Schedule

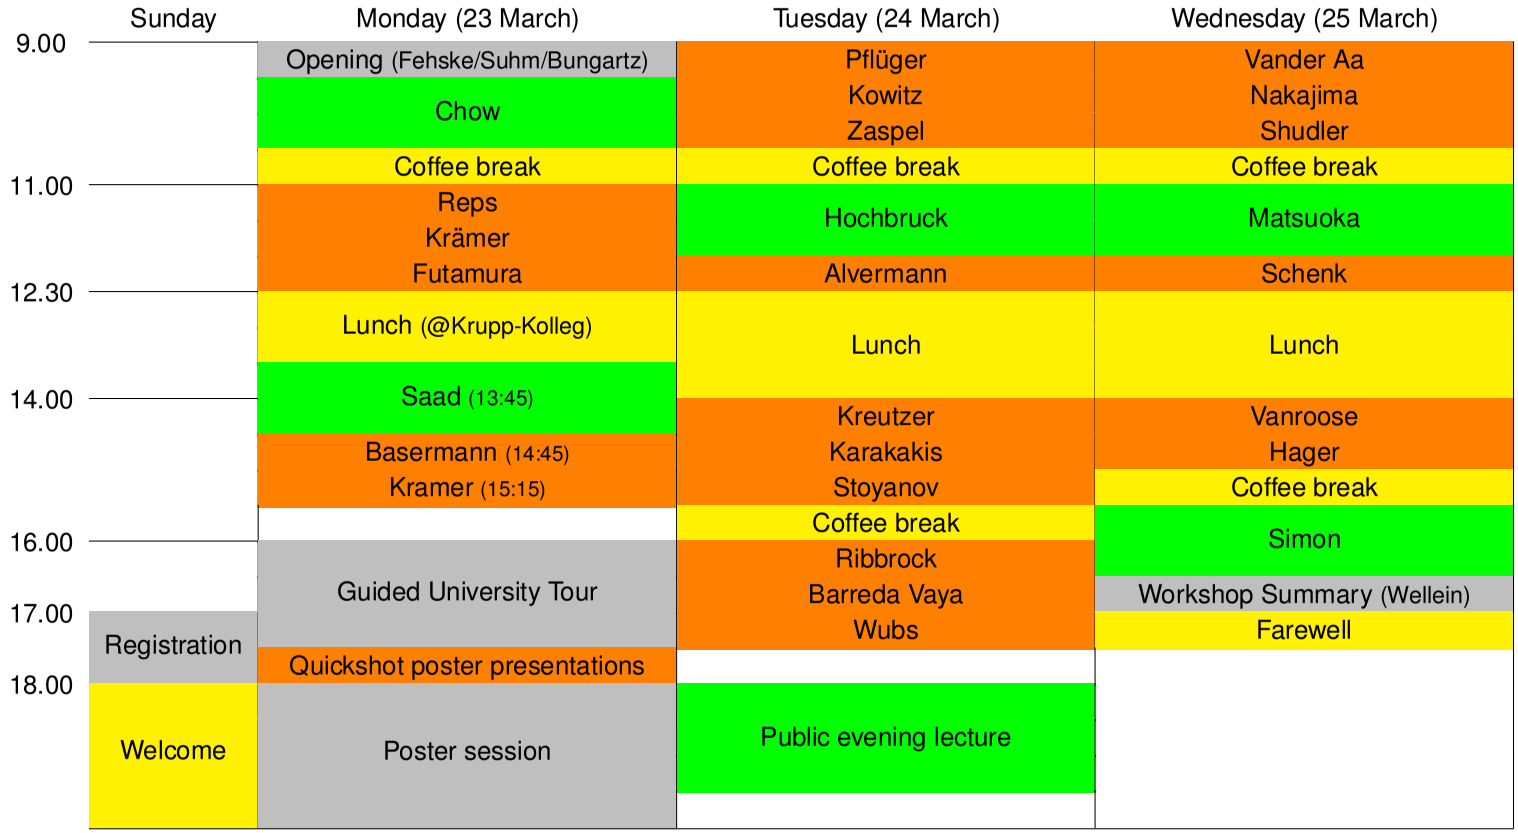

Registration:	Sunday, 17.00-19.00 in the Alfried-Krupp-Wissenschaftskolleg or during the workshop
Welcome:	Sunday, from 18.00 (Brasserie Hermann)
Public evening lecture:	Tuesday, 18.00 in the Alfried-Krupp-Wissenschaftskolleg
Workshop start:	Monday, 23 March, 9.00
Workshop end:	Wednesday, 25 March, 17.00

Monday Sessions

Workshop opening

9:00–9:10 **Welcome address**
Holger Fehske (Greifswald), Christian Suhm (WiKo Greifswald)

9:10–9:30 **SPPEXA: Software for Exascale Computing**
Hans-Joachim Bungartz (München)

Morning sessions

9:30–10:30 **Quantum Chemistry on Large-scale Heterogeneous Clusters**
Edmond Chow (Georgia Tech Atlanta)

10:30–11:00 **Coffee break** 

11:00–11:30 **Pipelined deflated Conjugate Gradients**
Bram Reps (Antwerpen)

11:30–12:00 **Adaptive choice of projectors in eigenvalue algorithms**
Lukas Krämer (Wuppertal)

12:00–12:30 **z-Pares: Software package for solving large sparse eigenvalue problems based on contour integral and complex moment**
Yasunori Futamura (Tsukuba)

12:30–13:45 **Lunch** 
Lunch is served in the Krupp-Kolleg

Afternoon sessions

13:45–14:45 **Divide and conquer algorithms for large Hermitian eigenvalue problems**
Yousef Saad (Univ. Minnesota)

14:45–15:15 **Scalable block methods and preconditioning for hardware efficient sparse eigensolutions**
Achim Basermann (Köln)

15:15–15:45 **GPU cloud computing for modeling photosynthetic energy transfer**
Tobias Kramer (Berlin)

Guided university tour

16:00–17:15 **Tour through the University main and historic buildings**
We start at 16:05 at the Rubenow Statue in the Domstraße

Poster session

17:30–18:00 **Quickshot poster presentations**

18:00–20:00 **Poster session**
Refreshments & Snacks

Tuesday Sessions

Morning sessions

9:00–9:30	**EXAHD – Simulation of Fusion Plasmas, Higher-Dimensional Problems, and Exa-Challenges** *Dirk Pflüger (Stuttgart)*
9:30–10:00	**EXAHD – Physics, Eigenproblems, and Faults** *Christoph Kowitz (München)*
10:00–10:30	**EXAHD – Combination Technique, Theory, and Preconditioning** *Peter Zaspel (Bonn)*
10:30–11:00	**Coffee break** 
11:00–12:00	**Time integration of evolution equations** *Marlis Hochbruck (Karlsruhe)*
12:00–12:30	**Polynomial expansions in quantum physics computations** *Andreas Alvermann (Greifswald)*
12:30–14:00	**Lunch** 

Afternoon sessions

14:00–14:30	**Optimizing the Performance of the Kernel Polynomial Method on Heterogeneous Systems** *Moritz Kreutzer (Erlangen)*
14:30–15:00	**SparseX: a library for high-performance Sparse Matrix-Vector multiplication on multicore platforms** *Vasileios Karakakis (München)*
15:00–15:30	**Task-based hybrid parallel sparse matrix-vector multiplication (SpMVM) with GASPI/GPI2** *Dimitar Stoyanov (Kaiserslautern)*
15:30–16:00	**Coffee break** 
16:00–16:30	**Integrating multi-threading and accelerators into DUNE-ISTL** *Dirk Ribbrock (Dortmund)*
16:30–17:00	**Leveraging Task-Parallelism in ILUPACK for many-core architectures** *Maria Barreda Vaya (Castellón de la Plana)*
17:00–17:30	**Eigenvalue computation with Trilinos and HYMLS for fluid flow problems** *Fred Wubs (Groningen)*

