ISTeC
The Information Science and Technology Center

COLORADO STATE UNIVERSITY

Distinguished Lectures
Fall 2022

Dr. Viktor K. Prasanna

Charles Lee Powell Chair in Engineering
Ming Hsieh Department of Electrical & Computer Engineering
Professor of Computer Science
University of Southern California

FPGA Accelerators in the Cloud

Monday, October 17, 2022
Reception with refreshments: 10:30 a.m.
Lecture: 11:00 a.m.-12:00 noon
Lory Student Center Ballroom-350D

Accelerating Graph Neural Networks

Tuesday, October 18, 2022
Lecture: 10:00-11:00 a.m.
Lory Student Center 308

Sponsored by
Colorado State University's Information Science and Technology Center (ISTeC)

In conjunction with the Department of Computer Science and Department of Electrical and Computer Engineering Seminar Series

Abstracts

FPGA Accelerators in the Cloud

With recent dramatic advances in Field Programmable Gate Arrays (FPGAs), these devices are being used along with multi-core and novel memory technologies to realize advanced platforms to accelerate complex applications in the Cloud. We will review advances in reconfigurable computing over the past 25 years leading up to accelerators for data science. We will illustrate FPGA-based parallel architectures and algorithms for a variety of data analytics kernels in streaming graph processing and graph machine learning. While demonstrating algorithm-architecture co-design methodology to realize high performance accelerators for graphs and machine learning, we demonstrate the role of modeling and algorithmic optimizations to develop highly efficient Intellectual Property (IP) cores for FPGAs. We show improved performance for two broad classes of graph analytics: iterative graph algorithms with variable workload (e. g., graph traversal, shortest paths, etc.) and machine learning on graphs (e. g., graph embedding). For variable workload iterative graph algorithms, we illustrate dynamic algorithm adaptation to exploit heterogeneity in the architecture. We conclude by identifying opportunities and challenges in exploiting emerging heterogeneous architectures composed of multi-core processors, FPGAs, GPUs and coherent memory.

Accelerating Graph Neural Networks

Recently, Graph Neural Networks (GNNs) have been used in many applications leading to improved accuracy and fast approximate solutions. Training as well as Inference in these networks is computationally demanding. Challenges include access to irregular data, large scale sparse as well as dense matrix computations, limited data reuse and heterogeneity in the various stages of the computation. This talk will review our recent work in the Data Science Lab (dslab.usc.edu) and FPGA/Parallel Computing Lab (fpga.usc.edu) at USC leading up to current trends in accelerators for data science. For graph embedding, we develop GraphSAINT, a novel computationally efficient technique using graph sampling and demonstrate scalable performance. We develop graph processing over partitions (GPOP) methodology to handle large scale graphs on parallel platforms. On a current FPGA device, we demonstrate up to 100X and 30X speed up for full graph GNN computations compared with state-of-the-art implementations on CPU and GPU respectively. We also demonstrate specific accelerators for two widely used GNN models: GraphSAGE and GraphSAINT. We conclude by identifying opportunities and challenges in exploiting emerging heterogeneous architectures towards a general framework for GNN acceleration.

Speaker Biography

Viktor K. Prasanna (sites.usc.edu/prasanna) is Charles Lee Powell Chair in Engineering in the Ming Hsieh Department of Electrical and Computer Engineering and Professor of Computer Science at the University of Southern California. He is the director of the Center for Energy Informatics at USC and leads the FPGA (fpga.usc.edu) and Data Science Labs (dslab.usc.edu). His research interests include parallel and distributed computing, accelerator design, reconfigurable architectures and algorithms and high performance computing. He serves as the Editor-in-Chief of the Journal of Parallel and Distributed Computing. Prasanna was the founding Chair of the IEEE Computer Society Technical Committee on Parallel Processing. He is the Steering Chair of the IEEE International Parallel and Distributed Processing Symposium. He is a Fellow of the IEEE, the ACM and the American Association for Advancement of Science (AAAS). He is a recipient of 2009 Outstanding Engineering Alumnus Award from the Pennsylvania State University and a 2019 Distinguished Alumnus Award from the Indian Institute of Science. He received the 2015 W. Wallace McDowell award from the IEEE Computer Society for his contributions to reconfigurable computing. He is a member of Academia Europaea.

To arrange a meeting with the speaker, please contact Prof. H.J. Siegel, HJ@ColoState.edu.

ISTeC (Information Science and Technology Center) is a university-wide organization for promoting, facilitating, and enhancing CSU's research, education, and outreach activities pertaining to the design and innovative application of computer, communication, and information systems. For more information please see ISTeC.ColoState.edu.