The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

Advanced Database System Architectures

Advanced Topics in Database Management (INFSCI 2711)

Distributed Databases (TELCOM 2326)

Textbook: Database System Concepts - 6th Edition, 2010

Vladimir Zadorozhny, GIST, University of Pittsburgh

Where we are now: Centralized Systems

- Run on a single computer system and do not interact with other computer systems.
- General-purpose computer system: one to a few CPUs and a number of device controllers that are connected through a common bus that provides access to shared memory.
- Single-user system (e.g., personal computer or workstation): desk-top unit, single user, usually has only one CPU and one or two hard disks; the OS may support only one user.
- Multi-user system: more disks, more memory, multiple CPUs, and a multi-user OS. Serve a large number of users who are connected to the system vie terminals. Often called *server* systems.

A Centralized Computer System

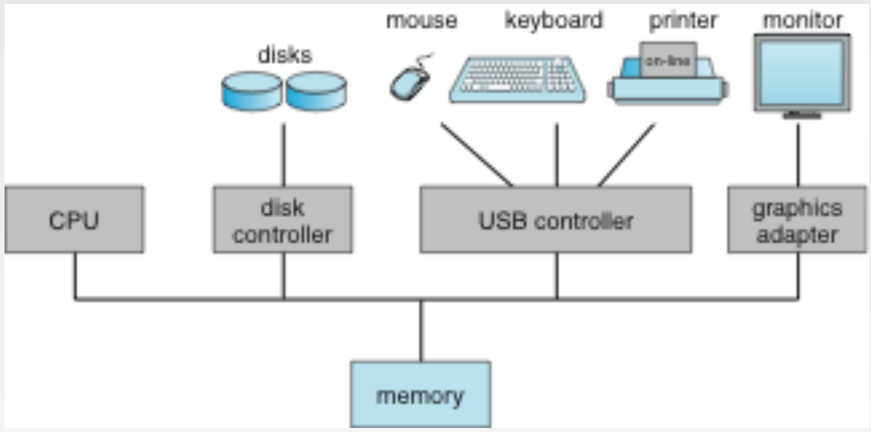

Next: Client-Server Systems

- Server systems satisfy requests generated at *m* client systems:

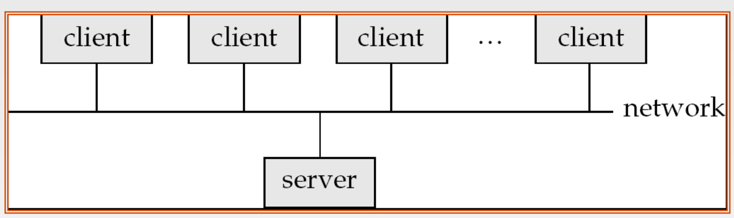

Client-Server Systems (Cont.)

- Database functionality can be divided into:
  - **Back-end**: manages access structures, query evaluation and optimization, concurrency control and recovery.
  - **Front-end**: consists of tools such as *forms*, *report-writers*, and graphical user interface facilities.
- The interface between the front-end and the back-end is through SQL or through an application program interface.

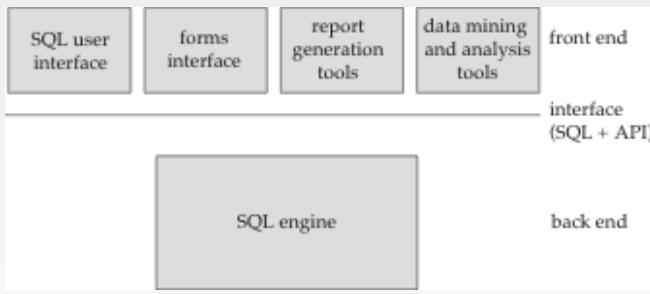

Server System Architecture

- Server systems can be broadly categorized into two kinds:
  - **transaction servers** which are widely used in relational database systems, and
  - **data servers**, used in object-oriented database systems

Transaction Servers

- Also called **query server** systems or SQL *server* systems
  - Clients send requests to the server
  - Transactions are executed at the server
  - Results are shipped back to the client.
- *Open Database Connectivity* (ODBC) is a C language application program interface standard from Microsoft for connecting to a server, sending SQL requests, and receiving results.
- JDBC standard is similar to ODBC, for Java

Data Servers

- Data are shipped to clients where processing is performed.
- This architecture requires full back-end functionality at the clients.
- Used in many object-oriented database systems
- Issues:
  - Page-Shipping versus Item-Shipping (tuple, or object)
  - Locking
  - Data Caching

Next: Distributed Systems

- Data spread over multiple machines (also referred to as **sites** or **nodes**).
- Network interconnects the machines
- Data shared by users on multiple machines