Public evening lecture

18:00–19:00	**What Supercomputers Do, and what Supercomputers Still Can't Do** *Horst Simon (Berkeley Lab)*

Wednesday Sessions

Morning sessions

9:00–9:30	**ExaShark: A Scalable Hybrid Array Kit for Exascale Simulation** *Tom Vander Aa (Leuven)*
9:30–10:00	**ppOpen-HPC: Open Source Infrastructure for Development and Execution of Large-Scale Scientific Applications on Post-Peta-Scale Supercomputers with Automatic Tuning (AT)** *Kengo Nakajima (Tokyo)*
10:00–10:30	*(Update)* **Exascaling Your Library: Will Your Implementation Meet Your Expectations?** *Sergei Shudler (Darmstadt)*
10:30–11:00	**Coffee break** 
11:00–12:00	**Harnessing the Deep Memory Hierarchy with Communication Reducing Algorithms for Exascale** *Satoshi Matsuoka (Tokyo)*
12:00–12:30	**Performance Engineering for Extreme-Scale Stochastic Optimizations on Petascale Architectures** *Olaf Schenk (Lugano)*
12:30–14:00	**Lunch** 

Afternoon sessions

14:00–14:30	**Raising the arithmetic intensity of Krylov solvers** *Wim Vanroose (Antwerpen)*
14:30–15:00	**The GHOST library: An overview of its functionality and a comparison to other approaches** *Georg Hager (Erlangen)*
15:00–15:30	**Coffee break** 
15:30–16:30	**Usable Exascale and beyond Moore's Law** *Horst Simon (Berkeley Lab)*

Closing & Farewell

16:30–17:00	**Workshop Summary** *Gerhard Wellein (Erlangen)*
17:00-17:30	**Farewell** Refreshments & Snacks

INVITED TALKS

Quantum Chemistry on Large-scale Heterogeneous Clusters

EDMOND CHOW

School of Computational Science and Engineering, College of Computing, Georgia Institute of Technology, 266 Ferst Drive, Atlanta, GA 30332-0765

echow@cc.gatech.edu

ABSTRACT

To efficiently utilize upcoming computer platforms, quantum chemistry algorithms and software must be able to use more and more parallelism. This is particularly challenging when the problem size is fixed, i.e., your molecule of interest and basis set are not increasing in size. We will discuss several new software components aimed at exascale quantum chemistry calculations. Features include 1) load balancing and reducing communication in the construction of Coulomb and exchange matrices, 2) using density matrix purification techniques in SCF iterations, 3) dynamic scheduling and heterogeneous calculation of electron repulsion integrals on CPUs and accelerators. Our algorithms have significantly lower parallel overhead than other commonly used algorithms, for example, those in NWChem. We will also discuss lessons learned and demonstrate our software components using more than 1 million cores on the Tianhe-2 supercomputer.

Time integration of evolution equations

MARLIS HOCHBRUCK

Institute for Applied and Numerical Analysis
Karlsruhe Institute of Technology
Kaiserstr. 89-93, 76133 Karlsruhe, Germany

marlis.hochbruck@kit.edu

ABSTRACT

In this talk we give an overview on recent advances in the time integration of evolution equations with special emphasis on the construction, analysis, implementation and application of various exponential integrators. Exponential integrators are time integration schemes, which involve the evaluation or approximation of the exponential (or related) function of a suitable matrix (e.g. the Jacobian of the differential equation). Such methods have been proposed about 60 years ago but for a long time they have been regarded as not practical. Significant advances on the approximation of the product of a matrix function with a vector including multiple time stepping approaches have renewed the interest in these integrators. By now it has been shown that such integrators are competitive or even outperform state of the art standard methods in certain applications.

We will discuss basic ideas to construct such integrators for different applications and show relations to standard explicit and implicit time integration schemes. Moreover, we consider the approximation of matrix functions using polynomial and rational Krylov methods and show how these methods can be implemented in problems arising in optics and photonics.

Harnessing the Deep Memory Hierarchy with Communication Reducing Algorithms for Exascale

SATOSHI MATSUOKA

Tokyo Institute of Technology
2-12-1 Oo-okayama, Meguro-ku Tokyo, Japan 152-8550

matsu@is.titech.ac.jp

ABSTRACT

Future exascale systems is predicted to be largely constrained by memory bandwidth and capacity, rather than FLOPS, in the context of real application performance. While it is plausible to accommodate either of the bandwidth or capacity, it would be difficult to satisfy both at the same time due to power and cost constraints. Although the new breed of non-volatile memory promises to provide lower power and higher density, due to their access characteristics they will likely supplement, not replace, DRAMs and associated caches in the processors. As a result, memory hierarchy will become increasingly deeper, and unless the algorithms will cope with them explicitly, it would be difficult to attain efficient performance. In order to satisfy both bandwidth and capacity requirements simultaneously, we are working on several breed of communication-reducing algorithms tailored to cope with the deep memory hierarchy. These include (1) partitioned variant of the recently proposed Cell-C-Sigma for sparse matrix operation on GPUs, (2) effective multi-level temporal blocking with programmability assist, and (3) offloading of sparse graph structures to NVM in graph algorithms. The speedup results are quite dramatic for each one, allowing large capacity problems to be resident in lower tiers of the memory hierarchy while preserving performance, without specialized hardware.

Divide and conquer algorithms for large hermitian eigenvalue problems

YOUSEF SAAD

Department of Computer Science and Engineering
University of Minnesota, Minneapolis, MN 55455

saad@cs.umn.edu

ABSTRACT

Algorithms based on Divide and conquer paradigms can lead to complex, but efficient and flexible algorithms for solving large hermitian eigenvalue problems. This talk will discuss how various such strategies can be combined to exploit both 'spectrum slicing', i.e., computing slices of the spectrum independently, and domain decomposition. These strategies are independent of each other but both are essential if one has to compute very large parts of the spectrum, as is the case in approaches that deal with excited states in solid state physics. The presentation will begin with an overview of polynomial filtering techniques. This general approach can be quite efficient in the situation where the matrix-vector product operation is inexpensive and when a large number of eigenvalues is sought, as is the case in electronic structure calculations for example. The method presented relies on a combination of the Lanczos algorithm with partial reorthogonalization and polynomial filtering. The second part of the talk will discuss our ongoing work on Domain-Decomposition (DD) type methods which rely on spectral Schur complements combined with Newton's iteration. This method is particularly appealing for interior eigenvalue problems. Finally, we will discuss methods based on what might be termed Rational Function Filtering (RFF) with good representatives like FEAST and the Sakurai-Sugiura approach. Here we will argue that DD should be an integral part of an approach of this type. We will also discuss a very important issue related to the choice of the filters in regards to the resulting linear systems to solve in the course of RFF algorithms.

Usable Exascale and beyond Moore's Law

HORST SIMON

Lawrence Berkeley National Laboratory
1 Cyclotron Road , Berkeley, CA 94720, USA
HDSimon@lbl.gov

ABSTRACT

It has been a truly remarkable feat of technological innovation that the number of components on an integrated circuit has been doubling about every 18 months since the introduction of this technology more than 50 years ago. Gordon Moore observed this trend in 1965, and today we are reaping the benefits of this continued exponential growth. High performance computing (HPC) has been a beneficiary of uninterrupted growth, and performance of the top HPC systems has been doubling about every year until 2004, based on about equal contributions of Moore's law and increasing parallelism in the highest end system. HPC performance growth has been well documented by the TOP500.

