August 03, 2017 | By Doug Ramsey

New Computer Engineering Professor Crosses Fields of Computer Architecture and System Software

She won't take up her new assistant professorship at the University of California San Diego until January, but Jishen Zhao is the latest hire into the university's Computer Science and Engineering (CSE) department. Currently an assistant professor of computer engineering at the University of California, Santa Cruz, Zhao intends to expand on her prior work in computer architecture and software systems, with a particular emphasis on memory and storage systems, domain-specific acceleration, high-performance computing and energy efficiency. Other research interests include electronic design automation and VLSI design for three-dimensional integrated circuits and nonvolatile memories.
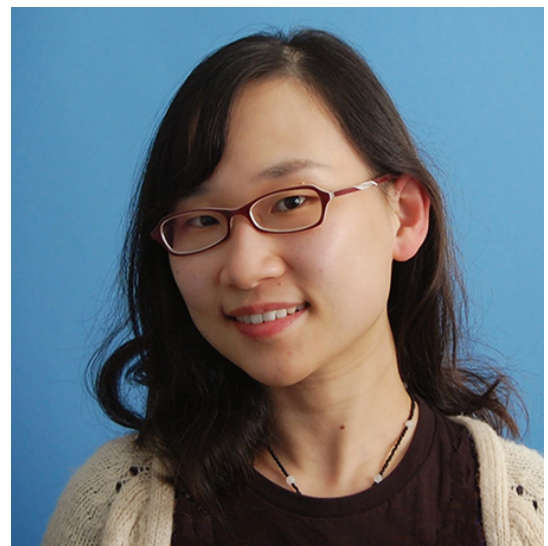
Incoming CSE assistant professor Jishen Zhao leaves UC Santa Cruz for UC San Diego. Her start date is January 1, 2018.

"The fields of computer architecture and systems have seen tremendous achievements in preserving continuous performance growth of computers," said Zhao. "Now two major challenges remain: the complexity and richness of application demands continue to overwhelm hardware capabilities, and new physical devices and circuit technologies are developed with properties that disrupt previous design assumptions – thus making current computer system designs suboptimal."

Prior to joining the UC Santa Cruz Baskin School of Engineering faculty in 2015, Zhao was a research scientist at HP Labs, where she worked in both the Systems Research Lab and HP's Intelligent Infrastructure Lab. She joined the company after completing her Ph.D. in Computer

Science and Engineering at Pennsylvania State University in 2014. She received her undergraduate and M.S. degrees in Electrical Engineering from China's Zhejiang University in 2003 and 2008, respectively.

Professor Zhao's honors include an NSF CAREER Award (2017), which came with nearly $500,000 over five years to pursue her research on enhancing data center storage systems with persistent memory. Persistent memory was also the focus of her seminar at UC San Diego in April, when she laid out a hardware-driven approach to "unlocking the full potential of persistent memory."Persistent memory was also the underlying topic that Zhao addressed in a paper to the Design Automation Conference (DAC 2017) earlier this summer in Austin, TX. Co-authored by Zhao with scientists from National Taiwan University and Academia Sinica, the joint paper explored a new persistent memory-accelerating approach.*
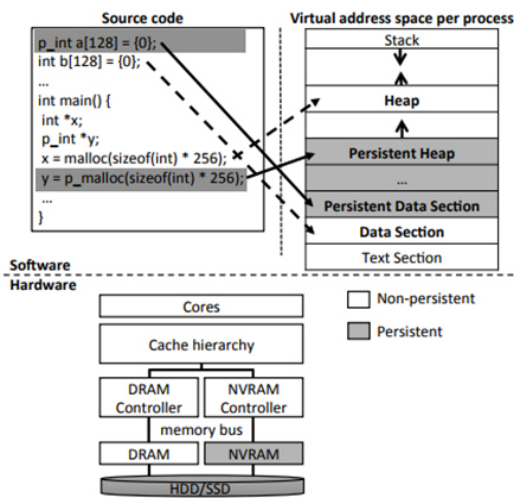
Overview of memory systems integrated with persistent memory [Source: DAC 2017 paper*]

Zhao's other honors include an honorable mention for the best-paper award at the International Symposium on Microarchitecture (MICRO 2013). The paper presents the first hardware-based persistent memory design that outperforms all previous software-based solutions.

