Foundations for Big Data Systems (B)

大数据系统基础(B)

Zhi Wang
wangzhi@sz.tsinghua.edu.cn

Course ID: 64100033

1

Lecturer and TA

Lecturer: 王智

F204, wangzhi@sz.tsinghua.edu.cn

External lecturers

Tencent,

HKUST, UCB

TA: 曾铖淋

F205, zephyr0703@qq.com

2

Background

大数据方向专业硕士

"π"型数据人才

大数据系统

数据系统决定数据处理和分析可能性

3

Big Data Trends

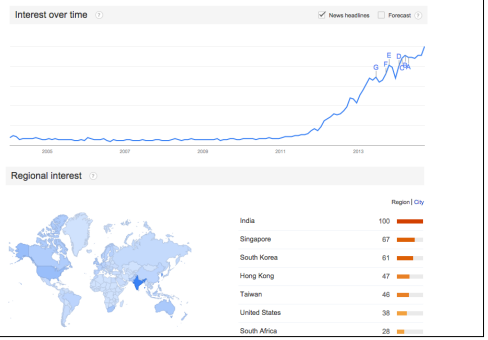

4

Big Data Trends

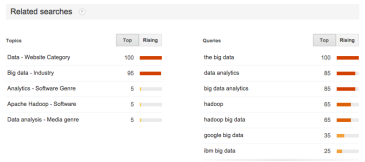

5

Big Data Trends

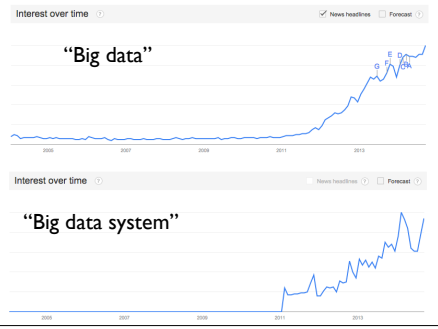

6

Big Data Trends

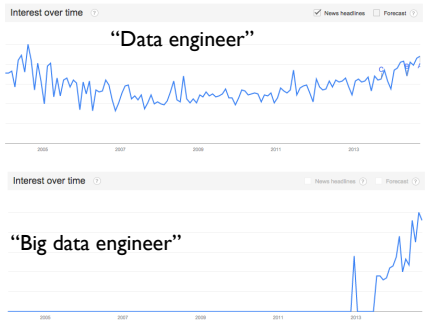

7

Outline

What is big data

What are big data systems

What can you learn in this course

Objectives and principles in big data systems

8

Data is Everywhere

Web data, e-commerce

Bank, credit cards

Online social networks

Sensors

Wearable devices

Robotics

9

How Much Data?

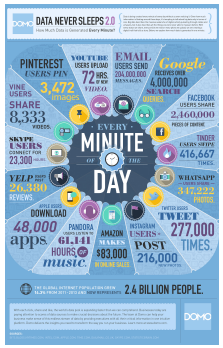

10

Types of Data

Relational Data (Tables/Transaction/Legacy Data)

Text Data (Web)

Semi-structured Data (XML)

Graph Data

Social Network, Semantic Web (RDF), …

Streaming Data

- You can only scan the data once

11

What do We do with the Data

Aggregation and Statistics

- Data warehouse and OLAP

Indexing, Searching, and Querying

- Keyword based search
- Pattern matching (XML/RDF)

Knowledge discovery

- Data Mining
- Statistical Modeling

12

To Whom does It Matter

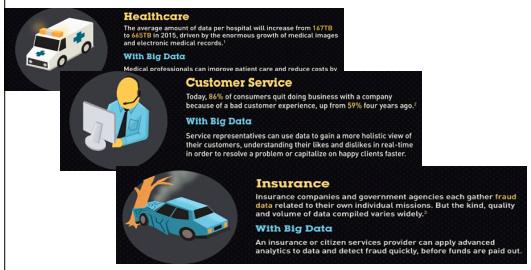

13

What does It Trigger

14

4V/3V view

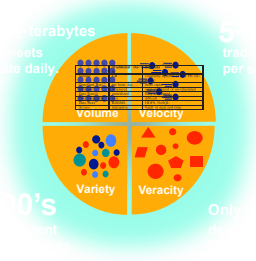

15

What is Big Data

The recognition that data is at the center of our digital world and that there are big challenges in collecting, storing, processing, analyzing, and making use of such data.

"Big" may refer to very large data volume, but not necessarily so.

