



北京大学高能效计算与应用中心学术报告

Invited Talk, Center for Energy-Efficient Computing and Applications

BIG DATA AND DARK SILICON: TAMING TWO IT INFLECTION POINTS ON A COLLISION COURSE

Prof. Babak Falsafi

École Polytechnique Fédérale de Lausanne

2014年8月6日 星期三 3:00pm

理科五号楼401会议室



ABSTRACT: Information technology is now an indispensable pillar of a modern-day society, thanks to the proliferation of digital platforms in the past several decades. We are now witnessing two inflection points, however, that are about to change IT as we know it. First, we are entering the Big Data era where demand on robust and economical data processing, communication and storage is growing faster than technology can sustain. Second, while forecasts indicate that chip density scaling will continue for another decade, the diminishing returns in supply voltage scaling and the impending "energy wall", is leading server designers towards energy-centric solutions and eventually Dark Silicon. In this talk, I will motivate these two IT trends and present promising research avenues for server chip design including specialized scale-out processors, on-chip networks and die-stacked DRAM cache hierarchies.

BIOGRAPHY: Babak is Professor in the School of Computer and Communication Sciences and the founding director of the EcoCloud research center pioneering future energy-efficient and environmentally-friendly cloud technologies at EPFL. He has made numerous contributions to computer system design and evaluation including a scalable multiprocessor architecture which was prototyped by Sun Microsystems (now Oracle), snoop filters and temporal stream prefetchers that are incorporated into IBM BlueGene/P and BlueGene/Q, and computer system simulation sampling methodologies that have been in use by AMD and HP for research and product development. His most notable contribution has been to be first to show that contrary to conventional wisdom, multiprocessor memory programming models -- known as memory consistency models -- prevalent in all modern systems are neither necessary nor sufficient to achieve high performance. He is a recipient of an NSF CAREER award, IBM Faculty Partnership Awards, and an Alfred P. Sloan Research Fellowship. He is a fellow of IEEE.