USC Viterbi School of Engineering
Ming Hsieh Department of Electrical Engineering

Course Announcement

EE 599: ACCELERATED COMPUTING USING FPGAS

TTH 330-450, LAB/DISCUSSION F 4-520

SPRING 2020

NO. OF UNITS: 2

INSTRUCTOR: VIKTOR K. PRASANNA

Recently, Field Programmable Gate Arrays have become a key computing platform to accelerate applications at data center, cloud and at the "edge". This course will review the technology and software tools from application acceleration perspective and discuss (application-specific) architectural, software and algorithmic innovations to realize the potential of this technology to optimize latency, throughput and energy efficiency.

Prerequisite: (EE 457 and CS 570) or consent of the instructor. Recommended preparation: EE 451. This course does NOT require any background in FPGA.

Text: Course will be based on recent research publications, survey articles and lecture notes by the instructor.

Course Grade: based on participation (10%), lab assignments (20%), project proposal (20%), presentation (20%), and final report (30%).

Project: The focus of the course is in designing accelerators using FPGAs. The project will be focused on specific application areas of interest to the students to identify a problem that needs acceleration, design an application specific architecture, develop scalable parallel algorithm and map it onto a target FPGA device. The project will consist of literature survey, problem definition, solution idea, hardware design and use of software tools to map the design to a FPGA. It will consist of proposal preparation, discussions with the instructor and the TA, present details of the design and implement it and report the resulting acceleration.
Sample project: Parallelizing LSTM models on FPGAs with coherent memory. Identifying opportunities for parallelism, surveying state of the art techniques for kernels and primitives, performance modeling and estimating projected performance. Implementation in VHDL or Verilog, synthesis, place and route results. Summarizing latency and throughput performance and energy dissipation.

Course Outline:

1. Introduction: Computing platforms and technology evolution
2. FPGA basics, architectural characteristics
3. FPGA abstractions and Computational models
4. Accelerating Dense Algebra
5. Accelerating FFT
6. Accelerating Networking (SDN)
7. Accelerating Networking (NFV)
8. Accelerating ML Kernels
9. Accelerating ML Kernels
10. Accelerating AI at the Edge
11. FPGAs in the Cloud

Professor Viktor K. Prasanna
Email: prasanna@usc.edu
Ext: 0-4483
Office: EEB-200C