

[LOCA SERIES: Addressing Challenges in Core Microarchitecture Research](#)

Abstract

Core microarchitecture research is the study of how to implement an Instruction Set Architecture for a general-purpose processor. The field has been studied for decades, but remains crucial due to the evolving demands of modern computing workloads. Growing instruction footprints, the influx of massive data into the processor, the overhead of modern programming languages, and the emphasis on productivity over performance all require innovative approaches. As Moore's Law reaches its end, the onus of improving performance and efficiency falls on microarchitecture research. Additionally, with more and more companies opting to design their own processors, academia is tasked not only with developing new processing technologies but also training the workforce to design these new chips.

In this talk, I will motivate the need for continued core microarchitecture research, give some recent examples of topics we study such as instruction fetch, address translation, and cache management, and give some insight into the challenges we face in this kind of work. For example, branch prediction has been a well-studied topic for decades, but recent trends in software design have caused huge growth in instruction footprints, putting pressure on other areas of instruction fetch as well as overwhelming the capacity of modern branch predictors and ultimately leading to performance degradation. We demonstrate a practical solution for this problem. I will go deep into a couple of microarchitecture research projects and describe the impact our work has had on thought in industry and academia over the years, in particular describing the path from my early work on branch prediction to the state of the art today. I will end by discussing our future work in this area.

Short bio

Daniel A. Jiménez is a Professor in the Department of Computer Science and Engineering at Texas A&M University. Previously he was Assistant and later Associate Professor in the Department of Computer Science at Rutgers University, then Associate Professor, then Professor and Chair of the Department of Computer Science at The University of Texas at San Antonio. Daniel received his Ph.D. in Computer Sciences from the University of Texas at Austin. Daniel's research is in microarchitecture, including microarchitectural prediction and cache management. He pioneered the development of neural-inspired branch predictors that have been implemented in millions of processors sold by IBM, AMD, Oracle, Samsung, and others.

Daniel designed the neural branch predictors used in the popular Samsung Galaxy S7/8/9/10/20. His 2001 paper on perceptron-based branch prediction won the "HPCA Test of Time Award" in 2019. Daniel won the 2021 IEEE CS B. Ramakrishna Rau Award for contributions to neural branch prediction. He is an IEEE Fellow, an ACM Distinguished Scientist, an NSF CAREER award winner, and member of the ISCA,

MICRO, and HPCA halls of fame. He is the Chair of the IEEE Computer Society Technical Committee on Computer Architecture (TCCA), co-Chair of the ISCA Steering Committee, and Vice-Chair of the ACM Special Interest Group on Microarchitecture (SIGMICRO). He was General Chair of IEEE HPCA in 2011, Program Chair for IEEE HPCA in 2017, and Selection Committee Chair for IEEE Micro "Top Picks" in 2020. He is Program Co-Chair for MICRO 2024.

Speakers

Speaker: Daniel A. Jiménez. Professor in the Department of Computer Science and Engineering at Texas A&M University.

Host: Xavier Martorell. Computer Sciences Department, BSC.
Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on 3 Abr 2025 - 13:45): <https://www.bsc.es/es/research-and-development/research-seminars/loca-series-addressing-challenges-core-microarchitecture-research>