Topic 10
Parallel Numerical Algorithms

Ian Duff, Michel Daydé, Matthias Bollhoefer, and Anne Trefethen

Topic Chairs

Efficient and robust parallel and distributed algorithms with portable and easy-to-use implementations for the solution of fundamental problems in numerical mathematics are essential components of most parallel software systems for scientific and engineering applications.

This topic provides a forum for the presentation and discussion of new developments in the area of parallel and distributed numerical methods. All aspects of the design and implementation of parallel algorithms will be addressed, ranging from discussion of the ideas on which they are based to analyses of their complexities and performance on current parallel and distributed architectures (including clusters and grids). Software design and prototyping in scientific computing or simulation of software environments are also covered.

Methods for the solution of large linear systems are of particular interest because of their widespread occurrence in many fields, particularly in the numerical solution of partial differential equations. However, contributions dealing with new and improved parallel and distributed algorithms for the solution of other problems in numerical linear algebra, linear and non-linear programming, numerical quadrature, differential equations, fast transforms and non-linear systems were also welcome.

Thirteen papers papers were submitted to the track and five were accepted. Three papers consider the parallel implementation of application-related software:

- In "An Efficient Parallel Particle Tracker for Advection-Diffusion Simulations in Heterogeneous Porous Media", the authors consider the contamination of groundwater by migration of pollutants.
- In "A Fully Scalable Parallel Algorithm for Solving Elliptic Partial Differential Equations", the authors compare the probabilistic domain decomposition (DD) method with deterministic methods for solving linear elliptic boundary-value problems.
- In "Locality Optimized Shared-Memory Implementations of Iterated Runge-Kutta Methods", the authors consider the sequential and parallel implementation of Iterated Runge-Kutta methods focusing on the optimization of the locality behaviour. They introduce different implementation variants for sequential and shared-memory computer systems and analyse their performance.

The two remaining papers focus on implementation issues (task scheduling, memory hierarchy, multi-core processor architectures,...) for high performance linear algebra computational kernels:

– In “Toward Scalable Matrix Multiply on Multithreaded Architectures”, the authors analyse the impact of SMP and future multicore architectures on efficient software design.
– In “Task Scheduling for Parallel Multifrontal Methods”, the authors present a new scheduling algorithm for task graphs arising from parallel multifrontal methods for sparse linear systems. duling algorithm for task graphs arising from parallel multifrontal methods for sparse linear systems.