ECE DISTINGUISHED SPEAKER SERIES

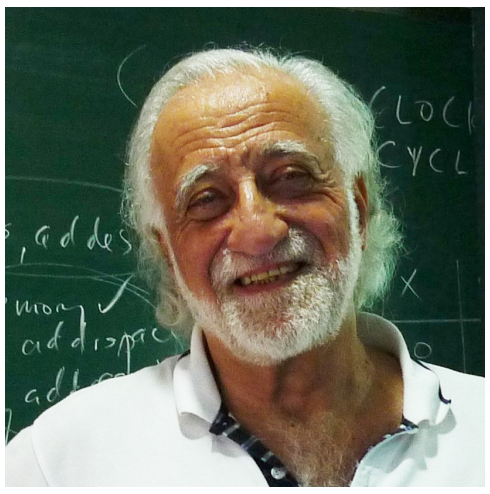

Yale Patt

The University of Texas at Austin

Host: Professor David Kaeli

The End of Von Neumann, The End of Moore's Law: What will the Microprocessor of 2025 look like

Friday, March 27

Room 378 • 140 The Fenway

3:30 pm

Reception to follow

Sponsored by the Department of Electrical and Computer Engineering

At Micro-47 in Cambridge, England, last December, we held a debate, the hypothesis being: it is the end of the Von Neumann architecture. Almost every speaker started by defining the Von Neumann Architecture! How can one debate whether something is over if we can not even agree on what it is? Another popular mantra, although one that I have been hearing for the past 40 years, is that Moore's Law is about to come to end, and even before it does, our continual growth in performance is already coming to an end. Third, we hear people trumpet the benefits of parallelism, whether or not they have any clue whatsoever as to what they are talking about. This is due in large part to the highly visible and well advertised continuing (temporarily, at least) benefits of more transistors on the chip, allowing designers to build the wildest of accelerators touted as non-Von Neumann architectures, unmindful of the two serious bugs that all those transistors come with. By 2025 we will have 50 Billion transistors on a chip, probably about 1000 cores and plenty of non-Von Neumann accelerators – whether we can effectively utilize them or not does not seem to curb the enthusiasm. What I would like to do today is discuss the Von Neumann architecture, why I think it is not at the end of the road and is in fact necessary for continued performance, and Moore's Law, what it has given us and what problems it provides. I would also like to examine parallelism, note that it did not start with the multicore chip, observe some of the silliness it has generated, and identify its fundamental pervasive element. Finally, I would like to discuss the microprocessor of 2025, and the transformation hierarchy, bereft of fanfare but essential to continuing our quest for higher and higher performance.

Yale Patt is the Ernest Cockrell, Jr. Centennial Chair in Engineering at The University of Texas at Austin, where he teaches the required intro to computing course to 400 freshmen every other fall, the advanced graduate course in microarchitecture to students wishing to do PhD work in the field or move onto work for leading chip manufacturers every other spring, and the senior course in computer architecture whenever they let him.

He started his career at Northeastern where he earned the BS in electrical engineering a long time ago, then moved on to get appropriate graduate degrees from another reputable university. He is responsible for leading the HPS work in the 1980s which provided out-of-order execution while maintaining precise exceptions and the two-level branch predictor work in the 1990s. These two concepts completely transformed our thinking when it came to designing high-end microprocessors.

He has won many awards for his research contributions, including the highest honor in computer architecture, the IEEE/ACM Eckert-Mauchly award in 1996 and the first IEEE Bob Rau Award in 2011. Notwithstanding, his first love is teaching, and importantly, how we introduce newcomers to computing. His freshman textbook, *Intro to Computing: from bits and gates to C and beyond*, co-authored with his former PhD student and current Professor at Illinois Sanjay Patel, eschews conventional wisdom in favor of his motivated bottom-up approach. Dr. Patt is a Fellow of the IEEE and the ACM, and a member of the National Academy of Engineering.

More detail can be found on his web site, at **www.ece.utexas.edu/~patt.**

Northeastern

ECE Distinguished Speaker Series Committee Members:

Prof. Edmund Yeh (eyeh@ece.neu.edu)
Prof. Mark Niedre (M.Niedre@neu.edu)
Prof. Matteo Rinaldi (rinaldi@ece.neu.edu)