

Introduction to Cloud Computing

Overview and Introduction

15-319, spring 2010

1st Lecture, Jan 12th

Majd F. Sakr

Why take 15-319?

- **Because you're cool!**
- **Because we're cool!**
- **Gain real world experience and learn new tools**
 - Emerging technology
 - New programming model
 - Could be the future of computing?
 - Running application on very large data-sets



Syllabus: Course Purpose

- **Introduce you to the basics of the emerging cloud computing paradigm**
 - learn how this paradigm came about
 - understand its enabling technologies
 - understand the computer systems constraints, tradeoffs and techniques in setting up and using the cloud
- **Teach you how to implement algorithms in the cloud**
 - gain competence in Hadoop/MapReduce as a programming model for distributed processing of large datasets.
 - understand how different algorithms can be implemented and executed in the Hadoop framework.
 - gain competence in evaluating the performance and identifying bottlenecks when mapping applications to the cloud.

Syllabus: Target Audience

- Juniors & Seniors

- Pre-requisites:
 - 15-213
 - 15-251
 - 15-212
 - Other equivalent courses

Syllabus: Instructor

- **Majd F. Sakr**
- **Office Hours:**
 - Tuesday 3-5pm
 - Welcome when my office door is open
 - By appointment

- **TA: Suhail Rehman**
- **Office Hours:**
 - To be decided NOW
 - By appointment

Syllabus: Course Components

What are we trying to answer?

What is Cloud Computing?


Why Cloud Computing?

What are its challenges and opportunities?



How does Cloud Computing work?

