

Robust Resource Management for Parallel Computing Systems

Prof. H. J. Siegel
Department of Electrical and Computer Engineering
and Department of Computer Science
Colorado State University
Fort Collins, Colorado, USA
E-mail: HJ@ColoState.edu
www.engr.colostate.edu/~hj

Abstract:

What does it mean for a computer system to be “robust”? How can robustness be described? How does one determine if a claim of robustness is true? How can one decide which of two systems is more robust? Parallel and distributed computing systems are often heterogeneous mixtures of machines, used to execute collections of tasks with diverse computational requirements. A critical research problem is how to allocate heterogeneous resources to tasks to optimize some performance objective. However, systems frequently have degraded performance due to uncertainties, such as inaccurate estimates of actual workload parameters. It is important for system performance to be robust against uncertainty. To accomplish this, we present a stochastic model for deriving the robustness of a resource allocation. This model assumes that stochastic (experiential) information is available for a parameter whose actual values are uncertain. The robustness of a resource allocation is quantified as the probability that a user-specified level of system performance can be met. We show how to use this stochastic model to evaluate the robustness of resource assignments and to design resource management heuristics that produce robust allocations. The stochastic robustness analysis approach can be applied to a variety of computing and communication system environments, including parallel, distributed, cluster, grid, Internet, cloud, embedded, multicore, content distribution networks, wireless networks, and sensor networks. Furthermore, the robustness model is generally applicable to design problems throughout various scientific and engineering fields.

Speaker:

H. J. Siegel is the George T. Abell Endowed Chair Distinguished Professor of Electrical and Computer Engineering at Colorado State University (CSU), where he is also a Professor of Computer Science. He is Director of the CSU Information Science and Technology Center (ISTeC), a university-wide organization for enhancing CSU's activities pertaining to the design and innovative application of computer, communication, and information systems. Before joining CSU, he was a Professor at Purdue University from 1976 to 2001. He received two B.S. degrees from the Massachusetts Institute of Technology (MIT), and the M.A., M.S.E., and Ph.D. degrees from Princeton University. He is a Fellow of the IEEE and a Fellow of the ACM. Prof. Siegel has co-authored over 380 published technical papers in the areas of parallel and distributed computing. He was a Coeditor-in-Chief of the *Journal of Parallel and Distributed Computing*, and was on the Editorial Boards of the *IEEE Transactions on Parallel and Distributed Systems* and the *IEEE Transactions on Computers*.