Introduction

Shirley Moore, Derrick Kondo, Brian Wylie, and Giuliano Casale

Topic chairs

Research on performance evaluation over the past several years has resulted in a range of techniques and tools for modeling, analyzing, and optimizing performance of applications on parallel and distributed computing systems. With the emergence of extreme-scale computing architectures, the need for tools and methodologies to predict and improve application performance and to adapt to evolving architectures will become even greater, due to increased complexity and heterogeneity of the systems. Furthermore, the scope of the term performance has expanded to include reliability, energy efficiency, scalability, and system-level context. This years conference topic Performance Prediction and Evaluation aims to bring together researchers involved with the various aspects of this broader scope of application and system performance modeling and evaluation on large-scale parallel and distributed systems.

The six papers accepted for this topic reflect a growing interest in end-to-end performance issues related to network performance, multicore architectures, resource allocation, scheduling, and energy usage. The first three papers focus on different aspects of communication modeling and performance. The fourth and fifth papers address cache partitioning and job scheduling, respectively. The sixth paper focuses on an application-level methodology for minimizing system energy consumption.

- The paper A contention-aware performance model for HPC-based networks: A case study of the Infiniband network presents a methodology for dynamically predicting communication times in congested networks and applies the methodology to an Infiniband network. The paper Using the last-mile model as a distributed scheme for available bandwidth prediction proposes decentralized heuristics for estimating the available bandwidth between nodes in a large-scale distributed system. The heuristics are based on the last-mile model, which characterizes each node by its incoming and outgoing capacity and uses this last-mile (end-host) bandwidth to predict overall performance of the end-to-end paths. The last paper in the network performance area, Self-stabilization versus robust self-stabilization for clustering in ad-hoc network, is an experimental comparison of the performance of four clustering protocols for maintaining a scalable hierarchical network routing scheme in the presence of topological changes due to failures and node motion in a mobile ad-hoc network. Two of the protocols are self-stabilizing, meaning that they converge in finite time to a state that provided optimum service, and two are robust self-stabilizing, meaning that they not only converge to

E. Jeannot, R. Namyst, and J. Roman (Eds.): Euro-Par 2011, LNCS 6852, Part I, pp. 77–78, 2011.

optimum service but also maintain minimal useful service during the stabilization period.
- The paper Multilayer cache partitioning for multiprogram workloads presents a coordinated cache partitioning scheme for multiprogram workloads on multicore systems that considers multiple levels of the cache hierarchy simultaneously. The scheme attempts to satisfy specified quality of service (QoS) values for all applications by partitioning the shared cache hierarchy across them and then distributes the remaining excess cache capacity (if any) across applications such that a global performance metric is maximized. In Backfilling with guarantees granted upon job submission, the authors present two scheduling algorithms based on conservative backfilling by adding prioritized compression and delayed prioritized compression. They use traces of actual workload data from the Parallel Workloads Archive to show that their algorithms generally perform better than normal conservative backfilling.
- The final paper addresses the growing challenge of reducing the energy consumption of high performance computing systems. Entitled Reducing energy usage with memory and computation-aware dynamic frequency scaling, it introduces a methodology that chooses fine-grained dynamic voltage frequency scaling (DVFS) settings, with potentially different setting for different parts of applications, with the goal of minimizing system-wide energy usage. The authors provide a set of automated tools that capture relevant application characteristics at the loop level, match these features using a database of benchmark results to determine the DVFS strategy, and insert the DVFS commands into the application using binary instrumentation.

We would like to take this opportunity to thank all authors who submitted a paper to this topic area, thank the reviewers for their careful evaluations, and finally thank the Euro-Par Organizing Committee for their outstanding management of this years conference.