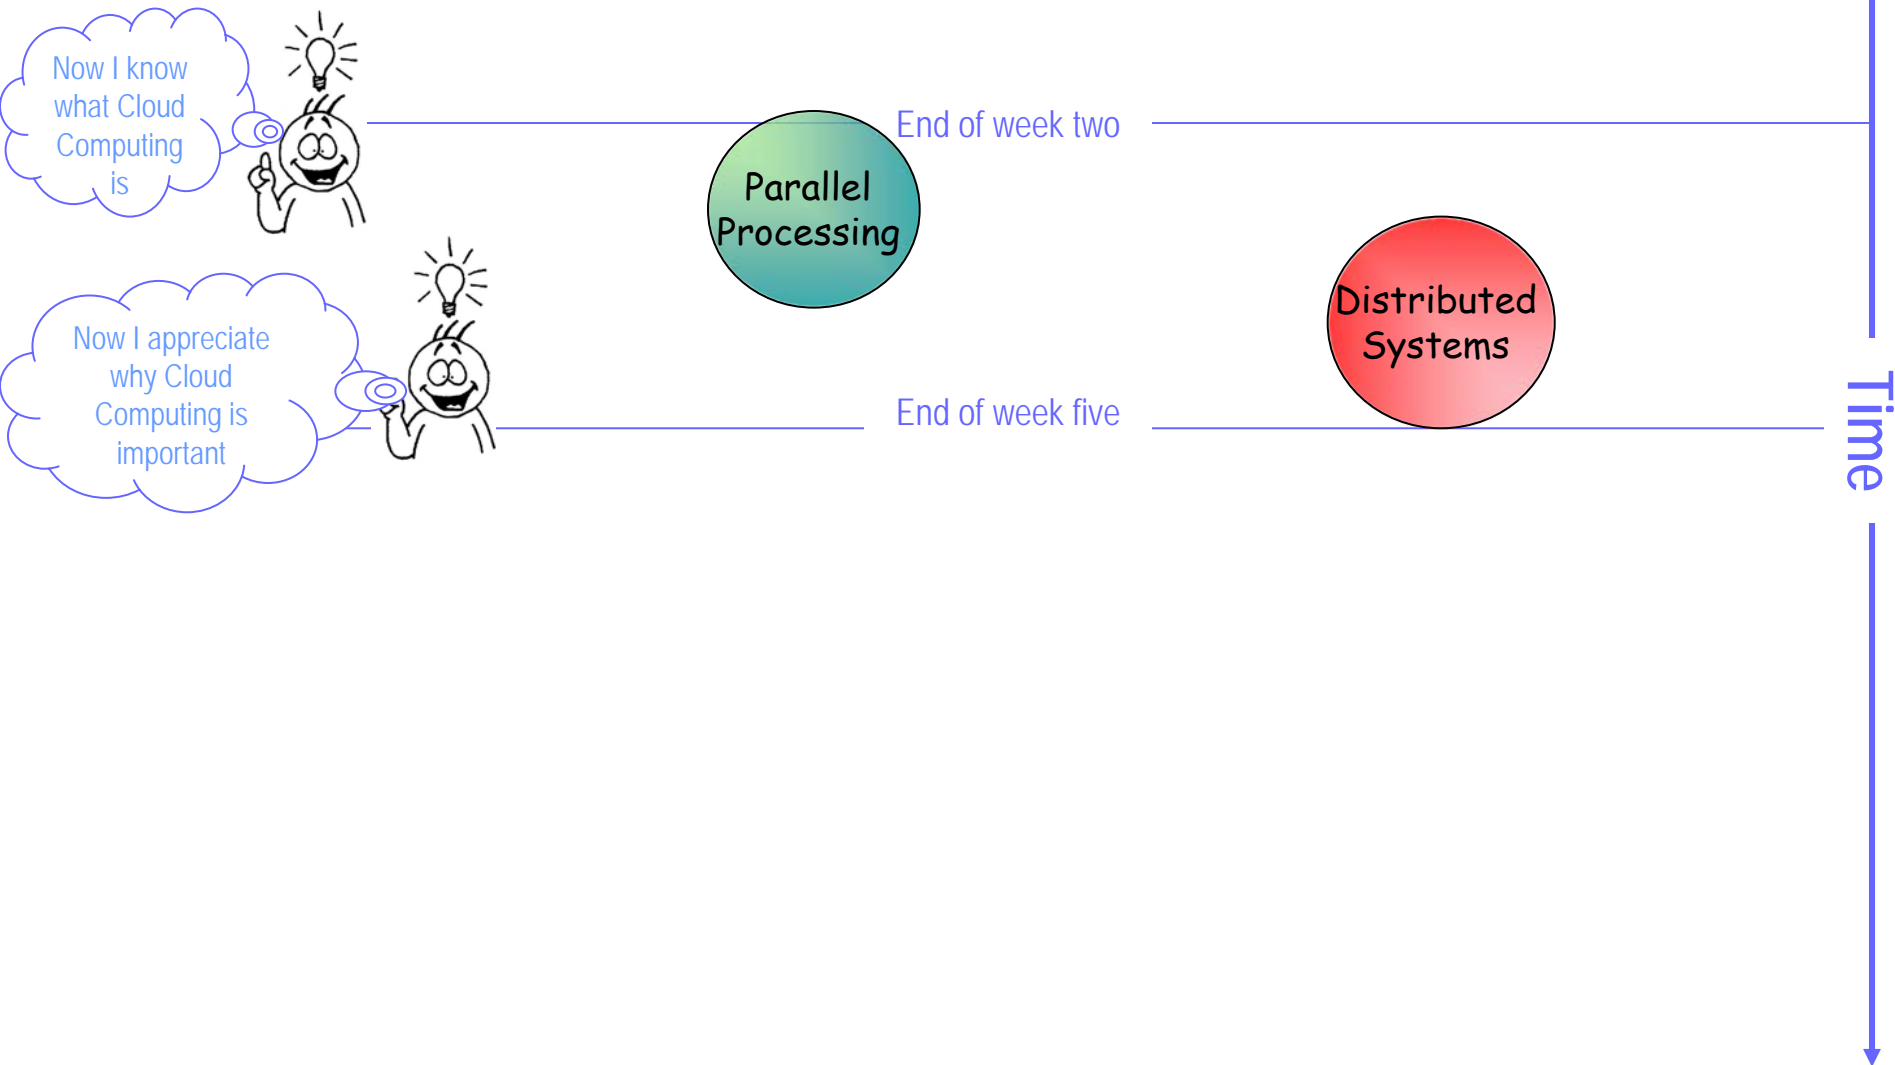
Syllabus: Course Components



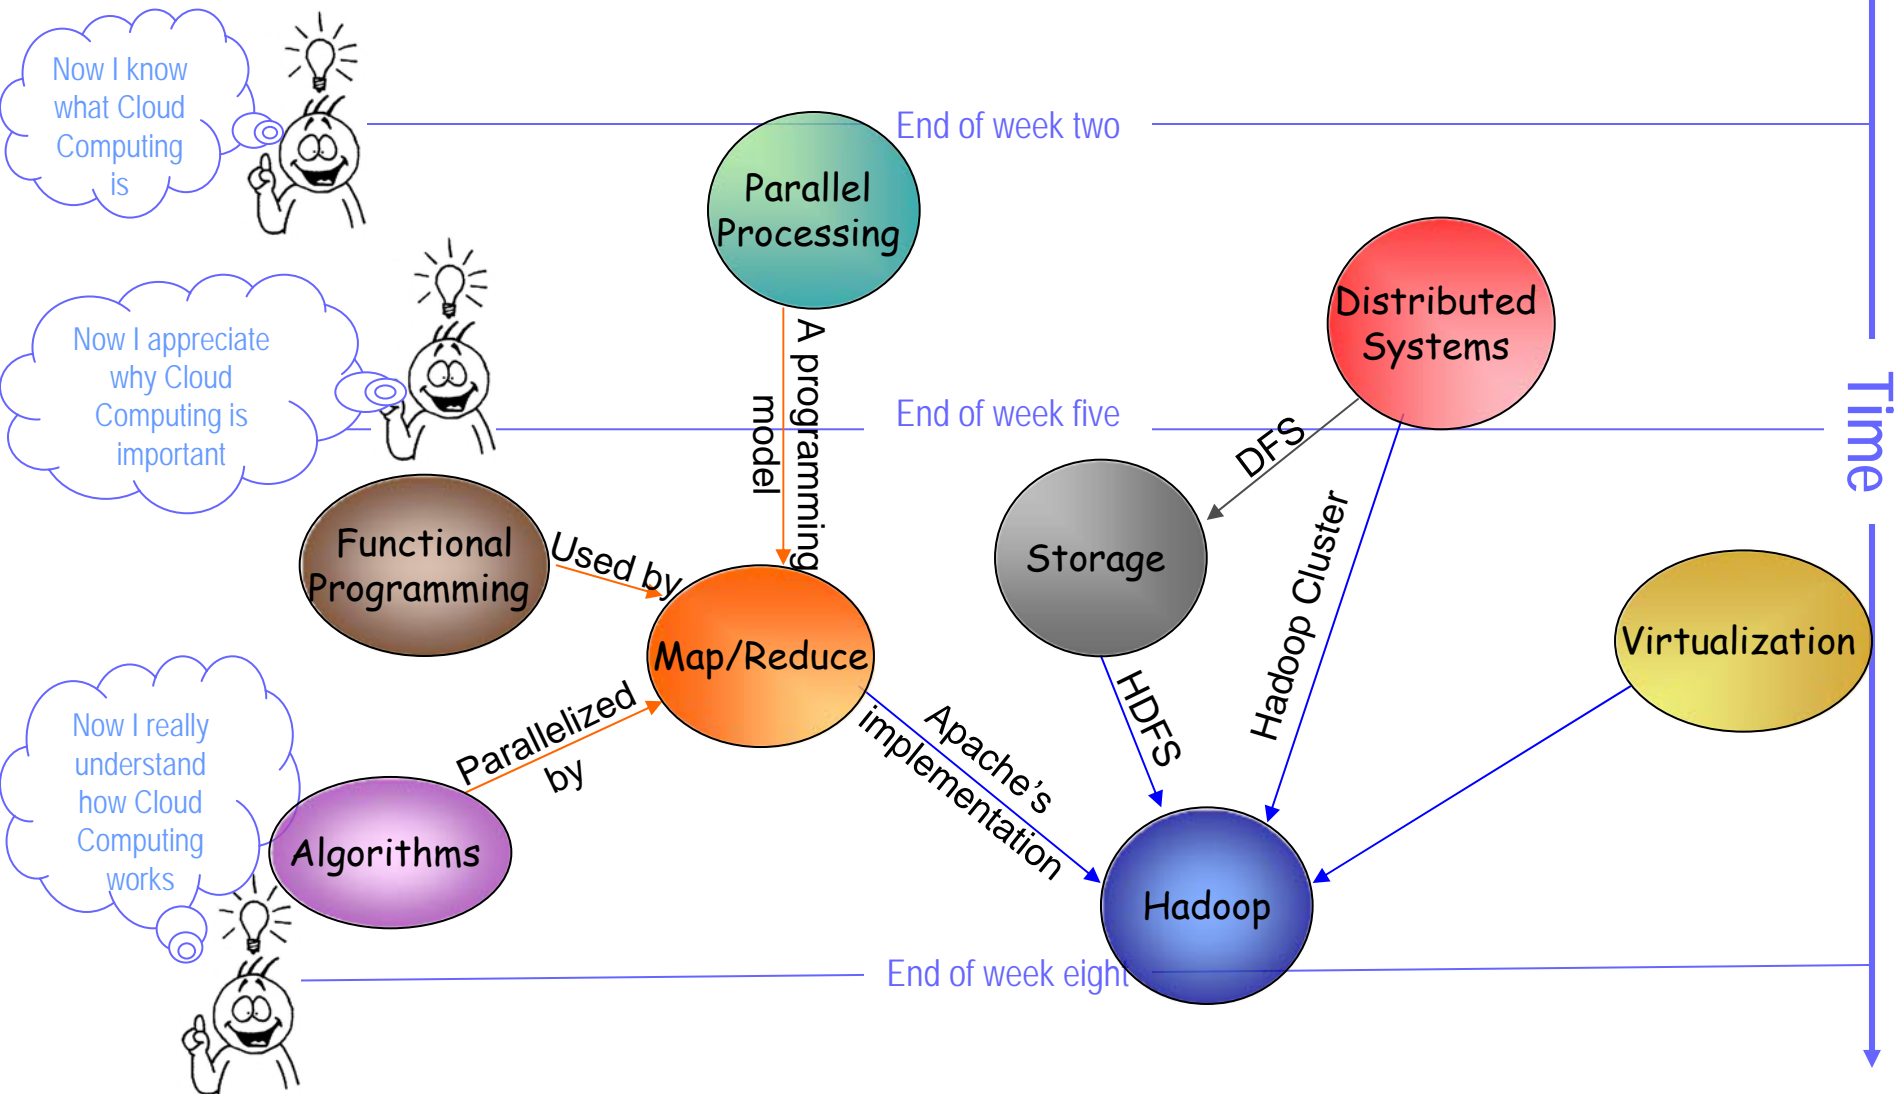
End of week two

Time


Syllabus: Course Components



Syllabus: Course Components




Syllabus: Course Components



Cloud
Computing

Syllabus: Course Components

Now I can play with



hadoop



Syllabus: Text Books

■ Primary Book:

- Tom White,
Hadoop: The Definitive Guide, O'Reilly Media, 2009.

■ Reference Books

- Tanenbaum and van Steen, *Distributed Systems: Principles and Paradigms*, Pearson, 2007.
- Jean Dollimore, Tim Kindberg, George Coulouris, *Distributed Systems: Concepts and Design, Fourth Edition*, Addison Wesley, 2005.
- Randal E. Bryant and David R. O'Hallaron,
Computer Systems: A Programmer's Perspective, Prentice Hall, 2003.
- Patterson and Hennessy, *Computer Organization and Design: The Hardware/Software Interface, Fourth Edition*, Morgan Kaufmann/Elsevier.
- Jason Venner, *Pro Hadoop*, Apress, 2009.


