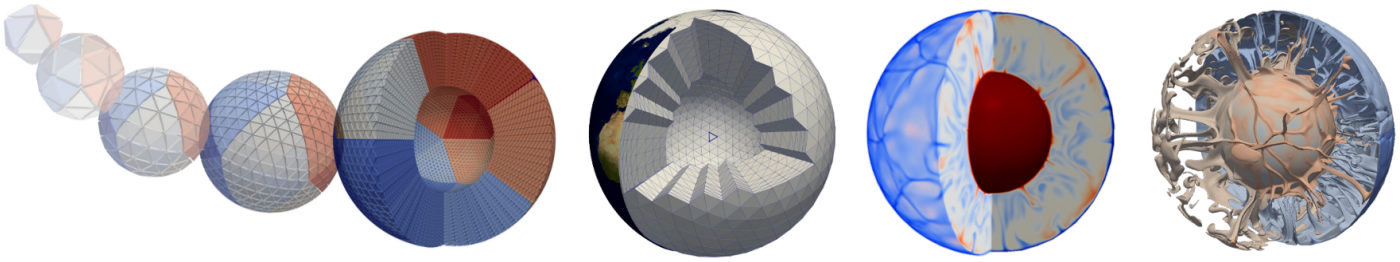

PARNUM 2017

International Workshop on Parallel Numerics
April 19 – 21, 2017
Fraunhofer Research Campus
Waischenfeld, Germany

BOOK OF ABSTRACTS

FRIEDRICH-ALEXANDER UNIVERSITÄT ERLANGEN-NÜRNBERG
UNIVERSITÄT SALZBURG
DFG Deutsche Forschungsgemeinschaft
SPPEXA

Program Committee

Hans-Joachim Bungartz, Technical University of Munich, Germany

Dietmar Fey, University of Erlangen-Nürnberg, Germany

Gundolf Haase, Karl-Franzens-University, Graz, Austria

Axel Klawonn, University of Cologne, Germany

Günter Leugering, University of Erlangen-Nürnberg, Germany

Miriam Mehl, University of Stuttgart, Germany

Gabriel Okša, Slovak Academy of Sciences, Bratislava, Slovakia

Roman Trobec, Jožef Stefan Institute, Ljubljana, Slovenia

Pavel Tvrdik, Czech Technical University, Prague, Czech Republic

Roman Wyrzykowski, Czestochowa University of Technology, Poland

Program Chairmen

Ulrich Rüde, University of Erlangen-Nürnberg, Germany

Marián Vajteršic, University of Salzburg, Austria &
Slovak Academy of Sciences, Bratislava, Slovakia

Organizing Committee

Dominik Bartuschat, University of Erlangen-Nürnberg, Germany

Julia Deserno, University of Erlangen-Nürnberg, Germany

Editors: Dominik Bartuschat, Ulrich Rüde, and Marián Vajteršic

April 2017, Erlangen, Germany

Speaker Index

DIRECT SOLUTION OF SPARSE LINEAR EQUATIONS ON PARALLEL COMPUTERS

Iain S. Duff
STFC Rutherford Appleton Laboratory,
Harwell Campus, Didcot, Oxfordshire, UK
and
CERFACS, Météo-France, Toulouse, France.
iain.duff@stfc.ac.uk

As part of the H2020 FET-HPC Project NLAFET, we are studying the scalability of algorithms and software for using direct methods for solving large sparse equations.

We briefly discuss the structure of NLAFET and the scope of the Project. We then discuss the algorithmic approaches for solving sparse systems: positive definite, symmetric indefinite, and unsymmetric. An important aspect of most of our algorithms is that although we are solving sparse equations most of the kernels are for dense linear algebra. We show why this is the case with a simple example before illustrating the various levels of parallelism available in the sparse case. We examine the benefits of using standard run time systems to assist us in developing codes for extreme scale computers.

For sparse matrices that are very unsymmetric inasmuch as their structure is quite different from the structure of $|A|+|A^T|$, we use sparse data structures. We discuss the design, coding, and performance of software for this case, including the development of a parallel threshold-Markowitz algorithm.

We illustrate our talk with runs of prototype codes that will be developed for inclusion in a library being developed in the context of the NLAFET Project.

The work described in this talk has been conducted by the STFC NLAFET Team who comprise: Florent Lopez, Stojce Nakov, and Vedran Novakovic.

