

# Networking in the Big Data Era

Nelson L. S. da Fonseca

Institute of Computing, State University of Campinas, Brazil

International Teletraffic Congress 27

Ghent, Belgium

Sep 10<sup>th</sup>, 2015

# Outline

- Big Data
- Networking for Big data
- Big Data for Networking
- Some Research Problems

What are big data?



# What Happens in an Internet Minute?




## And Future Growth is Staggering



# Big Data Size: The Volume Of Data Continues To Explode

The Digital Universe 2009 - 2020






"If a tree falls in a forest and no one is around to hear it, does it make a sound?" --- George Berkeley

**Data ≠ Knowledge**

# Processing Big Data






**1%**

of the world's data is analyzed  
today



# Big Data



# Volume

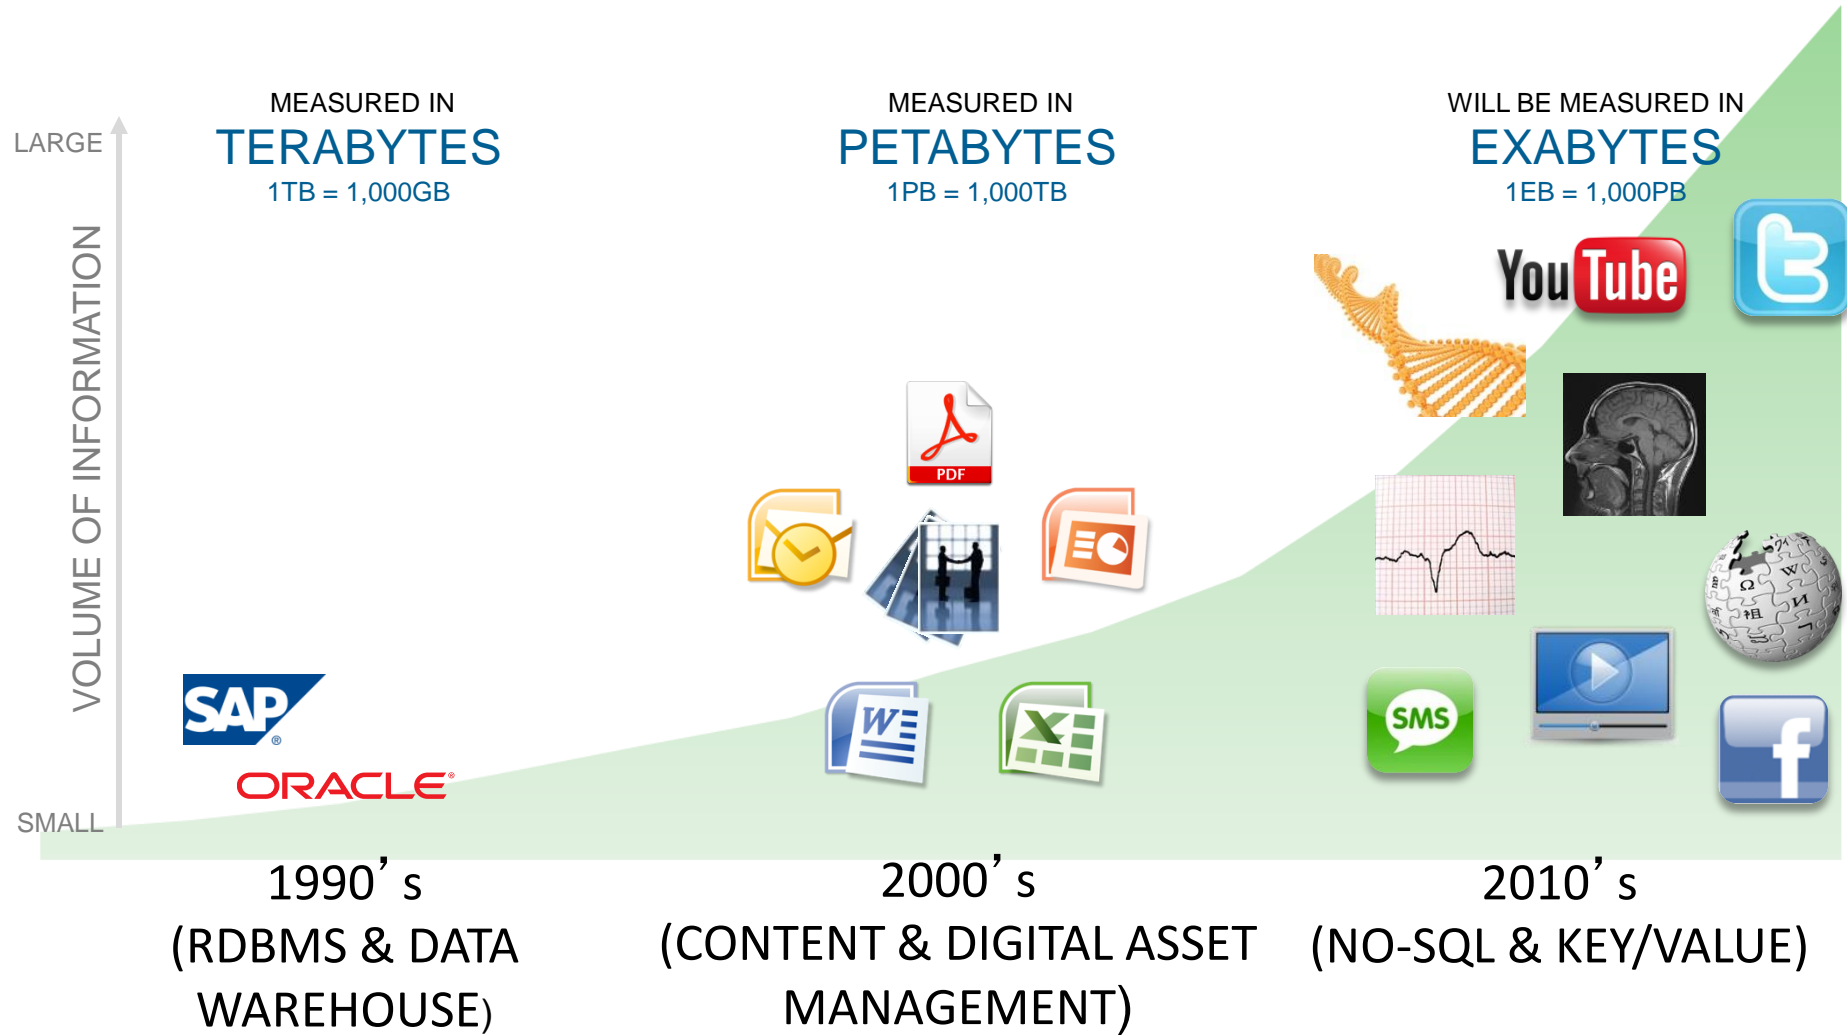
- Data set to be processed at a time is too large
- Data set is not too large but the collection of data set is large
- Volume of data set too large per se, but processing is time consuming perhaps due to too many IO operations



## Present of Big Data

Too big to handle

# New Applications Driving Data Volume



# Velocity

- Data arrive is faster than the processing capacity
- Results must be produced with certain delay bound, processing is limited by disk I/O throughput



## Future of Big Data

Drinking from a firehose

# Variety


## Structured Data

SUMMER FOOD SERVICE PROGRAM 1]				
(Data as of August 01, 2011)				
Fiscal Year	Number of Sites	Peak (July) Participation	Meals Served	Total Federal Expenditures 2]
	-----Thousands-----		--Mil.--	--Million \$--
1969	1.2	99	2.2	0.3
1970	1.9	227	8.2	1.8
1971	3.2	569	29.0	8.2
1972	6.5	1,080	73.5	21.9
1973	11.2	1,437	65.4	26.6
1974	10.6	1,403	63.6	33.6
1975	12.0	1,785	84.3	50.3
1976	16.0	2,453	104.8	73.4
TQ 3]	22.4	3,455	198.0	88.9
1977	23.7	2,791	170.4	114.4
1978	22.4	2,333	120.3	100.3
1979	23.0	2,126	121.8	108.6
1980	21.6	1,922	108.2	110.1

## Semi-Structured Data



View → Source

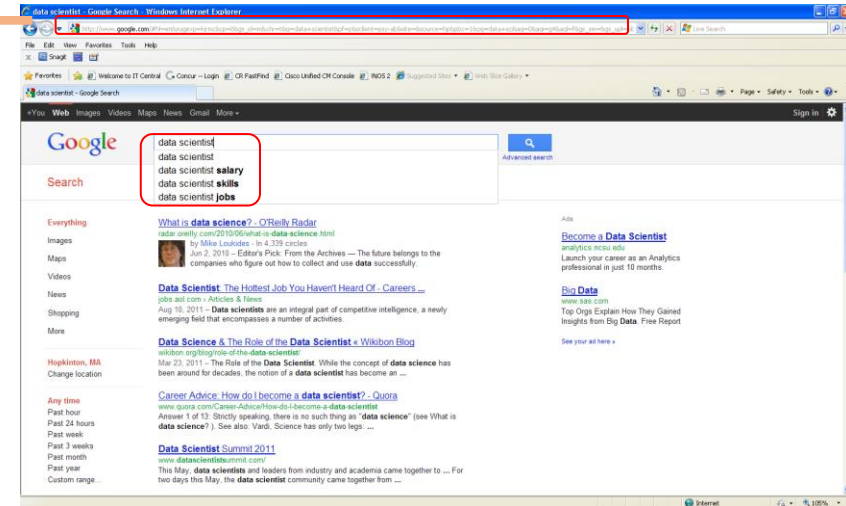


```

1 <!DOCTYPE html PUBLIC "-//W3C//DTD XHTML 1.0 Transitional//EN" "http://www.w3.org/TR/xhtml1/DTD/xhtml1-trans
2 <html xmlns="http://www.w3.org/1999/xhtml">
3
4 <head>
5
6 <meta http-equiv="Content-Type" content="text/html; charset=UTF-8" />
7 <META name="y_key" content="859b402e1c9acc">
8 <link rel="canonical" href="http://www.emc.com/index.htm" />
9 <META NAME="verify-v1" CONTENT="yi2t9VOP4eV0jFdiPeVVIfrP32g4qWFE0I2UvTfMSU
10 <title>EMC - Data Recovery, Cloud Computing, and Storage Hardware</title>
11 <META NAME="description" CONTENT="EMC is a leading provider of storage hardware solutions th
12 data recovery and improve cloud computing." />
13 <META NAME="keywords" CONTENT="emc, network storage, data recovery, information manage
14 software, nas storage, information protection, information management" />
15 <!-- Start: stylesheet includes -->
16 <link rel="stylesheet" href="/_admin/css/styles.css" />
17 <link rel="stylesheet" href="/_admin/css/styles_nav.css" />
18 <!--[if IE]>

```

## Quasi-Structured Data

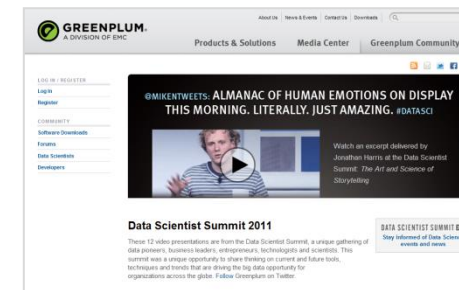


[http://www.google.com/#hl=en&sugexp=kjrmc&cp=8&gs\\_id=2m&xhr=t&q=data+scientist&pq=big+dat&a&pf=p&scient=psyb&source=hp&pbx=1&oq=data+sci&aq=0&aqi=g4&aql=f&gs\\_sm=&gs\\_upl=&bav=on.2,or.r\\_gc\\_r\\_pw,.cf.osb&fp=d566e0fbd09c8604&biw=1382&bih=651](http://www.google.com/#hl=en&sugexp=kjrmc&cp=8&gs_id=2m&xhr=t&q=data+scientist&pq=big+dat&a&pf=p&scient=psyb&source=hp&pbx=1&oq=data+sci&aq=0&aqi=g4&aql=f&gs_sm=&gs_upl=&bav=on.2,or.r_gc_r_pw,.cf.osb&fp=d566e0fbd09c8604&biw=1382&bih=651)


## Unstructured Data

*The Red Wheelbarrow*, by William Carlos Williams

so much depends  
upon  
a red wheel  
barrow  
glazed with rain  
water  
beside the white  
chickens.



