


CS 553

Cloud Computing

<http://www.cs.iit.edu/~iraicu/teaching/CS553-S12/>



Spring 2012

Monday/Wednesday, 11:25AM - 12:40PM, Stuart Building 106

Dr. Ioan Raicu

DETAILED COURSE TOPICS:

- **Distributed System Models**
- **Parallel Computing**
- **Virtualization**
- **Cloud Platform Architectures**
 - * Amazon AWS
 - * Microsoft Azure
 - * Google App Engine
 - * Google MapReduce / Yahoo Hadoop
 - * Eucalyptus, Nimbus, OpenStack
- **Service-Oriented Architectures**
- **Cloud Programming**
- **Grid Computing**
- **Peer-to-Peer Computing**



Dr. Ioan Raicu is an assistant professor in the Department of Computer Science (CS) at Illinois Institute of Technology (IIT), as well as a guest research faculty in the Math and Computer Science Division (MCS) at Argonne National Laboratory (ANL). He is also the founder and director of the Data-Intensive Distributed Systems Laboratory (DataSys) at IIT. He has received the prestigious NSF CAREER award (2011 - 2015) for his innovative

work on distributed file systems for exascale computing. He was a NSF/CRA Computation Innovation Fellow at Northwestern University in 2009 - 2010, and obtained his Ph.D. in Computer Science from University of Chicago under the guidance of Dr. Ian Foster in March 2009. He is a 3-year award winner of the GSRP Fellowship from NASA Ames Research Center. His research work and interests are in the general area of distributed systems. His work focuses on a relatively new paradigm of Many-Task Computing (MTC), which aims to bridge the gap between two predominant paradigms from distributed systems, High-Throughput Computing (HTC) and High-Performance Computing (HPC). His work has focused on defining and exploring both the theory and practical aspects of realizing MTC across a wide range of large-scale distributed systems. He is particularly interested in resource management in large scale distributed systems with a focus on many-task computing, data intensive computing, cloud computing, grid computing, and many-core computing. His work has been funded by the NASA Ames Research Center, DOE Office of Advanced Scientific Computing Research, the NSF/CRA CIFellows program, and the NSF CAREER program. He is a member of the IEEE and ACM.

Overview

Cloud Computing is “A large-scale distributed computing paradigm that is driven by economies of scale, in which a pool of abstracted, virtualized, dynamically-scalable, managed computing power, storage, platforms, and services are delivered on demand to external customers over the Internet.” It has become a driving force for information technology over the past several years, and it is hinting at a future in which we won’t compute on local computers, but on centralized facilities operated by third-party compute and storage utilities. Governments, research institutes, and industry leaders are rushing to adopt Cloud Computing to solve their ever-increasing computing and storage problems arising in the Internet Age. There are three main factors contributing to the surge and interests in Cloud Computing: 1) rapid decrease in hardware cost and increase in computing power and storage capacity, and the advent of multi-core architecture and modern supercomputers consisting of hundreds of thousands of cores; 2) the exponentially growing data size in scientific instrumentation/simulation and Internet publishing and archiving; and 3) the wide-spread adoption of Services Computing and Web 2.0 applications. This course is a tour through various topics and technologies related to Cloud Computing. We will explore solutions and learn design principles for building large network-based systems to support both compute and data intensive computing across geographically distributed infrastructure. Topics include resource management, programming models, application models, system characterizations, and implementations. Our discussions will often be grounded in the context of deployed Cloud Computing systems, such as Amazon EC2 and S3, Microsoft Azure, Google AppEngine, Eucalyptus, Nimbus, OpenStack, Google’s MapReduce, Yahoo’s Hadoop, Microsoft’s Dryad, Sphere/Sector, and many other systems. The course involves lectures, outside invited speakers, discussions of research papers, programming assignments, and a major project (including both a written report and an oral presentation).

We will be using the textbook *Distributed and Cloud Computing: Clusters, Grids, Clouds, and the Future Internet* by Kai Hwang, Jack Dongarra & Geoffrey C. Fox.