TEXAS TECH UNIVERSITY
DEPARTMENT OF COMPUTER SCIENCE
WHITACRE COLLEGE OF ENGINEERING

U-REASON SEMINAR SERIES FALL - 2013

An Introduction to Goblin-Core64: A Massively Parallel Processor Architecture Designed for Complex Data Analytics

By John Leidel
Texas Tech University

Date: November 26th, 2013 (Tuesday)
Time: 3:40pm-4:40pm
Venue:ECE 226 (Bullen Room)

Faculty Coordinator: Dr. Yong Chen (yong.chen@ttu.edu)
Student Coordinators: Navaneeth Thiagarajan, Dan Ferguson, Lakhan Jhawar

Abstract:

Given the decrease in sustainable manufacturing techniques, recent core processor architectures have relied upon increasing degrees of parallelism in order to maintain peak theoretical performance scalability. However, many of these architectures maintain the same core architectural techniques as were applied over the past two decades. The use of multi-level cache hierarchies, strong memory ordering and complex CISC [VLIW] techniques have failed to deliver scalable performance, especially for applications associated with large-scale analytics. In this talk, we introduce Goblin-Core64, a massively parallel processor architecture designed to efficiently solve non-deterministic or sparse algorithmic workloads commonly referred to as Big Data problems. The end result is a highly scalable, parallel processor infrastructure that natively maps modern parallel programming constructs, compilation techniques and low-latency runtime methodologies in order to efficiently solve problems in large-scale analytics.

Speaker Bio:

John Leidel is a PhD student under the advisement of Dr. Yong Chen in the Data Intensive Scalable Computing Lab. His research interests include massively parallel processor architectures, runtime models and compiler optimizations designed for non-linear or sparse applications. He holds patents and publications in the areas of multithreaded processor architecture, low-latency runtime systems and high bandwidth memory architecture. In his previous role as Chief Software Architect at Convey Computer Corp, he lead development efforts of the MX-100 platform which still remains the highest single node performance result on the Graph500 list.