A first change in the dynamics of the semiconductor market happened in 2004, when Dennard scaling become no longer feasible, clock rates leveled off, and single processor performance stalled. Continued HPC system performance increases were obtained by doubling parallelism. However, over the last five years HPC performance growth has been slowing measurably, and in this presentation several reasons for this slowdown will be analyzed.

It will be shown that in order to reach usable exascale performance over the next decade some fundamental changes will have to occur in HPC systems architecture, in particular a transition from a compute centric to a data movement centric point of view. The goal of usable exascale is understood to mean the ability to obtain a factor of 1000 improvement of sustained application performance over current petaflops level applications.

Moore's Law scaling of semiconductor devices will come to an end in about 10 years. The HPC community will have to begin exploring alternatives to current CMOS based von Neumann architectures. Explorations of quantum computing and neuromorphic computing have just been reported with early results. The prospects of these technologies for post-Moore's Law supercomputing will be explored.

TALKS

Polynomial expansions in quantum physics computations

ANDREAS ALVERMANN

Institut für Physik, Ernst-Moritz-Arndt-Universität, 17487 Greifswald, Germany

alvermann@physik.uni-greifswald.de

ABSTRACT

Polynomial expansions can play an important role in numerical quantum physics, especially for few particle problems or continuous degrees of freedom. They can be used for two complementary purposes: First, to set up the matrix representation of the physical Hamiltonian, using polynomials as the building blocks of the computational Hilbert space. Second, to compute the properties of interest, e.g., with time-propagation techniques based on Chebyshev and Faber polynomials.

The talk will address specific aspects of these two purposes: First, how to use polynomials on sparse grids for bosonic degrees of freedom. Second, how to understand three key algorithms for quantum physics computations (Kernel Polynomial Method, Chebyshev time propagation, Krylov eigensolvers) in terms of the underlying polynomial approximations. Although my emphasis will be on the second aspect, I will shortly discuss the relevance of these techniques for current physics research on non-Markovian dynamics in open quantum systems.

Leveraging Task-Parallelism in ILUPACK for many-core architectures

MARIA BARREDA VAYA

Universitat Jaume I, Departamento de Ingeniería y Ciencia de los Computadores, Av. de Vicent Sos Baynat, s/n, 12071 Castellón de la Plana, Spain

mvaya@uji.es

ABSTRACT

The ILUPACK library uses a preconditioned iterative method to solve sparse linear systems, incorporating the inverse-based approach to the ILU factorization to obtain multilevel preconditioners. The main objective of this work is to describe the efficient task-parallel implementation of ILUPACK on multithreaded processors by using OmpSs for the solution of SPD linear systems. We capture the task dependencies of the two most challenging operations in the method (calculation of the preconditioner and its application). From this information, the OmpSs runtime can implement a correct and efficient schedule of the entire solver. We test this implementation on multicore platforms with a high number of cores, and also in accelerators like Intel Xeon Phi. To improve the performance in these platforms, we have introduced several changes. We have assigned priorities to the tasks in order to execute the larger tasks before, if possible, reducing the execution time. Also, we have merged some small tasks to decrease the overhead of the runtime, and we have incorporated nested tasks and specific options of the runtime to improve the performance in NUMA architectures. For the Intel Xeon Phi, we have analyzed the thread affinity, implementing the best option to map the threads to the cores efficiently. Our results with high-end multicore platforms equipped with Intel, Intel Xeon Phi and AMD processors report significant performance gains, demonstrating that OmpSs provides an efficient and close-to-seamless means to leverage the concurrency in a complex scientific code like ILUPACK.

Scalable block methods and preconditioning for hardware efficient sparse eigensolutions

ACHIM BASERMANN[1], MELVEN RÖHRIG-ZÖLLNER, AND JONAS THIES

[1] *German Aerospace Center (DLR)*
Simulation and Software Technology
Linder Höhe, 51147 Cologne, Germany

achim.basermann@dlr.de

ABSTRACT

Block variants of the Jacobi-Davidson method for computing a few extreme eigenpairs of a large sparse matrix are known to improve the robustness of the standard algorithm, but are generally shunned because the total number of floating-point operations increases. We present a scalable implementation of a block Jacobi-Davidson solver with blocked GMRES preconditioning. By detailed performance engineering and numerical experiments we demonstrate that the increase in operations is typically more than compensated by the performance gains of the block operations on modern parallel hardware architectures. Thus the block solver is both more efficient and robust than its single vector counterpart.
Moreover, we discuss a row projection method for the solution of ill-conditioned linear systems encountered in the FEAST algorithm for computing inner eigenpairs. We introduce a novel hybridly parallel and fully iterative implementation of the eigenvalue solver which exploits parallelism on several levels and present numerical and performance results.

z-Pares : Software package for solving large sparse eigenvalue problems based on contour integral and complex moment

YASUNORI FUTAMURA[1] AND TETSUYA SAKURAI

[1] *University of Tsukuba, Sogokenkyutou B 1012, 1-1-1 Tennohdai, Tsukuba, Ibaraki, Japan*

futamura@cs.tsukuba.ac.jp

ABSTRACT

In this talk, we show an implementation and performance evaluation of a software package named z-Pares for solving large sparse generalized eigenvalue problems. z-Pares implements so-called the Sakurai-Sugiura (SS) method which computes eigenvalues located in a specified domain by a numerical contour integration. As in other contour integral based methods, independent linear systems with respect to each quadrature point are solved in the method. The key difference between the SS method and other contour integral based methods is the use of the complex moment which enables users to reduce the number of right hand sides of the linear systems. Thus the SS method is more beneficial when it is used in conjunction with iterative linear solvers. Since the linear systems can be solved independently of each other and each linear system can be solved in parallel, contour integral based methods have a nested parallelism. z-Pares implements this nested parallelism by introducing hierarchically defined MPI communicators. Currently, z-Pares is the only contour integral based eigensolver package which has this feature. By this hierarchical distributed parallel implementation, z-Pares has a capability to handle problems that are so large that all elements of vectors cannot be stored in one node. To show the capability and the performance of z-Pares, we present use cases on large-scale problems come from electronic structure calculations and vibration analyses. Parallel performances are evaluated on state-of-art supercomputers such as K-computer.

The GHOST library: An overview of its functionality and a comparison to other approaches

GEORG HAGER

Erlangen Regional Computing Center (RRZE)
Martensstr. 1, 91058 Erlangen, Germany

georg.hager@fau.de

ABSTRACT

If there is one definite consent about the architecture of future high performance computers, it is that they will be strongly heterogeneous. Moreover we expect topology effects to become even more important than they are today. On the other hand, MPI+X-style programming is here to stay. The GHOST library (General Hybrid Optimized Sparse Toolkit) aims at helping programmers to leverage parallelism without abstracting performance-critical hardware features. It provides low-level building blocks for constructing highly parallel algorithms dealing with sparse matrices.
The talk will give an overview of the functionality of GHOST and provide instructive examples for how to deal with heterogeneity and parallelism while still retaining a well-known programming model.