16

(Size, Activity) view

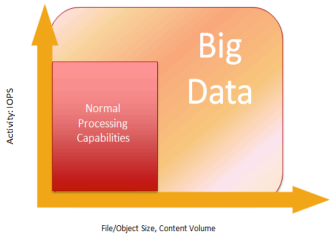

17

Examples

2顾客：那你们有什么可以推荐的？

1客服：您可以试试我们的低脂健康比萨。

2顾客：你怎么知道我会喜欢吃这种的？

1客服：您上星期一在国家图书馆借了一本《低脂健康食谱》。

2顾客：好。那我要一个家庭特大号比萨，要付多少钱？

1客服：99元，这个足够您一家六口吃了。但您母亲应该少吃，她上个月刚刚做了心脏搭桥手术，还处在恢复期。

18

19

Video

20

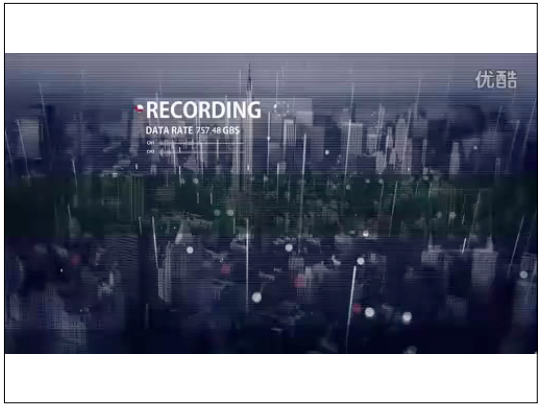

21

- *Multi-source, Correlation, Fuzzy, etc.*

22

Big Data is...

It is all about *better analytic* on a *broader* spectrum of data, and therefore represents an opportunity to create even more *differentiation* among industry (data) peers.

23

Big-Data Systems

Big data is a broad concept that covers many aspects of computer science.

We focus on the computer systems aspect, for instance,

- How various parts of a big data computer system (hardware, system software, and applications) are put together?
- What are the appropriate approaches to realize high performance, scalability, reliability, and security in practical big data computer systems?

Probably not the right course if you are hoping to learn about algorithmic design and theoretical/mathematical foundations for machine learning and data mining.

24

Simple to Start

What is the maximum download speed you get?

What is the maximum file size you have downloaded/uploaded?

What are the data types you have processed so far?

25

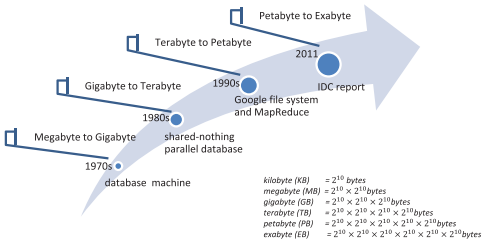

FIGURE 2. A brief history of big data with major milestones. It can be roughly split into four stages according to the data size growth of order, including Megabyte to Gigabyte, Gigabyte to Terabyte, Terabyte to Petabyte, and Petabyte to Exabyte.

Hu, Han, et al. "Toward scalable systems for big data analytics:A technology tutorial." Access, IEEE 2 (2014): 652-687.

26

Comparison

	Traditional Data	Big Data
Volume	GB	constantly updated (TB or PB currently)
Generated Rate	per hour, day, ...	more rapid
Structure	structured	semi-structured or un-structured
Data Source	centralized	fully distributed
Data Integration	easy	difficult
Data Store	RDBMS	HDFS, NoSQL
Access	interactive	batch or near real-time

Hu, Han, et al. "Toward scalable systems for big data analytics:A technology tutorial." Access, IEEE 2 (2014): 652-687.

27

Framework of Big-data Industry

28

Topics in Big-data Systems

Collection: missing information, dummy data, organization

Data Transfer: Limitations of current systems, CPU intensive

Storage: Data sets beyond relational database, clusters, data centers, distributed data

Processing: Software, processing power, parallel and distributed computing

User Interaction: Non-programmers need to perform complex information, real time GUI interfaces, visualization of data

29

Big-data Systems in this Course

Big Data - Infrastructure (Cloud)

Big Data - Storage (HDFS, GFS)

Big Data - Computing (MapReduce, Spark)

Big Data - Database (HBase)

Big Data - Graph DB (GraphLab)

Big Data - Streaming (Storm)

Big Data - Tool (R)

30

Big-data Systems in this Course

Not limited to these systems

We focus on the fundamental design principles, and performance issues

31

General Course Information

Online info:

