

HICOMB 2018 Keynote Speaker

Accelerating Genome Analysis: A Primer on an Ongoing Journey

Onur Mutlu, ETH Zurich

Abstract: Genome analysis is the foundation of many scientific and medical discoveries as well as a key pillar of personalized medicine. Any analysis of a genome fundamentally starts with the reconstruction of the genome from its sequenced fragments. This process is called read mapping. One key goal of read mapping is to find the variations that are present between the sequenced genome and reference genome(s) and to tolerate the errors introduced by the genome sequencing process. Read mapping is currently a major bottleneck in the entire genome analysis pipeline because state-of-the-art genome sequencing technologies are able to sequence a genome much faster than the computational techniques that are employed to reconstruct the genome. New sequencing technologies, like nanopore sequencing, greatly exacerbate this problem while at the same time making genome sequencing much less costly.

This talk describes our ongoing journey in greatly improving the performance of genome read mapping. We first provide a brief background on read mappers that can comprehensively find variations and tolerate sequencing errors. Then, we describe both algorithmic and hardware-based acceleration approaches. Algorithmic approaches exploit the structure of the genome as well as the structure of the underlying hardware. Hardware-based acceleration approaches exploit specialized microarchitectures or new execution paradigms, like processing in memory. We show that significant improvements are possible with both algorithmic and hardware-based approaches and their combination. We conclude with a foreshadowing of future challenges brought about by very low cost yet highly error prone new sequencing technologies.

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. A variety of techniques he, along with his group and collaborators, has invented over the years have influenced industry and have been employed in commercial microprocessors and memory/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, VMware, and Google. He received the inaugural IEEE Computer Society Young Computer Architect Award, the inaugural Intel Early Career Faculty Award, faculty partnership awards from various companies, a healthy number of best paper or "Top Pick" paper recognitions at various computer systems and architecture venues, and the ACM Fellow recognition "for contributions to computer architecture research, especially in memory systems." 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/>.