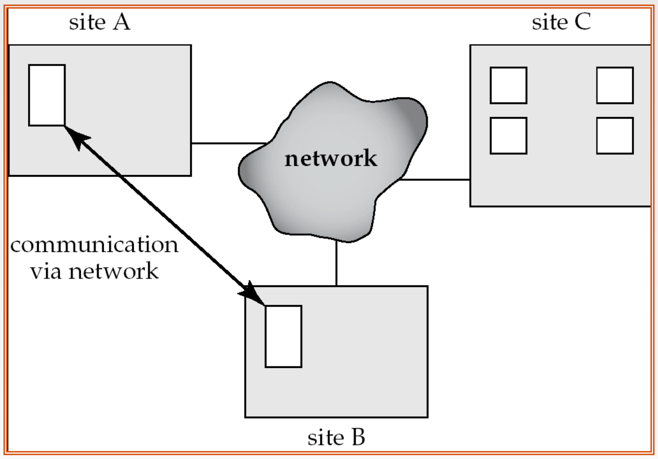

Distributed Databases

- Homogeneous distributed databases
  - Same software/schema on all sites, data may be partitioned among sites
  - Goal: provide a view of a single database, hiding details of distribution
- Heterogeneous distributed databases
  - Different software/schema on different sites
  - Goal: integrate existing databases to provide useful functionality
- Differentiate between *local* and *global* transactions
  - A local transaction accesses data in the *single* site at which the transaction was initiated.
  - A global transaction either accesses data in a site different from the one at which the transaction was initiated or accesses data in several different sites.

Trade-offs in Distributed Systems

- Sharing data – users at one site able to access the data residing at some other sites.
- Autonomy – each site is able to retain a degree of control over data stored locally.
- Higher system availability through redundancy — data can be replicated at remote sites, and system can function even if a site fails.
- Disadvantage: added complexity required to ensure proper coordination among sites.
  - Software development cost.
  - Greater potential for bugs.
  - Increased processing overhead.

Heterogeneous Distributed Databases

- Different software/schema on different sites
- Goal: integrate existing databases to provide useful functionality

Querying Web Data from a DB Perspective

- Manual navigation over multilevel links: inefficient

  Find the top selling book on C++ at Amazon ?

- Objective: database-like declarative queries:

  select bookTitle
  from Amazon
  where bookTopic = "C++" and
  bookSalesRank > all (select bookSalesRank
  from Amazon
  where bookTopic = "C++")

Structured vs Unstructured Data: 1996-2006

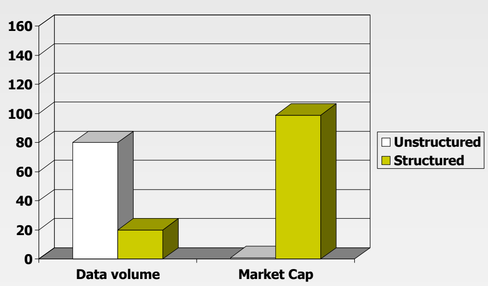

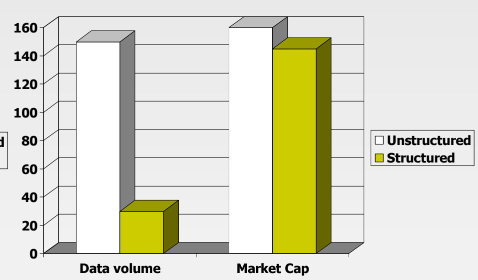

Database Management vs Information Retrieval

- Data: DB: Set of Tables with well defined schema
  IR: Set of (text) documents
- Goal: DB: Find an accurate response to a user query
  IR: Retrieve documents with information that is relevant to user's information need

15

Querying unstructured data

- Which plays of Shakespeare contain the words ***Brutus*** *AND* ***Caesar*** but *NOT* ***Calpurnia***?
  - One could grep all of Shakespeare's plays for ***Brutus*** and ***Caesar,*** then strip out lines containing ***Calpurnia***?
    - Slow (for large corpora)
    - *NOT* ***Calpurnia*** is non-trivial
    - Other operations (e.g., find the word ***Romans*** near ***countrymen***) not feasible
    - Ranked retrieval (best documents to return)

16

Term-document incidence matrix

	Antony and Cleopatra	Julius Caesar	The Tempest	Hamlet	Othello	Macbeth
Antony	1	1	0	0	0	1
Brutus	1	1	0	1	0	0
Caesar	1	1	0	1	1	1
Calpurnia	0	1	0	0	0	0
Cleopatra	1	0	0	0	0	0
mercy	1	0	1	1	1	1
worser	1	0	1	1	1	0

1 if play contains word, 0 otherwise

Brutus *AND* ***Caesar*** but *NOT* ***Calpurnia***

Query evaluation and optimization

- 0/1 vector for each term.
- To answer query: take the vectors for ***Brutus, Caesar*** and ***Calpurnia*** (complemented) ➔ bitwise *AND*.
- 110100 *AND* 110111 *AND* 101111 = 100100.