Syllabus: Projects

■ Five assignments

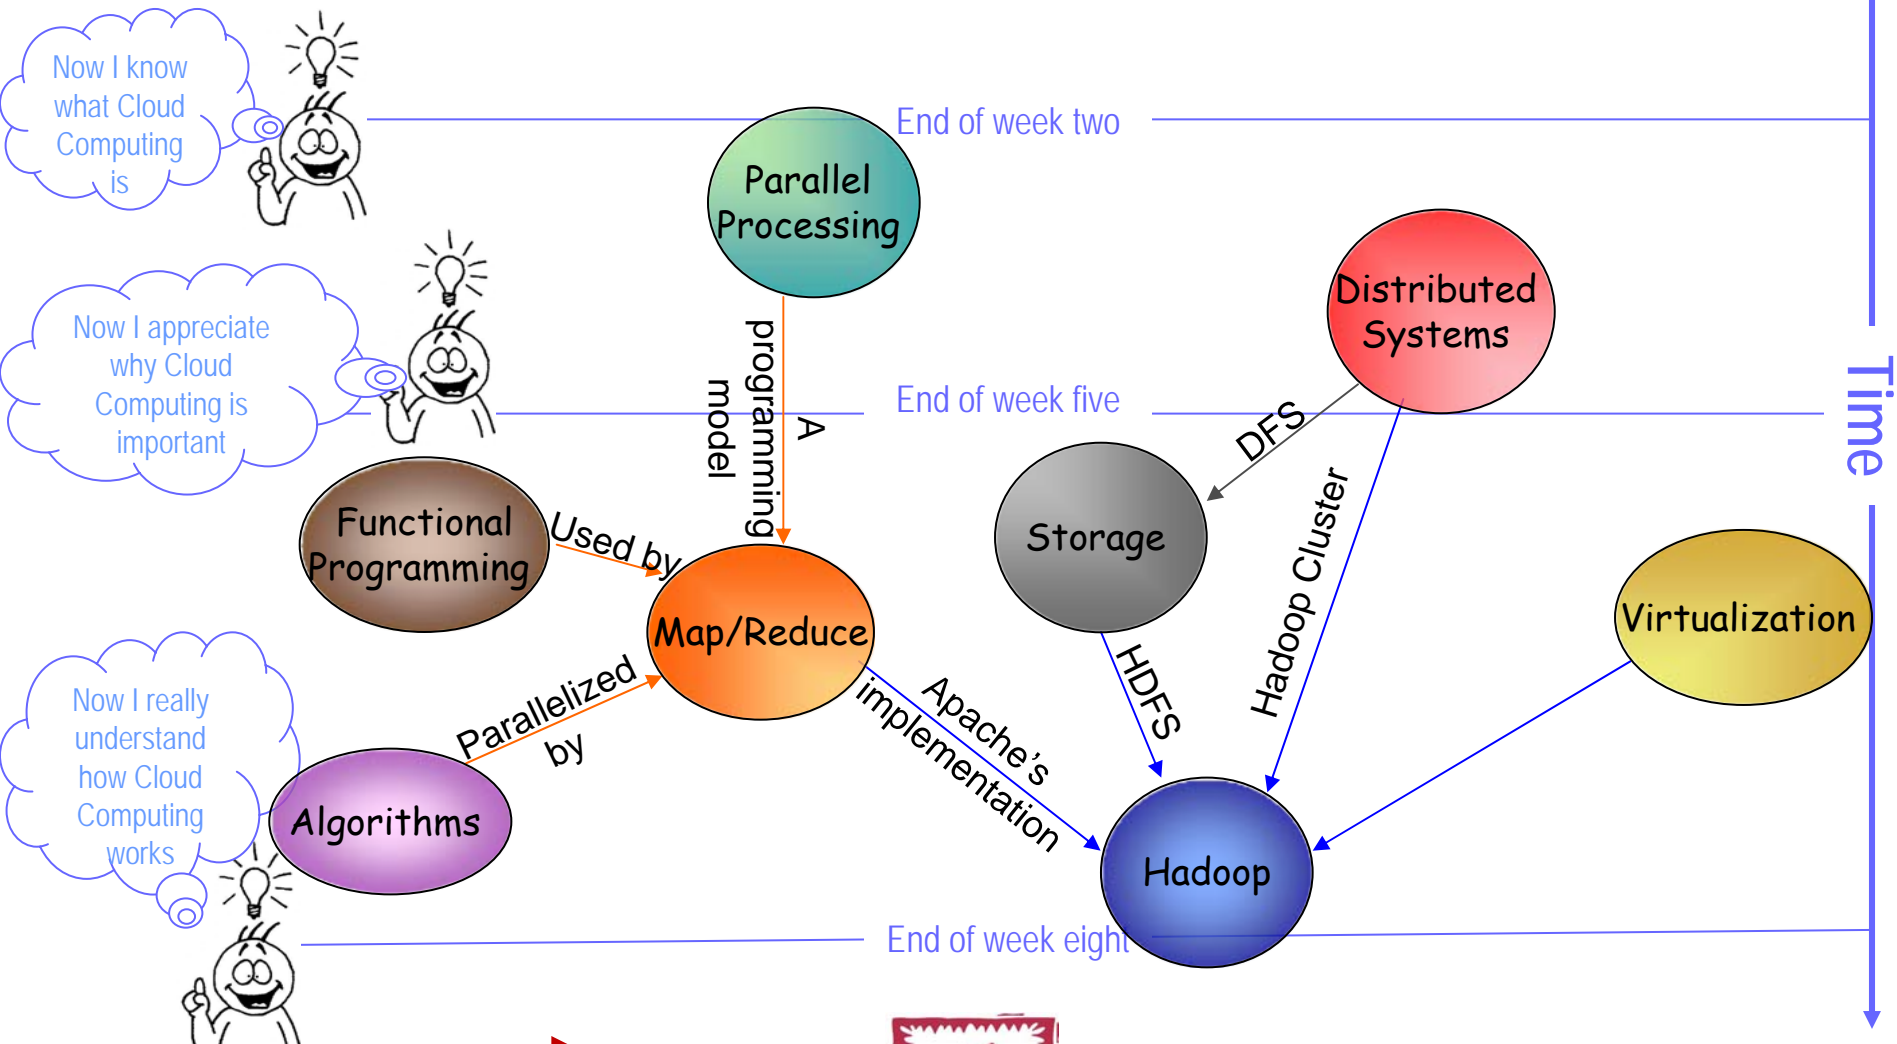
- 2 weeks per assignment
- Technical papers and case studies
- Short write-up
- In-class presentations and discussions

■ Four Projects:

- 4-5 weeks per project
- Final project includes a paper and presentation



Syllabus: Exam



End of week NINE



ONLY 1 EXAM



Syllabus: Grading

- Attendance/Participation 10%
- Assignments 15%
- Projects 60%
 - Project 1: 10%
 - Project 2: 10%
 - Project 3: 15%
 - Project 4: 25%
- Exam 15%

Final Thoughts



What Computing Paradigms Are There?

