Topic 2: Performance Prediction and Evaluation

Allen D. Malony, Helen Karatza, William Knottenbelt, and Sally McKee

Topic Committee

In recent years, a range of novel methodologies and tools have been developed for the purpose of evaluation, design, and model reduction of existing and emerging parallel and distributed systems. At the same time, the coverage of the term 'performance' has constantly broadened to include reliability, robustness, energy consumption, and scalability in addition to classical performance-oriented evaluations of system functionalities. Indeed, the increasing diversification of parallel systems, from cloud computing to exascale, being fueld by technological advances, is placing greater emphasis on the methods and tools to address more comprehensive concerns. The aim of the Performance Prediction and Evaluation topic is to bring together system designers and researchers involved with the qualitative and quantitative evaluation and modeling of large-scale parallel and distributed applications and systems to focus on current critical areas of performance prediction and evaluation theory and practice.

The five papers selected for the topic area reflect the broadening perspective of parallel performance involving alternative evaluation techniques (measurement, simulation, analytical modeling), different systems components (file systems, GPUs, I/O, multicore processors), and multiple targeted metrics (execution time, energy, network latency). The two papers based on modeling methodologies looked at two different systems aspects:

The paper "Energy Consumption Modeling for Hybrid Computing" by Marowka presented analytical models based on an energy consumption metric and used the model to analyze different design options for hybrid CPU-GPU chips. They studied the joint effect of performance and energy consumption to understand their relationship, particularly with respect to greater parallelism. The paper "HPC File Systems in Wide Area Networks: Understanding the Performance of Lustre over WAN" by Aguilera et al., also looked at performance interactions, but with respect to network system design. They evaluated the performance of the Lustre file system and its networking layer, with the goal of understanding the impact that the network latency has on Lustre's performance and deriving useful "rules of thumb" to help predict Lustre performance variation.

The next two papers share a common thread of how to capture a more comprehensive evaluation of performance for predictive purposes when the environment itself is complex and difficult to study directly. While sharing this theme, the two papers look at two distinct problems areas and take two different approaches:

"Understanding I/O Performance using I/O Skeletal Applications" by Logan et al. attempts to get a handle on the causes of I/O bottlenecks in HPC applications for purposes of guiding scaling optimization. By combining an approach to generate I/O benchmark codes from a high-level description with low-level

C. Kaklamanis et al. (Eds.): Euro-Par 2012, LNCS 7484, pp. 52–53, 2012.

performance characterization of I/O components, a more complete and representative picture of application I/O behavior is obtained. The methodology enables more meaningful I/O performance testing, improved prediction of I/O performance, and more flexible evaluation of new systems and I/O methods. The paper "ASK: Adaptive Sampling Kit for Performance Characterization," by Castro et al., addresses the complexity of a large optimization design space for performance tuning. Their approach measures and understands performance tradeoffs by applying multiple adaptive sampling methods and strategies with the goal of considerably reducing the cost of performance exploration. The outcome are precise models of performance, created with a small number of measures, that can be used for prediction of performance for specific features, such as memory stride accesses.

The last paper of the topic area "CRAWP: A Workload Partition Method for the Efficient Parallel Simulation of Manycores" by Jiao et al. points out the interesting aspect that the behavior of manycore execution affects the performance of parallel discrete event simuation (PDES) system used to study it. Thus, by altering how workload is partitioned in the PDES simulator, it is possible to achieve improved speed and accuracy. The authors propose an adaptive workload partition method – Core/Router-Adaptive Workload Partition (CRAWP) – that make the simulation of on-chip-network independent of the cores. Significant improvements are achieved for the simulation of the SPLASH2 benchmark applications.