Workshop on High-Performance Computing
ETH Zürich, September 2006

Talks of Monday 4

Where is HPC headed?
A. Gunzinger, ETH and SCS Supercomputing Systems, AG

Most of todays supercomputers are based on standard PCs and networks. What are the limitations of these systems? Next generation game consoles (XBox, PlayStation 3, etc.) contain increasingly powerful processors. What role will these processors play in the HPC market? Modern FPGAs can be used to build extremely powerful parallel processors. Will the supercomputers of the future be based on FPGAs?

Efficient Fully-coupled Newton-Krylov Solution of Transport/Reaction Systems
John N. Shadid, Computational Sciences R&D Group, Sandia National Laboratories, Albuquerque, NM

A current challenge before the computational science and numerical mathematics community is the efficient computational solution of multiphysics systems. These systems are strongly coupled, highly nonlinear and characterized by multiple physical phenomena that span a very large range of length and time scales. The myriad of complex, interacting physical mechanisms can balance to produce steady-state behavior, nearly balance to evolve a solution on a dynamical time scale that is long relative to the component time-scales, or they can be systems that are dominated by one (or a few) process(es). These characteristics make the computational solution of these systems, over relevant dynamical time scales of interest (or to steady-state solutions), extremely challenging.

In this presentation I will overview a number of the important solution methods that our research group has applied in the large-scale parallel simulation of such systems with a focus on transport/reaction applications. The solution methods that we employ include, fully-implicit time integration, direct-to-steady-state solution methods, continuation, bifurcation, and optimization techniques. The resulting large sparse linear systems that are generated by these methods are solved by the application of parallel preconditioned Krylov methods. The preconditioners include additive Schwarz domain decomposition (DD) and multi-level preconditioners. The multi-level preconditioners are based on geometric and algebraic methods as well as approximate block factorization techniques.

To demonstrate the capability of these methods I will present simulation results for representative low heat release and high heat release transport / reaction simulations. In this context I will discuss robustness, efficiency, and the parallel and algorithmic scaling of solution methods.

Parallel Computing Problems in Cosmological N-body Simulations
Joachim Stadel

Parallel computers have been very effective tools in the study of cosmological structure formation and our understanding of the dynamical evolution of the Universe. The simulations use particles which represent packets of the phase fluid which for dark matter behaves as an incompressible fluid in the 6-dimensional phase space. Accuracy of such simulations is primarily dictated by the number of particles used since the methods are based on Monte-Carlo integration. As more particles are used the densities that can be resolved in the simulations increases. This greatly increases the difficulties associated with load balancing on a parallel computer and current simulations set an ever increasing challenge to scaling on parallel architectures. I will discuss these problems and some of the strategies used to alleviate them. New machines with up to 1 million processors seem to require a new approach and I will discuss ways in which one can achieve efficiency at even this scale of parallelism.

Precision Cosmology and High Performance Computing
S. Habib, Los Alamos National Laboratory

A revolutionary change in cosmology has been underway for the last fifteen years – the transition to precision cosmology. It is now common to discuss cosmological measurements at the percent level of accuracy or better – in certain cases, achievable target accuracies as low as 0.1%! A powerful and diverse suite of cosmological observations has led to this remarkable transformation. Taken together, results from cosmic microwave background observations, studies of the large-scale distribution of matter in the Universe, and of the redshift distribution of supernovae, yield an impressively consistent picture of the fundamental make-up of the Universe, the so-called Concordance Model which requires the existence of a mysterious Dark Energy. The establishment of the concordance model has in turn ushered in the era of precision cosmology. Forecasts for determination of cosmological parameters from next-generation observations require calibration against large-scale cosmological simulations accurate to the level of one percent or better. The petabyte scale of simulation and observation datasets raises yet another set of challenges including unprecedented levels of automation. In this talk I will discuss where we stand today with regard to these challenges for next-generation computational cosmology. As a primary application, I will focus on how progress in high performance computing will impact dark energy studies, the primary target of future cosmological surveys.

High-performance Computing of Bone Structure and its Strength
Harry van Lenthe (ETH Zurich)

On a Finite Element based Grid-enabled HPC Toolchain for Patient Specific Surgery Simulations
Dr. Jens Georg Schmidt, C&C Research Laboratories, NEC Europe Ltd.

We are dealing with the simulation of maxillo-facial surgeries, especially with distraction osteogenesis. During this treatment the surgeon cuts free the patient's upper jaw (maxilla), which is subsequently relocated into a new position in the course of several weeks, using a distraction device.

