Date/Time: Friday, May 31, 2013
Location: Donald Bren Hall, Room 4011

Committee Members:
Professors Alexander V. Veidenbaum (Chair),
Professor Alexandru Nicolau,
Professor Nikil Dutt

Abstract
Maintaining and improving program performance in the multi-core era requires a large engineering effort, e.g., a large number of time consuming trials & tests. It involves finding, as efficiently as possible, a combination of attributes that characterize (i.e., expose similarity in) algorithmic optimizations, compiler optimizations, execution environment settings and hardware configurations to reach certain performance goals. This problem is currently unsolved in its entirety because the number of attributes involved is very large. Thus programs are optimized based on a very limited number of such attributes at a time.

This thesis investigates the construction of empirical performance models that provide program performance prediction across system configurations, where the term system includes the development environment, e.g., compilers, libraries and their settings, and the execution environment, e.g., operating system, run-time environment and their settings. Predictions have to be sufficiently accurate to reduce the engineering effort by replacing trials & tests with prediction.

Two new characterization approaches are proposed, that exploit similarity between programs and systems to predict performance of programs on unknown systems. The results demonstrate that proposed characterization approaches and the corresponding performance models are sufficiently accurate and equally applicable to specific performance evaluation studies, including characterization, comparison and tuning of hardware configurations, compilers, run-time environments or any combination thereof.