TEMPLE UNIVERSITY®

Spring 2013 Colloquium

Center for Networked Computing
Computer and Information Sciences

Performance Tuning in Clouds: A Multidimensional Approach

Distinguished Speaker

Chita Das
Distinguished Professor
Department of Computer Science and Engineering
Pennsylvania State University

Monday March 4, 2pm, Wachman 1015D

Abstract:
Cloud computing is an emerging paradigm that has the potential to transform the future computing landscape for many application domains. With cloud computing, both hardware and software resources can be leased from cloud providers, thereby avoiding the capital investment on procuring and maintaining such infrastructures. Therefore, various pay-as-you-go models like Saas, PaaS, and IaaS are becoming attractive options to achieve better economy of scale, elasticity and availability. However, a major roadblock for wide-scale adoption of clouds is the performance unpredictability of applications due to factors such as system scale, workload dynamism, virtualization overheads, and resource sharing and contention.
In this talk, I will discuss several issues that are critical for performance enhancement in clouds and summarize some of our current efforts for alleviating the performance problem. In this context, two complementary scheduling techniques for enhancing cloud performance will be discussed. The first technique will highlight the importance of an important cloud workload property, called task placement constraints and the second technique will summarize the impact of a fine-grained resource scheduling for data analytic systems like MapReduce. Next, I will summarize an end-to-end fault diagnosis framework for virtualized cloud platforms. The talk will conclude with a simple performance model, called D-factor, for estimating application performance.

Bio:
Chita Das is a Distinguished Professor of Computer Science and Engineering at the Pennsylvania State University. His main areas of interest include parallel and distributed computer architectures, multi-core architectures, mobile computing, performance evaluation, and fault-tolerant computing. In particular, he has worked extensively in the area of design and analysis of interconnection networks/on-chip interconnects. He has published more than 150 papers in the above areas, has received several best paper awards, and has served on many program committees, and editorial boards. He is a Fellow of the IEEE.