

## COLLOQUIUM

### *Scalable and Elastic Data Management in the Cloud*

Amr El Abbadi

University of California, Santa Barbara

**Abstract:** Over the past two decades, database and systems researchers have made significant advances in the development of algorithms and techniques to provide data management solutions that carefully balance the three major requirements when dealing with critical data: high availability, reliability, and data consistency. However, over the past few years the data requirements, in terms of data availability and system scalability, from Internet scale enterprises that provide services and cater to millions of users has been unprecedented. Cloud computing has emerged as an extremely successful paradigm for deploying Internet and Web-based applications. Scalability, elasticity, pay-per-use pricing, and autonomic control of large-scale operations are the major reasons for the successful widespread adoption of cloud infrastructures. In this talk, we analyze the design choices that allowed modern scalable data management systems to achieve orders of magnitude higher levels of scalability compared to traditional databases. With this understanding, we highlight some design principles for data management systems that can be used to augment existing databases with new cloud features such as scalability, elasticity, and autonomy. We then analyze several state of the art systems and discuss our proposed system, G-Store, which provides transactional guarantees on data granules formed on-demand while being efficient and scalable.

Bio: Amr El Abbadi is currently a Professor in the Computer Science Department at the University of California, Santa Barbara. He received his B. Eng. in Computer Science from Alexandria University, Egypt, and received his Ph.D. in Computer Science from Cornell University in August 1987. Prof. El Abbadi is an ACM Fellow. He has served as a journal editor for several database journals, including, currently, The VLDB Journal. He has been Program Chair for multiple database and distributed systems conferences, most recently SIGSPATIAL GIS 2010 and ACM Symposium on Cloud Computing (SoCC) 2011. He has also served as a board member of the VLDB Endowment from 2002-2008. In 2007, Prof. El Abbadi received the UCSB Senate Outstanding Mentorship Award for his excellence in mentoring graduate students. He has published over 250 articles in databases and distributed systems.

Friday, February 10, 2012, 3:00 pm  
Mathematics and Science Center: W301

MATHEMATICS AND COMPUTER SCIENCE  
EMORY UNIVERSITY