Topic 2
Performance Prediction and Evaluation

Allen D. Malony, Thomas Fahringer, Allan Snavely, and Luís Silva

Topic Chairs

Performance is the reason for parallel computing. Achieving high performance on parallel computer systems is the product of an intimate combination of hardware architecture (processor, memory, interconnection network), system software, runtime environment, algorithms, and application design. Performance evaluation is the science of understanding these factors that contribute to the overall expression of parallel performance on real machines and on systems yet to be realized. Benchmarking and performance characterization methodologies and tools provide an empirical foundation for performance evaluation. Performance prediction techniques provide a means to model performance behaviors and properties as system, algorithm, and software features change, particularly in the context of large-scale parallelism. These two areas are closely related since most prediction requires data to be gathered from measured runs of a program, to identify application signatures or to understand the performance characteristics of current machines.

A total of twenty-nine papers were submitted to the performance prediction and evaluation topic area. The submissions covered a broad range of prediction and evaluation topics, and reflect a high level of current interest in the parallel computing community. The eleven papers accepted (38%) represent state-of-the-art results from leading parallel performance researchers in the field today. The papers cover four general themes in performance prediction and evaluation.

The first theme considers methods to explore performance properties from different evaluation contexts: data access, processor, and interconnect. The understanding gained from looking at these different performance contexts is valuable to forming a more complete performance assessment. The second theme concerns advances in measurement infrastructure for performance analysis at the application level. In particular, the three tools reported illustrate techniques for instrumenting events closely tied to parallel program operation and for capturing performance data needed to correctly interpret performance behavior. Techniques for performance prediction for large-scale parallel systems is the third theme in the topic. The contributions here on performance extrapolation from traces, performance modeling and sensitivity analysis, and performance prediction using machine learning, are especially strong and are important contributions to the field. Lastly, we consider the connection of performance evaluation in tools for performance tuning in the fourth theme. Graphical user interface support for integrated performance environments and automatic tuning for parallel program archetypes are described.

J.C. Cunha and P.D. Medeiros (Eds.): Euro-Par 2005, LNCS 3648, p. 93, 2005.