Existing Computing Paradigms

Personal
Computing

Reconfigurable
Computing

Parallel
Computing

Distributed
Computing

Ubiquitous
Computing

Autonomic
Computing



Super
Computing

Grid
Computing

Cluster
Computing

Utility
Computing

Cloud
Computing

Pervasive
Computing

Mobile
Computing

Existing Computing Paradigms

Personal
Computing

Reconfigurable
Computing

Parallel
Computing

Distributed
Computing

Ubiquitous
Computing

Autonomic
Computing



Super
Computing

Grid
Computing

Cluster
Computing

Utility
Computing


Cloud
Computing

Pervasive
Computing

Mobile
Computing

Personal Computing

- Personal computing system
- Local software installation, maintenance
- Local system maintenance
- Customizable to user needs
- Very low utilization
- High up-front cost



Existing Computing Paradigms

Personal Computing

Reconfigurable Computing

Parallel Computing

Distributed Computing

Ubiquitous Computing

Autonomic Computing



Super Computing

Grid Computing

Cluster Computing

Utility Computing

Cloud Computing

Pervasive Computing

Mobile Computing

Reconfigurable Computing

- **Field Programmable Gate Arrays (FPGAs)**
 - Reprogrammable Hardware
 - Can exploit embarrassingly parallel code
 - Slow programming time (ms)
 - Power hungry

Existing Computing Paradigms

Personal Computing

Reconfigurable Computing

Parallel Computing

Distributed Computing

Ubiquitous Computing

Autonomic Computing



Super Computing

Grid Computing

Cluster Computing

Utility Computing

Cloud Computing

Pervasive Computing

Mobile Computing

Autonomic Computing

- Motivation: rapidly growing complexity of integrating, managing and operating computer systems
- introduced by **IBM** in 2001
- Inspired by Human ANS
- Self-management includes: **self-**

configuration

organization

protection

governing

recovery


healing

autonomy


adaptation

optimization

diagnosis of faults



Complex yet
self-managing



Existing Computing Paradigms

Personal Computing

Reconfigurable Computing

Parallel Computing

Distributed Computing

Ubiquitous Computing

Autonomic Computing



Super Computing

Grid Computing

Cluster Computing

Utility Computing

Cloud Computing

Pervasive Computing

Mobile Computing

Mobile Computing

- You can use computing technology on the move
- Since 1990s
- Intermittent connectivity
- Limited Bandwidth
- Mobile device maturity



Existing Computing Paradigms

~~Personal Computing~~

~~Reconfigurable Computing~~

Parallel Computing

Distributed Computing

Ubiquitous Computing

~~Autonomic Computing~~



Super Computing

Grid Computing

Cluster Computing

Utility Computing


Cloud Computing

Pervasive Computing

~~Mobile Computing~~

Utility Computing

- Water, gas, and electricity are provided to every home and business as commodity services
 - You get connected to the utility companies' "public" infrastructure
 - You get these utility services on-demand
 - And you pay-as-you use
- **Utility Computing** is doing same for **computing resources** (processing power, bandwidth, data storage, and enterprise software services)
- Thought of by 1960s but re-surfaced late 90s
 - "If computers of the kind I have advocated become the computers of the future, then computing may someday be organized as a public utility just as the telephone system is a public utility... The computer utility could become the basis of a new and important industry.,,"
—John McCarthy, MIT Centennial in 1961



Existing Computing Paradigms

~~Personal Computing~~

~~Reconfigurable Computing~~

Parallel Computing

Distributed Computing

Ubiquitous Computing

~~Autonomic Computing~~



Super Computing

Grid Computing

Cluster Computing

~~Utility Computing~~

Cloud Computing

Pervasive Computing

~~Mobile Computing~~

Existing Computing Paradigms - Blue Group



Distributed Computing

Parallel Computing

Super Computing

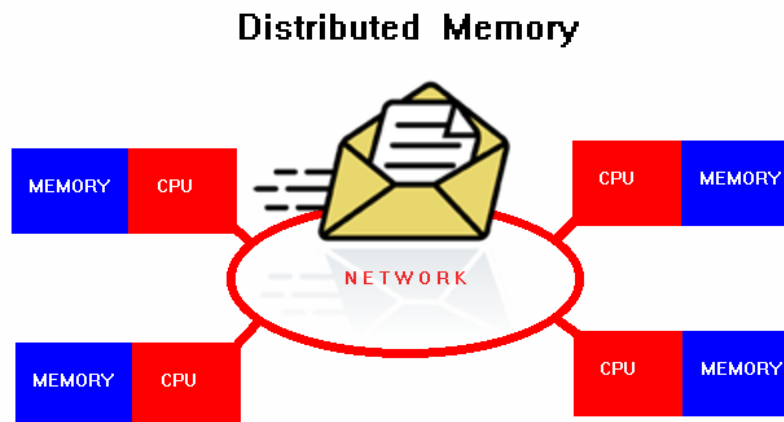
Grid Computing

Cluster Computing

Blue Group

■ Distributed Computing

- Using **distributed systems** to solve large problems.
- **Distributed System:** multiple autonomous computers connected through a communication network
- The system has a **distributed memory** where each processor has its private memory.
- Information exchanged using communication models, ex: **MPI**



Blue Group

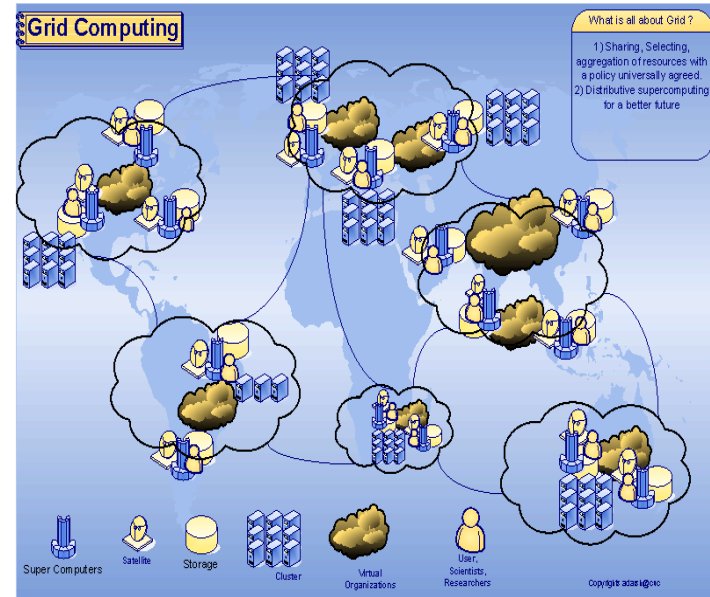
■ Distributed Computing

■ Cluster Computing:

- Characteristics:
 - tightly coupled computers
 - single system image
 - Centralized Job management & scheduling system
- Better performance and availability and more cost-effectiveness over single computer with same capabilities
- Since 1987

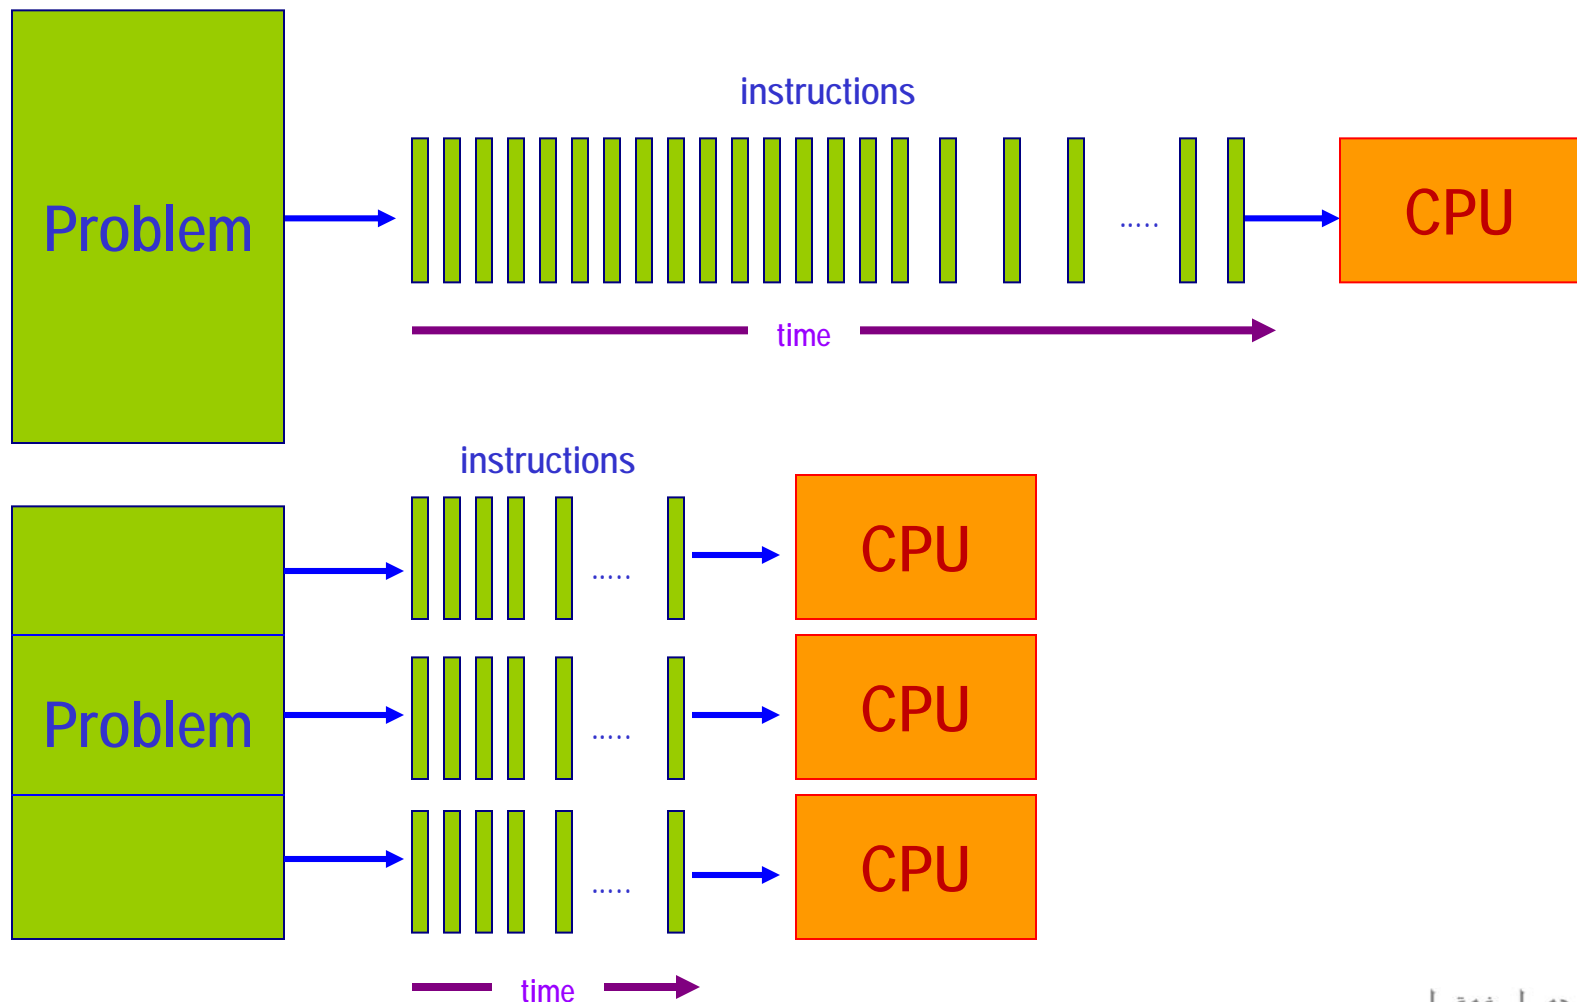
■ Grid Computing:

- According to Gartner, "a grid is a collection of resources owned by multiple organizations that is coordinated to allow them to solve a common problem."
- Characteristics:
 - loosely coupled
 - no Single System Image
 - distributed Job Management & scheduling
- Originated early 1990s



What is Parallel Computing

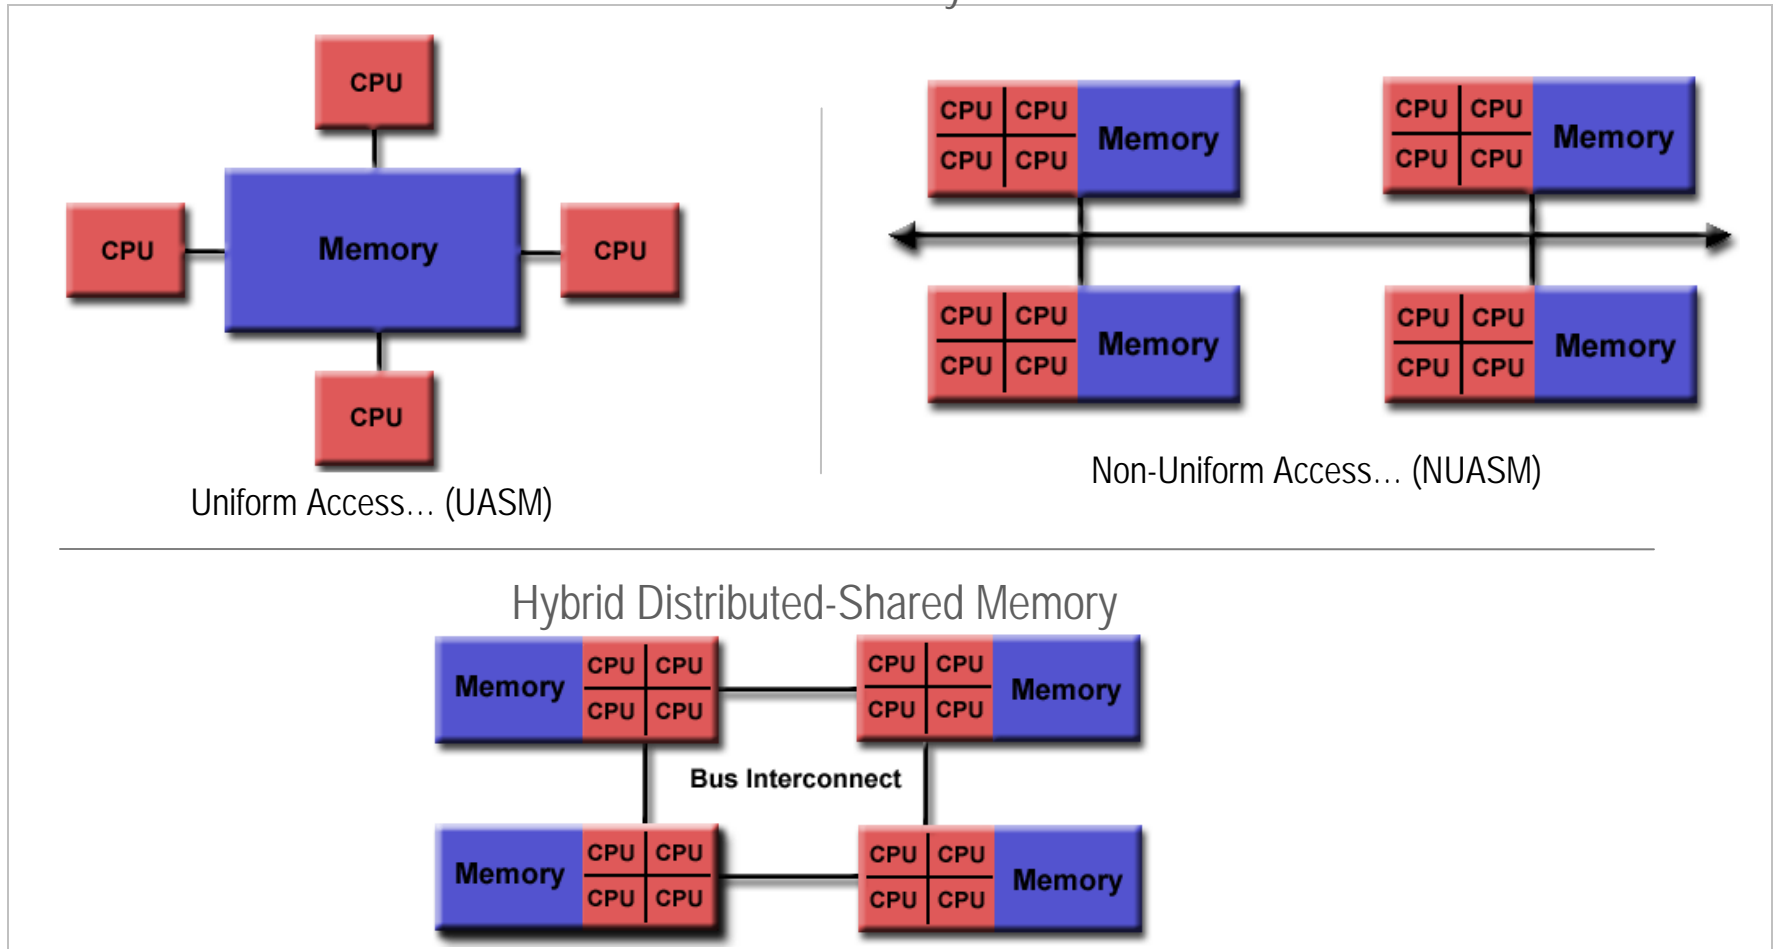
- Calculations of large problems are divided into smaller parts and carried out simultaneously/concurrently on different processors.