SparseX: a library for high-performance Sparse Matrix-Vector multiplication on multicore platforms

VASILEIOS KARAKASIS

Leibniz-Rechenzentrum (LRZ)
Boltzmannstr. 1, 85748 Garching b. München (Germany)

karakasis@lrz.de

ABSTRACT

The Sparse Matrix-Vector Multiplication (SpMV) kernel ranks among the most important and thoroughly studied linear algebra operations, as it lies at the heart of many iterative methods for the solution of sparse linear systems and often poses a sever performance bottleneck. Its optimization, which is intimately associated with the data structures used to store the sparse matrix, has always been of particular interest to the applied mathematics and computer science communities and has attracted further attention since the advent of multicore architectures. In this talk, we will present SparseX, a software package for performing sparse matrix-vector multiplications on multicore platforms that embraces the state-of-the-art CSX sparse matrix storage format in order to deliver high efficiency, while providing a highly usable interface that requires limited or no tuning. Performance results indicate that SparseX library achieves superior performance on large scale problems.

The paper for the CSX storage format appeared in IEEE Transactions on Parallel and Distributed Systems (TPDS), 24(10):1930–1940, 2013.

As of this writing, the SparseX library paper is under review by the ACM Transaction on Mathematical Software Journal, while SparseX is available online at http://research.cslab.ece.ntua.gr/sparsex

EXAHD – Physics, Eigenproblems, and Faults

CHRISTOPH KOWITZ

Technische Universität München, Fakultät für Informatik, Boltzmannstraße 3, 85748 Garching, Germany

kowitz@in.tum.de

ABSTRACT

The second EXAHD talk addresses the exascale challenge resilience. There, the combination approach introduces algorithm-based fault tolerance. Additionally, we present further details of the underlying physical problem. The simulation of hot fusion plasmas in plasma physics with the software GENE allows us to study several application scenarios. We will illustrate the use of the combination technique and necessary adaptions for eigenvalue computations, where we are interested in the modes that drive the turbulences.

GPU cloud computing for modeling photosynthetic energy transfer

TOBIAS KRAMER

Institut für Physik, Humboldt-Universität Berlin,
Newtonstr. 15, D 12489 Berlin (Germany)

tobias.kramer@physik.hu-berlin.de

ABSTRACT

We share our experience with developing the first GPU-compute enabled nanoHUB.org cloud-computing tool ("GPU-HEOM") for studying quantum-mechanical effects in photosynthesis.
NanoHUB is a NSF (USA) supported project administered by Purdue University. Registration is open to all academic users for free.
Using the nanohub GPU-HEOM tool, users without hpc experiences have GPU and conventional computing clusters available to solve computationally demanding problems.
We discuss how this low-entrance barrier to hpc computing helps to foster collaborations on computationally demanding problems. For reproducible results and citation purposes each tool at nanohub has a digital object identifier.

Adaptive choice of projectors in eigenvalue algorithms

LUKAS KRÄMER

Bergische Universität Wuppertal, Gaußstraße 20, 42119 Wuppertal

lkraemer@math.uni-wuppertal.de

ABSTRACT

We discuss algorithms for the Hermitian eigenvalue problem that are based on the approximate computation of spectral projectors. Such methods include the FEAST method or methods based on polynomial approximation. Typically, the approximation scheme is fixed over the iterations of the eigenvalue algorithm. Choosing the approximation dynamically has some advantages. It may even decrease the overall runtime.

We present two approaches of introducing adaptivity in eigensolvers that are based on projectors. First, polynomial approximation with dynamically chosen polynomial degree is presented. Second, we present techniques involving numerical integration in combination with an iterative linear solver. Here, the adaptivity is controlled via the accuracy of the linear solver. To conclude the talk we present numerical results.

Optimizing the Performance of the Kernel Polynomial Method on Heterogeneous Systems

MORITZ KREUTZER

Erlangen Regional Computing Center, Martensstraße 1, 91058 Erlangen

moritz.kreutzer@fau.de

ABSTRACT

In this work we present a highly optimized and fully heterogeneous implementation of the Kernel Polynomial Method (KPM, which is a well-established scheme used for the determination of the eigenvalue density and spectral properties of large sparse matrices in quantum physics and quantum chemistry. At the algorithmic level we apply optimizations which eventually decouple the sparse matrix problem posed by the KPM from main memory bandwidth both on CPU and GPU. By using block vectors we can alleviate the effects of scattered data access and enable pure data streaming. Using optimized device-specific implementations of this algorithm within our hybrid-parallel framework GHOST enables large scale heterogeneous computations of which we will show performance data on a petascale-class Cray XC30 system.

ppOpen-HPC: Open Source Infrastructure for Development and Execution of Large-Scale Scientific Applications on Post-Peta-Scale Supercomputers with Automatic Tuning (AT)

KENGO NAKAJIMA

The University of Tokyo, 2-11-16 Yayoi, Bunkyo-ku, 1138658 Tokyo

nakajima@cc.u-tokyo.ac.jp

ABSTRACT

ppOpen-HPC is an open source infrastructure for development and execution of optimized and reliable simulation code on post-peta-scale (pp) parallel computers based on many-core architectures, and it consists of various types of libraries, which cover general procedures for scientific computation. This is a part of JST/CREST "Development of System Software Technologies for post-Peta Scale High Performance Computing" project, and supported by JST during FY.2011-2015.

In ppOpen-HPC, we are focusing on five types of discretization methods for scientific computing, which are FEM, FDM, FVM, BEM, and DEM. ppOpenHPC provides a set of libraries covering various types of procedures for these five methods, such as parallel I/O of data-sets, assembling of coefficient matrix, linear-solvers with robust and scalable preconditioners, adaptive mesh refinement (AMR), and dynamic load-balancing. The target post-peta-scale system is the Post T2K System of the University of Tokyo and University of Tsukuba based on many-core architectures, such as Intel MIC/Xeon Phi. It will be installed in FY.2015-2016 and its peak performance is expected to be 20-30 PFLOPS. Automatic tuning (AT) enables a smooth and easy shift to further development on new/future architectures through the use of ppOpen-AT. Directive-based special AT languages for specific procedures in scientific computing, focused on optimum memory access, are being developed.

Source codes of ppOpen-HPC with documents are available at http://ppopenhpc.cc.u-tokyo.ac.jp/. In this presentation, recent progresses of the "ppOpen-HPC" project are overviewed. Especially, achievements in the development of ppOpen-MATH/MG, which is a geometric multigrid solver in ppOpen-HPC, are described in detail. Furthermore, future development issues of ppOpen-HPC, and collaborations with the ESSEX project will be addressed.

EXAHD – Simulation of Fusion Plasmas, Higher-Dimensional Problems, and Exa-Challenges

DIRK PFLÜGER

University of Stuttgart, Institute for Parallel and Distributed Systems, Universitätsstr. 38, 70569 Stuttgart

Dirk.Pflueger@ipvs.uni-stuttgart.de

ABSTRACT

Higher-dimensional problems (beyond the classical four dimensions: space and time) are classical candidates for current and future HPC. A prominent example are turbulence simulations of hot fusion plasmas, a core task on the way to fusion power plants. The project EXAHD within DFG's SPPEXA program approaches higher-dimensional problems with the sparse grid combination technique. The solution of a single problem is numerically decoupled into a superposition of many, but much smaller and anisotropic problems. On the one hand, this allows to reduce the number of degrees of freedoms to a feasible range: classical approaches suffer from an exponential dependency on the number of dimensions. On the other hand, this hierarchical splitting provides a great handle on the exascale challenges scalability, load balancing, and resilience. In three talks, we show different aspects of the EXAHD project.