# Other V's



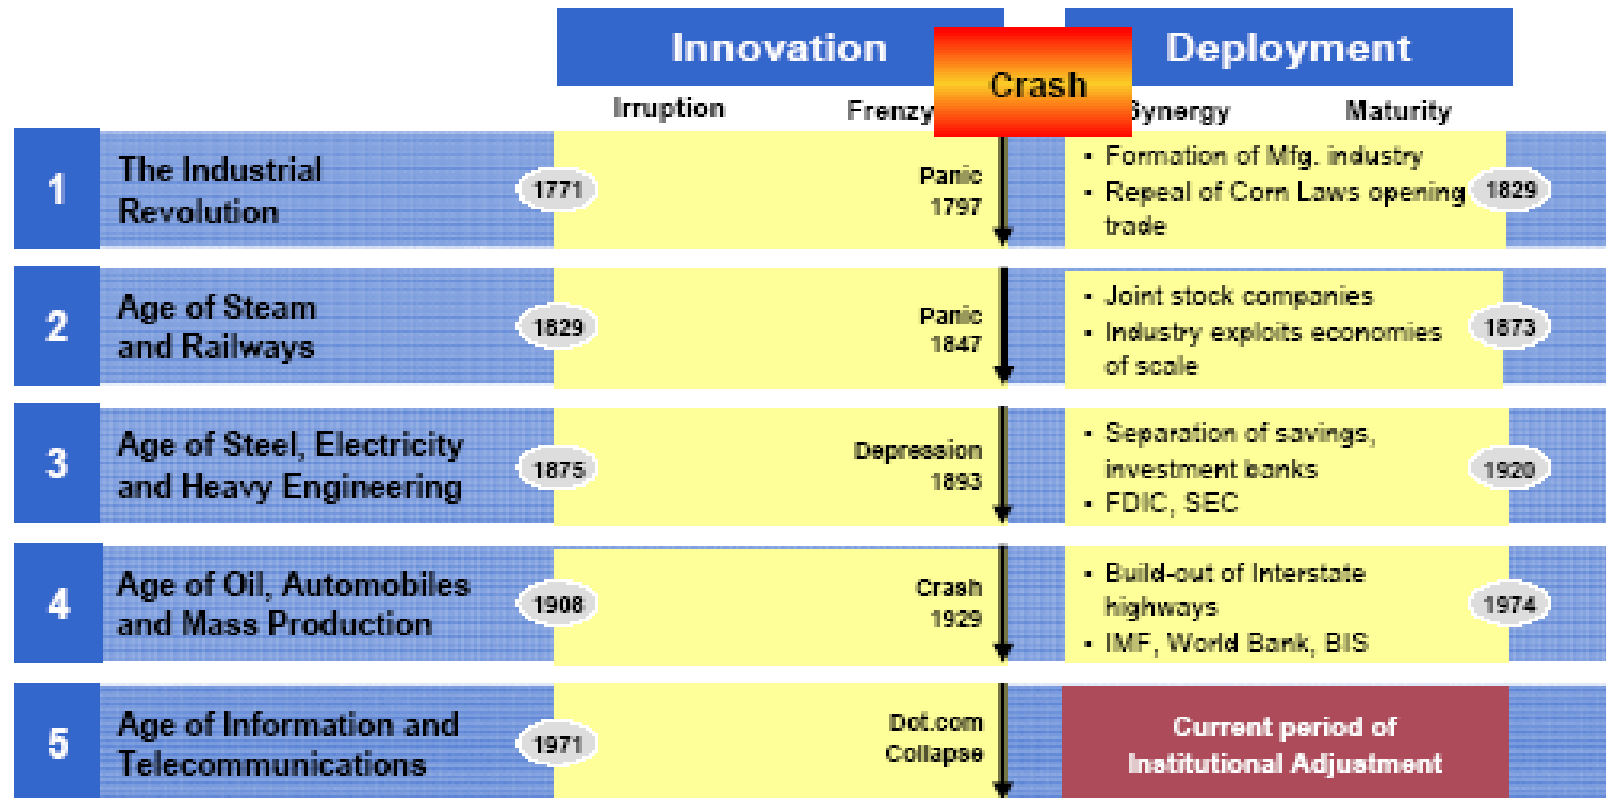
- Valence: Non-trivial inter-relatedness of data
- Veracity: The degree of certainty in data
- Variability: variable interpretations

# Why Big Data?



# Revolutions needs Innovation...

## Five historical cycles ...

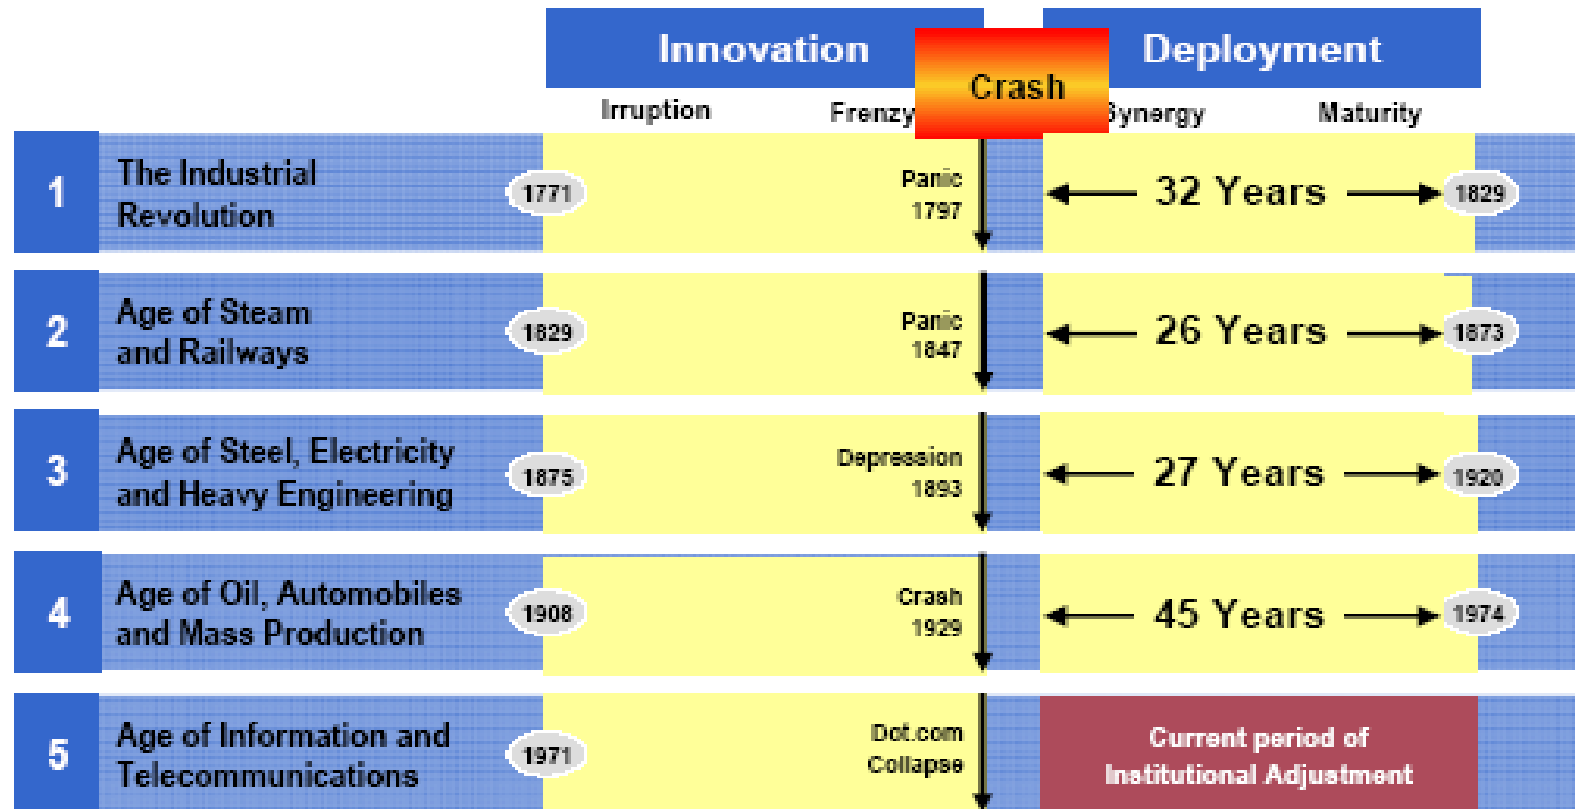


Source: "Technological Revolutions and Financial Capital, Carlota Perez, 2002



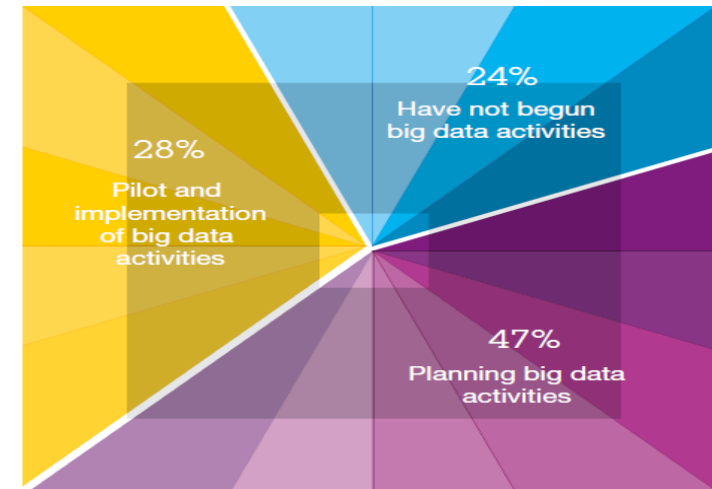
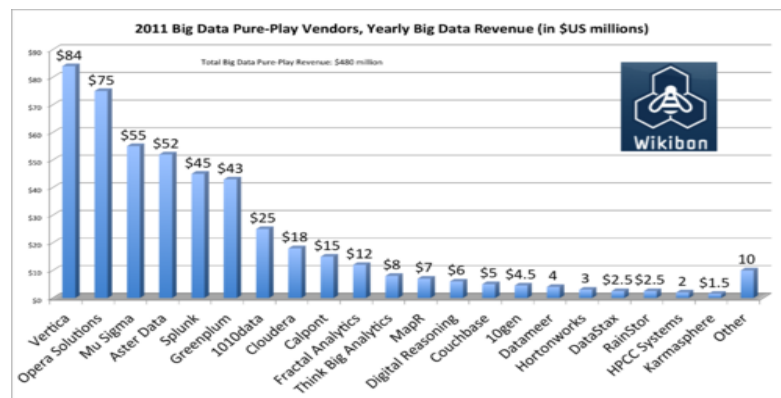

# Takes time to deploy...

The deployment phase lasts 26 to 45 years ...



