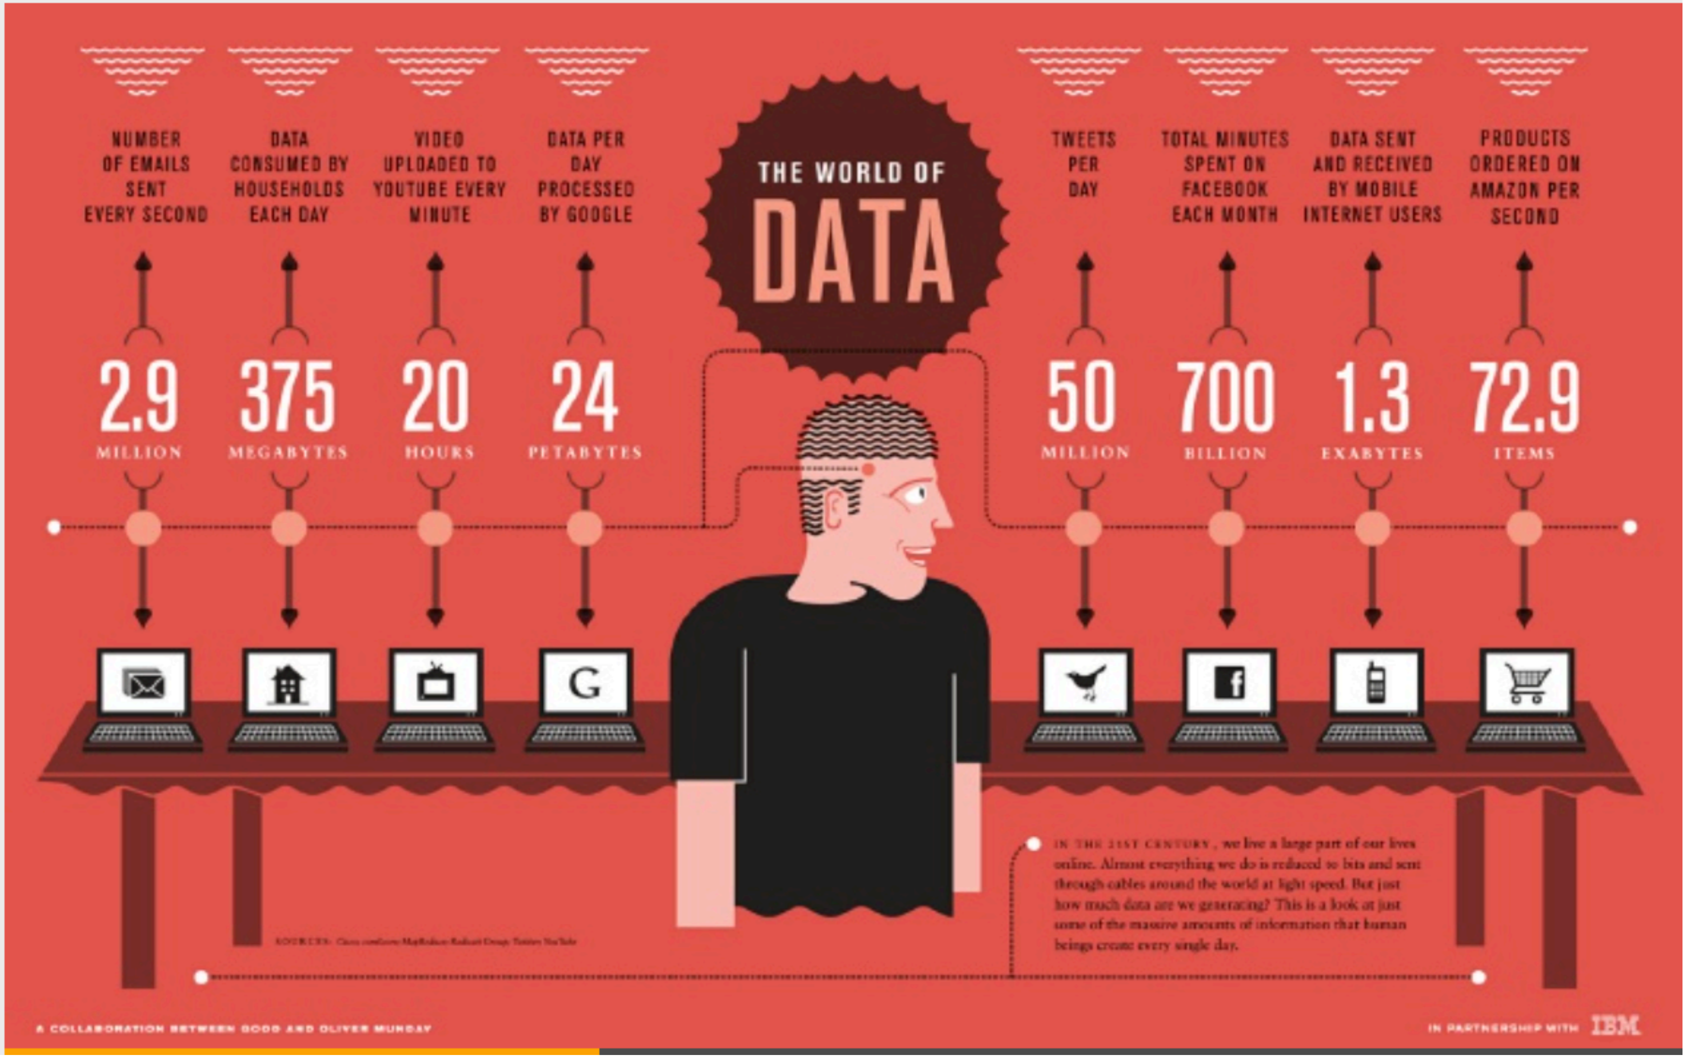
THE WORLD OF
DATA
NUMBER OF EMAILS SENT EVERY SECOND
2.9
MILLION
DATA CONSUMED BY HOUSEHOLDS EACH DAY
375
MEGABYTES
VIDEO UPLOADED TO YOUTUBE EVERY MINUTE
20
HOURS
DATA PER DAY PROCESSED BY GOOGLE
24
PETABYTES
TWEETS PER DAY
50
MILLION
TOTAL MINUTES SPENT ON FACEBOOK EACH MONTH
700
BILLION
DATA SENT AND RECEIVED BY MOBILE INTERNET USERS
1.3
EXABYTES
PRODUCTS ORDERED ON AMAZON PER SECOND
72.9
ITEMS
IN THE 21ST CENTURY, we live a large part of our lives online. Almost everything we do is reduced to bits and sent through cables around the world at light speed. But just how much data are we generating? This is a look at just some of the massive amounts of information that human beings create every single day.
A COLLABORATION BETWEEN GOOD AND OLIVER MUNDAY
IN PARTNERSHIP WITH IBM

Massive data

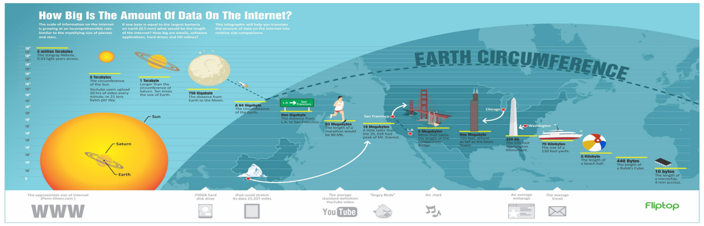

Data sources

- Information-sensing mobile devices, aerial sensory technologies (remote sensing)
- Software logs, posts to social media sites
- Telescopes, cameras, microphones, digital pictures and videos posted online
- Transaction records of online purchases
- RFID readers, wireless sensor networks (number of sensors increasing by 30% a year)
- Cell phone GPS signals (increasing 20% a year)

Data collections

- Sloan Digital Sky Survey (2000-) its archive contains 140 terabytes. Its successor: Large Synoptic Survey Telescope (2016-), will acquire that quantity of data every five days !
- Facebook hosts 140 billion photos, and will add 70 billion this year (ca. 1 petabyte). Every 2 minutes today we snap as many photos as the whole of humanity took in the 1800s !
- Wal-Mart, handles more than 1 million customer transactions every hour, feeding databases estimated at more than 2.5 petabytes (10^{15})
- The Large Hadron Collider (LHC): nearly 15 million billion bytes per year -15 petabytes (10^{15}). These data require 70,000 processors to be processed!

http://blog.websourcing.fr/infographie-la-vrai-taille-dinternet/

Big Data: las primeras 3 V's

- Collection of data sets so large and complex that it becomes difficult to process using on-hand database management tools or traditional data processing applications.

- Challenges include capture, curation, storage, search, sharing, analysis, and visualization **within a tolerable elapsed time**

- Data growth challenges and opportunities are three-dimensional (**3/4Vs model**)
  - increasing **volume** (amount of data)
  - **velocity** (speed of data in and out)
  - **variety** (range of data types and sources)
  - **veracity** (data consistency)

"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."

Characteristics of Big Data:
1-Scale (Volume)

- **Data Volume**
  - 44x increase from 2009 2020
  - From 0.8 zettabytes to 35zb
- Data volume is increasing exponentially

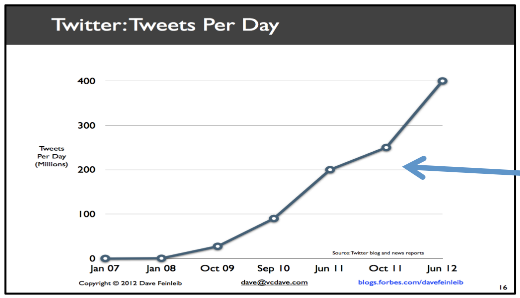

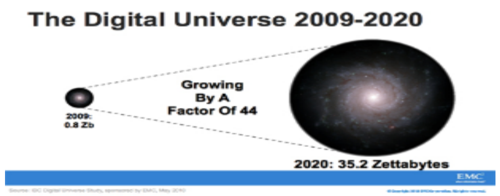

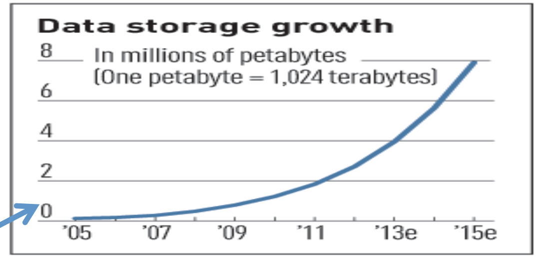

Exponential increase in collected/generated data

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

Characteristics of Big Data:
2-Complexity (Variety)

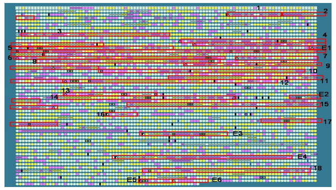

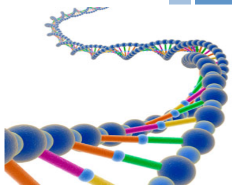

- Various formats, types, and structures
- Text, numerical, images, audio, video, sequences, time series, social media data, multi-dim arrays, etc…

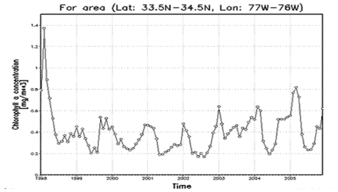

- Static data vs. streaming data
- A single application can be generating/collecting many types of data