- Webpage: http://mmlab.top:38080/bigdatasys/

Text and references

No official textbook, will use online resources and papers

Acknowledgements:

Several Internet material – Thanks to "Internet"

32

Our Tasks in this Course

Knowing

- How big-data systems work
- Limitations
- Performance evaluation
- Potential improvement

Engineering

- Helloworld in these systems
- Demo-level development
- Real job implementation

33

Assignment and Final Projects

Three programming assignments (individual)

- First on data collection (10%)
- The other two on data processing and analysis (2 x 10%)

A survey on big-data system subject to your choice (team, 20%)

A final course project on a topic of your choice (team, 30%)

- Proposal
- Presentation
- Demo

Exam (20%)

34

Pre-requisite

No formal prerequisite

Desire good programming skills

- Know C++/Java/Python
- Ability to learn new programming languages

35

Interaction, Please!

This is a new course; I am learning along with you

We encourage discussions and interactions

Extra credits for strong participation

36

Introduction for Making Groups

Group size: [2, 3]

Submit on 网络学堂

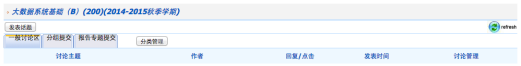

37

Update Your Info

38

Course Outline

Introduction to Big Data and Big-data Systems

Cloud Computing and Virtualization

Big-data System Management

Data Preprocessing (1st assignment)

Data Storage Systems

MapReduce (2nd assignment)

In-memory Data Processing (survey assignment)

39

Course Outline

Big-data vs. Database

Graph Computing (3rd assignment)

Streaming Big Data Processing

Data Analysis Tools (project assignment)

Big-data Applications

13. Survey & Project Presentation

阳光课堂

40

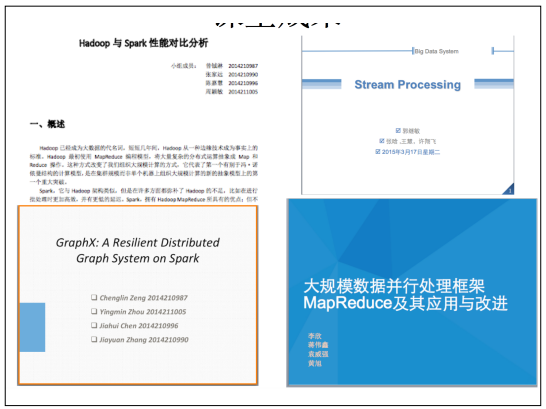

41

Framework View

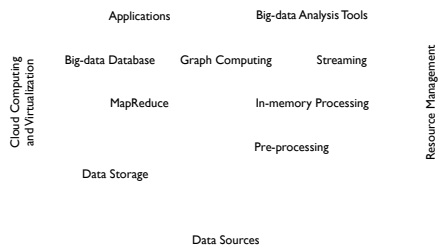

42

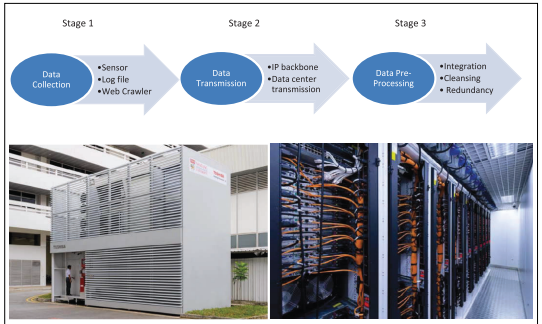

FIGURE 1. **A modular data center was built at Nanyang Technological University (NTU) for system/testbed research. The testbed hosts 270 servers organized into 10 racks.**

43

Computer Storage

44

Data Processing

RecordID	UserID	Sex	BirthYear	Age	星座	Major(计算机,电子,自动化,挖掘机)	one5s,小米4,锤子)	PrefWakeupTime(Morning,Noon,Evening)
0	2014210750	M	1988	26	处女座	挖掘机	iPhone5s	Morning
1	2014210751	M	1994	20	双子座	挖掘机	iPhone5s	Morning
2	2014210752	F	1998	16	金牛座	挖掘机	iPhone6	Noon
3	2014210753	F	2003	11	射手座	挖掘机	iPhone5s	Noon
4	2014210754	M	1994	20	白羊座	挖掘机	小米4	Afternoon
5	2014210755	F	2004	10	处女座	挖掘机	iPhone5s	Afternoon
6	2014210756	F	1991	23	处女座	挖掘机	锤子	Noon
7	2014210757	F	1992	22	双子座	挖掘机	小米4	Morning
8	2014210758	M	1990	24	白羊座	挖掘机	iPhone6	Afternoon
9	2014210759	M	2000	14	白羊座	挖掘机	小米4	Morning
10	2014210760	F	1996	18	双鱼座	挖掘机	iPhone5s	Morning
11	2014210761	F	1999	15	双鱼座	挖掘机	iPhone5s	Afternoon
12	2014210762	M	1990	24	金牛座	挖掘机	小米4	Afternoon
13	2014210763	M	2004	10	天秤座	挖掘机	小米4	Afternoon
14	2014210764	F	1989	25	双鱼座	挖掘机	小米4	Afternoon