Source: "Technological Revolutions and Financial Capital, Carlota Perez, 2002

# Big Data and Enterprise



[http://wikibon.org/wiki/v/Big\\_Data\\_Market\\_Size\\_and\\_Vendor\\_Revenues](http://wikibon.org/wiki/v/Big_Data_Market_Size_and_Vendor_Revenues)

# THE BIG DATA LANDSCAPE

JANUARY 2014

## Apps

### Vertical

KNEWTON ellucian practice fusion  
SurveyMonkey PROACTIVE PUBLISHING

### Operational Intelligence

splunk New Relic PopDynamics  
sumologic VITRIA

### Data As A Service

factust FICO Kaggle INRIX  
GNIP Opigive Placed LOGATE  
DATASET TOPSY LexisNexis

### Consumer

Google amazon  
Walmart Labs ebay BETAFUN in

### Ad/Media

METAMARKETS collectiveIQ  
rocketHub FLURBY  
collective DataXu  
Media Science TURN  
bloomreach

### Business Intelligence

ORACLE | Hyperion  
SAP Business Objects IBM  
Microsoft | Business Intelligence  
IBM TIBCO  
pentaho RECOMMIND  
Autonomy blime GoodData  
Chart.io WATTIVO Recorded Future

### Analytics and Visualization

tableau QlikView  
Palantir QUBA TRIFACTA  
Knoema  
centrifuge  
SAS TIBCO  
panopticon UFDRA  
Datameer mindata  
platform Near Story  
alteryx visual.ly AVATA  
metaLayer Adigo Alpine STScore

## Infrastructure

### Analytics

cloudera Hortonworks  
MAER Gobon HARAPT  
Pivotal INFOBRIGHT  
NETEZZA VERTICA  
ORACLE Kognitio

### Operational


COUCHBASE mongoDB  
AEROSPIKE splice  
DATASTRX VoltDB  
TERRACOTTA INFORMATICA  
MarkLogic

### As A Service

Qubole amazon web services  
Windows Azure MORTAR  
CSC Google BigQuery

### Structured DB

ORACLE MySQL  
SQL Server PostgreSQL  
IBM DB2 SYBASE  
memsql TERADATA



Technologies



HBASE



# Big Data and the Government

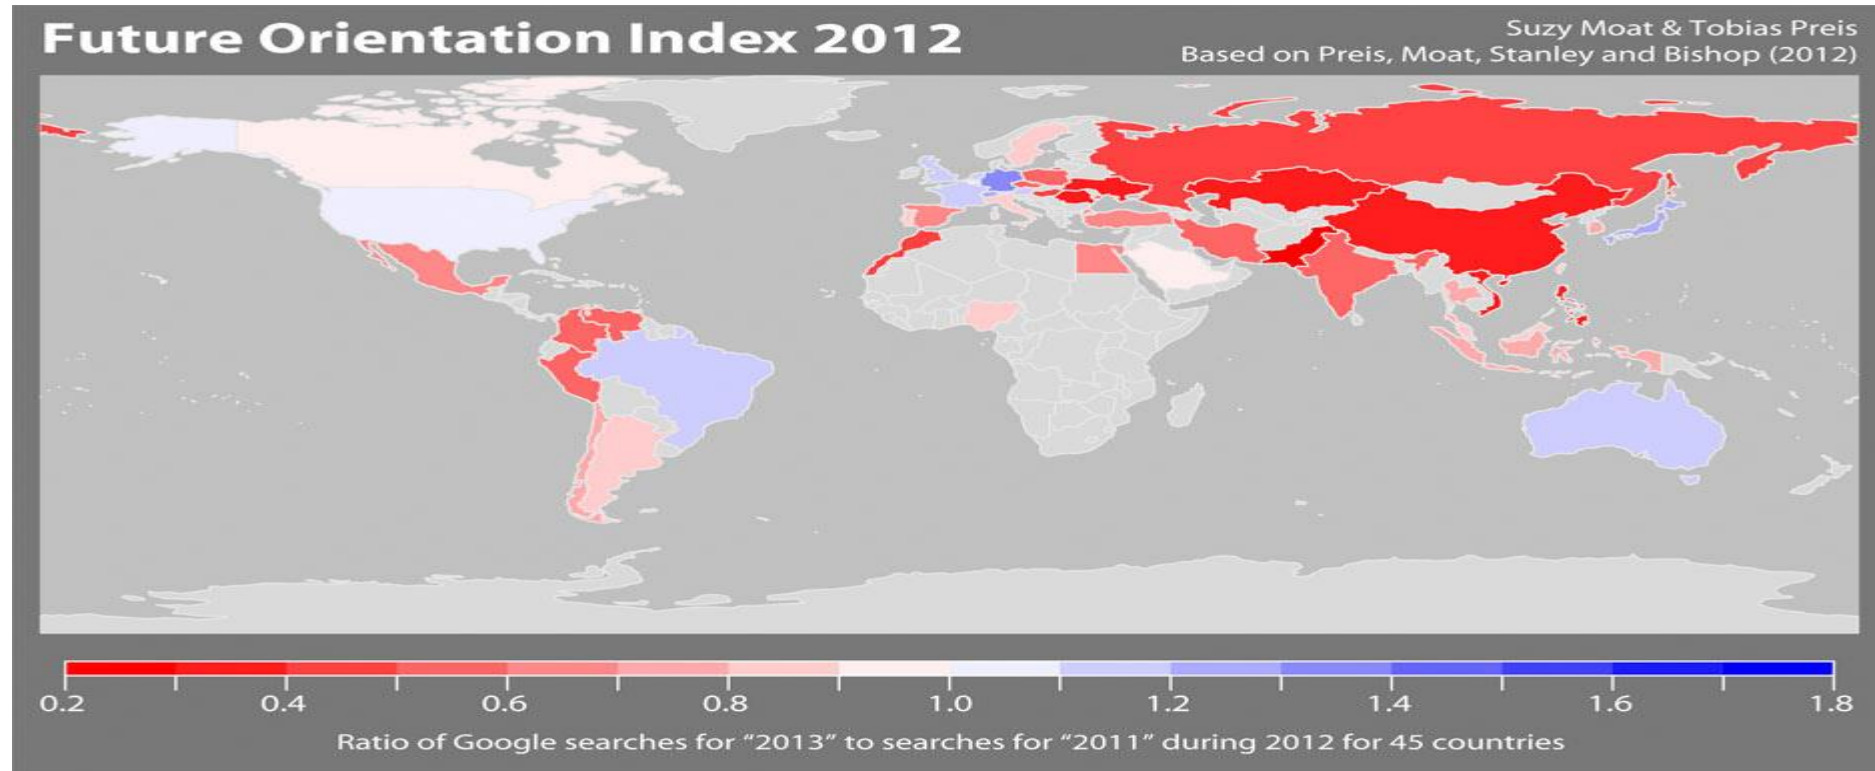


# Big Data and the Government



*“Can you explain all the emails you’ve received from Russia and Iran?”*

# Big Data and Economy





"We see two leading explanations for this relationship between search activity and GDP. Firstly, these findings may reflect international differences in attention to the future and the past, where a focus on the future supports economic success. Secondly, these findings may reflect international differences in the type of information sought online, perhaps due to economic influences on available Internet infrastructure."


What are the sources of data ?



# Sensors



During 2008, the number of things connected to the internet exceed the number of **people** on earth




By 2020 there will be 50 billion ***things***






# Mobile Traffic Growth (in ExaBytes)



# 5G Requirements




# You are What you Eat and Drink




- Food and drink became also a strong cultural aspect, being able to describe strong differences
- Foursquare, created in 2009, registered 5 million users in December 2010 and 45 million users in January 2014
- Possibility to sense human activities related to food and drink practices in large geographical areas
- Delineate and describe regions that have common cultural elements, defining signatures that represent cultural differences between distinct areas around the planet

# You are What you Eat and Drink




- 4.7 million tweets containing check-ins were gathered, each one providing a URL to the Foursquare website (one-week dataset same order of magnitude of the number of interviews performed in World Values Survey in almost three decades)
- Locationbased social networks (LBSNs)
- Location identified by free reverse geocoding API offered by Yahoo (<http://developer.yahoo.com>)

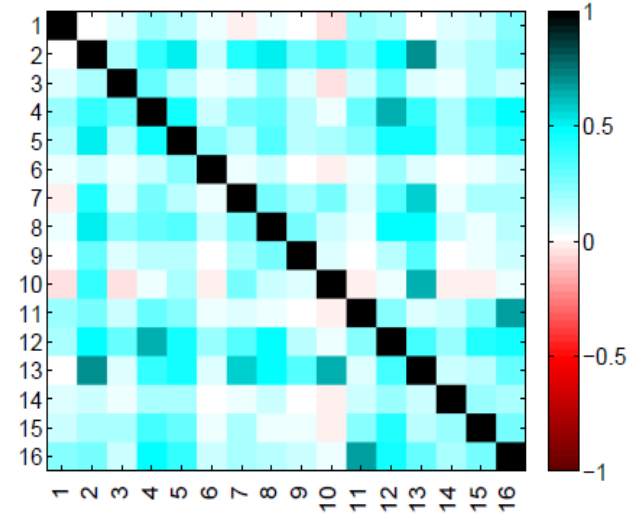
# You are What you Eat and Drink



(a) Drink




(b) Fast Food




(c) Slow Food

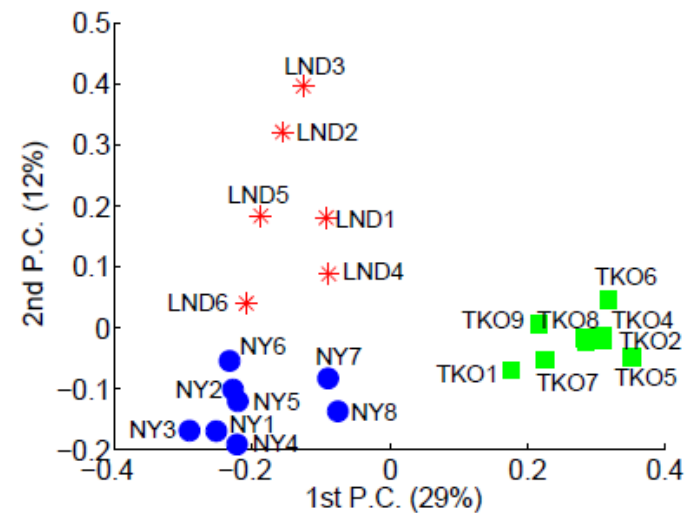
Figure 3: Correlation of preferences between countries.




(a) Countries



(b) Cities



(c) Regions



What is the role of networking in Big Data?



***Big Data ...  
and the Next Wave of **InfraStress*****  
***John R. Mashey***  
***Chief Scientist, SGI***

***Technology Waves:  
NOT technology for technology's sake  
IT'S WHAT YOU DO WITH IT  
But if you don't understand the trends  
IT'S WHAT IT WILL DO TO YOU***

Uh-oh!