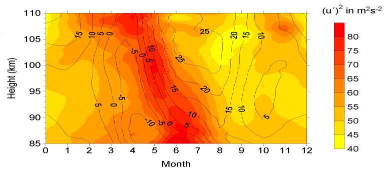

To extract knowledge➔ all these types of data need to linked together

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

Characteristics of Big Data:
3-Speed (Velocity)

- Data is begin generated fast and need to be processed fast
- Online Data Analytics
- Late decisions ➔ missing opportunities

- **Examples**
  - **E-Promotions:** Based on your current location, your purchase history, what you like ➔ send promotions right now for store next to you
  - **Healthcare monitoring:** sensors monitoring your activities and body ➔ any abnormal measurements require immediate reaction

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

Big Data: 3V's

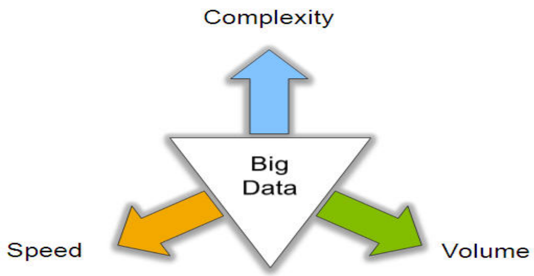

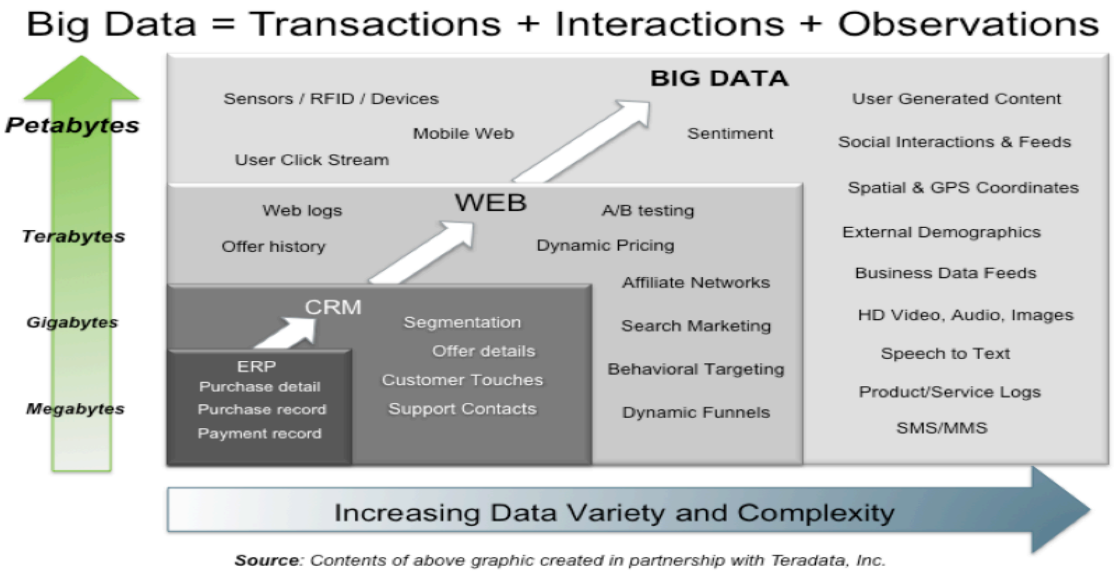

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

Some Make it 4V's

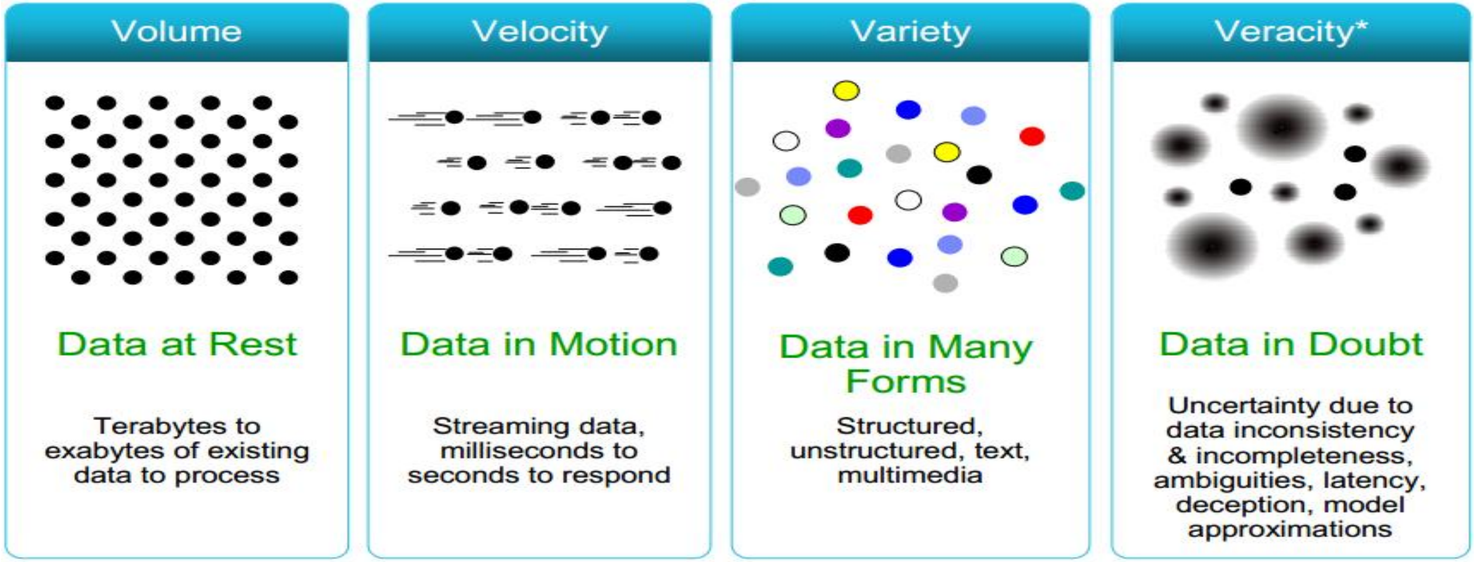

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

Why do we have to care about Big Data

Digital shadow

75% of the information is generated by individuals — writing documents, taking pictures, downloading music, etc. — but is far less than the amount of information being created about them in the digital universe

- 3,146 billion mail addresses, of which 360 million Hotmail
- 95.5 million domains.com
- 2.1 billion Internet users, including 922 million in Asia and 485 million in China
- 2.4 billion accounts on social networks
- 1000 billion videos viewed on YouTube, about 140 average per capita on the planet!
- 60 images posted on Instagram every second
- 250 million tweets per day in October

Data access requirements

- Capture, search, discovery, and analysis tools for unstructured data (90% of the digital universe)
- Metadata management
  - Most of tools can create data about data automatically.
  - Metadata, is growing twice as fast as the digital universe as a whole
- Business intelligence tools (also with real-time data)
- Storage management tools (duplication, auto-tiering, and virtualization, what exactly to store)
- Security practices and tools
  - identify the information that needs to be secured and at what level of security
  - specific threat protection devices and software, fraud management systems, or reputation protection services

Big **Value** from big data

Big data is not a "thing" but instead a dynamic/activity that crosses many IT borders

- Big data is not only about the original content stored or being consumed but also about the information around its consumption.
  - Gigabytes of stored content can generate a petabyte or more of transient data that we typically don't store digital TV signals we watch but don't record voice calls that are made digital in the network backbone during the duration of a call
- Big data technologies describe a **new** generation of technologies and architectures, designed to **economically** extract value from very large volumes of a wide variety of data, by enabling **high-velocity** capture, discovery, and/or analysis

http://www.emc.com/collateral/analyst-reports/idc-extracting-value-from-chaos-ar.pdf

Big Data 5V's model

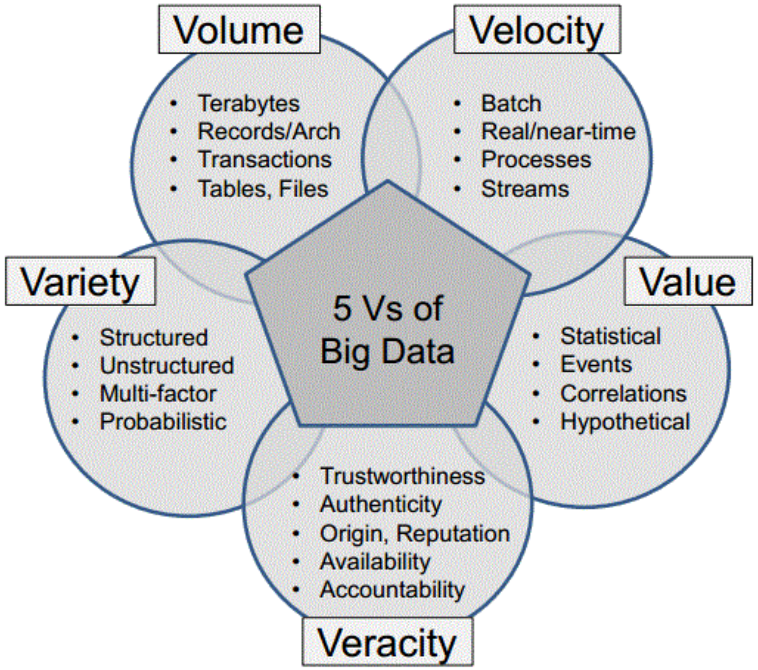

What's driving Big Data

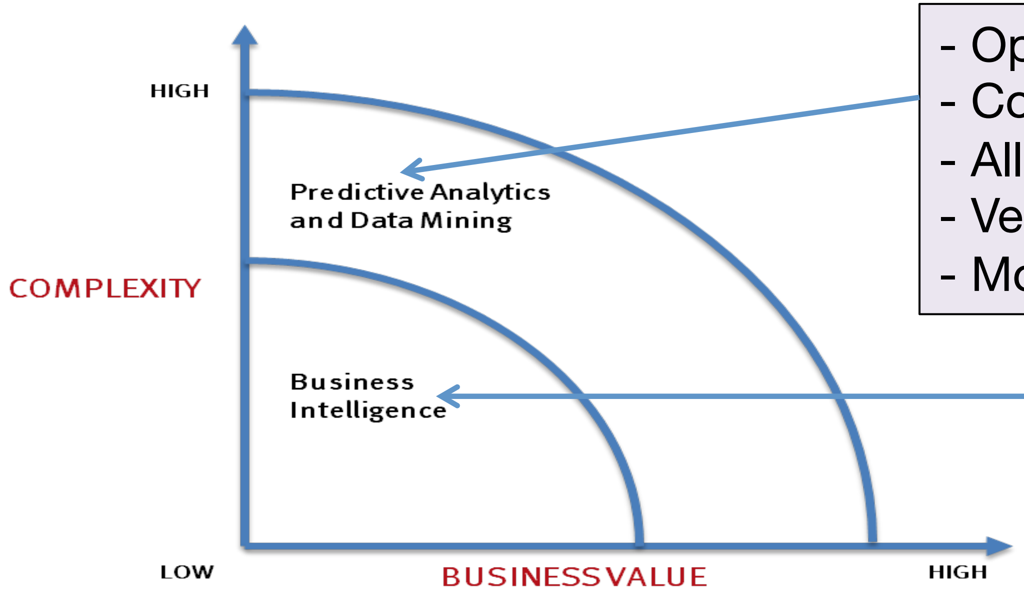

- Optimizations and predictive analytics
- Complex statistical analysis
- All types of data, and many sources
- Very large datasets
- More of a real-time

- Ad-hoc querying and reporting
- Data mining techniques
- Structured data, typical sources
- Small to mid-size datasets

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

The Model Has Changed…

- **The Model of Generating/Consuming Data has Changed**

Old Model: Few companies are generating data, all others are consuming data

New Model: all of us are generating data, and all of us are consuming data

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

Big data life cycle

It is not only the analytics results that matter, the whole processing phases are challenging:

- what to keep and what to throw?
- In which conditions is it acquired and cleaned?
- What is volatile or persistent, for how long?
- Are results volatile or should they be used in further analysis?