The presented simulation toolchain is set up to predict the displacements of the facial tissues during and after the pulling process and is based on individual CT images of the patient's head before treatment. Its purpose is to support the surgeon in optimizing the treatment plan and avoiding additional post operative plastic surgeries.

The input data for the simulation toolchain is generated by adding the suggested cuts, the geometry of the distraction device and the suggested forces to the CT data of the patient's head. From these data we generate a Finite Element mesh of the head and perform a Finite Element analysis of the distraction process. In order to achieve sufficient accuracy we have to resolve most of the geometrical features of the human head, which leads to a large number of unknowns, typically several millions. In addition to that the computational costs are significantly increased by the size of the displacements and the visco-elastic behaviour of the materials, which can only be properly approximated by non-linear modelling. The image processing, the meshing and the analysis of the model are run via a grid service on an HPC resource.

Focusing on the efficiency of the simulation, the linear solvers used to solve the arising systems of equations play the most crucial role. In our case standard iterative solvers like Krylov methods or ILU methods fail, as we will show in our presentation. Therefore we will focus on the use of Multigrid solvers.

But the complex geometry of the human head in combination with large jumps of the material parameters – Young's modulus jumps about 5 orders of magnitude between bone and soft tissues – prevents standard multigrid approaches from converging at a sufficient rate. In theory convergence can be dramatically improved by computing the "near null space" of the systems, consisting of quasi-rigid body modes, and treating it separately. We will show the limitations of this approach for highly complex geometries, like the human head.

In our presentation we will demonstrate the performance of the only two solvers we have found to work on our problems so far, which are BoomerAMG from LLNL's hypre package and ML from Sandia's Trilinos package. We will show the results of our intensive parameter studies and discuss the extensibility of our perfomance results for elasto-mechanical Finite Element simulations in general.

Cray Update

Roberto Ansaloni, Cray Italy

Unfolding Blue Gene

Michael Hennecke, HPC Technical Architect, IBM Germany

The presentation describes the system architecture of the IBM System Blue Gene Solution, including the microarchitecture of the Blue Gene compute chip, how systems are build from these compute chips, and the role of the different Blue Gene networks. The talk will also outline areas where the application programmer can perform Blue Gene specific optimization, covering both single-node performance and considerations for MPI programming on Blue Gene. Finally, some current customer usage areas are presented.

Software Strategies for Flexible High-performance Libraries
Roscoe A. Bartlett, Sandia National Laboratories

The Trilinos Project is an effort to facilitate the design, development, integration and ongoing support of mathematical software libraries within an object-oriented framework for the solution of large-scale, complex multi-physics engineering and scientific problems. Trilinos addresses two fundamental issues of developing software for these problems: (i) Providing a streamlined process and set of tools for development of new algorithmic implementations and (ii) promoting interoperability of independently developed software.

Trilinos uses a two-level software structure designed around collections of packages. A Trilinos package is an integral unit usually developed by a small team of experts in a particular algorithms area such as algebraic preconditioners, nonlinear solvers, etc. Packages exist underneath the Trilinos top level, which provides a common look-and-feel, including configuration, documentation, licensing, and bug-tracking.

Here we present the aspects of Trilinos that relate to performance. While there are several different aspects of performance addressed in Trilinos, this presentation will focus on the interaction between performance and the design of general and flexible object-oriented interfaces. A set of interfaces that are not carefully designed will result in undue constraints on possible implementations and result in sub par performance. Specifically, we will focus on the design of the Thyra operator/vector interfaces and on the design of the multi-vector interface.

Advanced Expression Template Concepts
Christoph Pflaum, Universitat Erlangen

One important difficulty concerning the implementation of numerical software for finite elements is the clear separation of the code in different modules which interact by clearly defined interfaces. In order to solve this problem, it is necessary to describe these interfaces by a mathematical language. To this end, it is very helpful to apply operator overloading in C++ and expression templates.

The first expression template implementations were based on non-trivial template constructions such that often programmers did not use expression templates for the implementation of mathematical software. Additionally, there exist applications, where the original expression template implementation does not lead to optimal performance. To avoid these problems we present expression template implementations which are more simple to implement and lead to an optimal performance.

Applications to different finite elements codes will be presented.

Generalizing Smoothed Aggregation Algebraic Multigrid
R. S. Tuminaro, Sandia National Laboratories

Smoothed aggregation is a well-known algebraic multigrid method that is based on minimizing the energy of grid transfer basis functions. It has been used on a wide variety of problems and has had particular success in elasticity.

