Spring 2011, EE 657, USC (updated Nov. 15, 2010)

Parallel and Distributed Computing

Prof. Kai Hwang, Email: kaihwang@usc.edu

Class Section 30994R, 10:00 – 11:45 am MW, Class Room VHE 214

Course Description:

This course covers the system architecture and enabling technologies of parallel and distributed computing systems and their innovative applications. We will cover scalable multicore or multiprocessor systems, distributed clusters of computers, P2P networks, computational Grids, virtual machines, virtual clusters, Internet Clouds, and the Internet of Things (IoT). Case studies include Cray XT-5, TeraGrids, Google AppEngine, Gnuttela, BitTorrent, Facebook, Twitter, content-delivery networks, VMWare Tools, Amazon Elastic Computing Cloud, Saleforce.com cloud, NASA and CERN clouds, Hadoop library, mobile and sensor networks, and some cloud and IoT projects, etc. The course will acquaint M.S. and Ph.D. students in computer science, electrical and computer engineering with Top-500 superclusters and various HPC and HTC systems for parallel and distributed computing in supercomputing, e-commerce, utility computing, cloud services, social networking, and emerging Internet applications.

Course Outline :

(In total, 24 lectures plus mid-term exam, quizzes, and project presentations)

1. **Advanced Processors and Interconnects** (3 lectures)
   Multicore Processors and High-bandwidth Networks
2. ***Clusters of Computers*****:** (3 lectures)
   Server Clusters, High Availability, and Disaster Recovery
3. **Virtualization for Distributed Computing** (3 lectures)
   Vertual Machines and Virtualized Datacenters
4. **Peer-to-Peer Computing Systems:** (3 lectures)
   P2P systems, Overlay networks, and Content Distribution.
5. **Computational Grids and Applications** (3 lectures)
   National or global computing Grids and Applications
6. **Cloud Computing Platforms and programming** (5 lectures)
   Upgraded Web services and Elastic Supercomputing
7. **The Internet of Things and Cyber-Physical Systems** (4 lectures)
   RFID, Sensors, GPS and Clouds for Future Internet

Prerequisite: **EE 557 or equivalent background approved by the Instructor**

Textbook: Hwang, Fox and Dongarra, ***Distributed Systems and Cloud Computing : Clusters, Grids, P2P and Clouds,*** published by Morgan Kaufmann, Publishers, 2011. The final manuscript of 9 Chapters will be handed out to registered students in the EE 657 class of Spring 2011**. (See the new Book Announcement for detailed contents on p.2**)

Grading Policy and Class Procedure:See

1. The course work consists of three parts :
   - **Quizzes in class (30%),** three quizzes will be given covering lecture material
   - **Mid-Term Exam (30%),** tentatively scheduled for late March 2010 after Spring Break.
   - **Final Research Project (40%)** : Project Proposal (5%) in early March, Oral Presentation (10%) during the last two weeks of April, and Final Project Report (25%) due early May 2010
2. **Possible Project Topics** include: multi-core processors, high-bandwidth networks, physical and virtual clusters, HPC clusters, MPP systems, virtual machine migration, disaster recovery, P2P networks, computational grids, network security, wireless sensor networks, cloud platforms, datacenter architecture, Internet of Things, public cloud services, data and copyright protection, reputation systems, trust management, performance modeling, MapReduce programming, benchmark experiments, and innovative applications, etc.

A New Book from Morgan Kaufmann Publishers, an imprint of Elsevier, in press to appear Spring 2011.

Distributed Systems and Cloud Computing : Clusters, Grids, P2P, and Clouds

Kai Hwang, Geoffrey Fox, Jack Dongarra

Book Summary: This new book presents design principles, systems architecture, strategic technologies, programming support, and innovative applications of parallel, distributed, and cloud computing systems over the Internet. This is a textbook, intended for use in a senior- or graduate-level course in computer science, computer engineering, scientific computation, information technology, and business computing services. The book was written by a group of leading scholars and IT experts in distributed systems and Internet computing. The lead authors selected the material and did most of the writing and enhancement. They integrated quite a few contributed sections or subsections from top frontier researchers in the US, Australia, and China. The book gives a balanced coverage and coherent treatment on a wide range of hot topics related to web-scale distributed computing systems.

The major emphases of this book lie in ubiquity, efficiency, scalability, availability, and programmability. Nine chapters feature the frontier research and development of *high-performance computing (*HPC) by scientific community and *high-throughput computing* (HTC) in business applications. Hot systems covered include *massively parallel processors* (MPP), supercomputing clusters, *service-oriented architecture* (SOA), computational grids, *peer-to-peer* (P2P) and social networks, virtualized datacenters, cloud computing platforms, *Internet of things* (IOT), and *cyber-physical systems* (CPS). Over 100 working examples are illustrated with 500 figures/tables plus 200 homework problems. The book serves well as a textbook for students and educators. It can be also used as a reference for scientists, engineers, information technologists, application developers, and computer professionals.

Chapter Contents and Highlights :

Part 1: Systems Modeling, Clustering and Service-Oriented Architecture

Chapter 1: Distributed System Models and Enabling Technologies
Chapter 2: Computer Clusters for Scalable Distributed Computing
Chapter 3: Service Oriented Architectures for Distributed Computing

Part 2: Grids, P2P Systems, Virtualization and Datacenter Design

Chapter 4: Grid Computing Systems and Resource Management
Chapter 5: Peer-to-Peer Computing and Overlay Networks
Chapter 6: Virtual Machines and Virtualized Clusters and Datacenters

Part 3: Computing Clouds, Social Networks and Internet of Things

Chapter 7: Architecture Design of Cloud Computing Platforms
Chapter 8: Cloud Computing Paradigms and Programming Support
Chapter 9: Ubiquitous Computing with Clouds and The Internet of Things

About The Authors and Contributors :

Kai Hwang is a Professor of Electrical Engineering and Computer Science, University of Southern California. He is also an IV-endowed visiting Chair Professor of Tsinghua University in China. He earned the Ph.D. in EECS from the University of California at Berkeley. An IEEE Fellow, he has served as the founding Editor-in-Chief of the *Journal of Parallel and Distributed Computing* for 26 years. Hwang is a world-renowned scholar for his contributions in computer architecture, parallel processing, distributed systems, Internet security, and cloud computing. Many of his Ph.D. graduates are now leaders in the academics or in IT industry in the USA, China, Korea, etc.

Geoffrey Fox is a Professor of Informatics, Computing and Physics and Associate Dean of Graduate studies and Research at the same College, Indiana University. He has taught and led many research groups at Caltech and Syracuse University, previously. A Ph.D. from Cambridge University, U.K., Fox is well known for his comprehensive work and extensive publications in parallel architecture, distributed programming, grid computing, web services, and Internet applications. He has produced 60+ Ph.D. students in computer science and engineering over the years.

Jack Dongarra is a University Distinguished Professor of Electrical Engineering and Computer Science, University of Tennessee and a Distinguished Research Staff, Oak Ridge National Lab. An ACM/IEEE/SIAM/AAAS Fellow, Dongarra pioneered the areas of supercomputer benchmarks, numerical analysis, PDE solvers, and high-performance computing and published extensively in these areas. He leads the Linpack benchmark evaluation of the Top-500 fastest computers every year. He is a Member of the National Academy of Engineering in the USA.

Points of Contact: Kai Hwang, **kaihwang@usc.edu,** Tel. 213 740 4470 at University of Southern California,