Spring 2013 Course (Grad/Undergrad)
CSE 489-02 / CSE 589-02 – Advances in Computing & Data Management

Instructor: A. Rezgui Email: rezgui@cs.nmt.edu Office: Cramer 224

Class Meeting Time: Monday/Wednesday, 9:30am-10:45am (Cramer 227)
Objective: The objective of this course is to explore recent advances in computing and data management through an in-depth study of recent literature. The course is offered for grad and undergrad students.
Assignments: Students will have to spend a significant amount of time reading recent papers presented at major conferences (e.g., ICDCS, PODC, Supercomputing, SIGMOD, VLDB, ICDE, PODS) or recently published in high impact journals. While **undergrad students will not be subject** to the same grading scheme as grad students, the assignments and requirements are the same. They include:

1. **Presentations:** Students will have to give one or more presentations throughout the semester.
2. **Projects/Papers:** Students will have to organize into groups of two or three. Each group will have to deliver two papers:
   (i) A survey paper: which will be a literature review with analysis (e.g., evaluation, comparison, etc.) of some relevant topic.
   (ii) A research paper: where students will have to demonstrate some degree of innovation, e.g., exploring a new idea, improving an existing algorithm/technique, etc. This second paper would normally be a continuation of the work done in the first paper. For example, based on their literature review, students could identify a significant problem and then work on designing, implementing, and evaluating a solution to that problem.

Organization & Content: Typically, lectures will be tutorial in nature, i.e., introducing topics that are then explored more in-depth by studying content from other sources. In addition to the lectures, some classes will be for presentations by students. With the instructor's permission, students may choose the topics that they desire to explore. Topics covered will be selected from a spectrum of topics that includes:

Advances in Computing	Advances in Data Management
Multicore/multiprocessor systems, accelerators	New database architectures
Memory, interconnection networks	Storage & indexing
Storage and file systems	Query processing & optimization, approximate query processing
Parallel and multicore programming languages and compilers	NoSQL databases
GPU computing, social computing, green computing	Big data, MapReduce
Cloud computing and virtualization	Replication and consistency
Cluster and grid computing, crowdsourcing	Streams, spatial, temporal, multimedia, scientific, and sensor databases
Communication networks, protocols, architectures, services, applications	Text databases and keyword search
Concurrent programming, synchronization, shared and transactional memory	Data security and privacy
Inter-cloud deployment of applications	Uncertainty and reasoning at scale
Multi-tenancy	Data cleaning
Applications of parallel and distributed computing	Database models, schema matching, data integration
Fault tolerance, reliability, availability	Data-center systems
Recommendation systems, P2P systems	Performance and evaluation, benchmarking, experimental methodology at scale

Grades: Grades will be based on the following distribution:

Participation/Quizzes: 20% Presentations: 30%
Midterm Exam/Survey Paper 25% Final Exam/Research Paper: 25%