Challenges in Handling Big Data

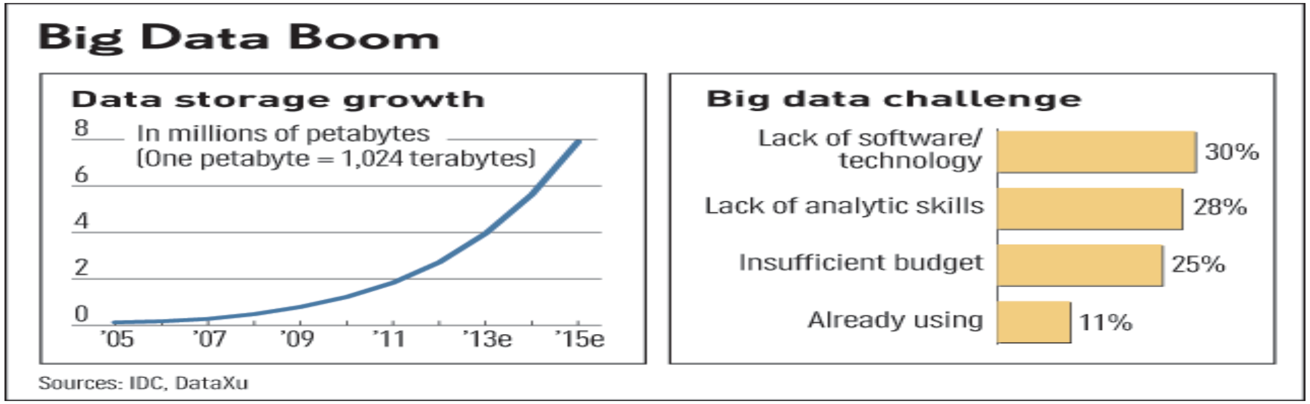

- **The Bottleneck is in technology**
  - New architecture, algorithms, techniques are needed
- **Also in technical skills**
  - Experts in using the new technology and dealing with big data

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

Big Data Landscape

Vertical Apps
Predictive Policing, bloomreach, Myrrix

Log Data Apps
splunk, loggly, sumologic

Ad/Media Apps
rocketfuel, collective[i], bluefin, Recorded Future, Lucky Sort, Media Science, Turn, DataXu

Data As A Service
kaggle, factual., knoema, GNIP, DataSift, Windows Azure Marketplace, INRIX, LexisNexis, Space Curve, LOQATE

Business Intelligence
Oracle | Hyperion, SAP Business Objects, RJMetrics, Microsoft | Business Intelligence, IBM, Cognos, birst, Autonomy, MicroStrategy, QlikView, bime, Domo, Chart.io, GoodData

Analytics and Visualization
tableau, Palantir, Opera, metaLayer, Metamarkets, dataspora, centrifuge, Teradata Aster, sas, TIBCO, Karmasphere, panopticon, pentaho, Datameer, ClearStory, platfora, Cirro, alteryx, visual.ly, Ayata

Analytics Infrastructure
Hortonworks, Vertica, MapR, cloudera, Infobright, ParAccel, EMC², Greenplum, Netezza, kognitio, DataStax, Exasol, calpont

Operational Infrastructure
Couchbase, 10gen | MongoDB company, Teradata, Hadapt, Terracotta, VoltDB, MarkLogic, Informatica

Infrastructure As A Service
amazon web services, Windows Azure, infochimps, Google BigQuery

Structured Databases
Oracle, MySQL, SQL Server, PostgreSQL, IBM DB2, Sybase, memsql

Technologies
hadoop, hadoop MapReduce, mahout, Apache HBase, Cassandra

dave@vcdave.com
blogs.forbes.com/davefeinleib

Big Data Technology

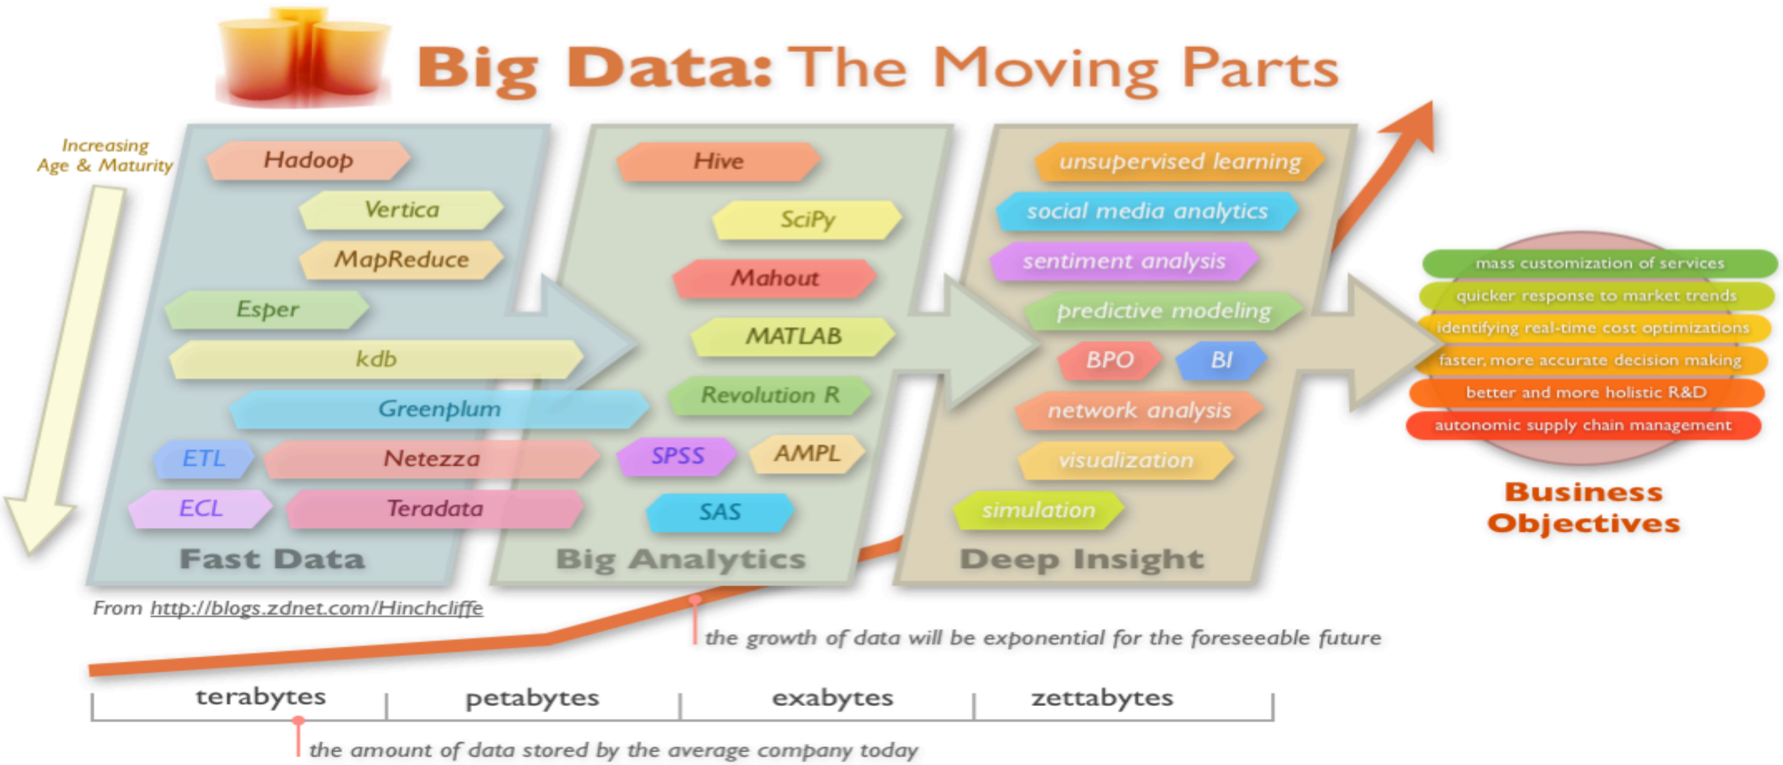

Mohamed Eltabakh, Special Topics in DBs Large-Scale Data Management, http://web.cs.wpi.edu/~cs525/s13-MYE/

Summer Big Data Fest

Menu & Method

Theme of this Course

Large-Scale Data Management

Big Data Analytics

Data Science and Analytics

- How to manage very large amounts of data and extract value and knowledge from them

Course objective

- Introduce challenges associated with Big Data
- Understand and apply data processing and analytics techniques to “big collections”
- **Methodology: theory & “hands on”**
  - **Lectures**
  - **Practical exercises & surprise trivia**
  - **Challenge:** *propose a front end for providing a data market showing 3 phases of the Big Data*

Course logistics

Course logistics

- Web site: all you need to follow lectures
  - http://vargas-solar.com/bigdata-fest/
- Facebook page: all crusty activities during and after lectures
- Teaching team:
  - genoveva.vargas@imag.fr
  - javiera.espinosa@imag.fr

What will you learn?

