

## Topic 2

# Performance Evaluation and Prediction

Jeff Hollingsworth, Allen D. Malony, Jesús Labarta, and Thomas Fahringer  
Topic Chairs

Performance is the reason for parallel computing. Despite years of work, many applications achieve only a few percentage of theoretical peak. Performance measurement and analysis tools exist to identify the problems with current programs and systems. Performance Prediction is intended to identify issues in new code or systems before they are fully available. These two topics are closely related since most prediction requires data to be gathered from measured runs of program (to identify application signatures or to understand the performance characteristics of current machines).

Several issues continue to confront tool builders. First, how can tools be scaled up to run on very large systems. There are a growing number of parallel and distributed systems with over 1,000 processors and machines with 10,000 to 100,000 processors are being considered. Many tools were designed to work with a few dozen processors. Second, memory system bottlenecks are often the largest source of the gap between theoretical peak and achieved performance. How can tools help programmers to understand where this loss of performance is coming from? Third, how can we achieve high quality predictions of the performance of large-scale applications. Many existing simulation tools are several orders of magnitude slower than running the application. How can we improve the running time of these tools without sacrificing too much accuracy?

Out of 16 submitted papers to this topic, 6 have been accepted as regular papers (37%), and two as short papers. Thanks to all of the contributing authors as well as all reviewers for their work. A special thanks to the program committee Thomas Fahringer, Jesus Labarta, and Allen Malony.