45

RDD Lineage

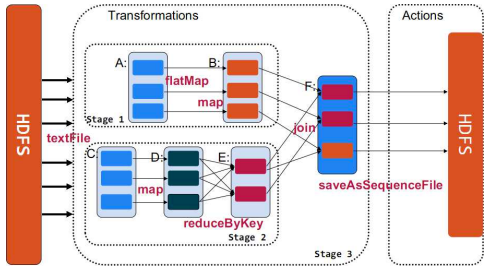

46

Google Bigtable

47

47

Storm

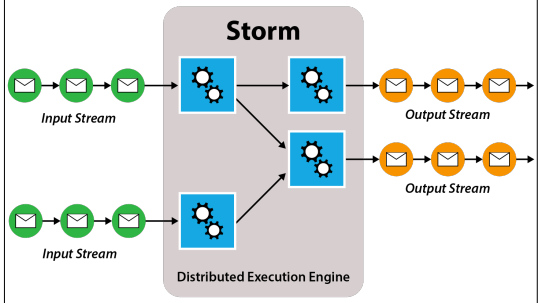

48

Everyone is talking about graphs...

49

Distributed Machine Learning System

mahout

50

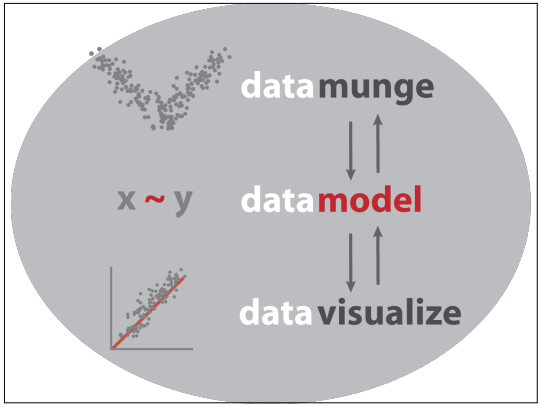

51

国内外相关参考课程

◆Introduction to Data Science (UC Berkeley)

- Instructor: Mike Franklin
- Pre-requisites: 61A, 61B, 61C, basic programming skill. Run VirtualBox on laptop.
- Grading: Class Participation and in-class labs:20%; Midterm:20%; Final Project (in groups):25%; Homeworks:30%; Bunnies: 5%.
- Final Project: Identify two or more data sets you would like to study, write the code to collect and integrate those data sets, then build two or three visualizations of data.
  - Keep diary of success and failures; final submission consists: paper document; presentation
  - Project proposals: problem intend to address; data intended to use…
  - Resources: Stanford 224w Page; Quandl-Find use and share numerical data
- Homework1: Text analysis and entity resolution
- Homework2: Introduction to machine learning: clustering & regression

52

国内外相关参考课程

◆Web Scale Data Management (Buffalo)

- Professor Oliver Kennedy, a seminar course
- Related to storing and querying large datasets, distributed systems and primitives, including data processing, synchronization, key-value stores, stream processors, as well as full SQL database system.
- Grading: S/U, all students submit a short, weekly abstract and critical analysis of the week's papers, and to participate in class discussion. Enrolled for 3 credits submit a simple experimental project.
  - Implement and compare two join algorithms(Nested Loop+Hash) on M/R
  - Implement and benchmark a ring DHT like Chord(Scaling Performance)
  - Implement and compare two distributed join algorithm as standalone processes.
- Project Resources: a 12-core development server; Amazon AWS.