CODE GENERATION FOR A HIGH ORDER ADER-DG SOLVER IN A HYPERBOLIC PDE ENGINE

Jean-Matthieu Gallard
Departments of Informatics, Technical University of Munich,
Boltzmannstr. 3, 85748 Garching, Germany
jean-matthieu.gallard@tum.de

In this talk the use of code generation to improve the performance and the energy efficiency of the solver engine ExaHyPE is discussed. ExaHyPE is an Horizon 2020 EU project to develop a high-performance engine to solve hyperbolic systems of PDEs using the high-order discontinuous Galerkin finite element method [1]. The engine will be flexible to support various applications and will be tailored towards expected exascale architectures. One of the main goals of the project is therefore to provide to the end-user an abstraction of the complicated algorithms to implement the ADER-DG numerical scheme and of the issues related to scalability and parallel adaptive mesh refinement, which are handled internally by the Peano framework [2].

Code generation within the engine produces optimized code that is tailored to the specific PDE problem, to the chosen polynomial order for the ADER-DG scheme, and especially to the compute architecture used. Compute kernels for the ADER-DG scheme exploit the high performance LIBXSMM library [3] for small matrix multiplications occurring in the element-local kernels, use tailored data layouts and support compiler auto-vectorization. The generated optimized kernels offer a speedup of a factor 2.5 when compared to a generic C++ implementation, and are currently benchmarked and improved in regards to performance and energy consumption. First results will be presented for benchmark scenarios in seismology and astrophysics.

References

[1] O. Zanotti, F. Fambri, M. Dumbser, and A. Hidalgo. Space–time adaptive ADER discontinuous Galerkin finite element schemes with a posteriori sub-cell finite volume limiting. *Comput. Fluids*, 118:204–224, 2015. DOI: 10.1016/j.compfluid.2015.06.020.

[2] T. Weinzierl. The Peano software - parallel, automaton-based, dynamically adaptive grid traversals. *CoRR*, abs/1506.04496, 2015. URL: http://arxiv.org/abs/1506.04496.

[3] A. Heinecke, G. Henry, M. Hutchinson, and H. Pabst. LIBXSMM: Accelerating Small Matrix Multiplications by Runtime Code Generation. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, SC '16. IEEE Press, 2016, pages 84:1–84:11. URL: http://dl.acm.org/citation.cfm?id=3014904.3015017.

PARALLELIZATION IN THE TIME DIMENSION OF GEOPHYSICAL DATA ASSIMILATION PROBLEMS

Mike Fisher
European Centre for Medium-Range Weather Forecasts, Reading, UK

Serge Gratton
INP-ENSEEIHT, Toulouse, France
serge.gratton@enseeiht.fr

Selime Gürol
CERFACS, Météo-France, Toulouse, France
gurol@cerfacs.fr

In this talk we will address the numerical solution of the saddle point system arising from four dimensional variational (4D-Var) data assimilation, including a study of preconditioning. This new saddle point formulation [1] of 4D-Var allows parallelization in time dimension. Therefore, it represents a crucial step towards higher computational efficiency, since 4D-Var approaches otherwise require many sequential computations.

In recent years, there has been increasing interest in saddle point problems which arise in many other applications such as constrained optimization, computational fluid dynamics, optimal control and so forth. The key issue of solving saddle point systems with Krylov subspace methods is to find efficient preconditioners. This talk will focus on the new low-rank limited memory preconditioners [2] exploiting the particular structure of the problem. Numerical experiments performed within the Object Oriented Prediction System (OOPS) are presented.

References

[1] M. Fisher and S. Gürol. Parallelization in the time dimension of four-dimensional variational data assimilation. *Q. J. R. Meteorol. Soc.*, 143(703):1136–1147, 2017. DOI: 10.1002/qj.2997.

[2] M. Fisher, S. Gratton, S. Gürol, Y. Trémolet, and X. Vasseur. Low rank updates in preconditioning the saddle point systems arising from data assimilation problems. *Optim. Meth. Softw*:1–25, 2016. DOI: 10.1080/10556788.2016.1264398.

EXPLORING PROGRAMMING MODELS FOR ACCELERATING SCIENTIFIC APPLICATIONS ON HYBRID CPU-MIC PLATFORMS

Roman Wyrzykowski, Lukasz Szustak, Kamil Halbiniak
Czestochowa University of Technology, Poland
roman@icis.pcz.pl

Modern heterogeneous computing platforms have become powerful HPC solutions, which could be applied to a wide range of real-life applications. In particular, the hybrid platforms equipped with Intel Xeon Phi coprocessors offer the advantages of massively parallel computing, while supporting practically the same (or similar) parallel programming model as conventional homogeneous solutions. However, there is still an open issue as to how scientific applications can efficiently utilize hybrid platforms with Intel MIC coprocessors.

In paper [1], we proposed a method for porting a real-life scientific application to computing platforms with Intel MICs. We focused on the parallel implementation of a numerical model of alloy solidification, which is based on the generalized finite difference method. We developed a sequence of steps that are necessary for porting this application to platforms with accelerators, assuming no significant modifications of the code. The proposed method considers not only overlapping computations with data movements, but also takes into account an adequate utilization of cores/threads and vector units. Using parallel resources of one Intel Xeon Phi coprocessor (KNC architecture), the developed approach allowed us to execute the whole application 3.45 times faster than the original parallel version running on two CPUs.

