SoC-HUST Summer School – 20-26 June 2012, Hanoi

CP3109: Introduction to Cloud Computing

Teo Yong Meng*

Department of Computer Science
National University of Singapore
Email: teoym@comp.nus.edu.sg
URL: www.comp.nus.edu.sg/~teoym

*Visiting Professor
Shanghai Advanced Research Institute
Chinese Academy of Sciences

What I do?

- **Teaching**
  - Parallel Computing
  - Performance Analysis of Computer Systems
  - Systems Modeling & Simulation
  - Applied Parallel Computing (co-teach with MIT)
  - Computer Systems Engineering (co-teach with MIT)
  - ....

- **Research**
  - parallel & distributed computing
  - performance evaluation

National University of Singapore

- 25K undergraduate + 8K graduate from 88 countries
- 14 faculties/schools

Faculty of Arts and Social Sciences
School of Business
School of Computing
Faculty of Dentistry
School of Design and Environment
Faculty of Engineering
Faculty of Law
Yong Loo Lin School of Medicine
Yong Siew Toh Conservatory of Music
Faculty of Science
University Scholars Programme
Lee Kuan Yew School of Public Policy
NUS Graduate School for Integrative Sciences & Engineering
Duke-NUS Graduate Medical School Singapore

National University of Singapore
School of Computing

- **Established July 1998 (formerly DISCS within FoS)**

- **Departments:**
  - Computer Science
  - Information Systems

- **Staff strength:**
  - 120 (academic staff)
  - 120 (research staff)

- **Student Population**
  ~ 2182 (total):
  - **1636 undergraduates**
  - **546 graduate students (350 PhD students)**

Computer Systems Group - Overview

Cloud Service Models

- **Software-as-a-Service** (SaaS)
- **Platform-as-a-Service** (PaaS)
- **Infrastructure-as-a-Service** (IaaS)

Virtualization Management (application, hardware, network, ..)

(Emerging) Technologies (virtualization, p2p, cloud, web services,..)

model of

fault tolerance

SNAP

CoDES

SkyBoxz
Elastic Computing on Multiple Clouds

STREAM
STraegic-proof REsource Allocation Mechanism

IRON
Idle Resource Overlay Network

TFTTP
Tit-for-Tat File Transfer Protocol

technologies

L0: Overview

"Cloud computing is cool technology,
but every time it rains I lose my data!"

buzzingup.com

Outline

0830-1130

Lecture 1: Principles of Cloud Computing

Lecture 2: Cloud Architecture and Systems

1330-1630

Lecture 3: Programming the Cloud

Lecture 4: Cloud Computing Demo

Course URL:
www.comp.nus.edu.sg/~teoym/CP3109/CP3109-Cloud-Computing.htm

Userid: **as announced** Password: **as announced**

L01: Principles of Cloud Computing

- What is Parallel Computing?
  - Motivation for Parallel Computing
- What is Cloud Computing?
  - Virtualization
  - Key Cloud Characteristics (Features)
  - Cloud Delivery Models
  - Cloud Services Model
  - Technical and Non-technical Challenges
  - Cloud Adoption and Barriers
  - Cloud Economics
- Summary

L02: Cloud Architecture and Systems

- Cloud reference architecture
  - Actors in cloud computing
  - Interactions between the actors
  - Usage scenarios
  - Cloud consumer: available services
  - Cloud provider: major activities
  - Cloud broker: key services
  - Scope of controls between provider and consumer
  - Service orchestration and management
  - Cloud use cases
  - Pros/Cons of service models
- Examples of Systems
  - Amazon Web Services: EC2 and S3
    - AWS ecosystems
    - Regions and availability zones
    - Amazon ‘s global datacenters
    - Amazon EC2
    - Amazon S3
    - Comparison of two leading cloud platforms
  - SkyBoxz: Elastic Computing with Multiple Clouds
- Summary

L03: Programming the Cloud

- Types of Parallel Applications
- Writing Parallel (cloud) Programs
- Parallel Programming Models
- Shared-memory Programming
  - Thread Model
  - What is OpenMP?
  - OpenMP Program to Calculate π
- Distributed-memory (message-passing) Programming
  - What is MPI?
  - MPI Program to Calculate π

L03: Programming the Cloud

- Data-intensive applications
  - What is MapReduce?
  - What is Hadoop?
  - MapReduce Framework
  - Structure of a MapReduce Program
  - High-level View of MapReduce
  - Example: Counting Words
  - Parallelism in MapReduce
  - Applications of MapReduce
- Comparison with Traditional Models
- Summary
- References

L04: Cloud Computing Demo

- Amazon EC2 and S3
  - Running serial, OpenMP and MPI programs
  - Summary

- SkyBoxz Federated Cloud
  - Running Hadoop program

Interesting Videos

- Cloud Computing
  http://www.youtube.com/watch?v=XdBd14rjcs0&NR=1

- SaaS
  http://www.youtube.com/watch?v=kGUPSvswmY0&feature=related

- Virtualization
  http://www.youtube.com/watch?v=p11lJOnALS4&feature=related