- Big Data characteristcs (5V model)
- Design & acquisition: different data models, acquisition techniques and data cleaning
- NoSQL storage techniques
- First steps to Map-reduce based data analytics

Some historical facts about data management

Sequential data processing

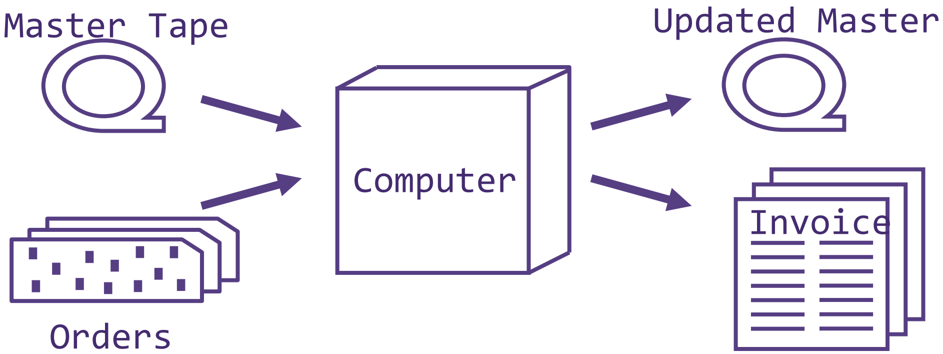

- External medium (tape or cards)
- Data dedicated to one application
- Data seen as flowing through a stationary processor

- Not much different from Hollerith's tabulator
- "Batch Processing"

Disks made modern databases possible

- The key property of disks:
  - Random access to stored data
- Records can contain pointers to other records
- Records can be indexed by their values and accessed directly using B trees

Prof. Rudolf Bayer, Tech. Univ Munich

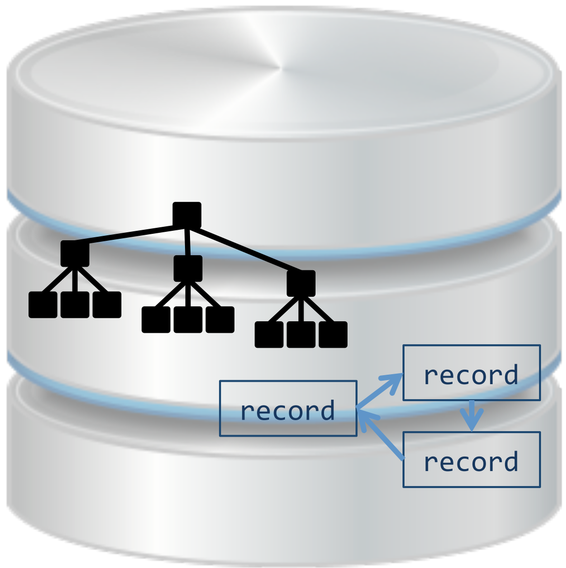

Don Chamberlain

Code and Data: Separated at Birth

COBOL

```
IDENTIFICATION: document

    AUTHOR, PROGRAM-ID, INSTALLATION,
    SOURCE-COMPUTER, OBJECT-COMPUTER,
    SPECIAL-NAMES, FILE-CONTROL, I-O-CONTROL,
    DATE-WRITTEN, DATE-COMPILED,
    SECURITY.

ENVIRONMENT: OS

    CONFIGURATION SECTION.
    INPUT-OUTPUT SECTION.

DATA: Files/Records

    FILE SECTION.
    WORKING-STORAGE SECTION.
    LINKAGE SECTION.
    REPORT SECTION.
    SCREEN SECTION.

PROCEDURE: code
```

"Us"

- **CODASYL – DBTG (1967)**
  COnference on DAta SYstems Languages
  Data Base Task Group
- Defined DDL for a network data model
  - Set-Relationship semantics
  - Cursor Verbs
- **Isolated from procedures**
- **No encapsulation**
- **DATA division is the Schema ancestor**

"them"

Jim Gray

Integrated databases (mid 60's)

- Many applications share data in common
- Data is managed by a centralized system
- Costs are shared
- Redundancy and inconsistency are minimized
- Control is improved
- Access language can be standardized
- Data becomes an enterprise resource
- Shared utilities: backup, recovery, replication, . . .

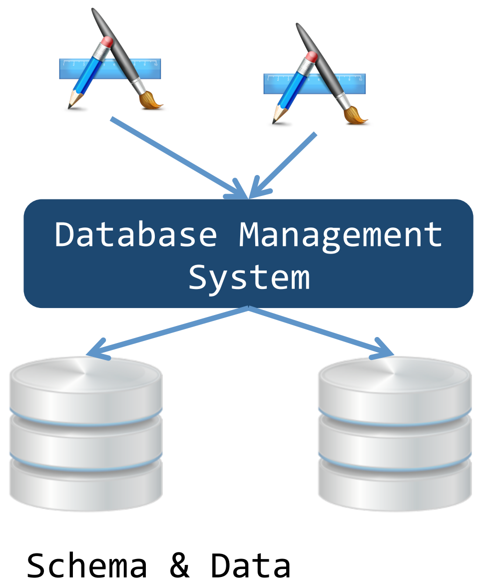

Data Stream Management System DSMS

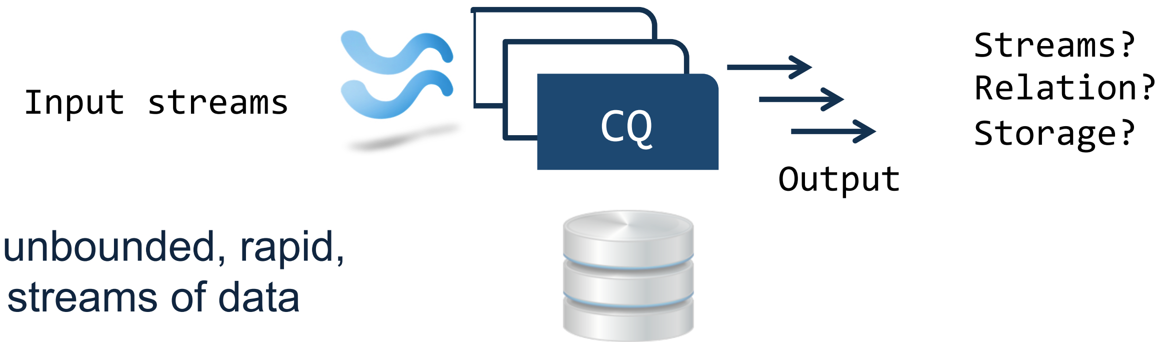

Continuous, unbounded, rapid, time-varying streams of data elements

- Network monitoring and traffic engineering
- Sensor networks, RFID tags
- Telecom call records
- Financial applications

DBMS architecture issues

3-Tier architecture

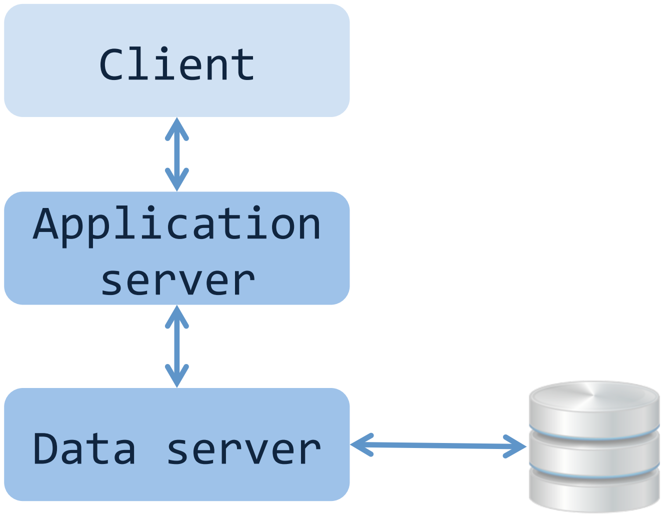

- Logical vs. Physical Architectures
- Components linked in different ways

Systems and architectures

- ***HW Influences*** : machine architecture (P, M, D), networks, data storage units, shared disk, shared nothing
- ***SW Influences*** : O.S. & Middleware

- Very large spectrum of DBMS
  - from smartcard to (paralell) DB Machine
  - Client(s) /Server (s)
  - Main Memory DBMS
  - Monolithic DBMS

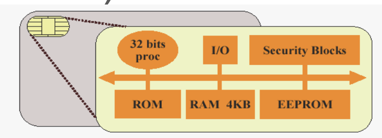

partition and/or replication

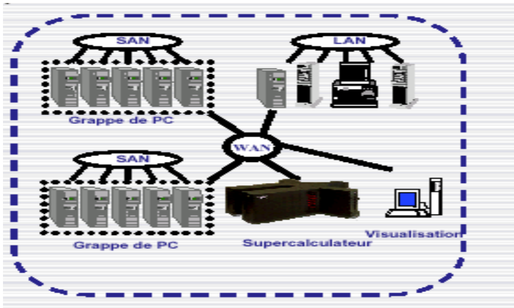

Parallel database architectures

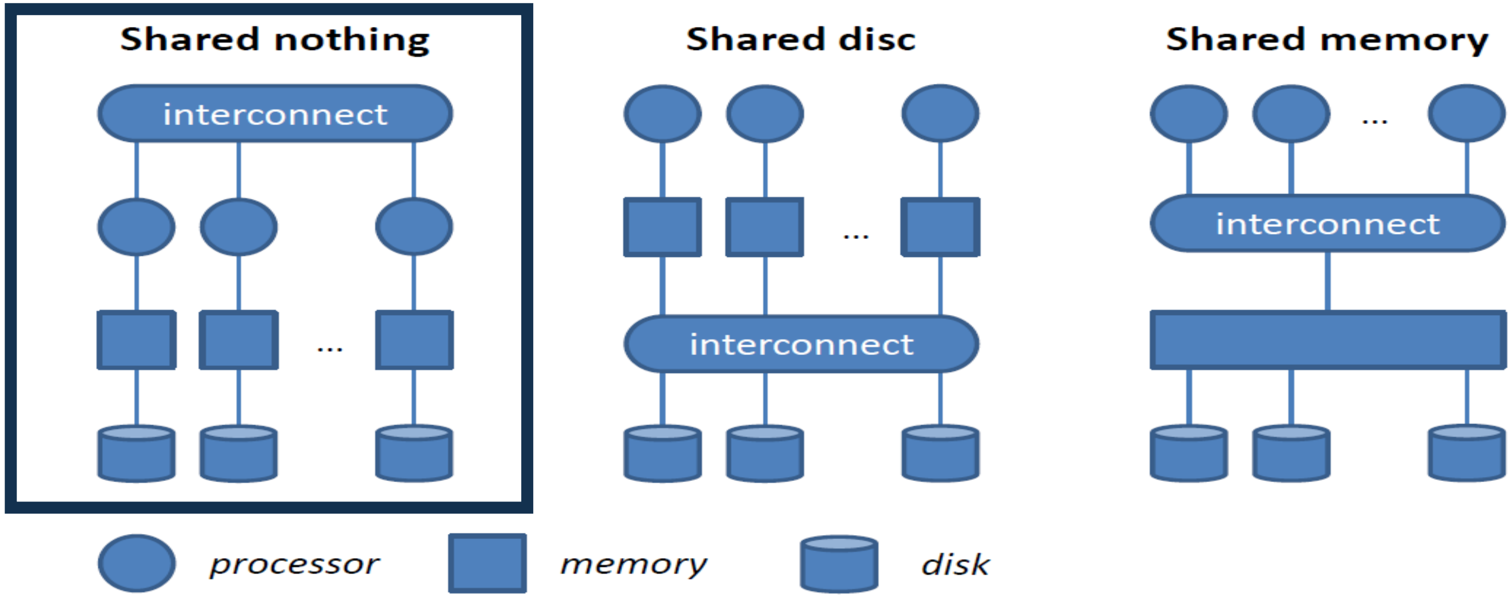

Source: D. DeWitt and J. Gray: "Parallel Database Systems: The Future of High Performance Database Processing", CACM 36(6), pp. 85-98, 1992.