In this work, we focus on studying various heterogeneous programming models for accelerating the solidification application on hybrid CPU-MIC platforms. We focus on two models: OpenMP 4.0 Accelerator Model and Hetero Streams Library (hStreams in short) [2]. Now the main challenge for achieving a desired high performance of computations is to take advantage of CPUs and coprocessors to work together, when all the available threads of CPUs and Intel MICs are utilized coherently to solve the modelling problem.

In the paper, we present the performance comparison of the above-mentioned models for various configurations of computing resources. In particular, using the hStreams library, our approach allows us parallelize efficiently the solidification application on hybrid platforms with two CPUs and two MICs, and accelerate computations about 10.5 times in comparison with the basic version for two CPUs. We also conclude that while OpenMP provides an unified directive-based programming model, the current stable version of this standard is not efficient in multi-device heterogeneous platforms. That is why, we plan to investigate new features available in version 4.5 of OpenMP, such as asynchronous offload mechanism.

References

[1] L. Szustak et al. Toward parallel modeling of solidification based on the generalized finite difference method using Intel Xeon Phi. *PPAM 2015, Part I. LNCS*, 9573:411–412, 2016.

[2] Ch. J. Newburn et al. Heterogeneous streaming. *IPDPSW, AsHES*, 2016.

ASYNCHRONOUS INTEGRATION OF CUDA/OPENCL WITHIN HPX FOR UTILIZING FULL CLUSTER CAPABILITIES

P. Diehl[1,5], T. Heller[2,5], L. Troska[1], H. Kaiser[3,5], D. Fey[2], M. A. Schweitzer[1,4]

[1] *Institute for Numerical Simulation, University of Bonn*
[2] *Computer Science 3, Computer Architectures, Friedrich-Alexander-University*
[3] *Center for Computation and Technology, Louisiana State University*
[4] *Meshfree Multiscale Methods, Fraunhofer SCAI, Schloss Birlinghoven*
[5] *The STE||AR Group (`http://stellar-group.org`)*

Experience shows that on today's high performance systems the utilization of different acceleration cards in conjunction with a high utilization of all other parts of the system is difficult. Future architectures, like exascale style clusters, are expected to aggravate this issue as the number of cores are expected to increase and memory hierarchies are expected to become deeper. One big aspect for distributed applications is to guarantee high utilization of all available resources, including local or remote acceleration cards on a cluster while fully using all the available CPU resources and the integration of the GPU work into the overall programming model. For the integration of CUDA and OpenCL code we extended HPX [1, 2], a general purpose C++ run time system for parallel and distributed applications of any scale, and enabled asynchronous data transfers from and to the GPU device and the asynchronous invocation of CUDA- and OpenCL kernels on this data. Both operations are well integrated into the general programming model of HPX which allows to seamlessly overlap any GPU operation with work on the main cores. Any user defined CUDA or OpenCL kernel can be launched on any (local or remote) GPU device available to the distributed application. We present asynchronous implementations for the data transfers and kernel launches for CUDA and OpenCL code as part of a HPX asynchronous execution graph. Using this approach we can combine all remotely and locally available acceleration cards on a cluster to utilize its full performance capabilities. Benchmarks show, that the integration of the asynchronous operations (data transfer + launches of the kernels) as part of the HPX execution graph imposes no additional computational overhead and significantly eases orchestrating coordinated and concurrent work on the main cores and the used GPU devices.

References

[1] H. Kaiser, T. Heller, D. Bourgeois, and D. Fey. Higher-level Parallelization for Local and Distributed Asynchronous Task-based Programming. In *Proceedings of the First International Workshop on Extreme Scale Programming Models and Middleware*, ESPM '15. ACM, 2015, pages 29–37. DOI: 10.1145/2832241.2832244.

[2] H. Kaiser, B. Adelstein-Lelbach, T. Heller, A. Bergé, and J. Biddiscombe. hpx: HPX V0.9.99: A general purpose C++ runtime system for parallel and distributed applications of any scale. 2016. DOI: 10.5281/zenodo.58027.

