

Hamed El Lozy, PhD, North Carolina State University, 1973  
Completeness Properties in m-Valued Logic

John Swensen, PhD, University of California, Berkeley, 1987  
High-Bandwidth/Low-Latency Temporary Storage for Supercomputers

Wen-mei Hwu, PhD, University of California, Berkeley, 1987  
HPSm: Exploiting Concurrency to Achieve High Performance in A Single-Chip Microarchitecture

Stephen Melvin, PhD, University of California, Berkeley, 1990  
Performance Enhancement through Dynamic Scheduling and Large Execution Units in Single Instruction Stream Processors

Ashok Singhal, PhD, University of California, Berkeley, 1990  
Exploiting Fine Grain Parallelism in Prolog

Chien Chen, PhD, University of California, Berkeley, 1991  
Scheduling Heuristics and Runtime Data Structures for the Parallel Execution of PROLOG Programs

Michael Butler, PhD, University of Michigan, 1993  
Aggressive Execution Engines for Surpassing Single Basic Block Execution

Tse-Yu Yeh, PhD, University of Michigan, 1993  
Two-Level Adaptive Branch Prediction and Instruction Fetch Mechanisms for High Performance Superscalar Processors

Michael Shebanow, PhD, University of California, Berkeley, 1994  
The Importance of Bus Transaction and Cache Consistency Protocol in the Design of High Performance Shared-Memory, Single-Bus Multiprocessor Systems

Robert Hou, PhD, University of Michigan, 1994  
Improving Reliability and Performance of Redundant Disk Arrays by Improving Rebuild Time and Response Time

Gregory Ganger, PhD, University of Michigan, 1995  
System-Oriented Evaluation of I/O Subsystem Performance

Bruce Worthington, PhD, University of Michigan, 1995  
Aggressive Centralized and Distributed Scheduling of Disk Requests

Po-Yung Chang, PhD, University of Michigan, 1997  
Classification-Directed Branch Predictor Design

Eric Hao, PhD, University of Michigan, 1998  
Block Enlargement Optimizations for Increasing the Instruction Fetch Rate in Block-Structured Instruction Set Architectures

Sanjay Patel, PhD, University of Michigan, 1999  
Trace Cache Design for Wide-Issue Superscalar Processors

Jared Stark IV, PhD, University of Michigan, 1999  
Out-of-Order Fetch, Decode, and Issue

Marius Evers, PhD, University of Michigan, 1999  
Improving Branch Prediction by Understanding Branch Behavior

Paul Racunas, PhD, University of Michigan, 2002  
Reducing Load Latency through Memory Instruction Characterization

Robert Chappell, PhD, University of Michigan, 2003  
Simultaneous Subordinate Microthreading

Mary Brown, PhD, University of Texas at Austin, 2005  
Reducing Critical Path Execution Time by Breaking Critical Loops

Onur Mutlu, PhD, University of Texas at Austin, 2006  
Efficient Runahead Execution Processors

Hyesoon Kim, PhD, University of Texas at Austin, 2007  
Adaptive Predication via Compiler-Microarchitecture Cooperation

Moinuddin Qureshi, PhD, University of Texas at Austin, 2007  
Adaptive Caching for High-Performance Memory Systems

Francis Tseng, PhD, University of Texas at Austin, 2007  
Braids: Out-of-Order Performance with Almost In-Order Complexity

Muhammad Aater Suleman, PhD, University of Texas at Austin, 2010  
An Assymmetric Multi-Core Architecture for Efficiently Accelerating Critical Paths in Multithreaded Programs

Chang Joo Lee, PhD, University of Texas at Austin, 2010  
DRAM-Aware Prefetching and Cache Management

Eiman Ebrahimi, PhD, University of Texas at Austin, 2011  
Fair and High Performance Shared Memory Resource Management

Veynu Narasiman, PhD, The University of Texas at Austin, 2014  
An Enhanced GPU Architecture for Not-So-Regular Parallelism with Special Implications for Database Search

Khubaib, PhD, The University of Texas at Austin, 2014  
Performance and Energy Efficiency via an Adaptive MorphCore Architecture

Rustam Miftakhutdinov, PhD, The University of Texas at Austin, 2014  
Performance Prediction for Dynamic Voltage and Frequency Scaling

Jose Joao, PhD, The University of Texas at Austin, 2014  
Bottleneck Identification and Acceleration in Multithreaded Applications

Milad Hashemi, PhD, The University of Texas at Austin, 2016  
On-Chip Mechanisms to Reduce Effective Memory Access Latency

Faruk Guvenilir, PhD, The University of Texas at Austin, 2019  
Scalable Virtual Memory via Tailored and Larger Page Sizes

Ben Ching-Pei Lin, PhD, The University of Texas at Austin, 2021  
Mitigating Bank Conflicts in Main Memory via Selective Data Duplication and Migration

Stephen Pruett, PhD, The University of Texas at Austin, 2022  
Maintaining High Performance in the Presence of Impossible-to-Predict Branches

Siavash Zangeneh Kamali, PhD, The University of Texas at Austin, 2022  
Using Convolutional Neural Networks to Improve Branch Prediction