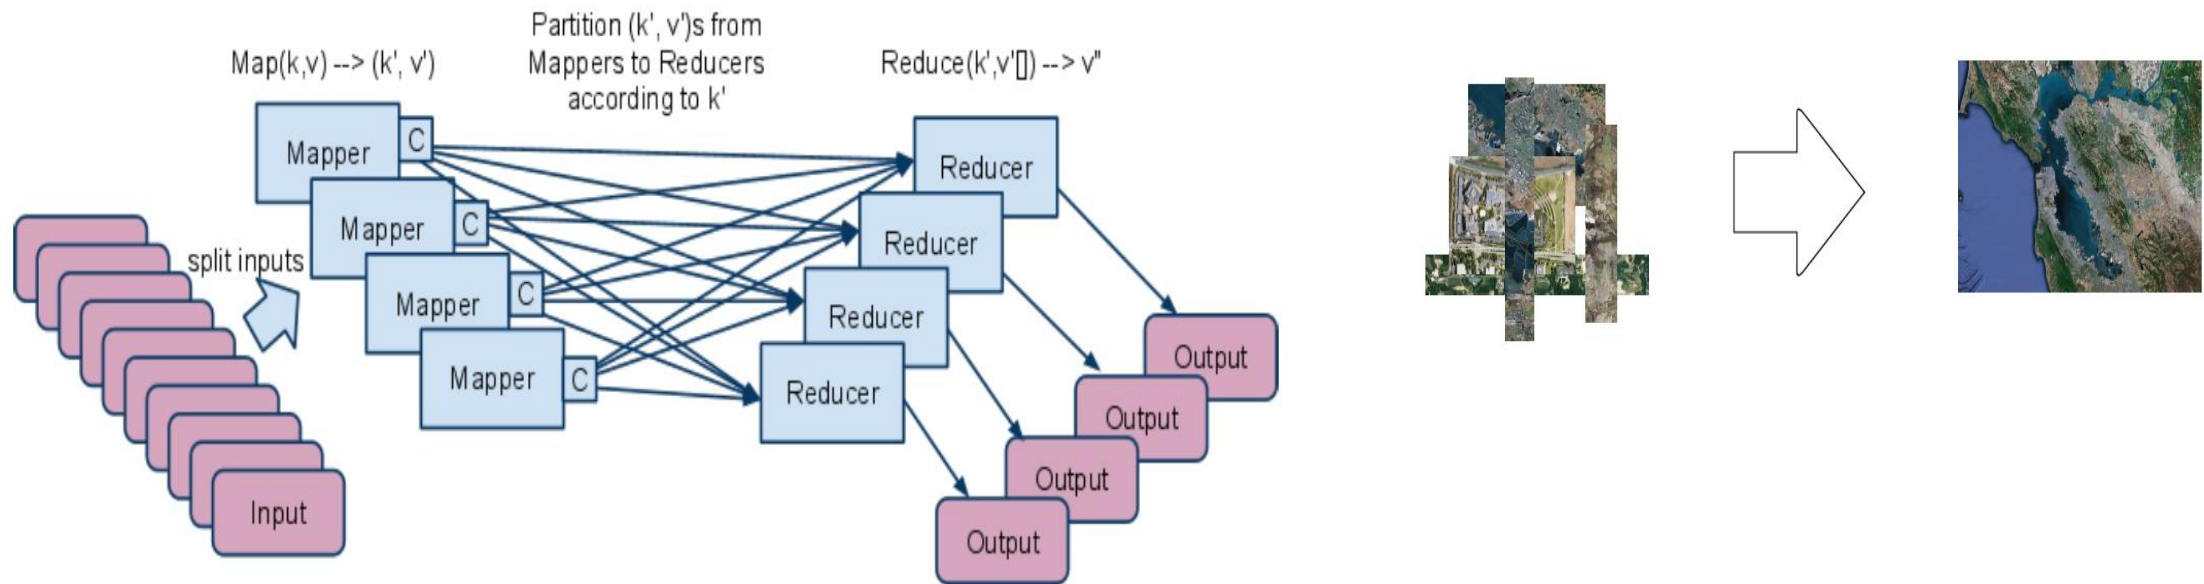


# “Infrastrass”



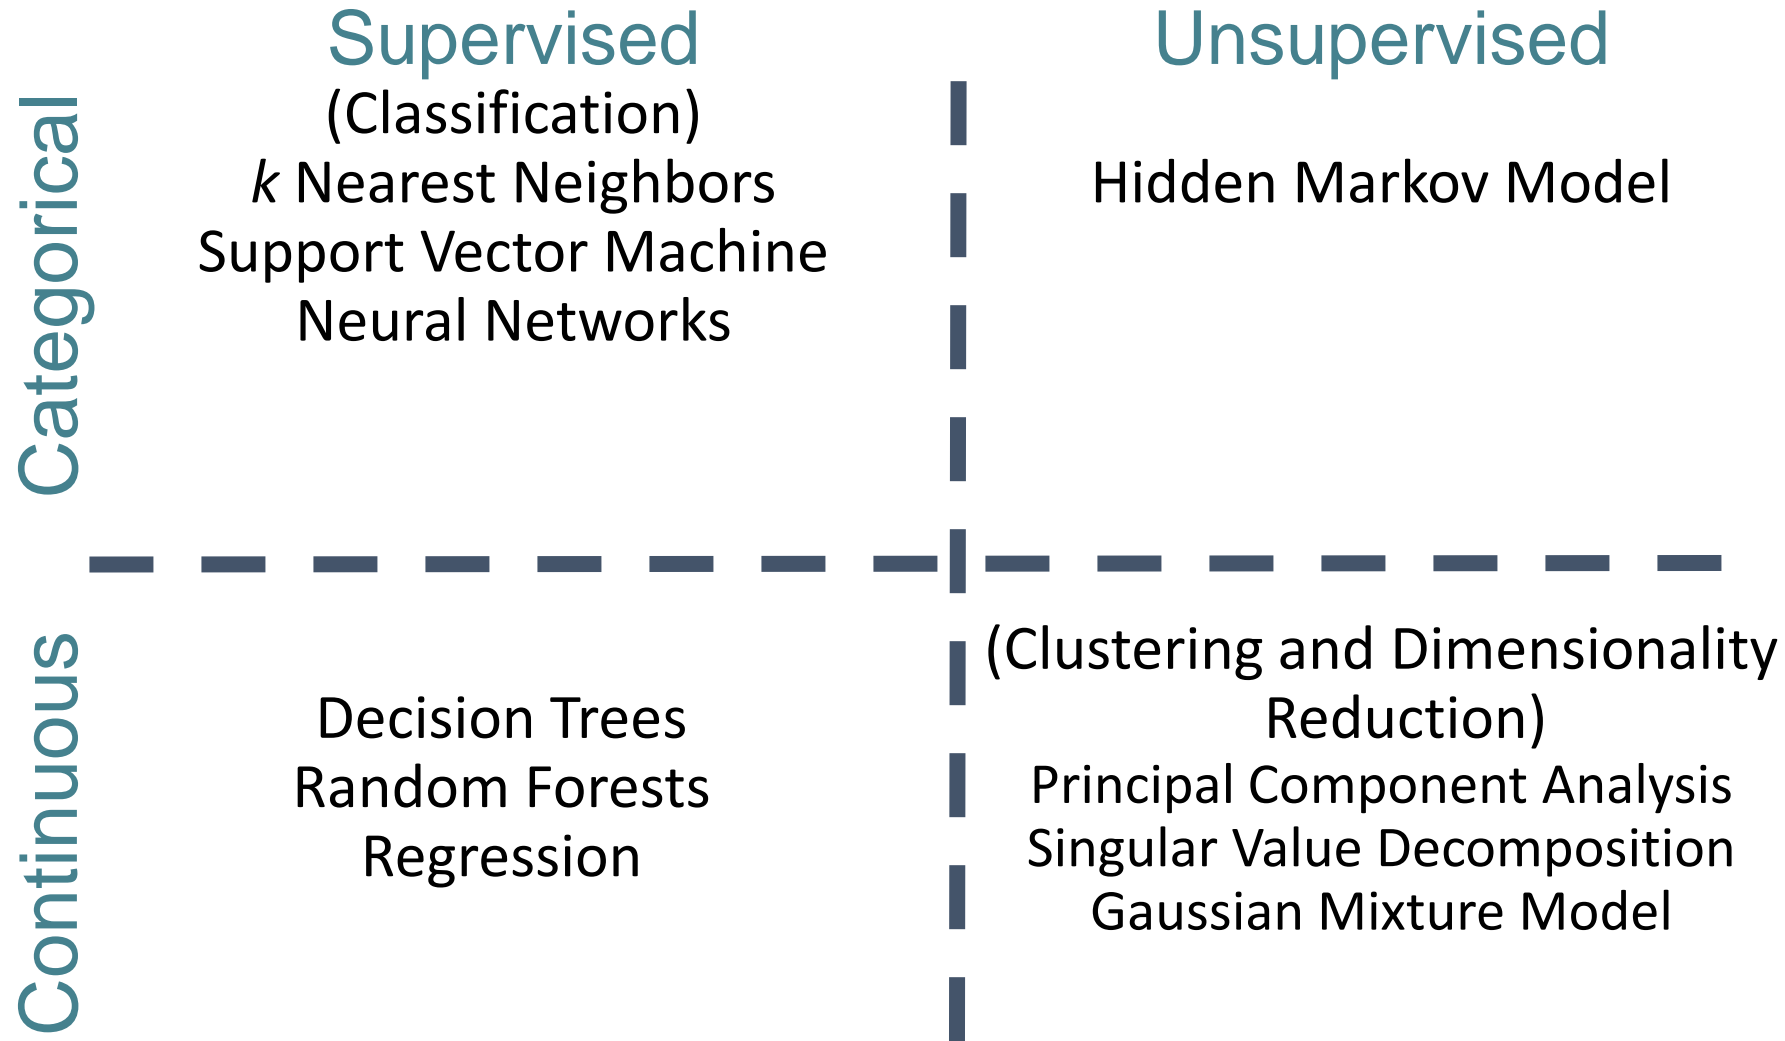
Alibaba Mall processes in a single day (Nov 11th, 2013) 105.8 million online transactions from 213 million users and 4.1 billion transactions

# Map-Reduce



Facebook Trace analysis: 30% to 50% of running time took up by communication phase

# Machine Learning Algorithms



# Machine Learning Algorithms

- The use of Latent Dirichlet Allocation algorithm for text mining requires all sampled topics to be multicasted interactively, which can exceed 1GB per interaction and over 1,000 interactions, leading to terabytes of data to be multicasted
- Logistic Regression algorithm for Twitter and the Alternating least Square algorithm for Netflix movie rating prediction take hundred of interactions. In each interaction roughly 300MB are distributed, leading to tens of gigabytes multicasted
- Multicast communication accounts for 30% to 45% of job completion time

# Networking for Big Data


# Networking Requirements for Big Data

- Elastic bandwidth: to match the variability of volume
- High Speed data transfer
- Security: Access control privacy, threat detection, all in real-time in a highly scalable manner
- Network partitioning to handle big data
- Network congestion control for big data applications
- Network service consistency

# Optical Multicast

- Data analytics applications routinely need to distribute terabytes of data from a central data source to hundreds of servers for processing
- In Hadoop Distributed File System(HDFS), multicast sender stores the data in HDFS and a multitude of receivers retrieve the data from a few data replicas, creating very high fan-out on the replicas. Bottleneck for over 40 recipients
- In Spark, a BitTorrent style P2P overlay among the recipient nodes, but BitTorrent suffers from suboptimal multicast trees that render high link stress, performs worse than HDFS


# Optical Multicast



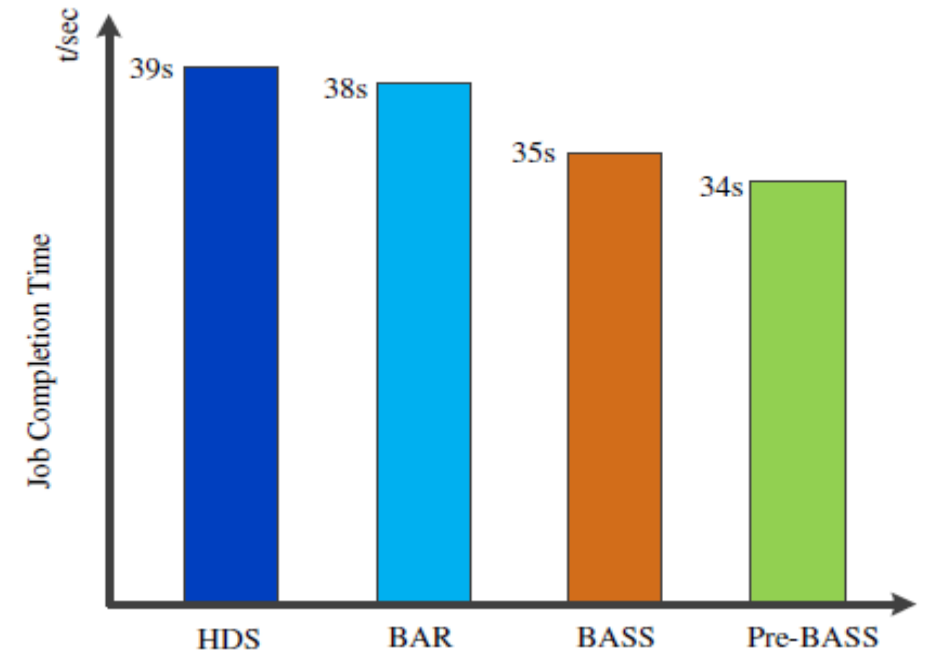
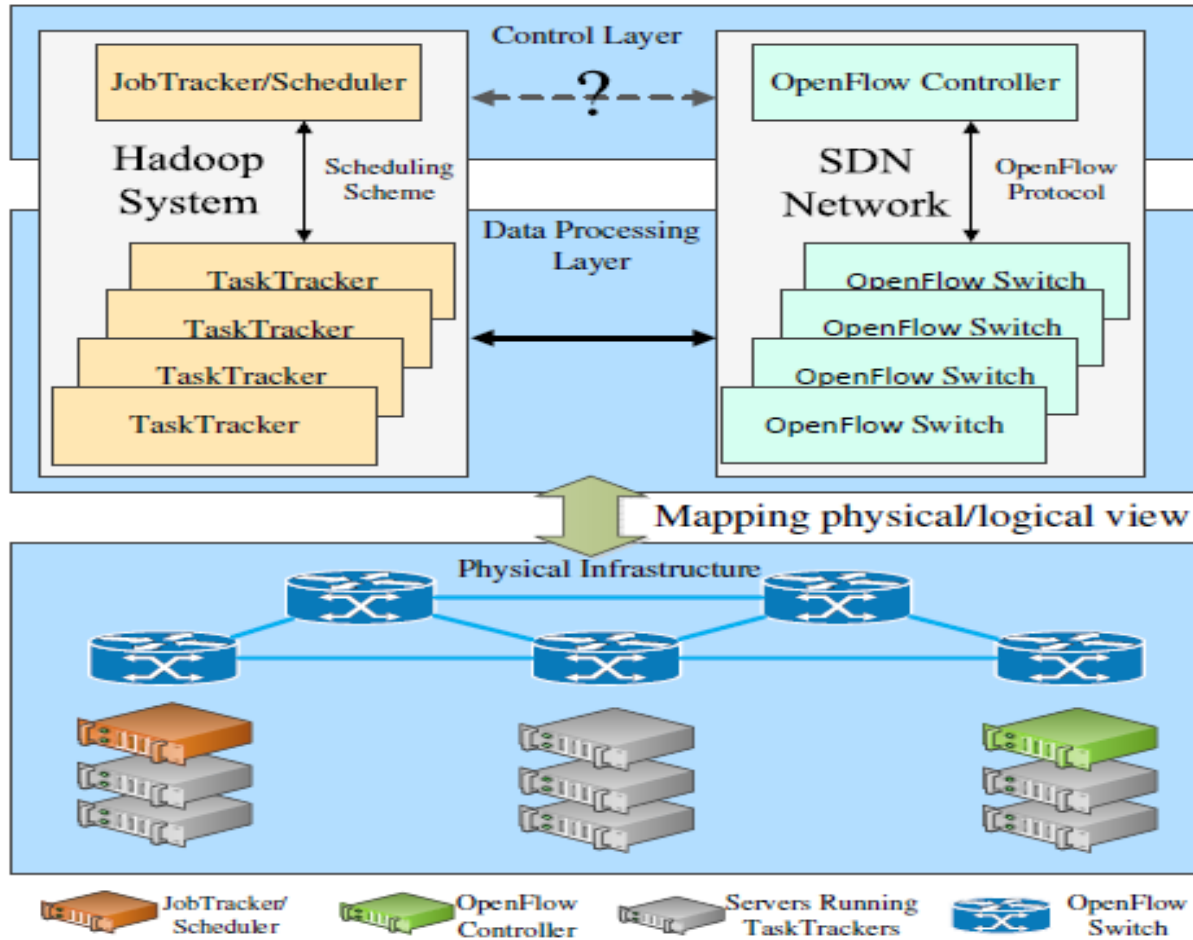
- Inherent performance limitations of application-layer overlays
- Difficult to perform TCP-friendly congestion control in network multicast traffic.
- Blast uses optical transmission to realize a physical-layer broadcast medium, via passive optical power splitting, to connect a data source to its receivers
- tailor-made control plane, capable of collaborating with data analytics applications interactively, making resource allocation decisions nearly optimally, and directing the data flows in optical and electrical components in the network



