The Department of Electrical and Computer Engineering (http://www.ece.uah.edu/) and the LaCASA Laboratory (http://lacasa.uah.edu) invite you to a research seminar.

Guest Speaker: **Dr. Martin Burtscher**
Texas State University, San Marcos, TX

Title: **An Efficient GPU Implementation of a Tree-based n-Body Algorithm**

Time&Place: **04/19/2013 (Friday), 11:00 - 12:00, EB 258, UAH**

Abstract
Compute GPUs offer an order of magnitude higher performance, energy efficiency, and price/performance than multicore CPUs, but it is substantially harder to write efficient code for GPUs. This is particularly true for irregular programs that operate on dynamic data structures, i.e., programs whose control flow and memory accesses are highly data dependent and statically unpredictable. Nevertheless, several irregular applications have already been accelerated using GPUs.
In this talk, I will describe how compute GPUs work, explain what makes them difficult to program, and present a complete CUDA implementation of the irregular Barnes-Hut n-body algorithm. This algorithm repeatedly builds an unbalanced tree data structure and performs complex traversals on it, making it difficult to map to GPUs. I will discuss how this irregular algorithm can be implemented to exploit the GPU's architecture, including features that are commonly viewed as hurdles to achieving high performance. On a five million star simulation running on a GTX 480, the CUDA code is over twenty times faster than an optimized OpenMP version running on a high-end hex-core Xeon X5690.

Speaker Biography: Martin Burtscher is Associate Professor in the Department of Computer Science at Texas State University-San Marcos. He received the combined BS/MS degree in computer science from the Swiss Federal Institute of Technology (ETH) Zurich in 1996 and the Ph.D. degree in computer science from the University of Colorado at Boulder in 2000. Martin's research interests include efficient parallelization of programs for GPUs and multicore CPUs, automatic performance assessment and optimization, and high-speed data compression. He is a senior member of the IEEE, its Computer Society, and the ACM. Martin has co-authored over 70 peer-reviewed publications, including a book chapter in NVIDIA's GPU Computing Gems, is the recipient of an NVIDIA Academic Partnership award, and is the PI of a CUDA Teaching Center.