An overview of the parallel smoothed aggregation multigrid method is first given. Then, several generalizations of the smoothed aggregation idea are discussed. These generalizations are intended to improve method robustness for anisotropic problems and highly variable coefficient problems as well as extend the applicability of smoothed aggregation to other problem domains such as the eddy current equations and the Navier-Stokes equations.

The central theme of this talk will be generalizations of the prolongator smoothing step which maintain the exact interpolation of null space components. Two particular ideas will be emphasized. The first concerns adapting the prolongator smoother to irregular aggregates. This involves shifting support between basis functions in a way that maintains the low energy nature of the basis functions while improving the sparsity of the resulting multigrid operators. The second idea considers generalizations of the notion of "energy" and the use of local damping parameters. This allows smoothed aggregation to be used on nonsymmetric systems and often improves the smoothing of individual basis functions (leading to faster rates of convergence).

We present numerical experiments that compare the new methods to traditional smoothed aggregation on a variety of problems.

An Algebraic Multilevel Approach for Highly Indefinite Systems of Equations
Matthias Bollhfer, TU Braunschweig

We will discuss algebraic multigrid techniques that address solving large sparse symmetric indefinite systems. In particular we consider the situation when the system is highly indefinite. Such cases arise e.g. from the discretization of the Helmholtz equation for high wave numbers or the Anderson model of localization.

As basis we mainly focus on three major aspects:

1. symmetric maximum weight matchings to increase the block diagonal dominance of the system,
2. inverse–based pivoting to drive the coarsening process and finally
3. filtering techniques to handle frequencies near zero eigenvalues.

These techniques are used within a multilevel framework and we will illustrate the resulting multilevel methods for selected numerical examples.

Joint work R. Rmer (U. Warwick) and M. Grote, O. Schenk (both U. Basel).

Scalable Parallel Performance of Molecular Dynamics Codes on Large Scale Cray Systems
Roberto Ansaloni, Cray Italy

Several molecular dynamics codes are routinely executed on large scale Cray vector and scalar MPP systems, with excellent scalability and performance results. Here we present some applications performance profiles and how the core algorithms involved can be optimally implemented in order to exploit the different architectural features of the Cray X1E and XT3.

New scalability frontiers in ab-initio Molecular Dynamics using the IBM BlueGene Supercomputer
Alessandro Curioni (IBM Research Lab Zurich)

Numerical solution of scattering and resonance problems by the pole condition method

F. Schmidt

We discuss idea and application of the pole condition method to solve time- harmonic scattering and resonance problems modeled by Helmholtz and Maxwell's equations on unbounded domains. A typical example for an electromagnetic scattering problem is the computation of transmission, re- flection and diffraction of light through photo-masks, a typical example for a resonance problem is the computation of modes of optical fibers.

This talk covers the essential aspects of the pole condition approach. First, the entire space is decomposed into an interior domain containing the ob ject of interest and an possibly heterogeneous exterior domain. The basic idea is to consider the Laplace transform of the field in the exterior domain.

We characterize the exterior fields by the poles of their Laplace transforms and say that a field satisfies the pole condition if its Laplace transform has no pole in the lower half of the complex plane. Fields satisfying the pole condition are outgoing fields. This enables us to characterize scattered fields based on the pole condition rather than Sommerfeld's radiation condition or the Silver-M uller radiation condition.

Moreover, the combination of finite element methods for the interior problem and the pole condition method for the exterior problem leads in a direct way to a number of new numerical algorithms. These algorithms allow to solve a number of practically relevant problems in 2D and 3D with variable coefficients and infinite obstacles in the exterior domain.

Particle Accelerator Science Development as a Driver for Large Scale Electromagnetic Simulations

Andreas Adelmann (PSI/AMAS)

A summary of activities that envisage a source to target simulation of the PSI cyclotron facility is presented. In the course of our ambitious high intensity upgrade, we aim to gain quantitative understanding of complex phenomena in our particle accelerators and beam lines. This ultimately leads to precise beam dynamics simulation. We discuss the implications of this undertaking with respect to physics and numerical models and highlight the importance of HPC. Examples will include three-dimensional space charge dominated beam transport in large structures and eigenmode calculations in large and complicated radio frequency cavities. Ideas for future development towards a time domain Maxwell solver for large structures will be presented.

In collaboration with P. Arbenz (ETH), R. Geus (PSI/Scientific Computing) and B. Oswald (PSI/AMAS)