# Optical Multicast



# SDN & Hadoop



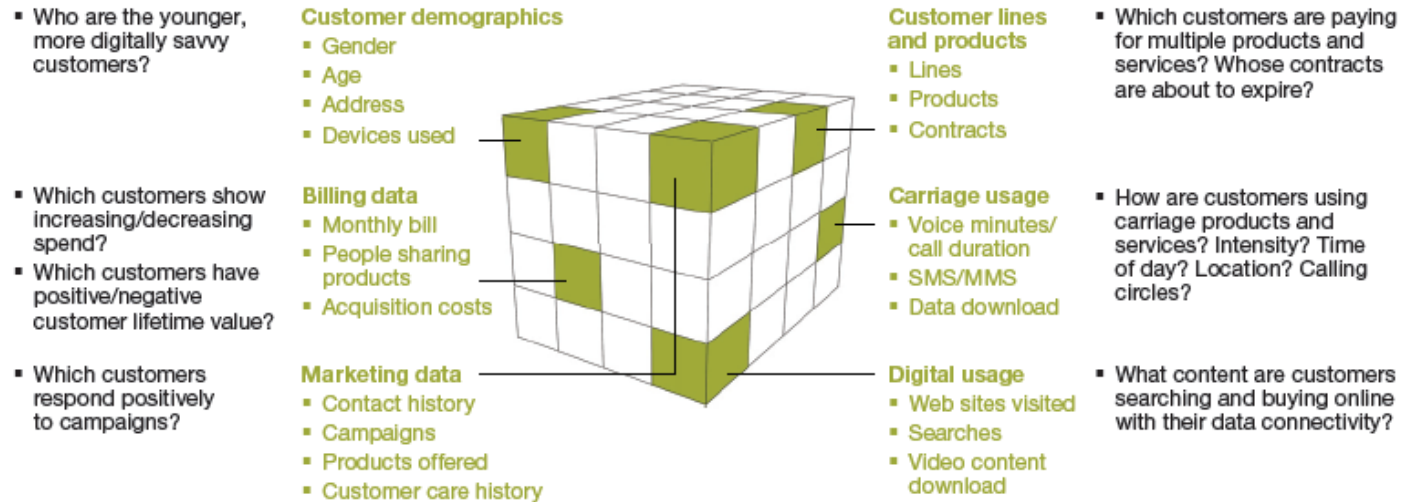
# Big Data for Networking

# Big data for Networking

- Traffic Classification
- Intrusion Detection
- Definition of new services and class of services
- Anomaly detection
- Fault detection
- Cognitive networks
- Network configuration
- Service level prediction

# Definition of new Services and Class of Services

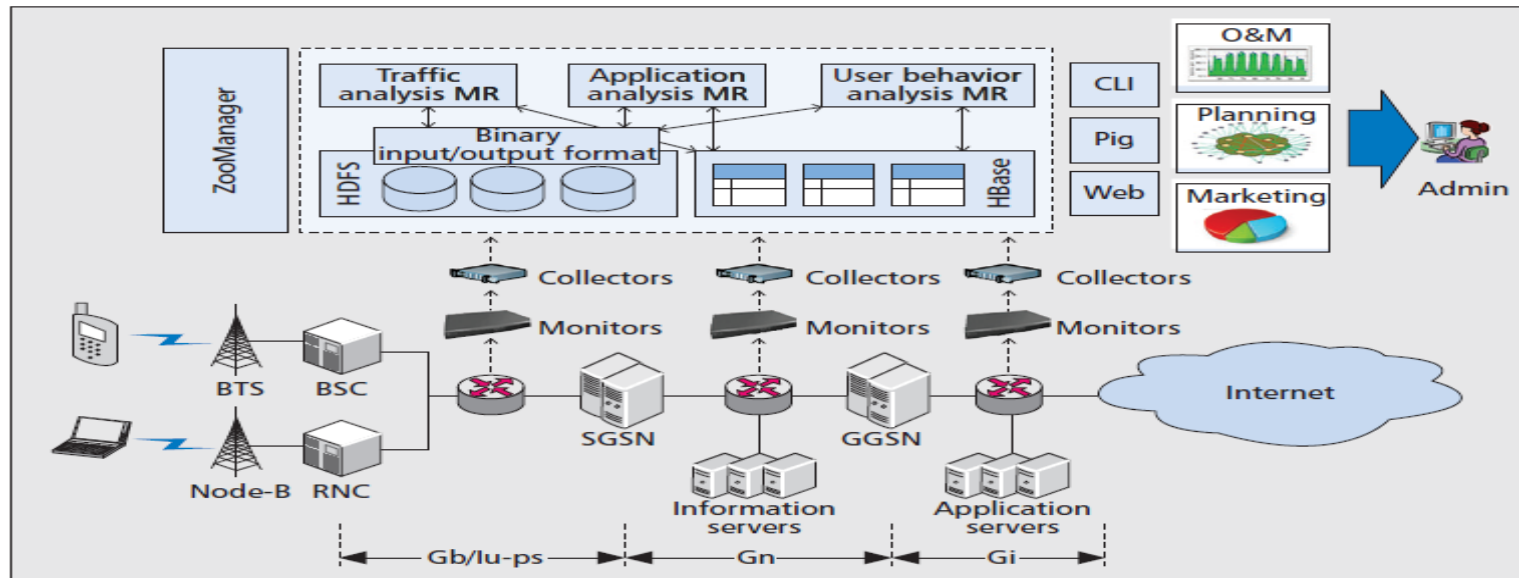
Based on detailed customer profiles, telcos can differentiate customer service models and develop individualized offer recommendations



SOURCE: McKinsey

# Identifying Communities of Website users

- Mobile operator in China
- Objective: to find website communities of users and identify their usage behavior



# Identifying Communities of Website users

- Traffic Monitor:
  - Line-speed packet parsing (PPPoE, GRE at various interface Gb, Gn, Gi)
  - Real-time traffic classification (19 sub-service classes)
  - Multi-level traffic statistics (packet, flow and aggregate level)
- Hbase in HDFS Hadoop, key/value stored in a columnar manner
- Mining logs of HTTP:
  - Affinity graphs models website usage
  - Sparsified affinity graph constructed by employing a scale free fitting index (node in-degree larger than a threshold)
  - Nodes are ranked by an influence score

# Identifying Communities of Website users

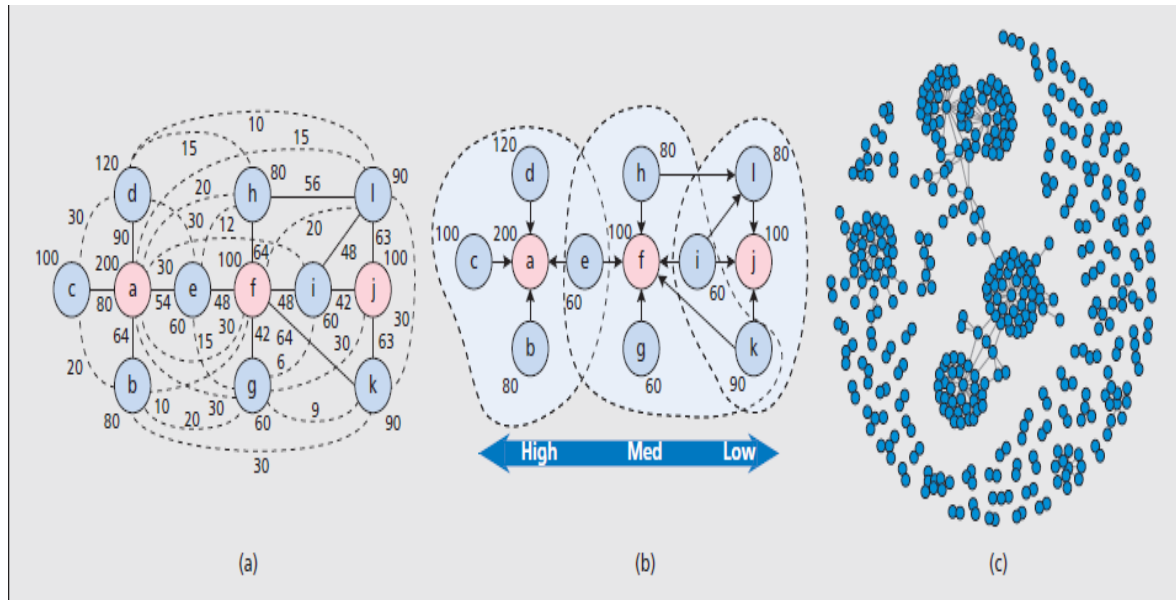


Figure 3. Website community identification: a) example of user distribution among websites; b) example sparsified affinity graph; c) identified website communities.