•Consider: 1M documents, each with about 1K terms.
- •6GB of data in the documents (avg 6 bytes/term incl spaces/punctuation)
- •Assume 500K <u>*distinct*</u> terms.

500K x 1M matrix has half-a-trillion 0's and 1's.

18

Information Integration from a DB Perspective

- Information Integration Challenge
  - *Given*: data sources S_1, ..., S_k (DBMS, web sites, ...) and user questions Q_1,...,Q_n that can be answered using the S_i
  - *Find*: the answers to Q_1, ..., Q_n
- The Database Perspective: source = "database"
  - ⇒ S_i has a schema
  - ⇒ S_i can be queried
  - ⇒ define virtual (or materialized) integrated views V over S_1,...,S_k using database query languages
  - ⇒ questions become queries Q_i against V(S_1,...,S_k)

Data Warehousing

- Integrated data spanning long time periods, often augmented with summary information.
- Several gigabytes to terabytes common.
- Interactive response times expected for complex queries; ad-hoc updates uncommon.

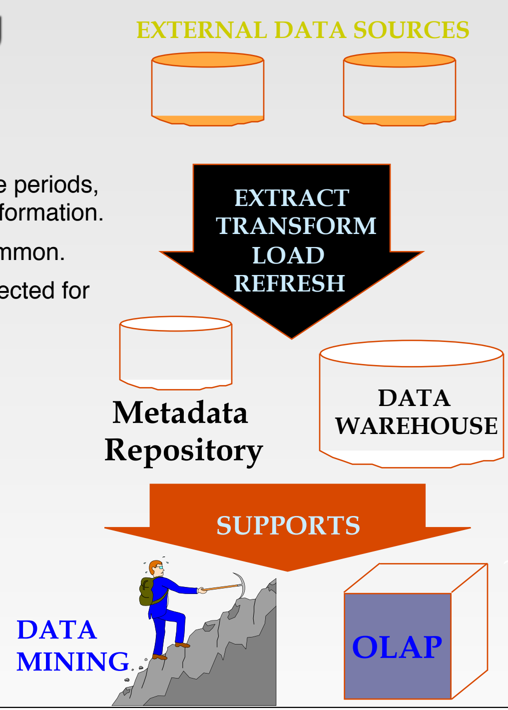

What Next?

More challenging network environments …

Wireless Sensors

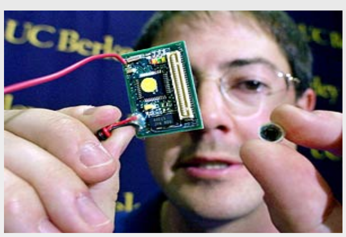

Courtesy: http://www.economist.com

- Small wireless devices (motes)
- Low cost, battery powered
- Sense physical phenomena
  - Light, temperature, vibration, acceleration, AC power, humidity.
- Process/aggregate data
- Communicate

Applications of Wireless Sensor Networks:

Information tracking systems (e.g., airport security);
Children monitoring in metro areas;
Product transition in warehouse networks;
Fine-grained weather measurements;
Structural Health Monitoring

Sensor Devices

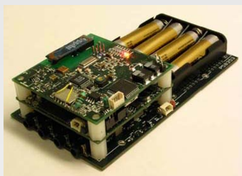

µAMPS-1
(MIT)

MICA-2 and MICA2Dot
(UC Berkeley, Crossbow)

Medusa MK-2
(UCLA)

Compare …

RAM: 4-128 KB
Flash: 32KB - 1 MB
Running at: 4 – 40 MHz

RAM: GBs
Disk: 100s of GBs
Running at: GHzs

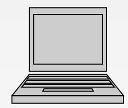

Wireless Network Standards

IEEE 802.15: wireless personal area networks

Suitable for low power and low data rate sensor networks

- 802.15.4: data rates: 20, 40 or 250 kbps in the 868, 915 or 2400 MHz bands respectively.
- Typical transmit power: 1 – 2 mW.

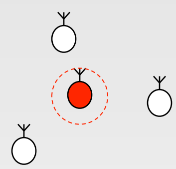

Compare …

IEEE 802.11: wireless local area networks

- 802.11b: 11 Mbps, operate at 2.4 GHz
- 802.11g: 54 Mbps, operate at 2.4 GHz
- 802.11a: 54 Mbps, operates at 5 GHz
- Typical transmit power: 250 mW.

