

Topic 02

Performance Evaluation and Prediction

Allen D. Malony, Graham D. Riley, Bernd Mohr,
Mark Bull, and Tomàs Margalef

Topic Chairpersons

The performance of parallel and distributed systems and applications — its evaluation, analysis, and optimization — is at once a fundamental topic for research investigation and a technological problem that requires innovations in tools and techniques to keep pace with system and application evolution. This dual view of performance “science” and performance “technology” jointly spans broad fields of performance modelling, evaluation, instrumentation, measurement, analysis, monitoring, optimization, and prediction.

Fourteen research papers were submitted to this workshop of which eight were selected for inclusion in the conference. The topics covered by the accepted papers have quite a broad range, which reflects the applicability of performance-related topics at all levels of computing systems from architecture and networking through to high level application development in the many programming paradigms which are in current use. A continuing challenge to the area is to seek abstractions and generic techniques which apply across these boundaries. Two themes particularly dominate the accepted papers and these form the basis of the two sessions of the Topic: first, the prediction of performance and performance modelling, and second, techniques to analyse actual performance and practical reports of analyses in specific cases.

Performance Prediction and Modelling

In the area of modelling, *Ponomarev et al.* use a theoretical approach to demonstrate how network queueing behaviour may be improved by using polling rather than more conventional interrupt driven servicing mechanisms. In the search for generic performance prediction methods, one approach is to seek methods which support particular programming paradigms, and this approach is seen in the paper, *Performance Prediction of Oblivious BSP Programs* by *González et al.* An alternative approach is to focus on classes of algorithms which occur in real applications. This approach is taken in the papers by *Gautama et al.* and *Moreno et al.*, each targeting a different algorithm type. In *Performance Prediction of Data-Dependent Task Parallel Programs*, recent new analysis techniques are used to overcome challenges to prediction for general task parallel programs with dependent behavior. In *The Tuning Problem on Pipelines*, focus is given on *pipeline algorithms* and it is shown that good prediction accuracy can be achieved with careful consideration of all algorithm features.

Performance Analysis and Support Tools

Performance tools for effective empirical performance analysis must address both technological and methodological concerns. The papers in the second session consider three different but related aspects of the problem. First, requirements for effective implementation and use of performance measurement tools is explored in *The Hardware Performance Monitor Toolkit* by Luiz A. DeRose. As the title implies, this paper focuses on the use of hardware monitoring facilities and what can be learned from hardware performance data. A second aspect is that of conducting and reporting the results of performance analysis on real systems. The papers by Baker *et al.* on *VIA Communication Performance on a Gigabit Ethernet Cluster* and Strey and Bange on *Performance Analysis of Intel's MMX and SSE: A Case Study* are two examples of quality performance analysis in practice. Third, questions of performance analysis methodology are addressed in the paper, *Group Based Performance Analysis for Multithreaded SMP Cluster Applications* by Brunst *et al.* This paper proposes innovative hierarchical viewing methods to support the exploration of larger-scale performance data.