The first part introduces the core problem, the solution of higher-dimensional PDEs at the example of the simulation of hot fusion plasmas, and briefly sketches the combination technique. We then address the second level of parallelism, introduced by the numerical decoupling of the single, full-resolution problem. We present how this tackles the exascale challenges scalability (by breaking the classical need for global communication) and load balancing (exploiting the two-level splitting).

Pipelined deflated Conjugate Gradients

BRAM REPS[1], PIETER GHYSELS, WIM VANROOSE, AND KARL MEERBERGEN

[1] *University of Antwerp, Middelheimlaan 1, B-2020 Antwerp, Belgium*

bram.reps@uantwerp.be

ABSTRACT

Key building blocks for iterative solvers based on Krylov subspaces are dot products that ensure orthogonality of the basis or make sure that the search directions point towards the solution. However, on a parallel machine, where solution vectors are distributed over many nodes, such a dot product is a global reduction and suffers from long latencies. Indeed, a global reduction requires synchronization of all participating nodes followed by a reduction tree and the time to complete this operation is susceptible to variability in the system.

Deflated Krylov subspace method are used when there are some troublesome eigenvectors with small eigenvalue in the system. These eigenvectors delay the convergence and by removing them through deflation the convergence is accelerated. However, deflation introduces additional global reductions in the iteration what leads to additional latencies.

In this contribution we present possible ways to reduce the cost of the global reductions in deflated Conjugate Gradients (CG). We explore fusion of the dot products into one instance of global communication per iteration. We also look at pipelining this communication phase with other useful work, similar to regular pipelined CG [2] and GMRES [1]. Furthermore, we introduce selective deflation where the solution is only deflated at selected steps in the iteration.

References

[1] P. GHYSELS, T.J. ASHBY, K. MEERBERGEN, W. VANROOSE, *Hiding Global Communication Latency in the GMRES Algorithm on Massively Parallel Machines*, SIAM J. Sci. Comput., 35, 2013.

[2] P. GHYSELS, W. VANROOSE, *Hiding global synchronization latency in the preconditioned Conjugate Gradient algorithm*, Parallel Computing, 2013.

Integrating multi-threading and accelerators into DUNE-ISTL

DIRK RIBBROCK[1], STEFFEN MÜTHING, MARKUS GEVELER, AND DOMINIK GÖDDEKE

[1] *Technische Universität Dortmund, Fakultät für Mathematik, Lehrstuhl LSIII, Vogelpothsweg 87, 44227 Dortmund*

dirk.ribbrock@mathematik.tu-dortmund.de

ABSTRACT

One important goal of the EXA-DUNE project within SPPEXA is to bring together the flexibility of the DUNE framework and the hardware-oriented numerics of the FEAST toolkit. In this talk, we discuss ongoing efforts in the iterative solvers subsystem, and demonstrate how its former MPI-only implementation is modified to support MPI+X, where X can be a threading implementation for multicore CPUs and the Xeon Phi, or a CUDA implementation for NVIDIA GPUs.

It is well known that the performance of sparse matrix operations, the key building block of our iterative solvers, varies strongly for different operations and hardware platforms. Here, we employ the unified SELL-C format recently developed by Kreutzer et al., which has been shown to efficiently exploits the SIMD features of all three target architectures listed above.

Additionally we exploit the fact, that our application uses discontinous Galerkin discretizations resulting in sparse matrices with small dense blocks, by extending the aforementioned SELL-C format with a block extension. This drastically improves the arithmetic intensity and reduces the memory consumption.

The efficiency of our approach is underlined by benchmark computations that exhibit reasonable speedups over the CPU-MPI-only case for complete preconditioned Krylov subspace methods on heterogeneous hardware.

Performance Engineering for Extreme-Scale Stochastic Optimizations on Petascale Architectures

OLAF SCHENK

Institute of Computational Science, Universita della Svizzera italiana, Via Giuseppe Buffi 13, 6900 Lugano, Switzerland

olaf.schenk@usi.ch

ABSTRACT

We present a scalable approach that computes stochastic optimization problems under uncertainty in operationally-compatible time. Several algorithmic and performance engineering advances are discussed for these stochastic optimization problems. The new developments include novel incomplete augmented multicore sparse factorizationss, multicore- and GPU-based dense matrix implementations, and communication-avoiding Krylov solvers. We also improve the interprocess communication on Cray systems to solve 24-hour horizon power grid problems from electrical power grid systems of realistic size with up to 1.95 billion decision variables and 1.94 billion constraints. Full-scale results are reported on Cray XC30 and BG/Q, where we observe very good parallel efficiencies and solution times within a operationally defined time interval. To our knowledge, "real-time"-compatible performance on a broad range of architectures for this class of problems has not been possible prior to present work.

(Update) Exascaling Your Library: Will Your Implementation Meet Your Expectations?

SERGEI SHUDLER

TU Darmstadt, Mornewegstraße 30, 64293 Darmstadt

s.shudler@grs-sim.de

ABSTRACT

Many libraries in the HPC field encapsulate sophisticated algorithms with clear theoretical scalability expectations. However, hardware constraints or programming bugs may sometimes render these expectations inaccurate or even plainly wrong. While algorithm engineers have already been advocating the systematic combination of analytical performance models with practical measurements for a very long time, we go one step further and show how this comparison can become part of automated testing procedures. The most important applications of our method include initial validation, regression testing, and benchmarking to compare implementation and platform alternatives. Advancing the concept of performance assertions, we verify asymptotic scaling trends rather than precise analytical expressions, relieving the developer from the burden of having to specify and maintain very fine-grained and potentially non-portable expectations. In this way, scalability validation can be continuously applied throughout the whole development cycle with very little effort. Using MPI as an example, we show how our method can help uncover non-obvious limitations of both libraries and underlying platforms.

Task-based hybrid parallel sparse matrix-vector multiplication (SpMVM) with GASPI/GPI2

DIMITAR STOYANOV

Fraunhofer-Institut für Techno- und Wirtschaftsmathematik ITWM
Fraunhofer-Platz 1, 67663 Kaiserslautern

stoyanov@itwm.fhg.de

ABSTRACT

We present a hybrid-parallel implementation of the sparse matrix-vector multiplication kernel (SpMVM) based on the communication library GASPI/GPI.
GASPI/GPI belongs to the partitioned global address space (PGAS) approaches, which are an alternative to the Message Passing (MPI). It is known that presently integrating multi-threading into message passing models is still more complicated. In this context, the key advantage of GASPI/GPI programming is that it is thread based and it offers full flexibility to follow the task model approach, being free from the restrictions of the standard hybridization with MPI and threads/OpenMP. Note also that such a task-based hybrid pattern is the appropriate choice to reveal the hierarchical parallelism of clusters with multi- and many- core nodes.
The SpMVM kernels are of key importance for the overall performance of Krylov subspace solvers and the hybrid-parallel implementation of such kernels is still a challenge nowadays. We present here a task-based SpMVM implementation over GPI. Our GPI-based SpMVM kernel includes asynchronous communication and differs in many aspects from the classical hybridization.
Further, we test our implementation on the numerical solution of a Poisson boundary value problem in a unit cube and we compare the performance of our GASPI-solver against PETSc. The Jacobi-preconditioned Richardson method is a good choice here because it allows for a clean comparison between the solvers: the calculations performed in the GPI-Richardson and PETSc-Richardson routines are identical. Furthermore, to guarantee a fair comparison we also employ two different types of Domain Decomposition (DD): cutting planes approach and graph partitioning (using MeTis).
Our tests show a significant performance advantage of the GASPI/GPI-implementation against (flat-MPI) PETSc. Moreover, in the case of the more appropriate MeTis-DD, our GASPI/GPI implementation shows also a much better scalability compared to PETSc.