53

国内外相关参考课程

◆Scalable and Data-Intensive Computing in the Cloud (Washington)

- Instructor: Bill Howe
- Explore the technology landscape at the intersection of big data and cloud computing. Each class consist of a 1-hour lecture, a 1-hour case study and demonstration of a specific system and a 1-hour of discussion and hands-on work.
- Student Assessment: Assignments:80%; Participation: 20%. All assignments will be due 1 week later by the start of class. Participation will be a combination of attendance and discussion involvement; in class and online involvement will both contribute.
- Websites:
  http://homes.cs.washington.edu/~billhowe/bigdatacloud/

54

MOOC

Q&A

55

Big-data Systems: Problems & Principles

How to make use of growing computer/network/ storage resource to deal with growing data, to satisfy growing user demands

Everything is unreliable, including user requirement

Objectives:

- Greedy: make use of all resource
- Adaptive: codes can change and can move
- Modular: let users manipulate data like playing a game

56

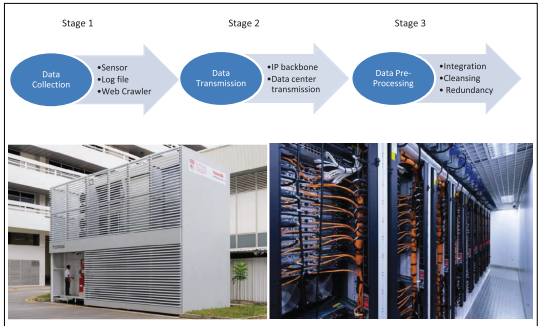

FIGURE 1. A modular data center was built at Nanyang Technological University (NTU) for system/testbed research. The testbed hosts 270 servers organized into 10 racks.

57

Data Collection

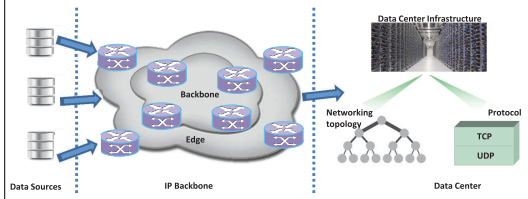

58

Data Storage

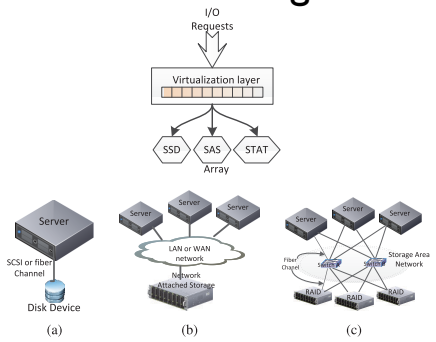

59

Data Processing

Programming Model	MapReduce	Pregel	Dryad	GraphLab
Database	BigTable	Dynamo	MongoDB	Spanner
File System	GFS	Haystack	TFS	

60

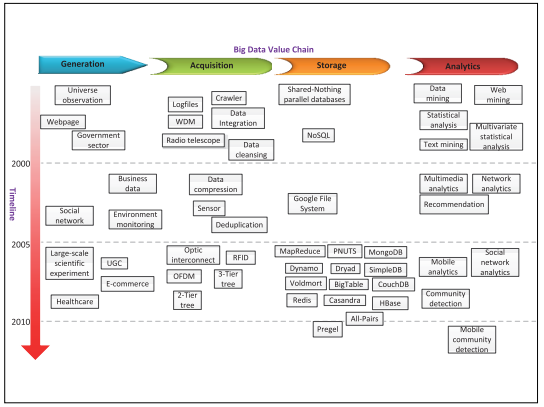

61

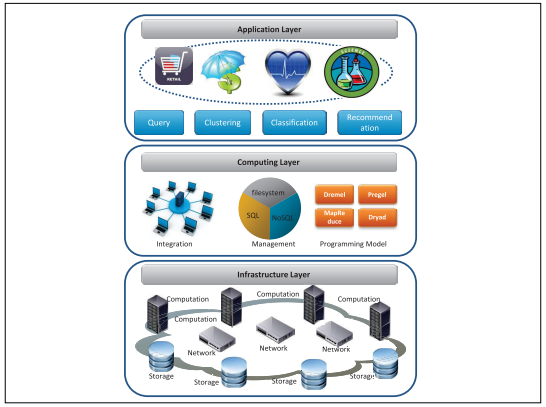

62

63

Divide and Conquer

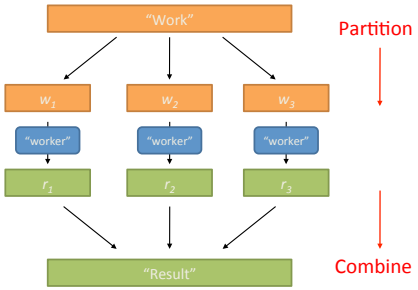

64

65

Parallelization Challenges

How do we assign work units to workers?