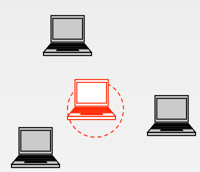

What to expect: 802.15.3 ultrawideband (UWB),

- very high data rates (110 - 400 Mbps) over short ranges,
- low power (order of) 0.1 mW;

<u>*Products are not yet available*</u>.

Sensor Databases_

Concept is Simple and Attractive

SELECT avg(*rainFallLevel*)
FROM Sensors;

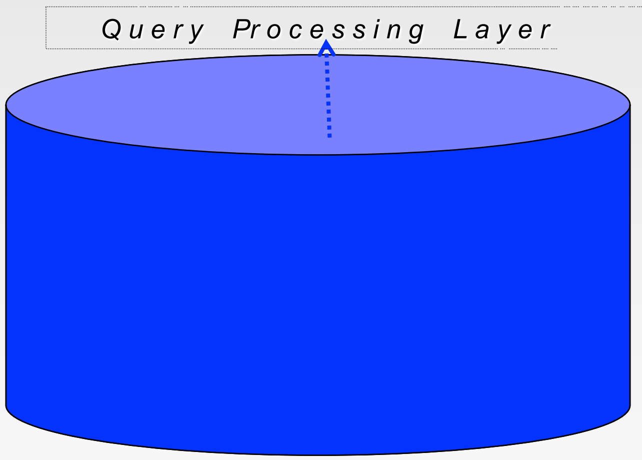

Network is a Database !

Network is a Blackbox ?

Database Query Evaluation Revisited

select A, B
from *r* ,*s*
where r.B=s.B

The image cannot be displayed. Your computer may not have enough memory to open the image, or the image may have been corrupted. Restart your computer, and then open the file again. If the red x still appears, you may have to delete the image and then insert it again.

$r \bowtie s$

Nested Loop Join:

for each tuple t_r **in** *r* **do begin**
  for each tuple t_s **in** *s* **do begin**
    test pair (t_r, t_s) if they satisfy the join condition
    if they do, add $t_r \cdot t_s$ to the *result*.
  end
end

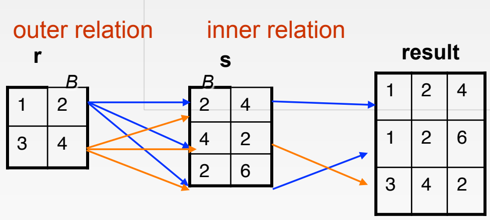

SQLQuery

Relational Algebra

Physical Query Execution Plan

Query Execution

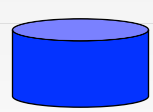

Disk Storage

Traditional Database Query Processing

select A, B
from *r* ,*s*
where r.B=s.B

SQLQuery

DBMS

Number of tuples in relations *r, s:* n_r =5000, n_s =1000
Number of disk blocks containing tuples of *r, s:* b_r =100, b_s = 400

r is outer:

$Cost(r \bowtie s) = n_r * b_s + b_r$

= 5000 * 400 + 100 = 2,000,100 disk accesses

s is outer:

$Cost(s \bowtie r) = n_s * b_r + b_s$

= 1000 * 100 + 400 = **1,000,400** disk accesses

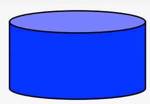

Disk Storage

Sensor Database Query Processing ?

SELECT *
FROM Sensors

SQLQuery

DBMS

?

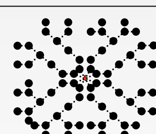

Sensor Network

Mobile Sensors

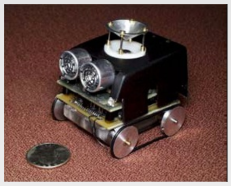

Milibot (CMU)

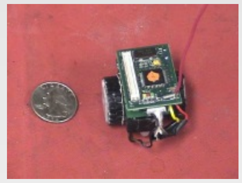

Robomote (USC)

COTS-BOTS
*(*Bergbreiter&Pister*)*

"Explore, discover, report ..."
Surveyor Wireless Mobile Robots:
video for telepresence;
autonomous and swarm Operation.

Mobility: Cool Applications

- E.g., a team of cooperative mobile robots can be considered as a wireless sensornet.
- Deployed in conjunction with stationary sensor nodes
- Acquire and process data for surveillance, tracking, environmental monitoring, or execute search and rescue operations.

