USC Viterbi
School of Engineering
Ming Hsieh Department of Electrical and Computer Engineering

Ming Hsieh Institute Seminar Series

Computer Engineering

Processing in Memory (PIM) – Power and Thermal Challenges and Opportunities

Mircea Stan
University of Virginia

Friday, November 15th, 2019
2:00-3:00pm
RTH 211

Abstract: Memory technology is a defining component of modern computing and has a strong impact on performance, power and cost of computing systems. However, the advances in memory performance have not been able to keep up with the performance advances for CPU's, thus leading to what is known as "the memory wall." Depending on the application, the memory wall manifests itself both in terms of memory latency, as well as memory bandwidth. An interesting solution to the memory wall problem is to bring memory closer to the processor, or vice-versa, to move some processing capability in the memory itself – this leads to variations of what is known as Near-Memory Processing, Processing in Memory (PiM), etc. This seminar will first introduce a PIM taxonomy along several dimensions of the PiM design space; it will then follow with a history of PIM, then go over several recent PIM examples. The seminar will then go in depth into the Thermal/Power delivery challenges for PIM that are a result of the increased switching activities inherent to the moving of processing into the memory fabric, and exacerbated by the evolution towards 3D integration due to the slow-down of traditional Moore's law methods. The seminar will conclude with some novel solutions that alleviate the Thermal/Power challenges for PiM.

Bio: Mircea R. Stan received the Ph.D. (1996) and the M.S. (1994) degrees from UMass Amherst and the Diploma (1984) from the Polytechnic Institute in Bucharest, Romania. Since 1996 he has been with the ECE Department at UVa, where he is now the Virginia Microelectronics Consortium (VMEC) Professor. Prof. Stan is teaching and doing research in the areas of high-performance low-power VLSI, temperature-aware circuits and architecture, embedded systems, spintronics, and nanoelectronics. He leads the High-Performance Low-Power (HPLP) lab, is an associate director of the Center for Automata Processing (CAP) and an assistant director of the Center for Research in Intelligent Storage and Processing-in-Memory (CRISP). He was a visiting faculty at UC Berkeley in 2004-2005, at IBM in 2000, and at Intel in 2002 and 1999. He received the 2018 Influential ISCA Paper Award (For 2003 paper "Temperature-aware microarchitecture"), the NSF CAREER award in 1997 and was a co-author on best paper awards at LASCAS19, SELSE17, ISQED08, GLSVLSI06, ISCA03 and SHAMAN02 and IEEE Micro Top Picks in 2008 and 2003. He gave keynotes at DCAS18, SOCC16, CogArch16, WoNDP15, iNIS15 and CNNA14. He was the chair of the VSA-TC of IEEE CAS in 2005-2007, general chair for ISLPED06 and GLSVLSI04, TPC chair for SOCC18, ISVLSI17, NanoNets07 and ISLPED05, and on technical committees for numerous conferences. He is Associate Editor-in-Chief for the IEEE TVLSI, Senior Editor for the IEEE TNano, AE for IEEE Design & Test, and was an AE for the IEEE TNano in 2012-2014, IEEE TCAS I in 2004-2008 and for the IEEE TVLSI in 2001-2003. Prof. Stan is a fellow of the IEEE, a member of ACM, and of Eta Kappa Nu, Phi Kappa Phi and Sigma Xi.