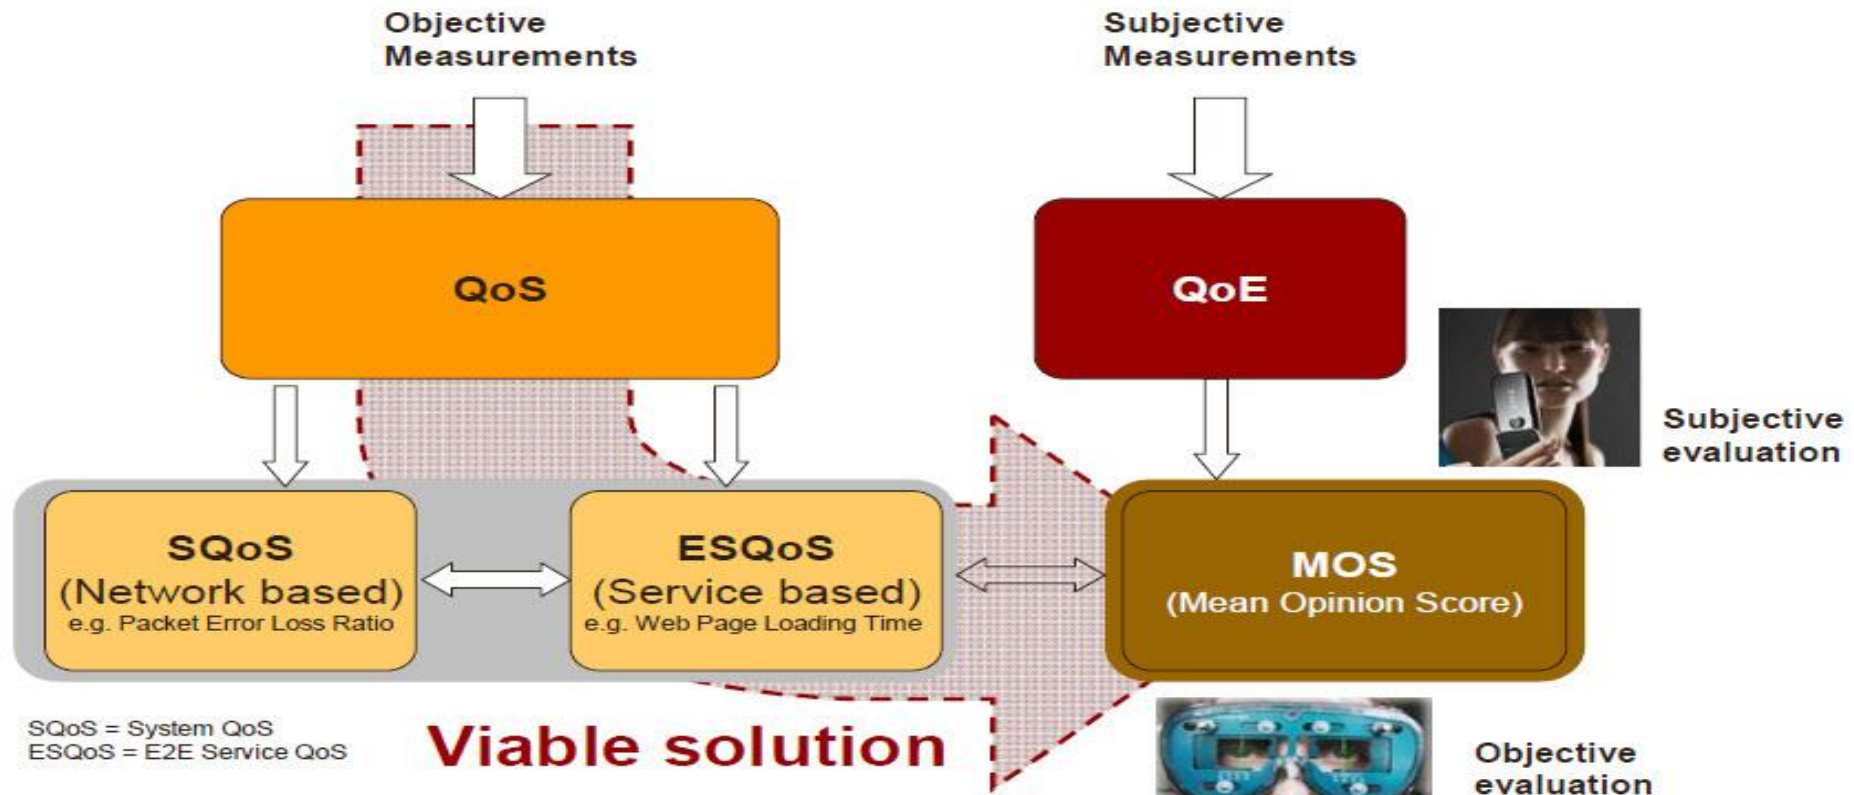


Figure 4. Mobile client model identification and user behavior analysis: a) example of identification results; b) accuracy of identification; c) density of network access time.




# Service Level Prediction

## QoS and QoE parameters – Mapping Model



# Service Level Prediction - VoD



- **Video streaming:** video frame rate, audio buffer rate, RTP packet rate

- CPU load, memory load, #network active sockets, #context switching, #processes, etc..
- raw data from /proc provided by Linux kernel


- Building block for real-time service assurance for a telecom cloud

# Service Level Prediction - VoD

- Statistical learning on low-level (OS-level) metrics, taking a large number of features ( $> 4000$ ) rather than few service-specific features ( $\leq 10$ )
- Measured metrics
  - Video frame rate (frames/sec)
  - Audio buffer rate (buffers/sec)
  - RTP packet rate (packets/sec)
- Linear regression, Regression tree, Random forest, Lasso regression
- Network statistics and client low-level metrics not considered
- Network and client machine are lightly loaded


# Service Level Prediction-VoD

- It is feasible to accurately predict client-side metrics based on low-level device statistics
  - Normalized Mean Absolute Error below 15% across service-level metrics and traces
- Preprocessing of X is critical
- Trade-off between computational resources vs. prediction accuracy
- No time dependence considered



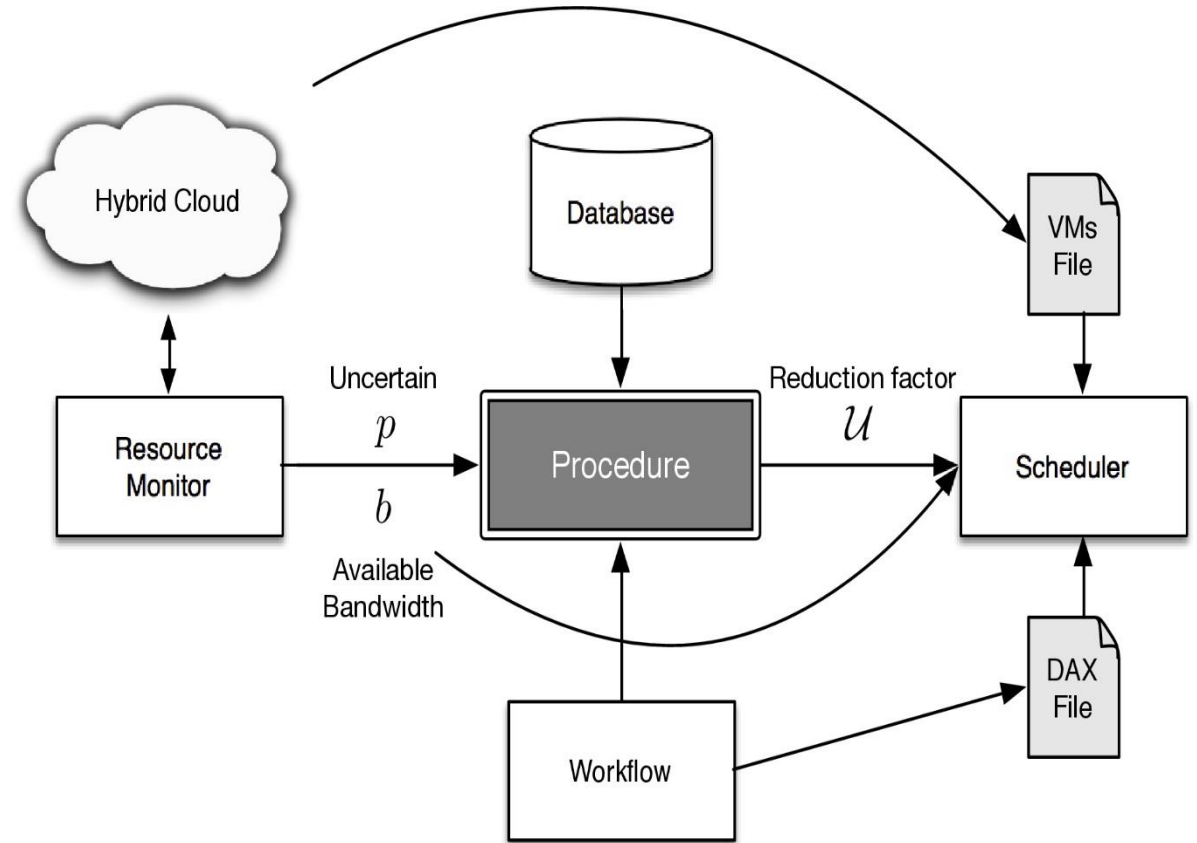
# Service Level Prediction- hybrid cloud

- Impact of intercloud link bandwidth on makespan of workflow execution
- Overestimation of available bandwidth can lead to increased makespan and higher cost
- Underestimation of the bandwidth leads to unnecessary leasing of resources
- Available bandwidth varies during execution of workflow