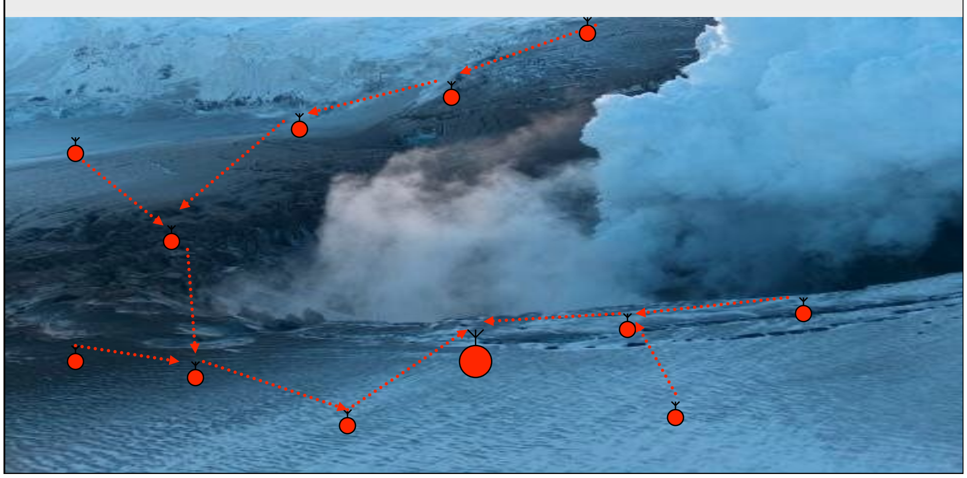

Application #2

• Large-scale human health monitoring with body sensors reporting critical health parameters (e.g., blood pressure) to a processing station.

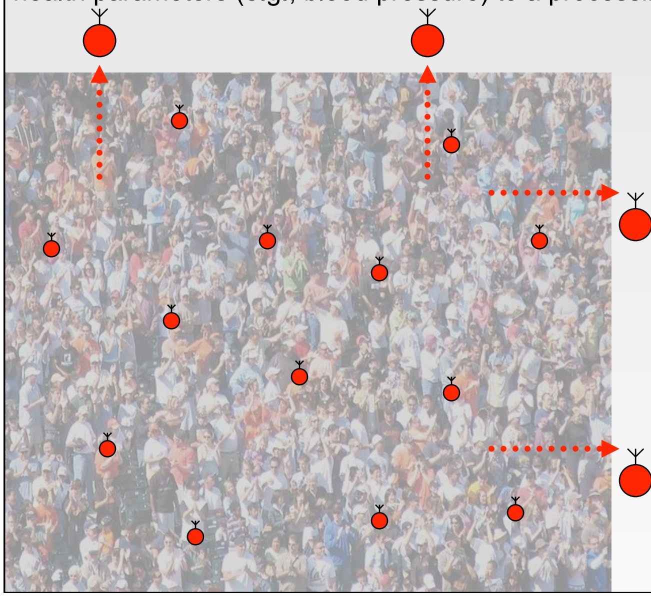

•More complicated version: monitoring the health of soldiers in a battlefield.

• Robomotes could deliver the information to a remote center that is away from the battlefield.

Application #3

• Discovering traffic conditions.

• Assuming that each vehicle is provided with a group of sensors that reports its local parameters (e.g., speed) and surrounding condition (e.g., snow, icy road, etc.).

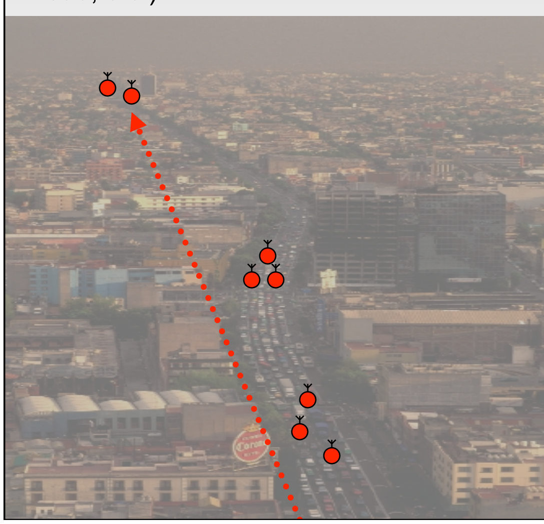

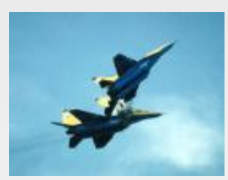

A complicated case: battlefield reports and extra speed (e.g., a swarm of jets).

Mobile Database Query Processing?