DBMS evolution

- No more monolithic DBMS
- Extensible, lightweight DBMS
- Unbundled technology*
- Component-based architectures* (thick-grain vs. fine-grain)
- OO Frameworks
- Components are providing Services
- Blur the boundaries between OS & DBMS
- Self-adaptive Systems
- Multi-tier architectures, Web, P2P, GRID, CLOUD,…

* See Dittrich, Geppert, Eds, "Component Database Systems", MK 2000

* Chaudhuri & Weikum, Rethinking Database System Architecture: Towards a Self-tuning RISC-style Database System, VLDB 2000

Service oriented DBMS[1]

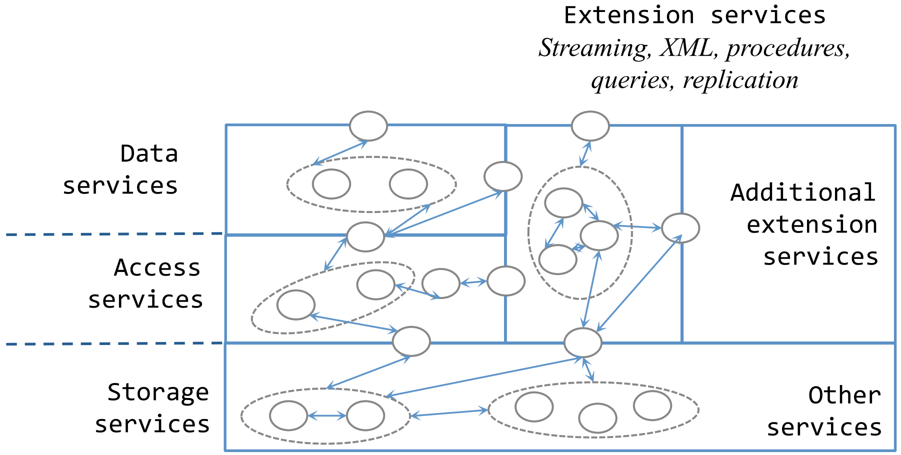

[1] Ionut Subasu, Patrick Ziegler, and Klaus R Dittrich. *Towards service-based data management systems*. In Workshop Proceedings of Datenbanksysteme in Business, Technologie und Web (BTW 2007)
Klaus R Dittrich and Andreas Geppert. *Component database systems*. Morgan Kaufmann, 2000.

Service oriented DBMS[1]

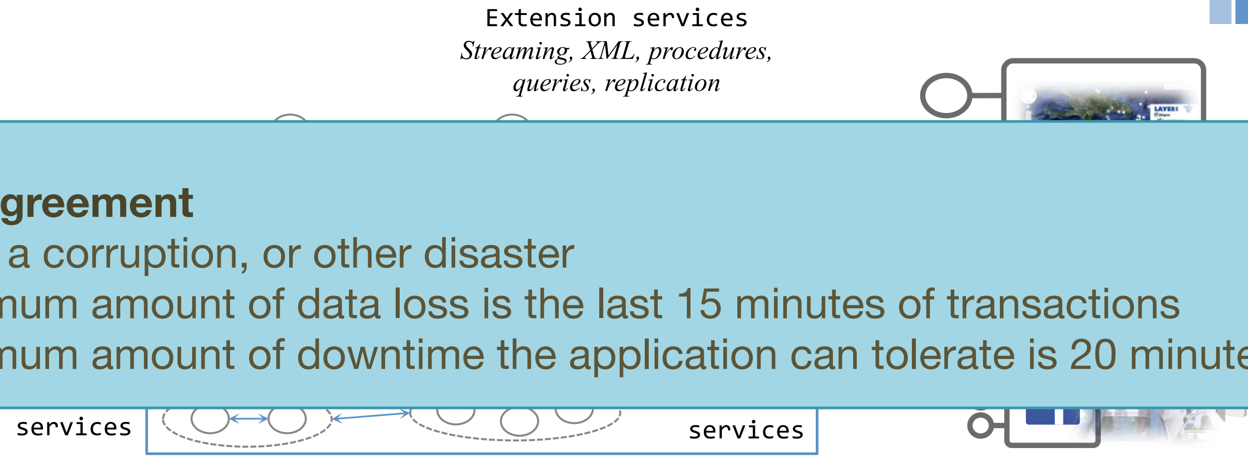

Service Level Agreement

- In the event of a corruption, or other disaster
  - the maximum amount of data loss is the last 15 minutes of transactions
  - the maximum amount of downtime the application can tolerate is 20 minutes

- *Service level agreement:* the contracted delivery time of the service or performance
- *Required SLA*: agreements between the user and SDBMS expressed as a combination of weighted measures associated to a query

[1] Ionut Subasu, Patrick Ziegler, and Klaus R Dittrich. *Towards service-based data management systems*. In Workshop Proceedings of Datenbanksysteme in Business, Technologie und Web (BTW 2007)
Klaus R Dittrich and Andreas Geppert. *Component database systems*. Morgan Kaufmann, 2000.

Service oriented DBMS[1]

Service oriented DBMS[1]

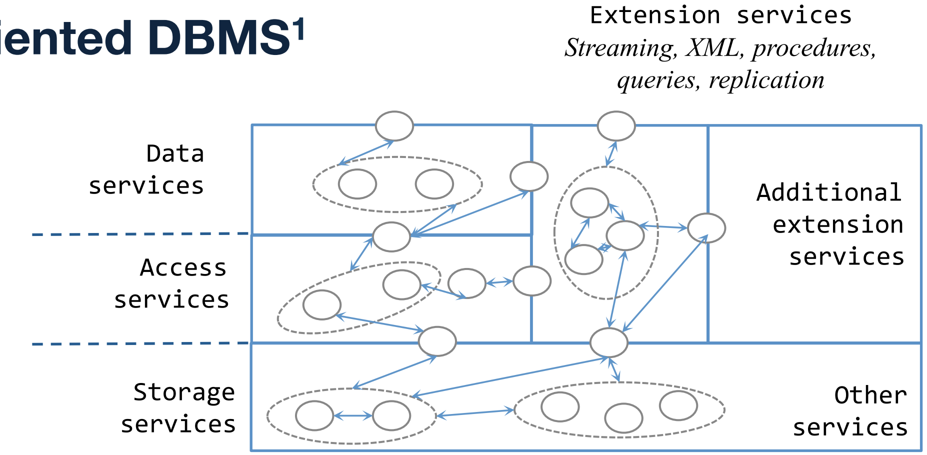

- *Service level agreement:* the contracted delivery time of the service or performance
- *Required SLA*: agreements between the user and SDBMS expressed as a combination of weighted measures associated to a query

[1] Ionut Subasu, Patrick Ziegler, and Klaus R Dittrich. *Towards service-based data management systems*. In Workshop Proceedings of Datenbanksysteme in Business, Technologie und Web (BTW 2007)
Klaus R Dittrich and Andreas Geppert. *Component database systems*. Morgan Kaufmann, 2000.

Challenges and objective

- How to combine, deploy, and deliver DBMS functionalities:
  - **Compliant** to application/user requirements
  - **Optimizing** the consumption of computing resources in the presence of **greedy** data processing tasks
  - Delivered according to **Service Level Agreement (SLA)** contracts
  - Deployed in **elastic** and distributed **platforms**

* See Dittrich, Geppert, Eds, "Component Database Systems", MK 2000

* Chaudhuri & Weikum, Rethinking Database System Architecture: Towards a Self-tuning RISC-style Database System, VLDB 2000

Challenges and objective

- How to combine, deploy, and deliver DBMS functionalities:

"*Design the* ***next generation*** *of* ***data management systems guided by SLA***"

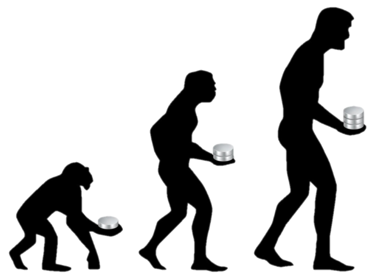

Open Source Big Data Stacks

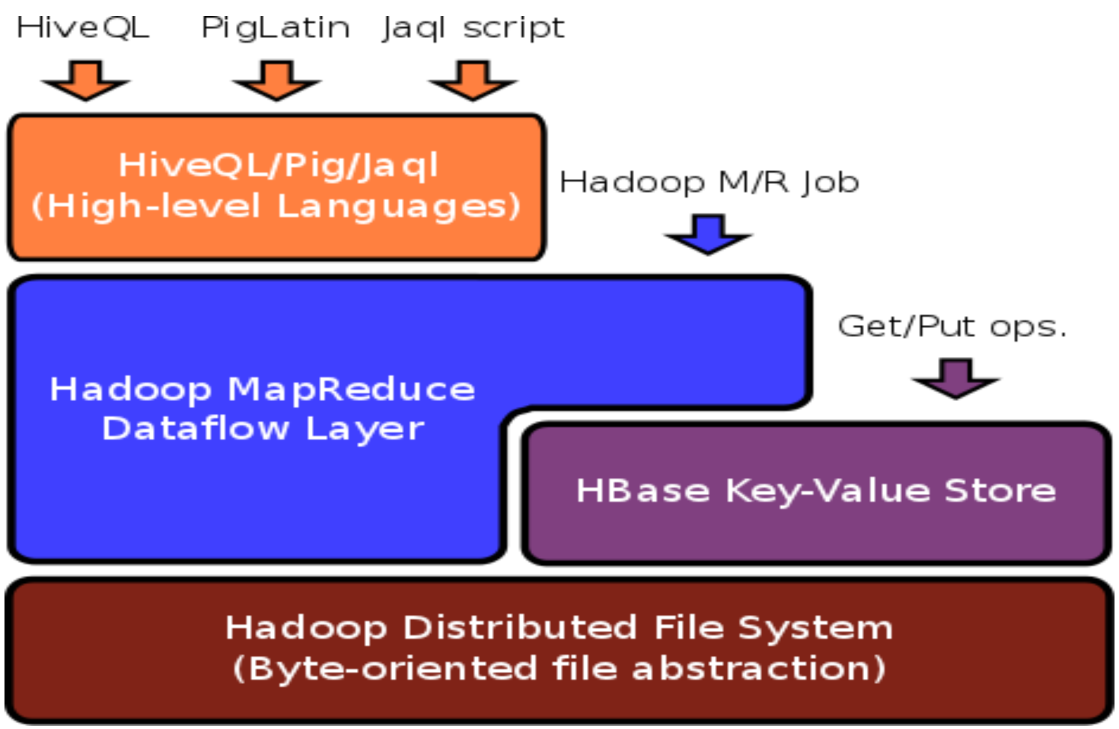

Notes:

- Giant byte sequence at the bottom
- Map, sort, shuffle, reduce layer in middle
- Possible storage layer in middle as well
- HLLs now at the top

From Mike Carey

ASTERIXDB Project @ UCI

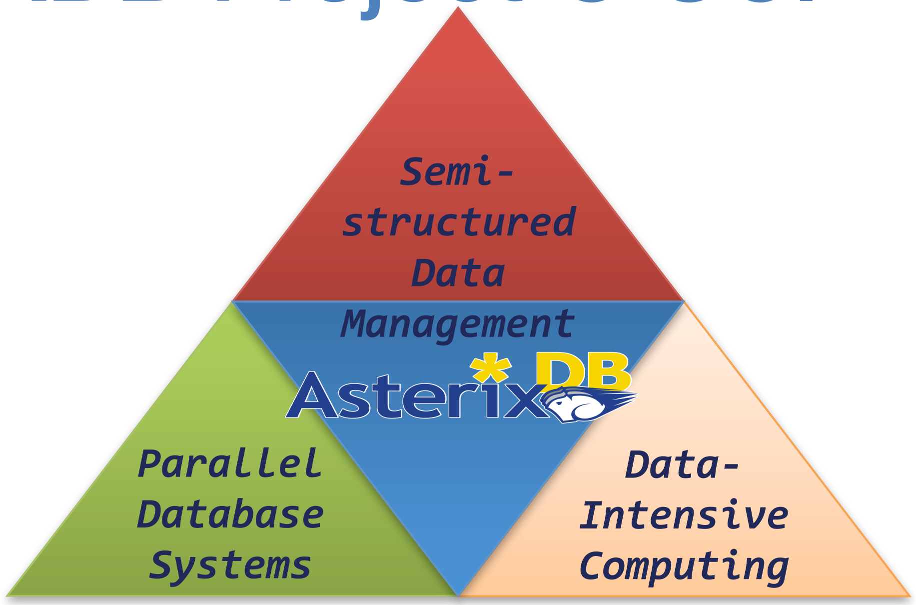

UCIRVINE
http://isg.ics.uci.edu

"One Size Fits a Bunch"

http://asterixdb.ics.uci.edu

•Inside "Big Data Management": Ogres, Onions, or Parfaits?, Vinayak Borkar, Michael J. Carey, Chen Li, EDBT/ICDT 2012 Joint Conference Berlin