ExaShark: A Scalable Hybrid Array Kit for Exascale Simulation

TOM VANDER AA

ExaScale Life Lab at imec, Kapeldreef 75, B-3001 Leuven, Belgium

Tom.VanderAa@imec.be

ABSTRACT

Design and implementation of solvers that scale to large HPC systems is not only a challenge from a numerical point of view, but also - and even more so - from an implementation and verification point of view. To estimate how a newly developed solver would behave on modern heterogeneous HPC architectures, knowledge of many different programming paradigms and libraries is needed: shared memory threading techniques like Pthreads, OpenMP or TBB, inter-node distribution techniques like MPI or GPI, and most probably a combinations of these.

This talk presents ExaShark: a library for handling n-dimensional distributed data structures that are at the core of these HPC applications. ExaShark's goal is to reduce the increasing programming burden while still offering good performance. It offers its users a global array like usability while its runtime can be configured to use any or a combination of the aforementioned programming models.

ExaShark has been used by to develop applications of different scale: ranging from standalone advanced pipelined conjugate gradient solvers, to a complete particle-in-cell simulator.

Raising the arithmetic intensity of Krylov solvers

WIM VANROOSE[1], BRAM REPS, PIETER GHYSELS, OLAF SCHENK, AND SIMPLICE DONFACK

[1] *University of Antwerpen, Middelheimlaan 1, 2020 Antwerpen*

wim.vanroose@uantwerpen.be

ABSTRACT

The performance of preconditioned Krylov solvers are severely hampered by the limited memory bandwidth. Each of the building blocks of a multigrid preconditioned Krylov solver is a low arithmetic intensity operation. We discuss how the algorithm can be organized such the arithmetic intensity is raised and that we can benefit from SIMD operations.

Eigenvalue computation with Trilinos and HYMLS for fluid flow problems

SVEN BAARS, JONAS THIES, AND FRED WUBS[1]

[1] *University of Groningen*
P.O. box 407, 9700 AK, Groningen, The Netherlands
f.w.wubs@rug.nl

ABSTRACT

For the determination of the stability of equilibria in flows, one needs to find the most critical eigenvalues of the associated generalized eigenvalue problem as a function of a parameter [3]. We performed such computations using the Anasazi package available in Trilinos. For this, one needs to activate the Cayley transformation available in the package and to provide a preconditioner to solve the systems it generates; for the preconditioner we use our in house developed multilevel incomplete LU factorization HYMLS [2].
HYMLS is a hybrid direct/iterative preconditioner for the Jacobian of the incompressible Navier-Stokes equations on structured grids. The method features a near grid-independent convergence rate and does not break down at high Reynolds numbers. The structure-preserving preconditioning technique allows recursive application of the algorithm resulting a a multilevel solver. The implementation of HYMLS is performed using the Epetra package of Trilinos.
Robustness and parallel performance will be demonstrated by directly computing highly unstable steady states with the associated eigenvalue problem of among others the lid-driven cavity problem in two and three space dimensions.

References

[1] G.L.G. Sleijpen and F.W. Wubs. Exploiting Multilevel Preconditioning Techniques in Eigenvalue Computations. SIAM Journal on Scientific Computing, 25(4):1249-1272, 2003.

[2] F.W.Wubs and J.Thies, A robust two-level incomplete factorization for (Navier-) Stokes saddle point matrices, SIAM J. Matrix Anal. Appl., 32:1475 - 1499, 2011.

[3] H.A. Dijkstra et al. Numerical Bifurcation Methods and their Application to Fluid Dynamics: Analysis beyond Simulation, *Communications in Computational Physics (CiCP)*, 2014.

EXAHD – Combination Technique, Theory, and Preconditioning

PETER ZASPEL

University of Bonn, Institute for Numerical Simulation,
Wegelerstrasse 6, 53115 Bonn

zaspel@ins.uni-bonn.de

ABSTRACT

The third EXAHD talk presents theoretical results for the combination technique. Our approach applies to the solution of general, higher-dimensional PDEs and is not restricted to our example application. We will show new developments and numerical results and explain where they can be applied. They include, but are not limited to, appropriate preconditioning techniques for the solution of time-dependent PDEs with the combination technique.

POSTERS

On the adaptive finite element analysis of the Kohn-Sham equations: Methods, algorithms, and implementation

DENIS DAVYDOV[1], TOBY D YOUNG, AND PAUL STEINMANN

[1] *Lehrstuhl für Technische Mechanik, Universität Erlangen-Nürnberg*
Paul-Gordan-Straße 3, 91052 Erlangen

denis.davydov@ltm.uni-erlangen.de

ABSTRACT

Details of an implementation of a numerical code for computing the Kohn-Sham equations are presented and discussed. A fully self-consistent method of solving the quantum many- body problem within the context of density functional theory using a real-space method based on finite element discretisation of real-space is considered. Various numerical issues are explored such as, (i) initial mesh motion aimed at co-aligning ions and vertices; (ii) a priori and a posteriori optimization of the mesh based on Kelly's error estimate; (iii) the influence of the quadrature rule and variation of the polynomial degree of interpolation in the finite element discretisation on the resulting total energy. Additionally, (iv) explicit, implicit and Gaussian approaches to treat the ionic potential are compared. A quadrupole expansion is employed to provide boundary conditions for the Poisson problem. To exemplify the soundness of our method, accurate computations are performed for hydrogen, helium, lithium, carbon, oxygen, neon and hydrogen molecule ion. Our methods, algorithms, and implementation are shown to be stable with respect to convergence of the total energy in a parallel computational environment.

The FEAST eigensolver

MARTIN GALGON

Bergische Universität Wuppertal, Gaußstraße 20, D-42119 Wuppertal

galgon@math.uni-wuppertal.de

ABSTRACT

A poster will be presented to give an overview of the current state of the FEAST eigensolver in the context of the ESSEX project, including improvements in terms of robustness and reliability. Also, approaches for the solution of the inner linear systems as well as polynomial approximation as alternative to integration of the resolvent will be detailed. Results for some of the benchmark problems from the ESSEX project will be shown.

PERMON toolbox

DAVID HORAK

VSB-Technical University of Ostrava, IT4Innovations
17.listopadu 15, Ostrava - Poruba, 70833, Czech Republic
david.horak@vsb.cz

ABSTRACT

Many problems are described by partial differential equations. To be solved with computers, they have to be discretized, e.g. with the popular Finite Element Method (FEM). We typically get large sparse linear systems of equations, but in case of constrained problems such as contact problems of mechanics, quadratic programming problems (QPs) arise. QPs also arise in other disciplines like least-squares regression, data fitting, data mining, support vector machines, control systems and many others. Large scale problems that are not solvable on usual personal computers can be solved only in parallel on supercomputers. Domain decomposition methods (DDM) come here into play. They solve an original problem by splitting it into smaller subdomain problems that are independent, allowing natural parallelization. Finite Element Tearing and Interconnecting (FETI) methods form a successful subclass of DDM. They belong to non-overlapping methods and combine sparse iterative and direct solvers. FETI methods allow highly accurate computations scaling up to tens of thousands of processors. Due to limitations of commercial packages, problems often have to be adapted to be solvable. This is an expensive process and results reflect less accurately physical phenomena. Moreover, it takes a long time before the most recent numerical methods needed for High Performance Computing (HPC) are implemented into such packages. These issues lead us to establish the PERMON (Parallel, Efficient, Robust, Modular, Object-oriented, Numerical) toolbox. It makes use of theoretical results in discretization techniques, quadratic programming algorithms, and domain decomposition methods. It incorporates our own codes, and makes use of renowned open source libraries. We focus on engineering applications (linear elasticity, contact problems with friction, elasto-plasticity, shape optimization and others) as well as altruistic ones (medical imaging, ice-sheet melting, climate changes modelling, etc.).