What if we have more work units than workers?

What if workers need to share partial results?

How do we aggregate partial results?

How do we know all the workers have finished?

What if workers die?

66

Current Tools

Programming models

- Shared memory (pthreads)
- Message passing (MPI)

Design Patterns

- Master-slaves
- Producer-consumer flows
- Shared work queues

Concurrency is difficult to reason about, not to mention debug

What is the Point

It's all about the right level of abstraction

- The von Neumann architecture has served us well, but is no longer appropriate for the big-data processing

Hide system-level details from the developers

- No more race conditions, lock contention, etc.

Separating the what from how

- Developer specifies the computation that needs to be performed
- Execution framework ("runtime") handles actual execution

Datacenter is the Computer!

Virtualization

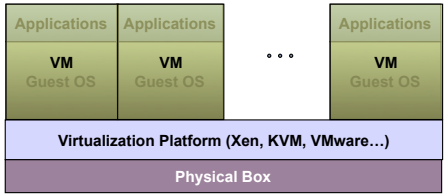

What is Cloud Computing?

- **Cloud Computing** is a general term used to describe a new class of network based computing that takes place over the Internet,
  - basically a step on from Utility Computing
  - a collection/group of integrated and networked hardware, software and Internet infrastructure (called a platform).
  - Using the Internet for communication and transport provides hardware, software and networking services to clients
- These platforms hide the complexity and details of the underlying infrastructure from users and applications by providing very simple graphical interface or API (Applications Programming Interface).

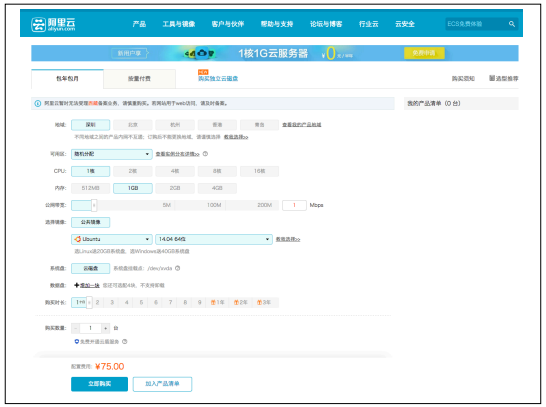

Cloud Computing Characteristics

Common Characteristics:

Massive Scale	Resilient Computing
Homogeneity	Geographic Distribution
Virtualization	Service Orientation
Low Cost Software	Advanced Security

Essential Characteristics:

On Demand Self-Service	
Broad Network Access	Rapid Elasticity
Resource Pooling	Measured Service

Cloud Service Models

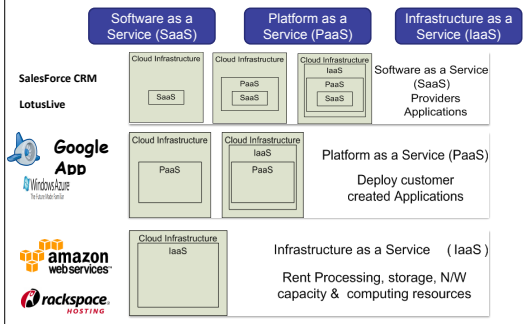

73

Virtualization Layers

Application Service (SaaS)	MS Live/ExchangeLabs, IBM, Google Apps; Salesforce.com Quicken Online, Zoho, Cisco
Application Platform	Google App Engine, Mosso, Force.com, Engine Yard, Facebook, Heroku, AWS
Server Platform	3Tera, EC2, SliceHost, GoGrid, RightScale, Linode
Storage Platform	Amazon S3, Dell, Apple, …

74

Virtual Machines

- VM technology allows multiple virtual machines to run on a single physical machine.

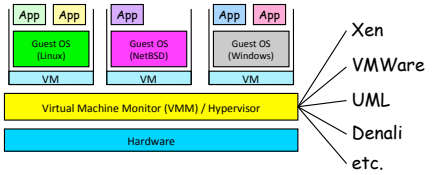

Performance: Para-virtualization (e.g. Xen) is very close to raw physical performance!

75

Summary: Principles for Big-data Systems

Scale "out", not "up"

- Limits of SMP and large shared-memory machines

Move processing to the data

- Cluster have limited bandwidth

Process data sequentially, avoid random access

- Seeks are expensive, disk throughput is reasonable

Seamless scalability

- From the mythical man-month to the tradable machine-hour

76