

Topic 2

Performance Prediction and Evaluation

Wolfgang Nagel, Bruno Gaujal, Tugrul Dayar, and Nihal Pekergin

Topic Chairs

Parallel algorithms used to be evaluated using some version of the PRAM model where actual execution platforms are abstracted as ideal parallel machines. On the other hand the performance of hardware is often given in terms of individual pick performances which can be useless for actual applications. The real challenge for performance predictions and evaluations of parallel systems is to combine the top and low layers points of view, where congestions, control mechanisms, failures or even breakdowns do alter the behavior of a large distributed platform running a parallel application. In many different ways, this challenge is addressed in all the papers presented in the track, using either new modelling techniques, or sophisticated measures, or experimental approaches.

This track is made of eight papers.

In “Decision Trees and MPI Collective Algorithm Selection Problem”, the authors study the applicability of decision trees to the MPI collective algorithm selection problem.

In “Profiling of Task-based Applications on Shared Memory Machines: Scalability and Bottlenecks”, new measures are introduced to measure the scalability of task pools algorithms using a non-intrusive profiling method.

In “Search Strategies for Automatic Performance Analysis Tools”, a set of agents distributed on the parallel machine are used to/ analyse the performance of parallel programs.

In “Experiences Understanding Performance in a Commercial Scale-Out Environment”, the new performance challenges faced in clusters of loosely connected machines are studied using a systematic experimental approach.

In “TAUoverSupermon: Low-Overhead Online Parallel Performance Monitoring”, two systems - Tuning and Analysis Utility (TAU) and Supermon are combined to perform online monitoring of large distributed applications with a very low overhead.

In “Practical Differential Profiling”, the authors present eGprof, a tool that facilitate performance profile comparisons.

In “Detecting Application Load Imbalance on High End Massively Parallel Systems”, a new set of indirect metrics are used to identify and measure application load imbalance. These metrics have been added to the Cray performance measurement and analysis infrastructure.

Finally, in “A First Step Towards Automatically Building Network Representations”, the authors present new algorithms to design tools which automatically build a topological network model through some measurements and show that they accurately predict the running time of simple applications over actual platforms.