■ Parallel Computing

- All have access to a **shared memory** that is used to exchange information between processors

Shared Memory



Source: https://computing.llnl.gov/tutorials/parallel_comp/

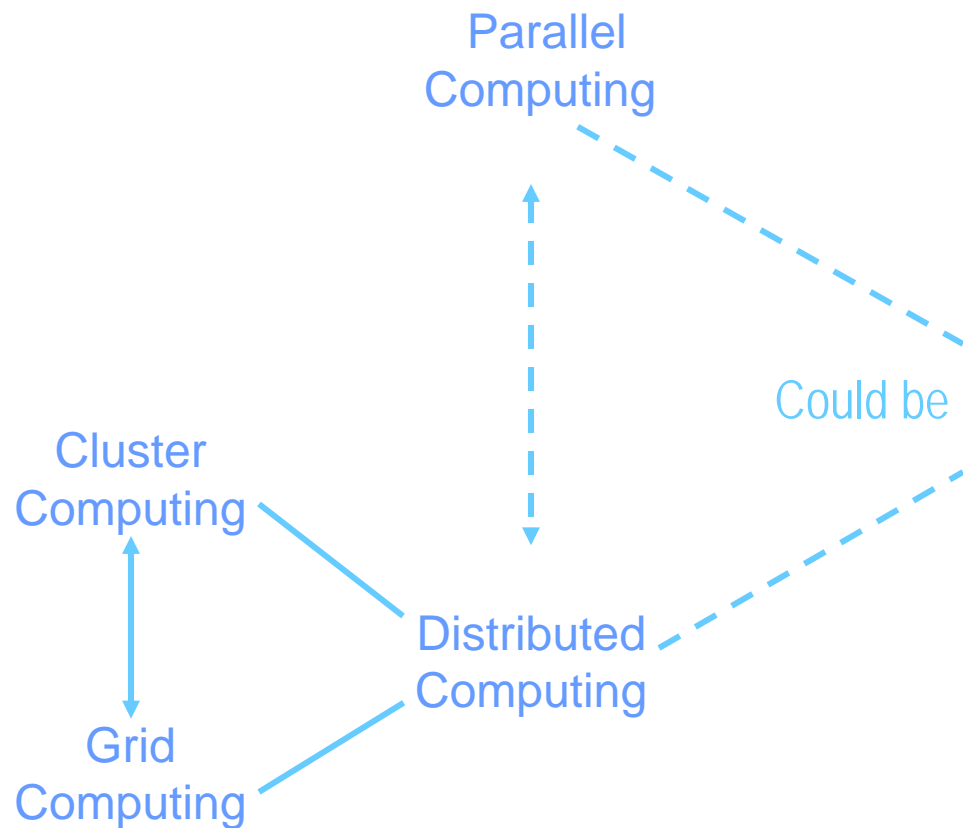
Blue Group

■ Super Computing

- Thousands of processors
- Used for compute-intensive problems
 - Days instead of Years!!!
- introduced in the 1960s



Blue Group



Super Computing



Existing Computing Paradigms

~~Personal Computing~~

~~Reconfigurable Computing~~

~~Parallel Computing~~

~~Distributed Computing~~

~~Ubiquitous Computing~~

~~Autonomic Computing~~



~~Super Computing~~

~~Grid Computing~~

~~Cluster Computing~~

~~Utility Computing~~

~~Cloud Computing~~

~~Pervasive Computing~~

~~Mobile Computing~~

Existing Computing Paradigms - Green Group



Ubiquitous
Computing

Pervasive
Computing

Green Group

- **Ubiquitous**= “seeming to be in all places”
- **Pervasive**= “present or noticeable in every part of a thing or place”
- Information processing engaged in everyday’s activities and objects.
- Term used since 1980s
- Different models but same vision:
 - Small, inexpensive, robust devices distributed throughout everyday’s life



Existing Computing Paradigms

~~Personal Computing~~

~~Reconfigurable Computing~~

~~Parallel Computing~~

~~Distributed Computing~~

~~Ubiquitous Computing~~

~~Autonomic Computing~~



~~Grid Computing~~

~~Super Computing~~

~~Cluster Computing~~

~~Utility Computing~~

~~Cloud Computing~~

~~Pervasive Computing~~

~~Mobile Computing~~

Existing Computing Paradigms



Cloud
Computing

Existing Computing Paradigms

Personal
Computing

Reconfigurable
Computing


Parallel
Computing

Distributed
Computing

Ubiquitous
Computing

Autonomic
Computing

How about cloud
computing?



Grid
Computing

Super
Computing

Cluster
Computing

Utility
Computing

Cloud
Computing

Pervasive
Computing

Mobile
Computing

Think of it this way ...

■ Banking

LIFE ON EARTH by Ham




Think of it this way ...

- Power/ heat/electricity/water supply to your home

Before



Now



Think of it this way ...

■ Transportation

- Which one should you pick?
- Should you buy/rent?



Cloud Computing

Think of it as Internet Computing

- Computation done over the internet
- **Enabled through:**
 - High Bandwidth and High Speed Internet
 - Utility Computing
 - Virtualization
 - ...

Cloud Computing Services

Three basic services:

- **Software as a Service (SAAS) model**
 - Apps through browser
- **Platform as a Service (PAAS) model**
 - Delivery of a computing platform for custom software development as a service
- **Infrastructure as a Service (IAAS) model**
 - Deliver of computer infrastructure as a service
- **XAAS, the list continues to grow...**

Interesting Videos

- SaaS:

<http://www.youtube.com/watch?v=kGUPSvswmY0&feature=related>

- Virtualization:

<http://www.youtube.com/watch?v=p11IJOALS4&feature=related>

- Cloud Computing:

<http://www.youtube.com/watch?v=XdBd14rjcs0&NR=1>