SELECT *Environmental_Conditions*
FROM Sensors

SQLQuery

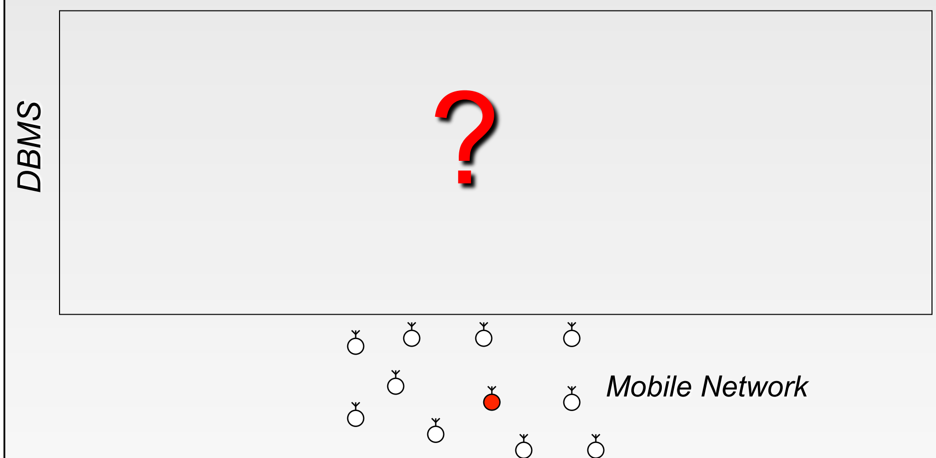

What Next?

Big Data Challenge

" … Researchers in a growing number of fields are generating extremely large and complicated data sets commonly referred to as "Big Data."

A wealth of information may be found within these sets with enormous potential to shed light on some of the toughest and most pressing challenges facing the nation.

To capitalize on this unprecedented opportunity to extract insights and make new connections across disciplines we need better tools and programs to access, store, search, visualize and analyze these data.

To maximize this historic opportunity -- and in support of recommendations from the President's Council of Advisors on Science and Technology -- the Obama Administration is launching a Big Data Research and Development Initiative, coordinated by the White House Office of Science and Technology Policy and supported by several federal departments and agencies. …"

http://www.cccblog.org/2012/03/27/white-house-to-announce-big-data-rlive-webcast-planned/

From Common Big Data Perception to 3D Big Data Challenge

- **Size:** 5PB (Facebook) , 6 PB (eBay): **STORAGE?**
- **Rate:** billions of records per day, billions of search impressions per day: **UTILIZATION ?**
- **Diversity:** formats, conflicting contents,… : **DATA INTEGRATION, FUSION, AND SENCEMAKING ?**

Challenge of Data Size

- 30 billion pieces of content shared on Facebook every month
- 40% projected growth in global data generated per year vs. 5% growth in global IT spending
- 235 terabytes data collected by the US Library of Congress by April 2011
- 15 out of 17 sectors in the United States have more data stored per company than the US Library of Congress
- MGI estimates that enterprises globally stored more than 7 exabytes of new data on disk drives in 2010, while consumers stored more than 6 exabytes of new data on devices such as PCs and notebooks.
- One exabyte of data is the equivalent of more than 4,000 times the information stored in the US Library of Congress

Big data: The next frontier for innovation, competition, and productivity. McKinsey Global Institute, June 2011