Fault tolerance inherent in a complex moment-based eigensolver

AKIRA IMAKURA[1], YASUNORI FUTAMURA, AND TETSUYA SAKURAI

[1] *University of Tsukuba, 1-1-1 Tennodai, Tsukuba, Ibaraki 305-8573, Japan*

imakura@cs.tsukuba.ac.jp

ABSTRACT

In this paper, we consider complex moment-based eigensolvers for computing all eigenvalues located in a certain region and their corresponding eigenvectors for a generalized eigenvalue problem. For solving the generalized eigenvalue problem, Sakurai and Sugiura have proposed a projection type method that uses certain complex moment matrices constructed by contour integral in 2003 [3]. Thereafter, several researchers have actively studied improvements and related eigensolvers based on the complex moment-based eigensolver.

Recently we focused on the Rayleigh–Ritz type complex moment-based eigensolver called the block SS–RR method [1], and analyzed the error bounds of the block SS–RR method [2]. In this talk, we show that the results of the analyses indicate a fault tolerance inherent in the block SS–RR method. We then provide an fault tolerant strategy which does not require checkpointing and replication techniques in the most time-consuming part of the eigensolver. Our numerical experiments suggest that our strategy recovers software faults like bit-flip with small additional costs.

References

[1] Ikegami, T., Sakurai, T.:, Contour integral eigensolver for non-Hermitian systems: a Rayleigh-Ritz-type approach. Taiwan. J. Math., **14** (2010), 825–837

[2] Imakura, A., Du, L. Sakurai, T.: Accuracy analysis on the Rayleigh-Ritz type of the contour integral based eigensolver for solving generalized eigenvalue problems, Technical Report of Department of Computer Science, University of Tsukuba (CS-TR), CS-TR-14-23.

[3] Sakurai, T., Sugiura, H.: A projection method for generalized eigenvalue problems using numerical integration. J. Comput. Appl. Math., **159** (2003), 119–128.

The EigenExa library : dense eigenvalue solver for post-petascale computing

TOSHIYUKI IMAMURA

RIKEN AICS, 7-1-26, Minatojima-minami-machi, Chuo-ku, Kobe, Hyogo, 650-0047, Japan

imamura.toshiyuki@riken.jp

ABSTRACT

Emerging exa-scale computational science technologies requires the progress of both sparse-matrix solver and dense-matrix solver (former for ultra-scale problems, and latter for medium and small ranged problems). Currently, Japan Science Technology agency has conducted development of the pos-petascale system software, where we organize a taskforce team to develop a high-performance dense eigenvalue solver EigenExa [1, 2]. We recognize that numerical linear algebra for exa-scale computing needs following issues to be refined: 1) Highly parallel algorithm, and 2) Hiding and avoiding technique for reducing overhead like memory transfer, synchronization and load imbalance, etc. As the modern eigensolvers do, EigenExa also employs the band-reduction algorithm [1] to refine the wall of memory bandwidth. However, EigenExa avoids the tridiagonal-banded back-transformation. We introduce a novel 1-stage scheme on which we compute eigenpairs without banded-tridiagonal transformation. In addition, asynchronous collective communication is adopted to conceal them behind computation. EigenExa successively diagonalizes a large-scale matrix by taking account of the full node of K computer. In the present, it is thought as a world record that one million matrix is diagonalized within one hour by EigenExa [2]. The result suggests that the employed algorithm performs on a 10 petaflops system, and such an ultra-scale diagonalzation can be done within a realistic time by dense matrix solver. It gives a big impact to the general user of a supercomputer to put such a large-scale eigenvalue calculation into effect aggressively.

References

[1] T. Imamura et al., "Development of a High Performance Eigensolver on the Peta-Scale Next Generation Supercomputer System", PNST, AESJ, Vol. 2, pp.643-650 (2011) .

[2] T. Imamura, The EigenExa Library – High Performance & Scalable Direct Eigensolver for Large-Scale Computational Science, HPC in Asia, ISC14.

Auto-tuning shared memory parallel Sparse BLAS operations in librsb-1.2

MICHELE MARTONE

Max-Planck-Institut für Plasmaphysik
Boltzmannstrasse 2, D-85748 Garching bei München, Germany

michele.martone@ipp.mpg.de

ABSTRACT

The Sparse BLAS standard specifies a programming interface (C and Fortran) for several 'sparse' computing kernels, foremost the performace critical ones used in iterative methods: multiply and triangular solve of a sparse matrix by a dense one. This includes support for real or complex numbers in single or double precision, eventual symmetry or hermitianness, eventual operands transposition, scalar factors, multiple right hand sides, and arbitrary vectors stride. Gearing a Sparse BLAS implementation to be efficient in all these cases is challenging.

Our experience specific to the RSB (Recursive Sparse Blocks) data structure suggests that attacking this problem by means of an auto-tuning based approach is possible and can be beneficial. At run time and with fairly simple heuristic rules we rearrange a sparse matrix instance to perform faster a given operation with specified operands. In case of success, subsequent invocations of the same operation may perform faster, leading to amortization of the auto-tuning time first and overall time/resources saving afterwards.

This poster will first sketch how an RSB structure can be rearranged and thereby 'tuned' via an automated procedure. Then, a quantitative overview of our techniques effectiveness on several matrices and variants of the sparse matrix-matrix operation will be given. To better highlight the strong and weak points of our approach, we will comment on the tuning time amortization and compare our performance to that of the proprietary, highly tuned 'Intel Math Kernels Library' CSR implementation.

Auto-tuning support can be added to the Sparse BLAS standard by means of minimal, compatible extensions, and independently of the sparse matrix format. The techniques we present here pertain the soon to be released librsb-1.2.

Dot-bound and dispersive states in disordered graphene quantum dot superlattices: a kernel polynomial method based approach

ANDREAS PIEPER[1], H. FEHSKE, AND G. WELLEIN

[1] *University Greifswald, Felix-Hausdorff-Str. 6, 17487 Greifswald*

pieper@physik.uni-greifswald.de

ABSTRACT

We use numerically exact Chebyshev expansion and kernel polynomial methods to study transport through circular gate-defined graphene quantum dots in the framework of a tight-binding honeycomb lattice model. Our focus lies on the regime where individual modes of the electrostatically defined dot dominate the charge carrier dynamics. In particular, we monitor the scattering of an injected Dirac electron wave packet for a single quantum dot, electron confinement in the dot, and the propagation of an electronic excitation along a linear array of dots. Moreover we consider a square lattice configuration of dots and calculate local density of states and the momentum resolved photoemission spectrum. Again we find clear evidence for a series of quasibound states at the dots. We further investigate the interplay of the superlattice structure with dot-localized modes on the electron energy dispersion. Effects of disordered dot lattices are discussed too.