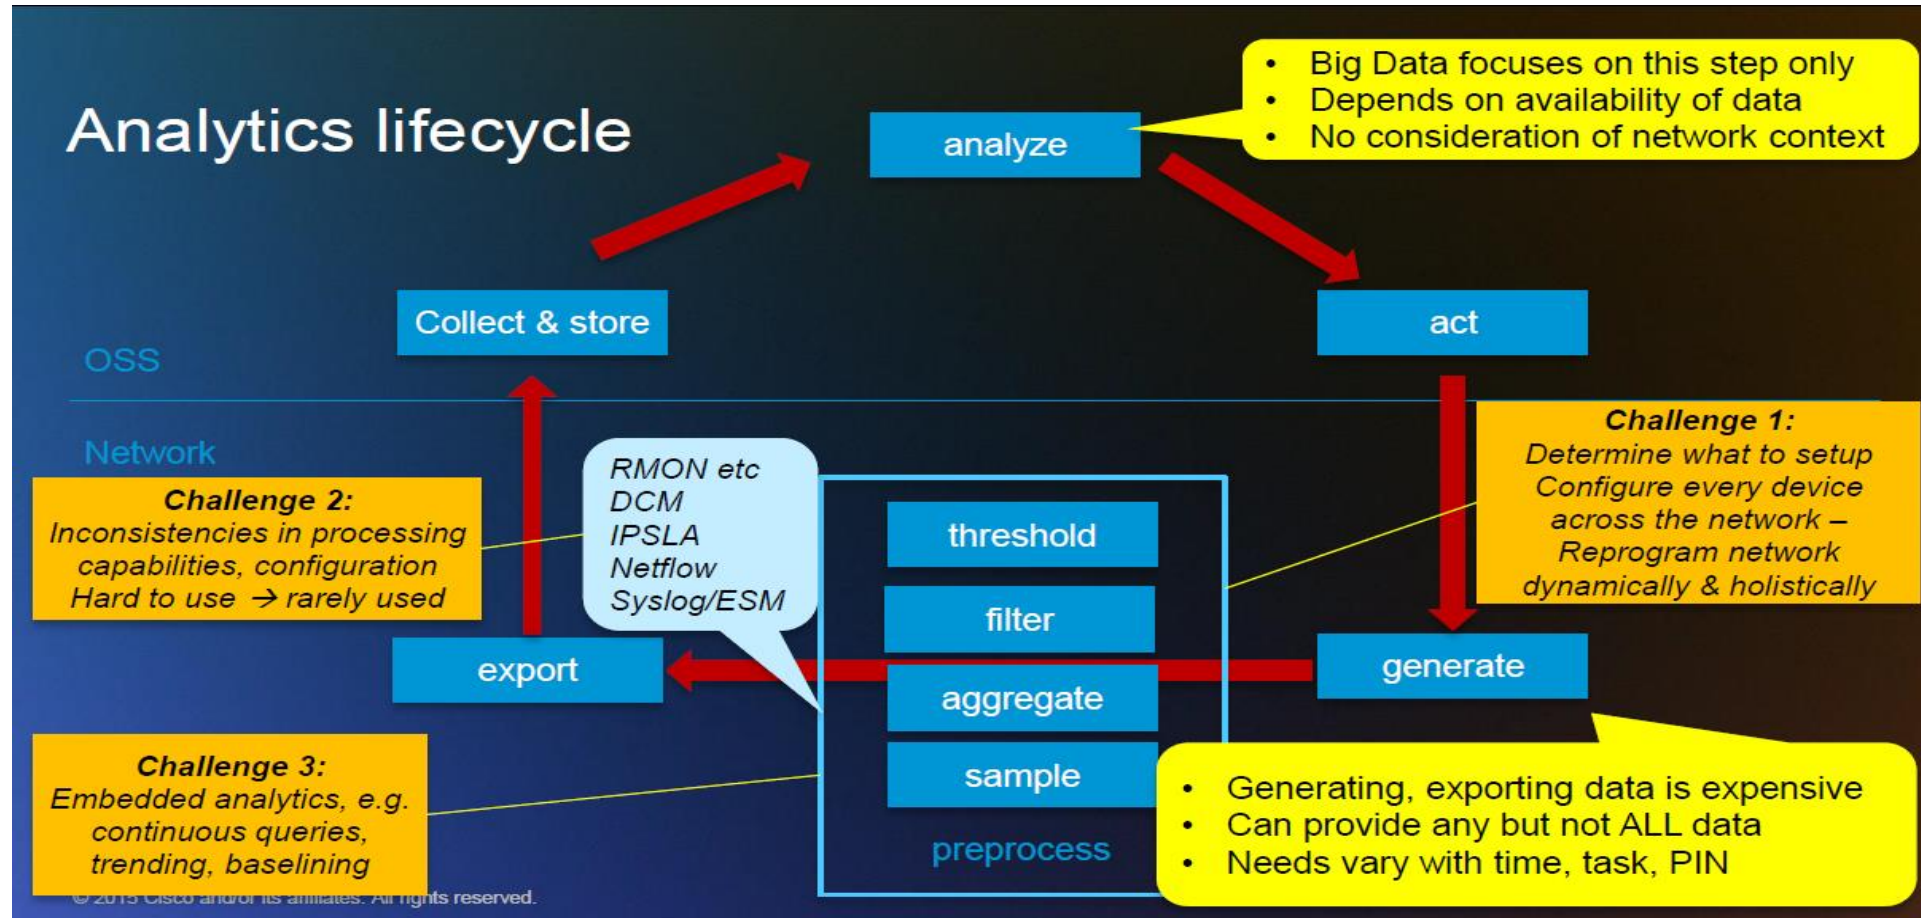


# Service Level Prediction- hybrid cloud

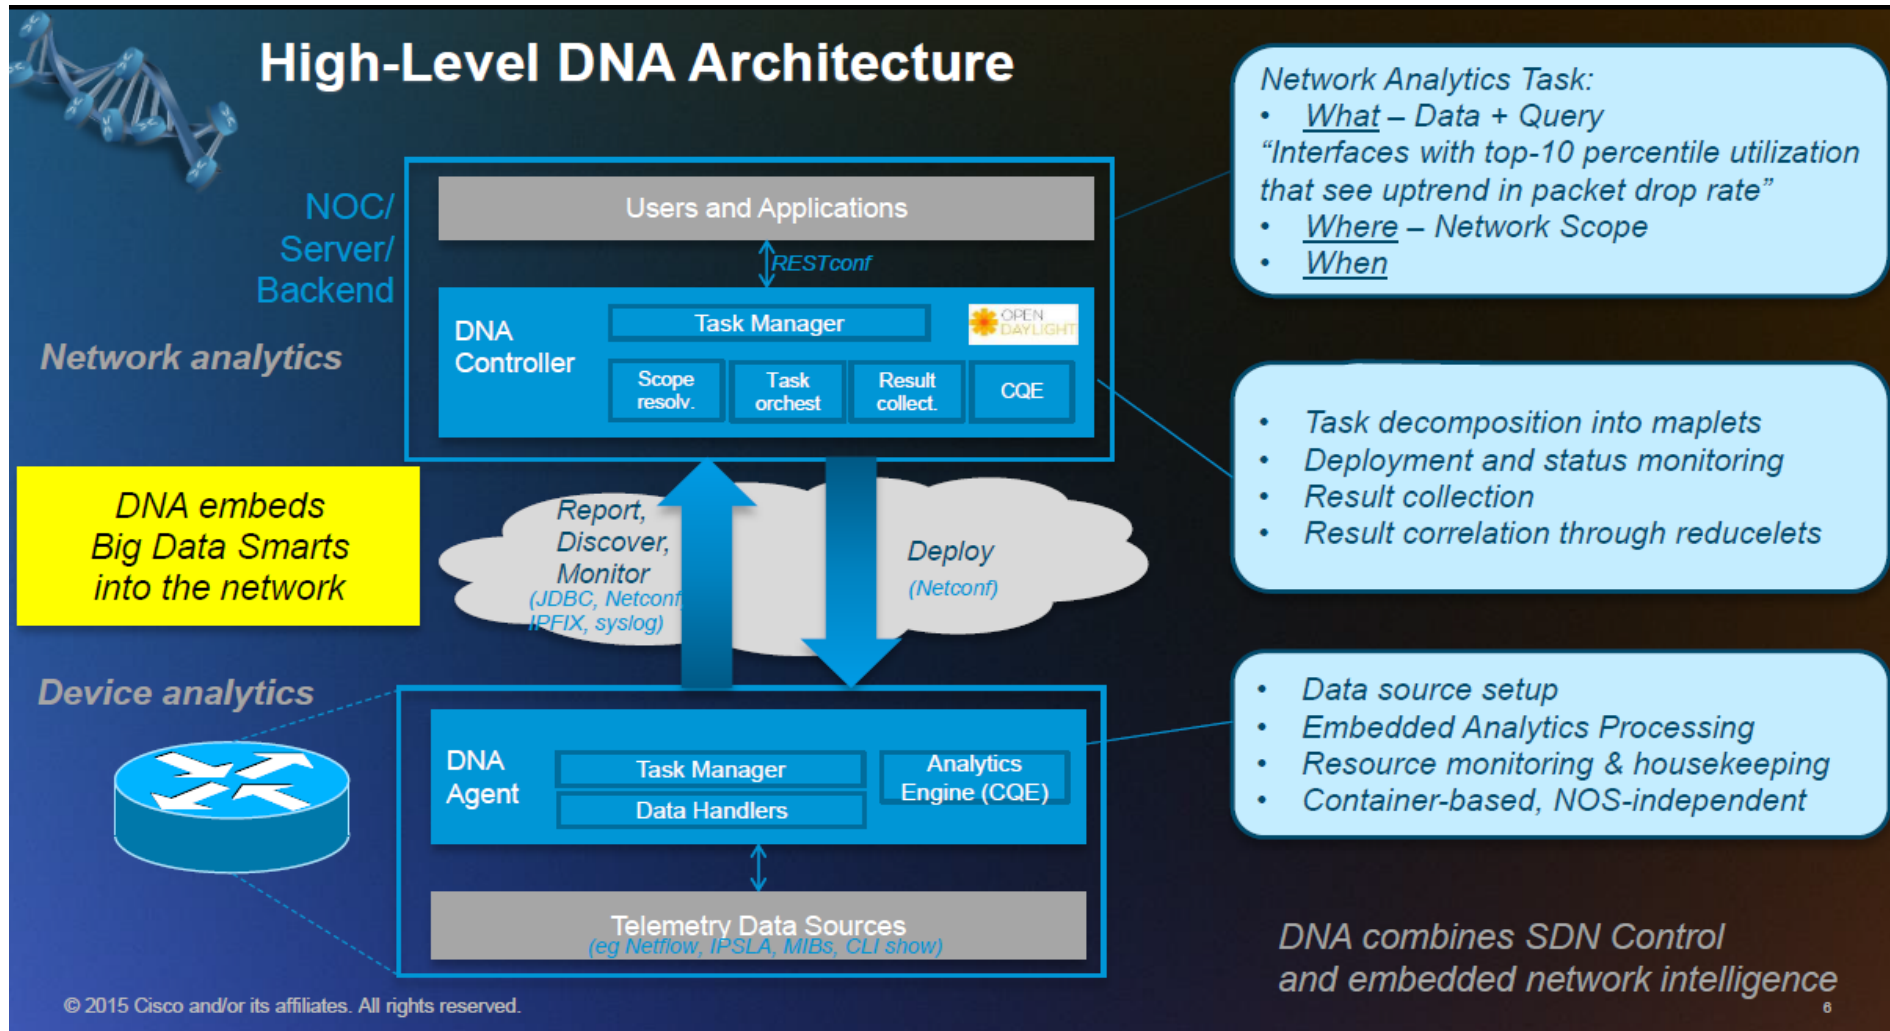
- Deflating fator to measured available bandwidth
- Database with historical data on performance relating available bandwidth, type of workflow, deflating fator, quality of information and estimation erros
- Multiple linear regression
- Reduced the number of disqualified solutions and reduced costs



# Big Data for Network Management and Operation



# Big Data for Network Management and Operation







# Verification of Forwarding Tables

- 10,000s switches in modern data centers
- Traditional approach of scaling-out and redundant design assumes correct reaction to errors
- Dormant bugs in routing systems triggers rare boundary conditions, some times perceived when benign event happens
- Two major deficiencies of available tools:
  - Assumption of consistent snapshot of forwarding tables
  - Not sufficiently fast to meet requirements of modern datacenters

# Verification of Forwarding Tables



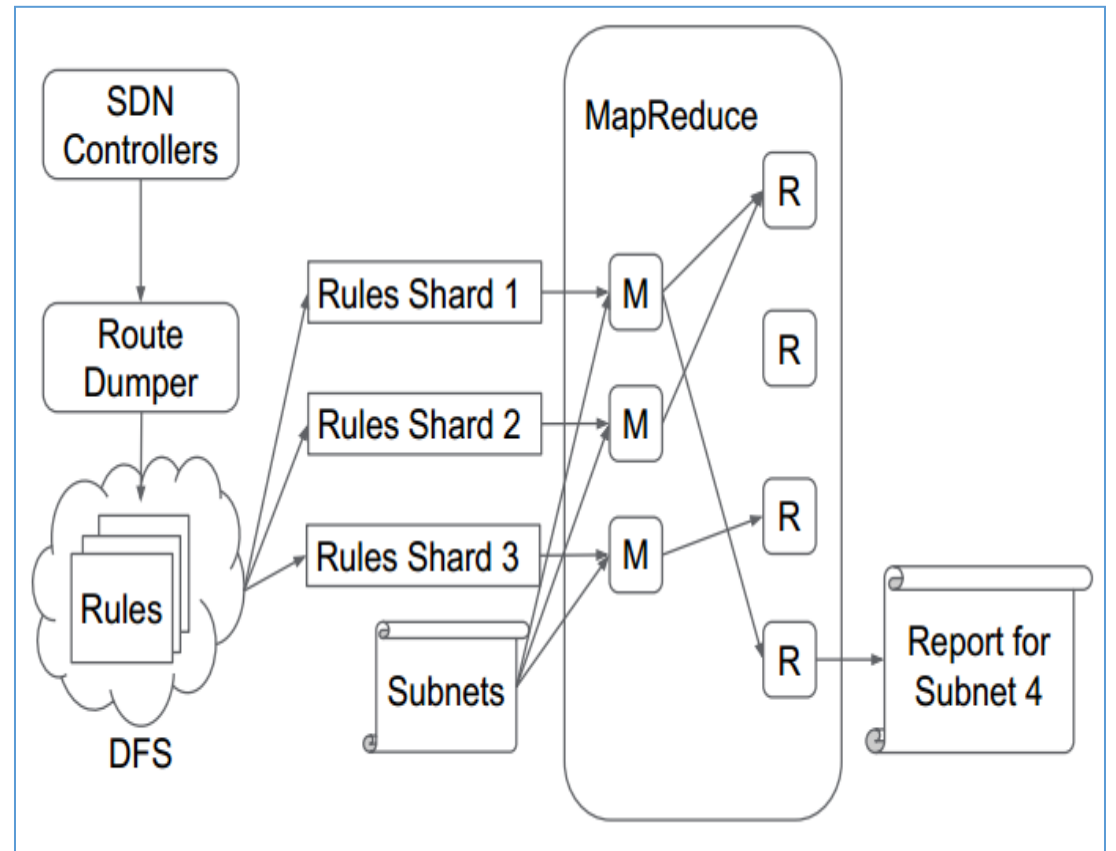
Data Plane Tickets in a Google Data Center



- Loops - usually caused by prefix aggregation
- Black-holes - lost BGP updating information
- Inconsistent snapshot - in an event of failure, changes to forwarding table may conflict with previous forwarding information

# Verification of Forwarding Tables

- Reachability - subnet can be reached from any switch; depth-first-search from the subnet switch, verify if all other switches are reachable
- Loop detection - strongly connected component
- Black-role - if a switch does not have a matching entry for the subnet



# Verification of Forwarding Tables

- Libra: Dive-and-Conquer approach

Data set	Switches	Rules	Subnets
DCN	11,260	2,657,422	11,136
DCN-G	1,126,001	265,742,626	1,113,600
INET	316	151,649,486	482,966