•Data Services, Michael J. Carey, Nicola Onose, Michalis Petropoulos
CACM June 2012, (Vol55, N.6)

The ASTERIX Software Stack

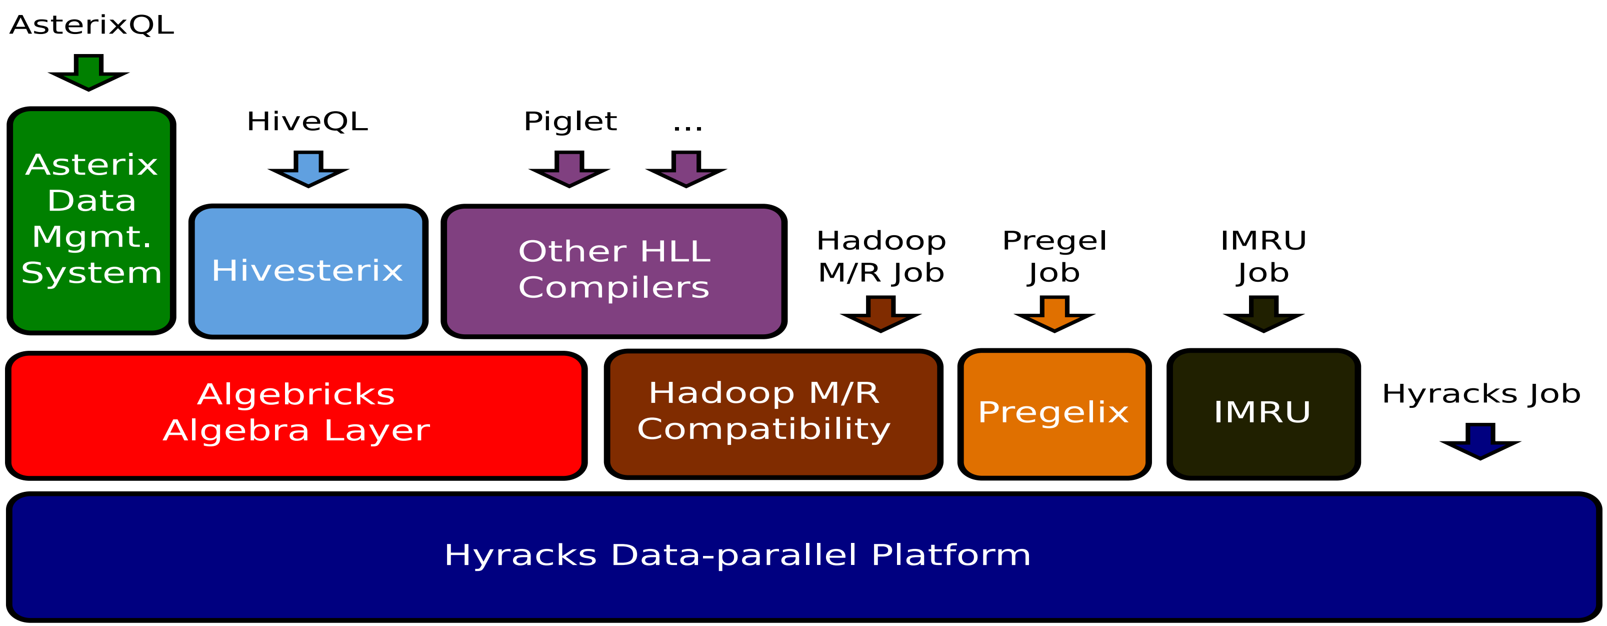

The “pay as you go” era

Economy oriented data and service delivery

What are the forces behind the explosive growth of the digital universe?

- Technology has helped by driving the cost of creating, capturing, managing, and storing information
- … Prime mover is financial
- Since 2005, the investment by enterprises in the digital universe has increased 50% — to $4 trillion
  - spent on hardware, software, services, and staff to create, manage, and store — and derive revenues from the digital universe
  - the trick is to generate value by extracting the right information from the digital universe

Big data tech economy

- Data-driven world guided by a rapid ROD (**Return on Dat**a)

  → reducing cost, complexity, risk and increasing the value of your holdings thanks to a mastery of technologies

- The key is how quickly data can be turned in to currency by:
  - Analysing patterns and spotting relationships/trends that enable decisions to be made faster with more precision and confidence
  - Identifying actions and bits of information that are out of compliance with company policies can avoid millions in fines
  - Proactively reducing the amount of data you pay ($18,750/gigabyte to review in eDiscovery) by identifying only the relevant pieces of information
  - Optimizing storage by deleting or offloading non-critical assets to cheaper cloud storage thus saving millions in archive solutions

Scientific challenges

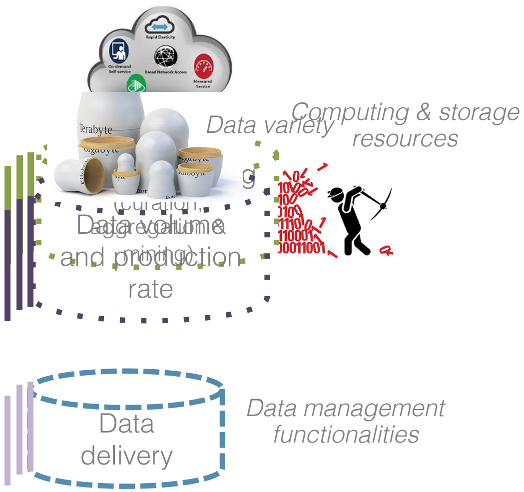

- Estimate the cost of data sanitation, curation, management, and delivery
- Design the cost and business models that can guide these tasks
- Determine the implications of considering the cost/price for implementing data delivery strategies

- Paraschos Koutris, Prasang Upadhyaya, Magdalena Balazinska, Bill Howe, and Dan Suciu. Toward practical query pricing with QueryMarket. In Proceedings of the 2013 international conference on Management of data - SIGMOD '13, page 613, New York, New York, USA, 2013. ACM Press
- Kien Le, Ricardo Bianchini, Jingru Zhang, Yogesh Jaluria, Jiandong Meng, and Thu D. Nguyen. Reducing electricity cost through virtual machine placement in high performance computing clouds. In International Conference for High Performance Computing, Networking, Storage and Analysis, pages 223–243, New York, New York, USA, 2011. ACM Press
- Massimo Mecella, Monica Scannapieco, Antonino Virgillito, Roberto Baldoni, Tiziana Catarci, and Carlo Batini. Managing data quality in cooperative information systems. On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE, pages 486–502, 2002

Scientific challenges

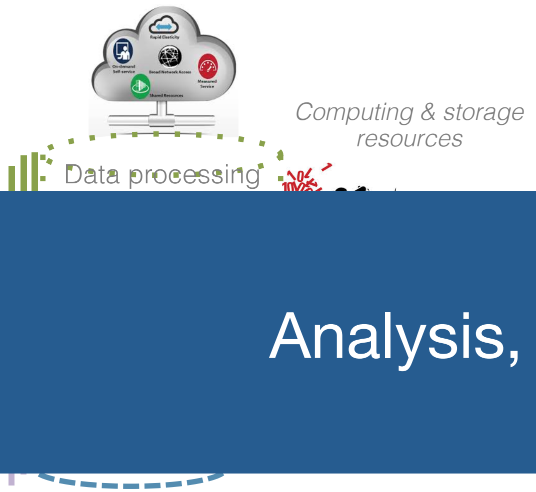

- Estimate the cost of data sanitation, curation, management, and delivery

Analysis, testing, algorithmic, architecture

the cost/price for implementing data brokering strategies

- Paraschos Koutris, Prasang Upadhyaya, Magdalena Balazinska, Bill Howe, and Dan Suciu. Toward practical query pricing with QueryMarket. In Proceedings of the 2013 international conference on Management of data - SIGMOD '13, page 613, New York, New York, USA, 2013. ACM Press
- Kien Le, Ricardo Bianchini, Jingru Zhang, Yogesh Jaluria, Jiandong Meng, and Thu D. Nguyen. Reducing electricity cost through virtual machine placement in high performance computing clouds. In International Conference for High Performance Computing, Networking, Storage and Analysis, pages 223–243, New York, New York, USA, 2011. ACM Press
- Massimo Mecella, Monica Scannapieco, Antonino Virgillito, Roberto Baldoni, Tiziana Catarci, and Carlo Batini. Managing data quality in cooperative information systems. On the Move to Meaningful Internet Systems 2002: CoopIS, DOA, and ODBASE, pages 486–502, 2002

Economy oriented data management

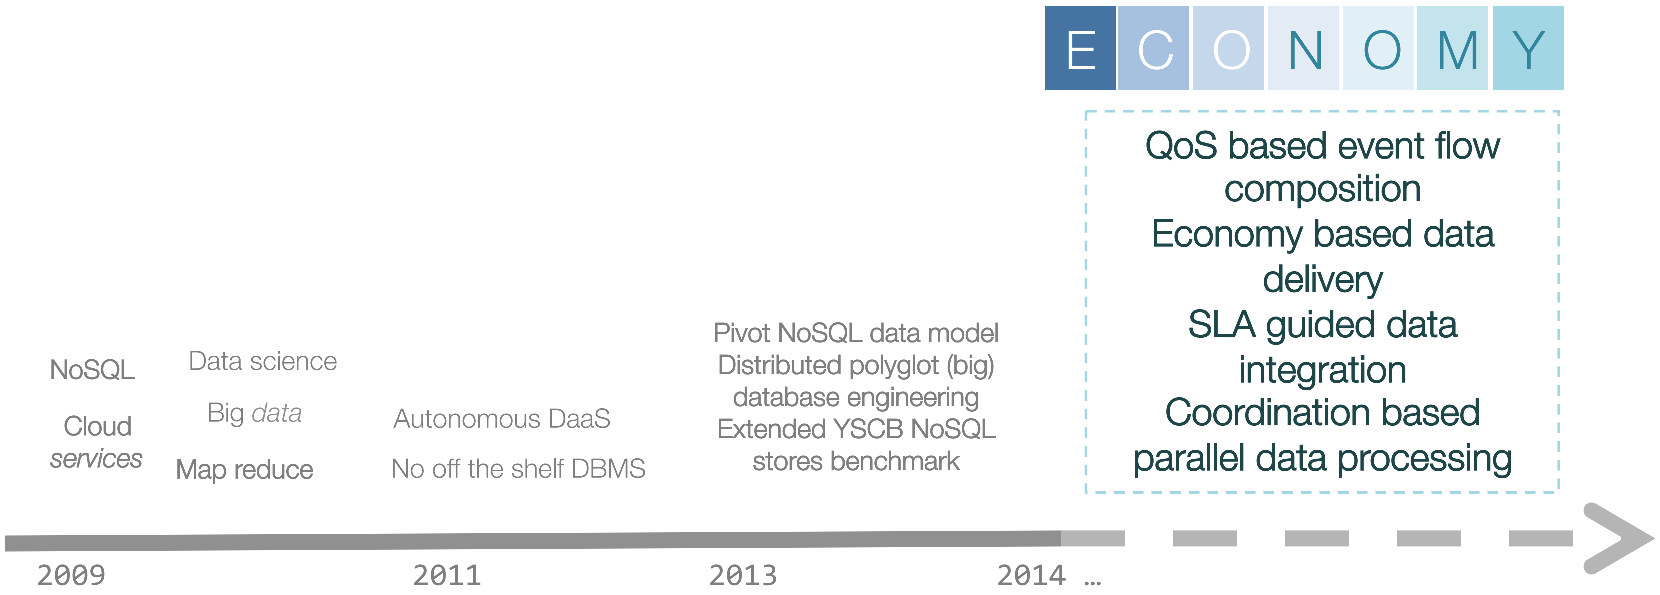

Economy oriented data management

Manage and **capitalize** on the **digital footprint** of the *homo interneticus* and *technologicus*