Coarse-Grained ILU Methods for Aggressive Smoothing in Stokes Preconditioners

PATRICK SANAN

Facoltà di scienze informatiche, Università della Svizzera italiana
Via Giuseppe Buffi 13, CH-6904 Lugano

sananp@usi.ch

ABSTRACT

Chow and Patel recently introduced a fine-grained parallel ILU factorization technique [1]. We use the package pTatin3d [2], built on top of PETSc [3], along with the ViennaCL library [4], to investigate applications and extensions of this technique to precondition large linear systems arising from discretization of nonlinear Stokes flow describing mantle convection. These problems are challenging to precondition due to large viscosity contrasts both over the computational domain and across sharp jumps with potentially complex geometric structure. "Heavy Smoothing" approaches are promising to maintain optimal scaling while mitigating the communication bottlenecks of standard multigrid methods. These employ more aggressive coarsening and smoothing, leveraging available accelerators with subdomain solvers using local preconditioners, including ILU. The novel factorization technique above is particularly valuable in this context due to its "warm start" capabilities; preconditioners can be cheaply updated across time steps in a simulation, and sparsity patterns can be adapted to provide stronger preconditioners in problematic physical subregions. We investigate the direct application of the new ILU algorithm to precondition Stokes sub-systems on the hybrid CSCS Piz Daint machine, assess the effectiveness of reusing factorizations, and discuss potential extensions of the methods to provide more effective multigrid smoothers.

References

[1] http://www.cc.gatech.edu/ echow/pubs/parilu-paper.pdf

[2] http://dl.acm.org/citation.cfm?id=2683624

[3] http://www.mcs.anl.gov/petsc/

[4] http://viennacl.sourceforge.net

Optimised collective communication operations for latency-bound applications

MICHAEL SCHLIEPHAKE

KTH Royal Institute of Technology
Teknikringen 14, SE-100 44 Stockholm, Sweden

michs@kth.se

ABSTRACT

There is a need to exchange short messages in many MPI applications, for example PDE solvers, simulations in neuroinformatics and programs using adaptive mesh refinements. The typical performance of collective operations for short messages is one the limiting factors for these applications that make it hard to achieve good scaling in large-scale simulations.

We developed optimisations for collective communiation operations in the context of the EU FP7 exascale project CRESTA. Starting point was the analysis of the superior communication efficiency in the application NEK500, a solver for incompressible Navier-Stokes equations that scales to more than one million processes. NEK5000 uses amongst other algorithms the crystal router.

We developed this algorithm further in order to use features of recent computer architectures. Our communication routines execute communication operations faster than the widely used MPI libraries for short messages up to a few kilobyte length. Their use is especially beneficial for operations with in place message exchange. We achieved here speedups with factors of several hundred times in benchmarks up to 8000 cores.

The poster will present the latest version of the algorithm and benchmarks from PDC's new Cray system "Beskow" that became available this month. The system is a Cray XC40 system with more than 50000 cores and an Aries interconnect.

References

[1] Michael Schliephake and Erwin Laure: Analysis of the Communication Performance of CRESTA's Co-Design Application NEK5000. In High Performance Computing, Networking, Storage and Analysis (SCC), 2012 SC Companion, p. 669-674.

[2] Michael Schliephake et.al.: Design, implementation and use of mampicl, the multialgorithm MPI Collective library. In Solving Software Challenges for Exascale, EASC 2013.

[3] Michael Schliephake and Erwin Laure: Performance Analysis of Irregular Collective Communication with the Crystal Router Algorithm. In Solving Software Challenges for Exascale, EASC 2014.

A fault tolerant application using GASPI communication layer

FAISAL SHAHZAD[1], MORITZ KREUTZER, MARKUS WITTMANN, THOMAS ZEISER, RUI MACHADO, ANDREAS PIEPER, GEORG HAGER, AND GERHARD WELLEIN

[1] *Erlangen Regional Computing Center (RRZE)*
Martensstr. 1, 91058 Erlangen, Germany

faisal.shahzad@fau.de

ABSTRACT

It is commonly agreed that highly parallel software on Exascale computers will suffer from many more runtime failures due to the decreasing trend in the mean time between failures (MTBF). This motivates substantial research in the area of fault tolerance (FT) and fault mitigation. Applica- tions should survive a failure and/or be able to recover with least effort. The MPI standard still does not contain any FT interface definition. The promising MPI User-Level Failure Mitigation (ULFM) activities are still at the proposal stage. As an alternative programming approach we use GASPI, a communication library based on the PGAS model that supports the design of fault tolerant applications. Integrating checkpointing techniques and low cost health checks into existing applications we demonstrate on the fly application recovery using pre-started spare processes. We choose the Lanczos algorithm as a prototype for iterative sparse eigenvalue solvers to show- case the implementation of such fault tolerance techniques. As a first step, the optimized check- point/restart functions are implemented. Dumping of the checkpoints is done via a separate library thread/process, which writes the checkpoints asynchronously while the application continues. This way, the application runtime overhead for checkpointing can be substantially reduced. Moreover, the library is able to store the checkpoints on a parallel file systems or on neighboring nodes' local storage. The underlying sparse matrix vector operation library is developed using GASPI and supplemented with fault tolerance features. This includes checkpoint functions relating to the communication infrastructure for the matrix algebraic operations. The recurring application checkpoints are algorithm aware, so that only the restart-relevant data is stored, i.e., only modified data is checkpointed. In a second step, health-check functions are implemented which depend on the return value of the timeout-based communication functions provided by GASPI. In case of a communication error, the health of the communicating partner is inquired. After the correct detection of any dead process, a globally consistent view of the healthy processes is formed. The application then recovers with help of pre-allocated spare processes. The poster presents implementation details about the fault tolerance method and its optimiza- tion techniques. We have thoroughly analyzed the overheads involved in the recovery mechanism and their dependence on the application parameters. In our observation, this recovery technique can enable an application to recover from the process failures at a reasonably low cost. We identify several possible improvements of the scheme.

Density-matrix renormalization group study of the metal-insulator transition in the Hubbard model on the triangular lattice

TOMONORI SHIRAKAWA

Computational Condensed Matter Physics Laboratory, RIKEN
Hirosawa 2-1, Wako, Saitama, 351-0198, Japan

t-shirakawa@riken.jp

ABSTRACT

Performing large-scale density-matrix renormalizaiton group calculations, we examine the ground state properties of the Hubbard model on the triangular lattice at half-filling. We find that the double occupancy exhibits two discontinuous changes with increasing the Hubbard interaction U, at $U_{c1} \sim 7.0t$ and $U_{c2} \sim 10.0t$, which indicates that there are three phases in this system. We also calculate the nearest neighbor bond order correlations and spin-spin correlations. Although these quantities generally oscillate due to the open boundary conditions, we find that the oscillation is strongly suppressed in the intermediate region ($U_{c1} < U < U_{c2}$). This suggests that neither bond ordering nor valence bond solid order is stable. We argue that the possible state in the intermediate region is a resonating valence bonds state.

High Performance Block Jacobi-Davidson

JONAS THIES

DLR, Simulations- und Softwaretechnik, Linder Höhe, 51147 Köln

jonas.thies@dlr.de

ABSTRACT

We present details of a new implementation of the Jacobi-Davidson method for computing exterior eigenvalues. The poster presents both algorithmic and implementation aspects of how to get the optimal performance and scalability out of a block variant of this popular algorithm. On the one hand, the basic building blocks need to be tuned to the needs of the algorithm, but on the other, the algorithm needs to be adapted to the specific properties of the machine, such as SIMD width or heterogeneity. The algorithm formulation is presented, along with details on the performance engineering for the key kernels (block orthogonalization, projection and sparse matrix-multiple-vector multiplication) Results are shown for a variety of modern architectures (CPUs, GPUs, hybrid nodes etc.) and a number of applications from theoretical physics and other areas.