Title: Statistics in Workload Characterization and Benchmarking

Abstract:
Computer architects and system designers have made tremendous advances in the performance of computer systems over the past several decades. Unfortunately, the performance of a computer system is impacted by many different components in extremely complex and nonlinear ways. For example, it is well understood that simply increasing the clock rate will not necessarily produce a proportionate increase in the overall performance. These complex interactions introduce uncertainty into the measurements of a system's performance, which makes it difficult to determine the impact any changes made to the system actually have on the overall performance. This measurement noise also makes it difficult to compare the performance of different systems. This tutorial provides a gentle introduction to some of the key statistical tools and techniques needed to interpret noisy performance measurements and to sort through large collections of simulation results. It also presents techniques that can be used to appropriately design experiments to obtain the maximum amount of information for a given level of experimental effort. . Finally, it discusses statistical data analysis techniques for constructing representative workloads.

Bio:
David J. Lilja received a Ph.D. and an M.S., both in Electrical Engineering, from the University of Illinois at Urbana-Champaign, and a B.S. in Computer Engineering from Iowa State University in Ames. He is currently the Louis John Schnell Professor of Electrical and Computer Engineering at the University of Minnesota in Minneapolis, where he also serves as the ECE department head, as a member of the graduate faculties in Computer Science and Scientific Computation, and as a Fellow of the Minnesota Supercomputer Institute. Previously, he worked as a research assistant at the Center for Supercomputing Research and Development at the University of Illinois, and as a development engineer at Tandem Computers Incorporated in Cupertino, California. He has chaired and served on the program committees of numerous conferences, and was a distinguished visitor of the IEEE Computer Society. He received a Fulbright Senior Scholar Award to visit the University of Western Australia in 2001, and was awarded a McKnight Land-Grant Professorship by the Board of Regents of the University of Minnesota. His main research interests include computer architecture, parallel processing, computer systems performance analysis, and high-performance storage systems. He has a special interest in the interaction of software and compilers with computer architecture, and the interaction of computer architecture and circuits. He is a Fellow of both the Institute of Electrical and Electronics Engineers (IEEE) and the American Association for the Advancement of Science (AAAS), is a Member of the ACM, and is a registered Professional Engineer.

Lieven Eeckhout is an Associate Professor in the Department of Electronics and Information Systems at Ghent University, Belgium. He received his PhD degree in computer science and engineering from Ghent University in Dec 2002. His main research interests include computer architecture, virtual machines, performance modeling and analysis, simulation methodology, and workload characterization. He has published papers in top conferences such as ISCA, ASPLOS, HPCA, OOPSLA, PACT, CGO, DAC and DATE; he has served on multiple program committees including ISCA, PLDI, HPCA and IEEE Micro Top Picks; and he was the program chair for ISPASS 2009 and is the general chair for ISPASS 2010. His work on hardware performance counter architectures was selected by IEEE Micro Top Picks from 2006 Computer Architecture Conferences as one of the "most significant research publications in computer architecture based on novelty and industry relevance". He graduated 5 PhD students, and currently supervises two postdoctoral researchers and 8 PhD students. He is a member of the IEEE and ACM.