2009
NoSQL
Cloud *services*
Data science
Big *data*
Map reduce

2011
Autonomous DaaS
No off the shelf DBMS

2013
Pivot NoSQL data model
Distributed polyglot (big) database engineering
Extended YSCB NoSQL stores benchmark

2014 …
SLA guided data integration
Coordination based parallel data processing

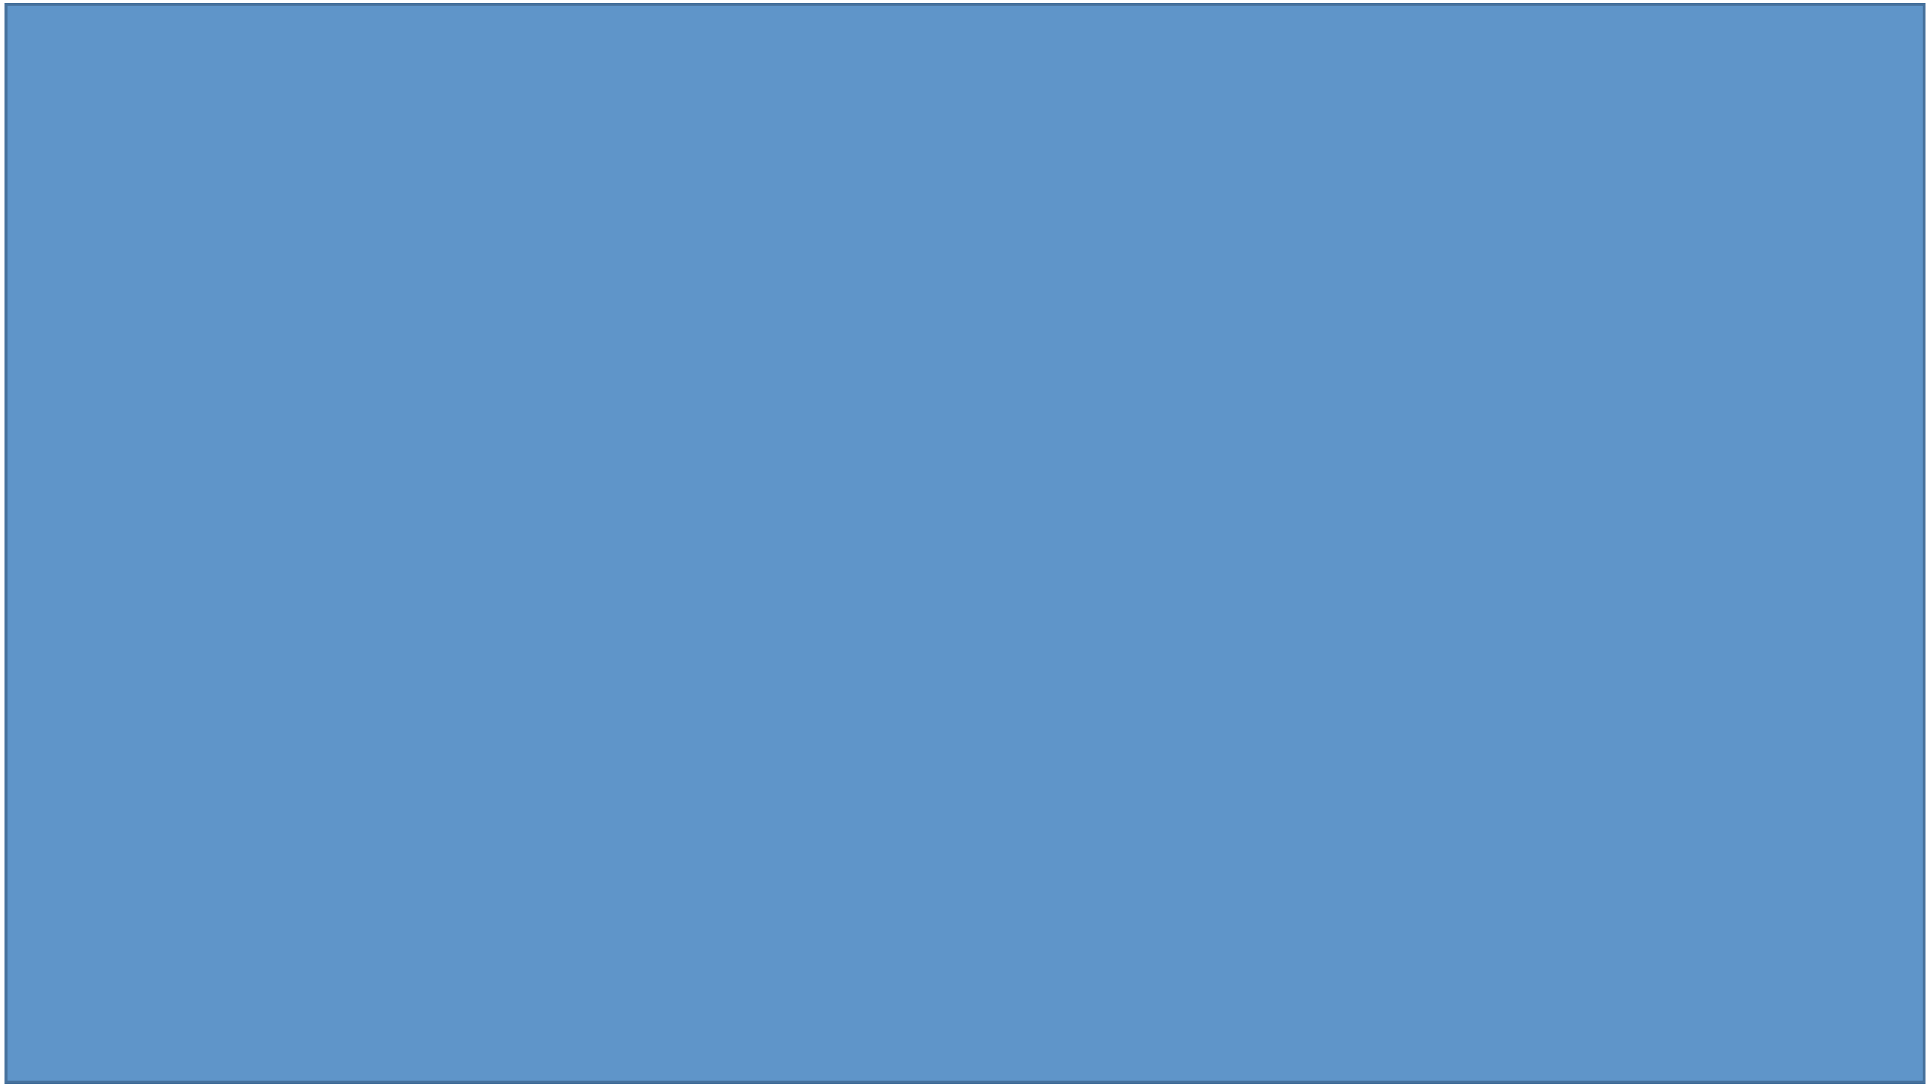

Big data variety: the right model according to data

1970 - 2000 Relational DB

More than 30 years: maturity!

- Theoretical & Practical aspects (DBMS)
- Domains & $R \subseteq D_1 \times D_2 \times \ldots D_n$, Algebra →
- 1st Order Predicate Logic
- Languages: SQL (wins), QUEL, QBE
- DBMS Prototypes (1975), Products (1980)
- A major improvement in DB: provide **data independence** & a simple, **tabular view** of data
- Normal Forms & Dependencies (DB design, **consistency**)
- Controversial: missing values, duplicates

R x S

R ∪ S

R - S

R[α]

R : φ

R * S

Data models

- Tuple
  - Row in a relational table, where attributes are pre-defined in a schema, and the values are scalar
- Document
  - Allows values to be nested documents or lists, as well as scalar values.
  - Attributes are not defined in a global schema
- Extensible record
  - Hybrid between tuple and document, where families of attributes are defined in a schema, but new attributes can be added on a per-record basis

Data stores

- Key-value
  - Systems that store values and an index to find them, based on a key
- Document
  - Systems that store documents, providing index and simple query mechanisms
- Extensible record
  - Systems that store extensible records that can be partitioned vertically and horizontally across nodes
- Graph
  - Systems that store model data as graphs where nodes can represent content modelled as document or key-value structures and arcs represent a relation between the data modelled by the node
- Relational
  - Systems that store, index and query tuples

Key-value stores

- "Simplest data stores" use a data model similar to the memcached distributed in-memory cache
- Single key-value index for all data
- Provide a persistence mechanism
- Replication, versioning, locking, transactions, sorting
- API: inserts, deletes, index lookups
- No secondary indices or keys

System	Address
Redis	`code.google.com/p/redis`
Scalaris	`code.google.com/p/scalaris`
Tokyo	`tokyocabinet.sourceforge.net`
Voldemort	`project-voldemort.com`
Riak	`riak.basho.com`
Membrain	`schoonerinfotech.com/products`
Membase	`membase.com`


```
SELECT   name, pic, profile_url
FROM     user
WHERE    uid = me()
```


```
SELECT   message, attachment
FROM     stream
WHERE    source_id = me() AND type = 80
```