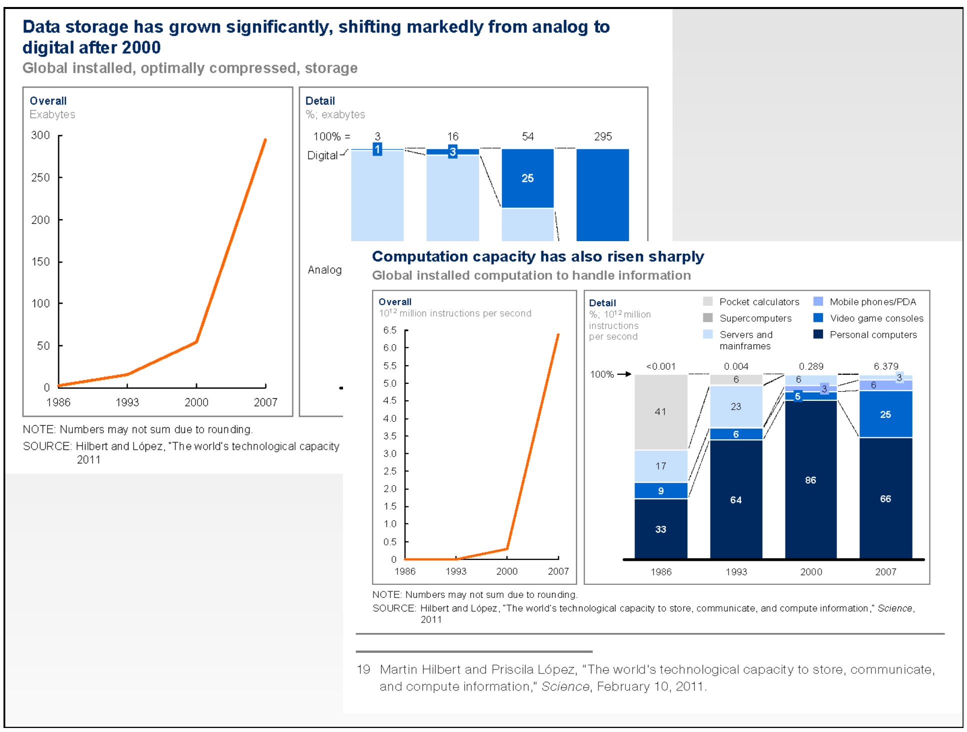

Amount of new data stored varies across geography

New data stored[1] by geography, 2010

Petabytes

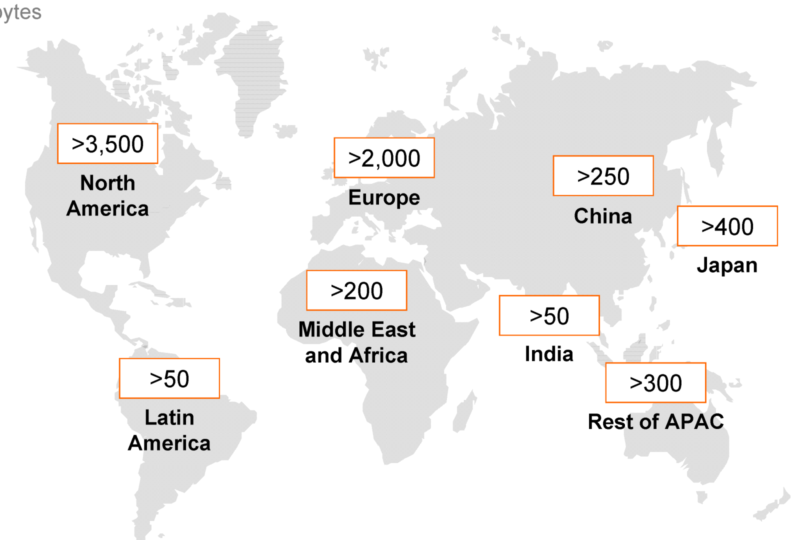

1 New data stored defined as the amount of available storage used in a given year; see appendix for more on the definition and assumptions.

SOURCE: IDC storage reports; McKinsey Global Institute analysis

Challenge of Data Rate

Across sectors and regions, several cross-cutting trends have fueled growth in data generation and will continue to propel the rapidly expanding pools of data.

These trends include growth in traditional transactional databases, continued expansion of multimedia content, increasing popularity of social media, and proliferation of applications of sensors in the Internet of Things.

Multimedia data already accounts for more than half of Internet backbone traffic and this share is expected to grow to 70 percent by 2013.

We are generating so much data today that it is physically impossible to store it all.

Health care providers, for instance, discard 90 percent of the data that they generate (e.g., almost all real-time video feeds created during surgery).

Case Study: CHP Bedside Medical Device Interface (BMDI)

- Automate data acquisition of device data at patient bedside *(e.g., blood pressure, heart rate, etc.).*
- Data automatically acquired via interfaces running on bedside PC or BMDI server
- Data written to an Oracle database replicated on network file servers
- Results reviewed/validated by CHP Critical Care Clinician and then written to patient's permanent medical record in the database

Device (Ventilator) Interface

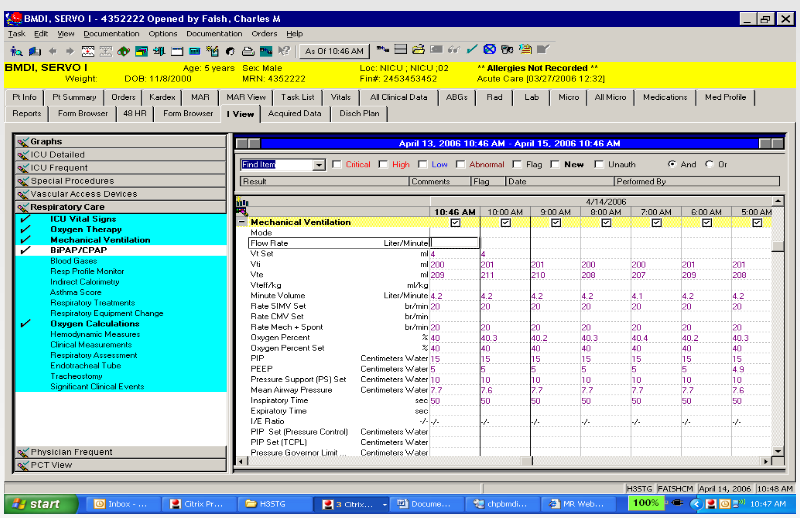

BMDI Infrastructure collects a lot of data points, but …

But … Data Utility is Low

•Amount of hospital operational data is considerable
50 Patients x 100 Data Parameters Per Minute
Equals 5000 Data Parameters/Minute or
300,000 Data Parameters/Hour or
7,200,000 Data Parameters/Day
2,160,000,000 Data Parameters/Month

•CHP expects to review (chart) no more than 20% of these Parameters Daily. This is only
1,440,000 Data Parameters/Day or
432,000,000 Data Parameters/Month

Do they chart most relevant data?