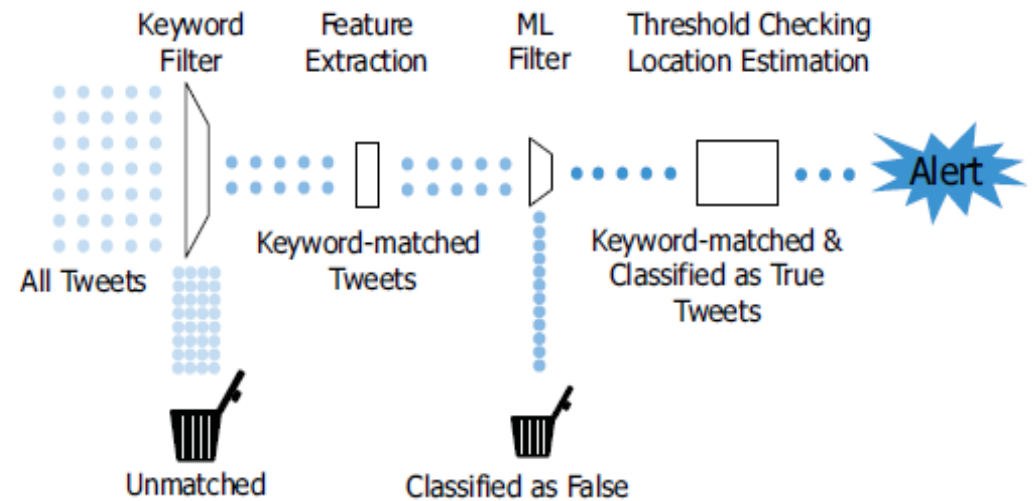
	<b>DCN</b>	<b>DCN-G</b>	<b>INET</b>
Machines	50	20,000	50
Map Input/Byte	844M	52.41G	12.04G
Shuffle Input/Byte	1.61G	16.95T	5.72G
Reduce Input/Byte	15.65G	132T	15.71G
Map Time/s	31	258	76.8
Shuffle Time/s	32	768	76.2
Reduce Time/s	25	672	16
<b>Total Time/s</b>	<b>57</b>	<b>906</b>	<b>93</b>

# Fault Identification and Location




- Some network failures become silent failures to mobile operators; difficult to establish rules for failure detection
- Most failures are not reported to call centers
- Monitoring Tweeter to detect failure in mobile services
  - ✓ "Why can't I text messages?"
- Most tweets are not related to network failure
  - ✓ I dropped my phone in toilet when I called my friend

# Fault Identification and Location




- Two requirements:
  - ✓ reduction of false positives and
  - ✓ location and when incident happened
- Imbalance between network failure tweets and others
- Closeness of network failure tweets and false positives
- Three stages: keyword filtering, machine learn filter and alert system

# Fault Identification and Location



- Keyword filtering: selection of features related to network failure
  - ✓ 4398 network failure and 6924 false positives
  - ✓ 10K tweets/sec for 40 keywords
- Four different techniques in machine learning filtering: SVM with and without Gaussian Kernel, Naive Bayes and Adaptive Regularization of weights
  - ✓ 1% of raw tweets,
  - ✓ Capacity of 1 K tweets per second
  - ✓ 1 hour to construct a model
- Location: name of city, station or landmark; use of gazetteer and kernel density estimation to estimate location of fault; known GPS location of tweets
  - ✓ Heavy processing load; 0.2 users/second

# Radio Environment Maps



- Spatial maps of received signal strengths can be used in dynamic spectrum access to, for example, discover coverage holes in cellular networks
- Perform drive test to collect data and perform spatial interpolation
- Methods from Spatial statistics, such as Kriging, are accurate and robust. However, its complexity is  $O(n^3)$  where  $n$  is the number of measurement points
- Current used approaches can only give a rough estimation since propagation simulations have inherent inaccuracies due to limited information on landscape data. Moreover, drive tests are too expensive



# Radio Environment Maps

- Minimization of drive tests developed by 3GPP makes every mobile phone a spectrum measurement device, making available a large number of path loss or received signal strength and GPS location information
- Operators can harvest unprecedented amount of data
- Employment of fixed rank kriging techniques with linear computational complexity to process hundreds of thousands of measurements, spatial estimate
- Model fitting takes 20 seconds of computation in a desktop and individual predictions less than milliseconds

# Radio Environment Maps


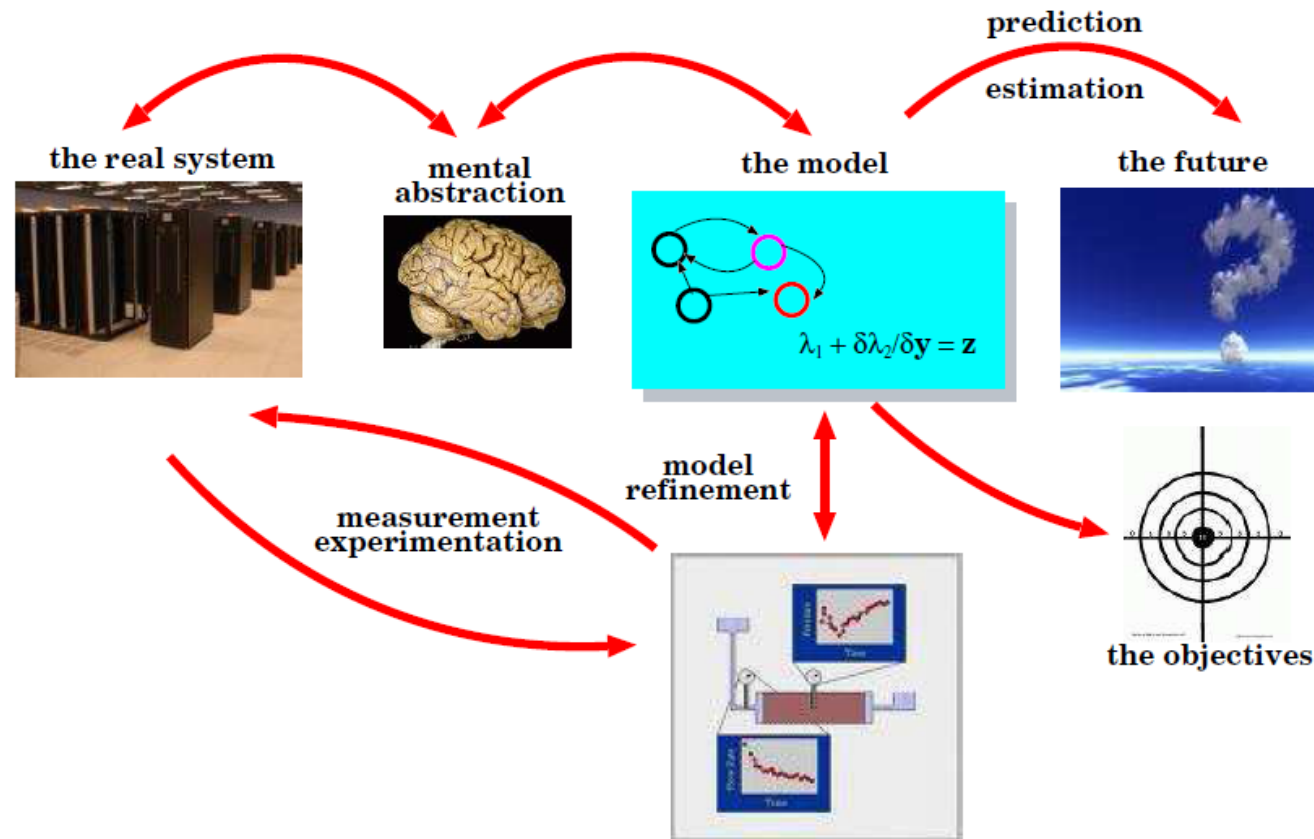


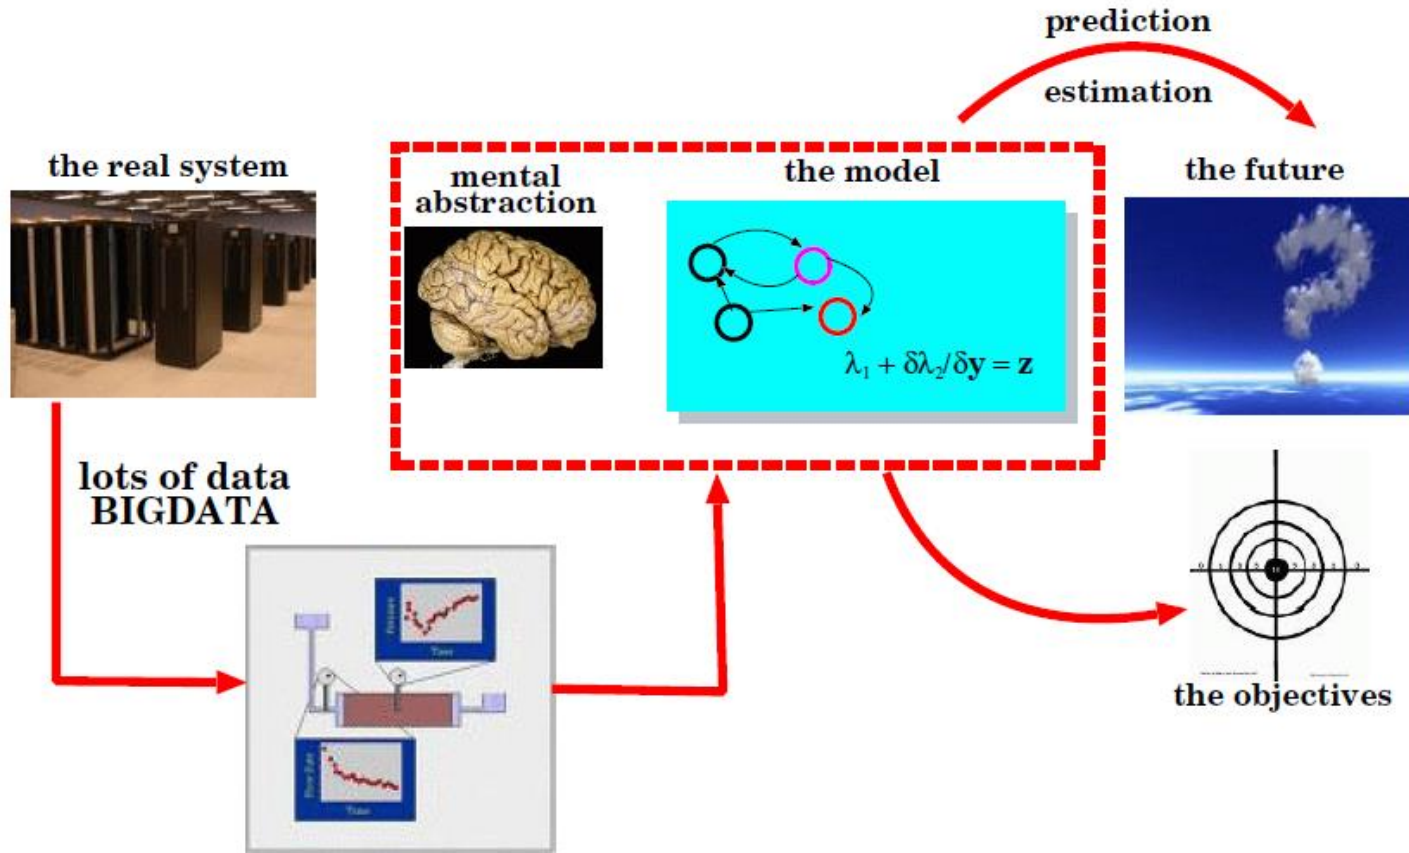
Figure 1. Example of a typical spatial estimation task in management of wireless networks. The map on the left shows the actual coverage of a digital TV network over Germany (total area of approximately 350,000 km<sup>2</sup>), whereas the middle and right figures show spatial estimates based on 10,000 and 27,000 distributed measurements, respectively.

# Performance Evaluation

# "Traditional" Analytical Modeling



# Analytics



# Some Research Problems

- Understand the traffic generated by Big Data processing to evaluate the need for novel congestion control mechanisms
- Process data at the edge of the network and transport partial results - use of distributed machine learning algorithms
- Investigate how software defined network and network function virtualization can be explored for application-aware networking, specially for congestion control
- Lavarage the integration of "traditional" analytical modelling and machine learning techniques for performance evaluation



- My hearty thank you for ITC27 organisers, especially to Michela Meo, Sabine Wittevrongel and Peter von Daele
- Alexander Clemm, Edmundo Sousa e Silva and Rolf Stadler for providing some slides
- ITC attendees for their attention and sincere apologies for not being able to make it to Ghent