Her latest publications include two papers accepted for presentation at the 2017 International Conference on Computer Aided Design (ICCAD 2017) in Irvine, CA. Both relate to magnetic random access memory (MRAM): "Approximate Image Storage with Multi-level Cell STT-MRAM Main Memory" and "PRESCOTT: Preset-based Cross-Point Architecture for Spin-Orbit-Torque Magnetic Random Access Memory."

Apart from her previous experience at HP Labs, Professor Zhao brings to UC San Diego a strong set of research relationships with industry, notably companies such as AMD, Intel, Micron, and Western Digital. "My collaborations exposed me to the technical challenges at the cutting edge of future computer systems but also promoted my research impact," explained Zhao. Her work on GPU system design using 3D+2.5D integration technology followed her collaboration with AMD, one of the top GPU vendors. Added Zhao: "Shortly after my work was published, both AMD and NVIDIA began to adopt such technology in their high-bandwidth, memory-based GPU products. My research also contributed nine patents to industry's intellectual property portfolios."

While her work addresses both technologies and applications roughly equally, Zhao admits that modern applications such as web search are usually executed on tens of thousands of interconnected processors. "This calls for a much broader research paradigm beyond stand-alone computers," she noted, "so I plan to explore digital ecosystems that are composed of interconnected sensors, mobile and cloud servers. I would like to research the scalability of my standalone computer designs as well as the problems posed by the heterogeneity of digital ecosystems."

In CSE, Jishen Zhao looks forward to collaborating with faculty in the Non-Volatile Systems Lab, including Steven Swanson. Continuing her dual focus on applications and technologies, she sees research directions driven by advancements on each count. With new technology (in the form of persistent memory) expected to take its place as a new tier in a data center's memory/storage stack because of their promising density and cost, Zhao plans to study next-generation designs that aim to use persistent memory in the cloud. "It's promising," she noted. "But this technical transition fundamentally changes current memory and storage system design assumptions… and introduces critical design challenges that demand a full-scale rethinking of data center memory and storage systems."

On the application side, Zhao's future research will hinge on machine learning's growth as a key enabler in a broad range of application domains (among them computer vision, autonomous navigation and social network analytics). In the past, the incoming UC San Diego professor has designed effective GPUs and custom accelerator application-specific integrated circuits (ASICs) and their clusters. "My team has optimized both architectures for neuromorphic computing, data mining and pattern recognition," recalled Zhao. "In the future, I want to investigate architectural support for reducing the performance constraints posed by specialization, and for increasing the programmability of special-purpose accelerators, while reducing the energy overhead for communication and moving data between accelerator clusters."

At UC San Diego, Zhao expects to teach both undergraduate and graduate courses, including introductory computer engineering courses such as computer architecture, computer organization and systems programming, digital systems and parallel computing. For the undergraduate courses, she will also design at least one project for each course to allow students to gain general skills that are usually not taught in the classroom. Said Zhao: "These skills are usually gained through immersively diving into every implementation detail – and finishing entire projects individually."

*C.-H. Lai, J. Zhao and C.-L. Yang, "Leave the Cache Hierarchy Operation As It Is: A New Persistent Memory Accelerating Approach (PDF)," *Design Automation Conference (DAC 2017)*, June 2017, Austin, TX.

MEDIA CONTACT

Doug Ramsey, , dramsey@ucsd.edu

UC San Diego's Studio Ten 300 offers radio and television connections for media interviews with our faculty, which can be coordinated via studio@ucsd.edu. To connect with a UC San Diego faculty expert on relevant issues and trending news stories, visit https://ucsdnews.ucsd.edu/media-resources/faculty-experts.