Challenge: more efficient way to utilize such data

Big Data: From In the Way to On the Way

Large Scale HC Infrastructures produce data flood ☹

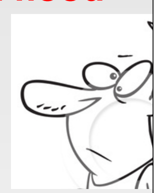

Large Scale HC Infrastructures should help utilize the data ☺

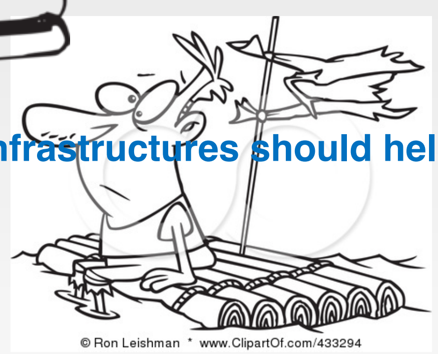

Major Challenge:
XLarge-Scale Human-Computer Synergy

Data Flood:
crowds are to blame

Global Data Refinery
Data Sensemaking
Data Curation
Data Reliability Assessment
Data Fusion
Data Integration
Data Storage
Data Collection

Data Utilization:
enlist crowds to he

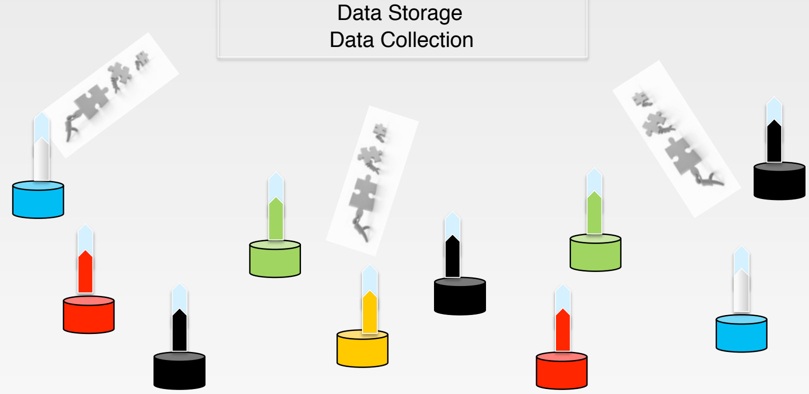

Big Picture:

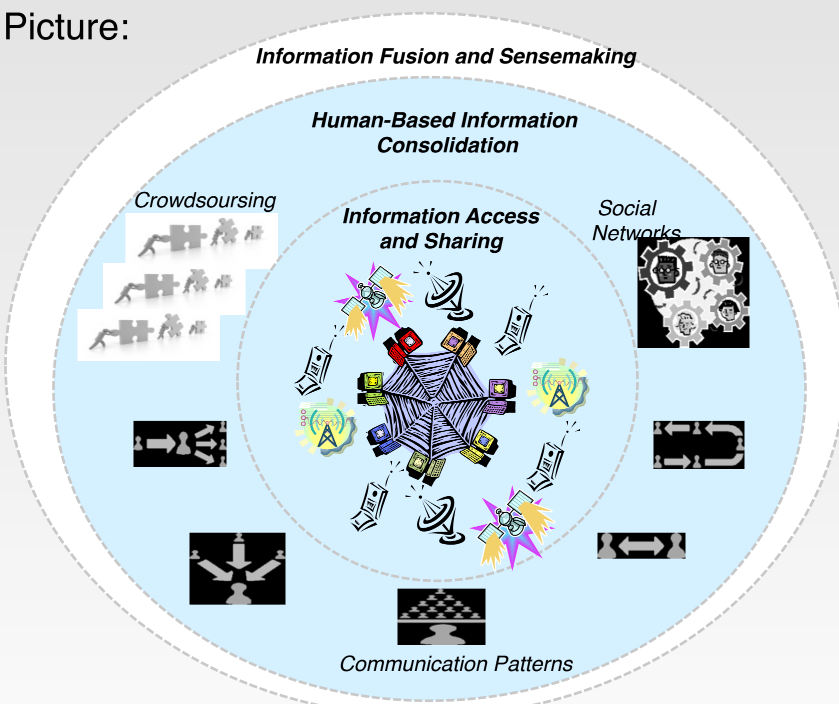

Big Data—capturing its value

- 60% potential increase in retailers' operating margins possible with big data
- 140,000–190,000 more deep analytical talent positions
- 1.5 million more data-savvy managers needed to take full advantage of big data in the United States

Big data: The next frontier for innovation, competition, and productivity. McKinsey Global Institute, June 2011

47

What Next?

Back to schedule …