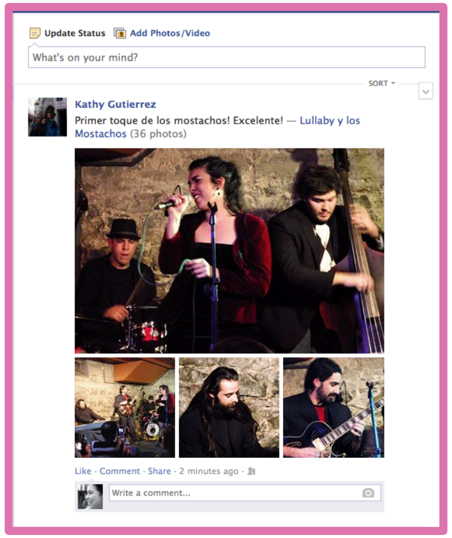

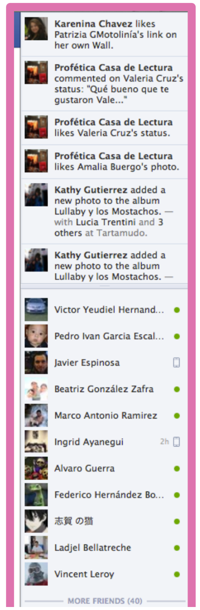


```
SELECT   name
FROM     friendlist
WHERE    owner = me()
```

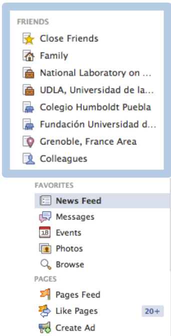


```
SELECT   name
FROM     group
WHERE    gid IN ( SELECT   gid
                  FROM     group_member
                  WHERE    uid = me() )
```

```
SELECT   name, pic
FROM     user
WHERE    online_presence = "active"
         AND
         uid IN ( SELECT   uid2
                  FROM     friend
                  WHERE    uid1 = me() )
```

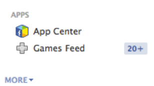

https://developers.facebook.com/docs/reference/fql/

<805114856

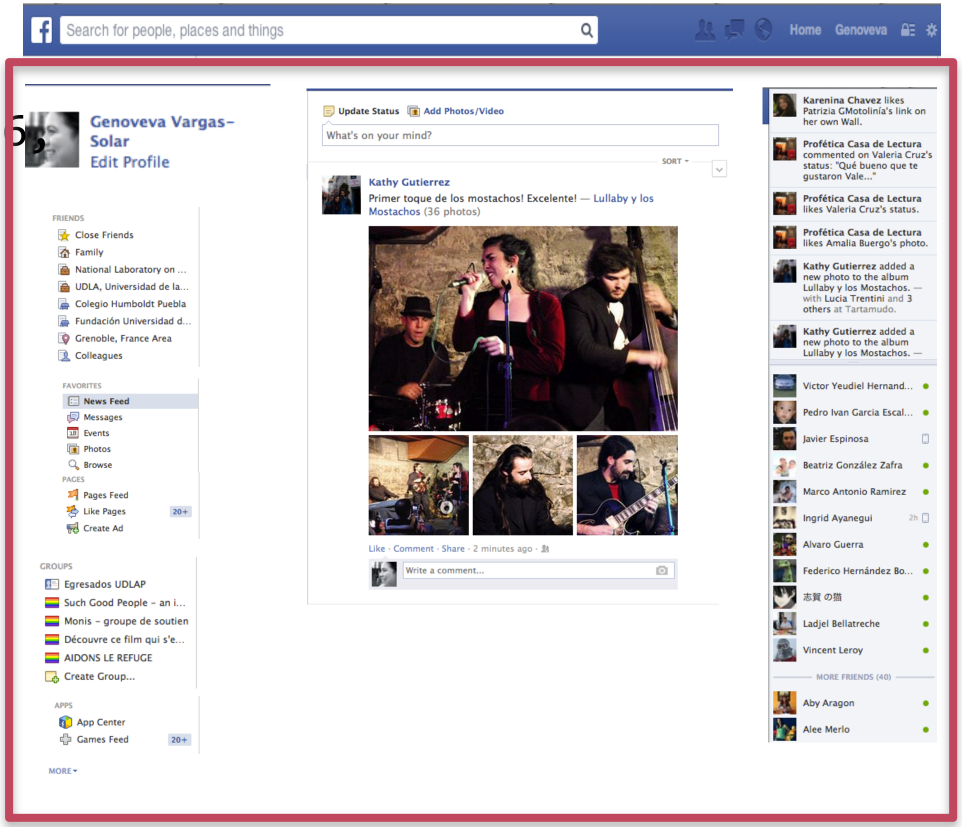

Search for people, places and things
Home
Genoveva
Genoveva Vargas-Solar
Edit Profile
FRIENDS
Close Friends
Family
National Laboratory on ...
UDLA, Universidad de la...
Colegio Humboldt Puebla
Fundación Universidad d...
Grenoble, France Area
Colleagues
FAVORITES
News Feed
Messages
Events
Photos
Browse
PAGES
Pages Feed
Like Pages
20+
Create Ad
GROUPS
Egresados UDLAP
Such Good People – an i...
Monis – groupe de soutien
Découvre ce film qui s'e...
AIDONS LE REFUGE
Create Group...
APPS
App Center
Games Feed
20+
MORE
Update Status
Add Photos/Video
What's on your mind?
SORT
Kathy Gutierrez
Primer toque de los mostachos! Excelente! — Lullaby y los Mostachos (36 photos)
Like · Comment · Share · 2 minutes ago
Write a comment...
Karenina Chavez likes Patrizia GMotolinía's link on her own Wall.
Profética Casa de Lectura commented on Valeria Cruz's status: "Qué bueno que te gustaron Vale..."
Profética Casa de Lectura likes Valeria Cruz's status.
Profética Casa de Lectura likes Amalia Buergo's photo.
Kathy Gutierrez added a new photo to the album Lullaby y los Mostachos. — with Lucia Trentini and 3 others at Tartamudo.
Kathy Gutierrez added a new photo to the album Lullaby y los Mostachos. —
Victor Yeudiel Hernand...
Pedro Ivan Garcia Escal...
Javier Espinosa
Beatriz González Zafra
Marco Antonio Ramirez
Ingrid Ayanegui
Alvaro Guerra
Federico Hernández Bo...
Ladjel Bellatreche
Vincent Leroy
MORE FRIENDS (40)
Aby Aragon
Alee Merlo

>

Document stores

- Support more complex data: pointerless objects, i.e., documents
- Secondary indexes, multiple types of documents (objects) per database, nested documents and lists, e.g. B-trees
- Automatic sharding (scale writes), no explicit locks, weaker concurrency (eventual for scaling reads) and atomicity properties
- API: `select, delete, getAttributes, putAttributes` on documents
- Queries can be distributed in parallel over multiple nodes using a map-reduce mechanism

SYSTEM	ADDRESS
SimpleDB	`amazon.com/simpledb`
Couch DB	`couchdb.apache.org`
Mongo DB	`mongodb.org`
Terrastore	`code.google.com/terrastore`

```
{
  "name": "Genoveva Vargas-Solar",
  "id": "805114856"
}
```



```
{
  "data": [
    {
      "name": "Genoveva Vargas-Solar",
      "pic": "https://fbcdn-profile-a.akamaihd.net/hprofile-ak-ash4/275915_805114856_16986061_s.jpg",
      "profile_url": "https://www.facebook.com/genoveva.vargas"
    }
  ]
}
```


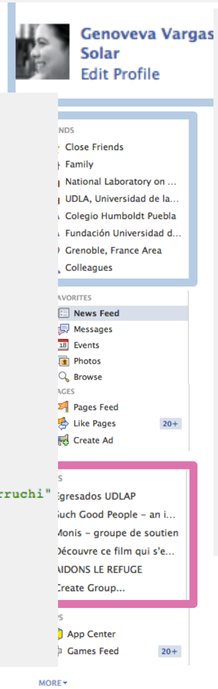


```
{
  "data": [
    {
      "name": "$$$ Se Vende Jeep Compass 2008 - 60,000kms. $$$"
    },
    {
      "name": "Découvre ce film qui s'engage pour le mariage pour tous"
    },
    {
      "name": "emepink"
    },
    {
      "name": "Such Good People - an indie screwball comedy"
    },
    {
      "name": "Comunidad Mexicana de Tecnologías Semánticas"
    },
    {
      "name": "TI-502 Administración de Datos"
    },
    {
      "name": "exaUDLAP Sistemas Computacionales"
    },
    {
      "name": "\"Hombre Nuevo\" artículos de valores humanos del P. Otaolaurruchi"
    },
    {
      "name": "LACCIR"
    },
    {
      "name": "Monis - groupe de soutien"
    },
    {
      "name": "Red Temática de las TIC"
    },
```



```
{
  "data": [
    {
      "message": "",
      "attachment": {
        "media": [
          {
            "href": "https://www.facebook.com/photo.php?fbid=10151871935952502&set=a.99396912501.109184.98871212501&type=1",
            "alt": "",
            "type": "photo",
            "src": "https://fbcdn-photos-e-a.akamaihd.net/hphotos-ak-ash3/1146527_10151871935952502_258686255_s.jpg",
            "photo": {
              "aid": "98871212501_109184",
              "pid": "98871212501_1073742168",
              "fbid": "10151871935952502",
              "owner": 98871212501,
              "index": 1,
              "width": 611,
              "height": 458,
              "images": [
                {
                  "src": "https://fbcdn-photos-e-a.akamaihd.net/hphotos-ak-ash3/1146527_10151871935952502_258686255_s.jpg",
                  "width": 130,
                  "height": 97
                }
              ]
            }
          }
        ],
        "name": "Timeline Photos",
        "href": "https://www.facebook.com/album.php?fbid=99396912501&id=98871212501&aid=109184",
        "caption": "El sutil arte de cantinflear.\r\n\r\nvía - Lectura Cinematográfica",
        "description": "",
```


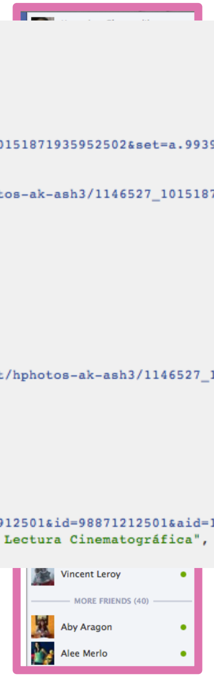

Extensible record stores

- Basic data model is rows and columns
- Basic scalability model is splitting rows and columns over multiple nodes
  - Rows split across nodes through sharding on the primary key
    - Split by range rather than hash function
    - Rows analogous to documents: variable number of attributes, attribute names must be unique
    - Grouped into collections (tables)
    - Queries on ranges of values do not go to every node
- Columns are distributed over multiple nodes using “column groups”
  - Which columns are best stored together
  - Column groups must be pre-defined with the extensible record stores

SYSTEM	ADDRESS
HBase	`hbase.apache.com`
HyperTable	`hypertable.org`
Cassandra	`incubator.apache.org/cassandra`

Scalable relational systems

- SQL: rich declarative query language
- Databases reinforce referential integrity
- ACID semantics
- Well understood operations:
  - Configuration, Care and feeding, Backups, Tuning, Failure and recovery, Performance characteristics
- Use small-scope operations
  - Challenge: joins that do not scale with sharding
- Use small-scope transactions
  - ACID transactions inefficient with communication and 2PC overhead
- Shared nothing architecture for scalability
- Avoid cross-node operations

SYSTEM	ADDRESS
MySQL C	`mysql.com/cluster`
Volt DB	`voltdb.com`
Clustrix	`clustrix.com`
ScaleDB	`scaledb.com`
Scale Base	`scalebase.com`
Nimbus DB	`nimbusdb.com`