

## Rethinking Memory System Design (and the Computing Platforms We Design Around It)

**Onur Mutlu**

ETH Zurich

**Monday, 27.02.2017, 14:50 – 16:20, TU Dresden, Fakultät Informatik  
Nöthnitzer Str. 46, Raum APB-3105**

### **Abstract:**

The memory system is a fundamental performance and energy bottleneck in almost all computing systems. Recent system design, application, and technology trends that require more capacity, bandwidth, efficiency, and predictability out of the memory system make it an even more important system bottleneck. At the same time, DRAM and flash technologies are experiencing difficult technology scaling challenges that make the maintenance and enhancement of their capacity, energy efficiency, and reliability significantly more costly with conventional techniques. In fact, recent reliability issues with DRAM, such as the RowHammer problem, are already threatening system security and predictability.

In this talk, we first discuss major challenges facing modern memory systems in the presence of greatly increasing demand for data and its fast analysis. We then examine some promising research and design directions to overcome these challenges and thus enable scalable memory systems for the future. We discuss three key solution directions: 1) enabling new memory architectures, functions, interfaces, and better integration of memory and the rest of the system, 2) designing a memory system that intelligently employs emerging non-volatile memory (NVM) technologies and coordinates memory and storage management, 3) reducing memory interference and providing predictable performance to applications sharing the memory system. If time permits, we will also touch upon our ongoing related work in combating scaling challenges of NAND flash memory.

An accompanying paper, slightly outdated (circa 2015), can be found here:

[http://people.inf.ethz.ch/omutlu/pub/memory-systems-research\\_superfri14.pdf](http://people.inf.ethz.ch/omutlu/pub/memory-systems-research_superfri14.pdf)

### **Bio:**

Onur Mutlu is a Professor of Computer Science at ETH Zurich. He is also a faculty member at Carnegie Mellon University, where he previously held the William D. and Nancy W. Strecker Early Career Professorship. His current broader research interests are in computer architecture, systems, and bioinformatics. He is especially interested in interactions across domains and between applications, system software, compilers, and microarchitecture, with a major current focus on memory and storage systems. He obtained his PhD and MS in ECE from the University of Texas at Austin and BS degrees in Computer Engineering and Psychology from the University of Michigan, Ann Arbor. His industrial experience spans starting the Computer Architecture Group at Microsoft Research (2006-2009), and various product and research positions at Intel Corporation, Advanced Micro Devices, and VMware. He received the inaugural IEEE Computer Society Young Computer Architect Award, the inaugural Intel Early Career Faculty Award, faculty partnership awards from various companies, and a healthy number of best paper or "Top Pick" paper recognitions at various computer systems and architecture venues. His computer architecture course lectures and materials are freely available on YouTube, and his research group makes software artifacts freely available online. For more information, please see his webpage at <http://people.inf.ethz.ch/omutlu/>