TRANSPARENT EXECUTION OF NUMERICAL LIBRARIES ON DISTRIBUTED HPC PLATFORMS

Michael Hofmann, Florian Polster, Riko Streller, Daniel Walther
Department of Computer Science, Chemnitz University of Technology
mhofma@cs.tu-chemnitz.de

The usage of numerical software libraries is well established in scientific computing as they can provide advanced methods and efficient implementations for solving common problems of scientific applications. One major goal of using these libraries is to improve the application performance by fully exploiting the available computational hardware. For example, BLAS libraries, such as OpenBLAS or cuBLAS provide efficient implementations of linear algebra operations that exploit modern multicore processors or graphics processing units. In this contribution, we propose a method for redirecting the execution of an existing numerical software library to a distributed HPC platform. The redirection is transparent in the sense that the application does not have to distinguish whether the utilized library functions are executed locally or distributed. Thus, the method allows to exploit the computational power of HPC platforms even in non-parallel application codes. Our proposed solution provides replacements of the utilized library functions, that can be used without additional programming efforts for adapting the application code. Furthermore, by providing the replacement functions as a shared library, the redirection can also be applied to applications that are only available as a binary executable. We demonstrate the approach for several numerical software libraries, such as BLAS and LAPACK libraries for linear algebra operations, the FFTW library for fast Fourier transforms, and the ScaFaCoS library for fast Coulomb interactions in particle systems. This includes sequential libraries as well as parallel libraries based on multi-threading or MPI. The implementation utilizes the Simulation Component and Data Coupling library [1] for performing the program interactions and the data transfers between the locally executed application and the numerical software library executed on a distributed HPC platform. Experimental results are presented to investigate the overhead of the required data transfers and the achieved performance improvements.

Acknowledgments This work was performed within the Federal Cluster of Excellence EXC 1075 "MERGE Technologies for Multifunctional Lightweight Structures" and supported by the German Research Foundation (DFG). Financial support is gratefully acknowledged.

References

[1] M. Hofmann and G. Rünger. Sustainability through flexibility: Building complex simulation programs for distributed computing systems. *Simul. Model. Pract. Theory*, 58(1):65–78, 2015. Special Issue on Techniques And Applications For Sustainable Ultrascale Computing Systems. DOI: 10.1016/j.simpat.2015.05.007.

PARALLEL SOLUTION OF TRIDIAGONAL MATRICES

Thomas Huckle
Fakultät für Informatik, TU München, Boltzmannstr. 3, D-85748 Garching
huckle@in.tum.de

For solving sparse linear systems of equations iteratively, an efficient preconditioner is necessary, that also allows fast solution in parallel. So ILU is a preconditioner that can be derived efficiently in parallel [1]. For eigenvalue algorithms like MRRR also twisted factorizations of tridiagonal matrices are used. But here the factorizations have to be computed to high accuracy. In these applications the convergence of the fixed point iteration of Chow for computing ILU takes too many iterations. Therefore, in this talk we will discuss fast iterative methods for computing (twisted) factorizations and for efficiently solving linear systems with (twisted) bidiagonal matrices in parallel.

References

[1] E. Chow and A. Patel. Fine-grained parallel incomplete LU factorization. *SIAM J. Sci. Comput.*, 37(2):C169–C193, 2015. DOI: 10.1137/140968896.

HPC AS A SERVICE FOR COMPUTATIONAL FLUID DYNAMICS PROBLEMS

Fayssal Benkhaldoun
LAGA - Institut Galilée, 99 avenue Jean Baptiste Clément, 93430 Villetaneuse
fayssal@math.univ-paris13.fr

Christophe Cérin
LIPN - Institut Galilée, 99 avenue Jean Baptiste Clément, 93430 Villetaneuse
christophe.cerin@lipn.univ-paris13.fr

Imad Kissami
LAGA & LIPN -Institut Galilée, 99 avenue Jean Baptiste Clément, 93430 Villetaneuse
imad@lipn.univ-paris13.fr

Walid Saad
LIPN & University of Tunis, LATICE, ENSIT,
5 av. Taha Hussein, B.P. 56, Bab Mnara, Tunis, Tunisia
walid.saad@lipn.univ-paris13.fr

Abstract

In this paper we decline the full cycle for transforming a parallel code for a Computational Fluid Dynamics (CFD) problem into a parallel version for the RedisDG workflow engine. This system is able to capture heterogeneous and highly dynamic environments, thanks to opportunistic scheduling strategies. It also captures multi-criteria approaches to decide the allocation of tasks to machines. We show how to move to the field of 'HPC as a Service' in order to use heterogeneous platforms and to also investigate other performance metrics than the makespan (the minimum completion time). We also provide an experimental evaluation of the implemented solution where we discuss of the accuracy of the multi-criteria approach. This paper states that new models for High Performance Computing are possible, under the condition we revisit our mind in the direction of the potential of new paradigms such as cloud, edge computing... New challenges are to aggregate resources, from anywhere, at any time under Service Level Agreements (SLAs) constraints.

Introduction

Our research focuses on the design of Systems for heterogeneous and highly dynamic environments, notably clouds, desktop grids and volunteer computing projects. The overall objective is to execute computational codes in such environments... and progressively moving from a traditional view for High Performance Computing (HPC) to service oriented and workflow oriented views. The context we consider is also of particular interest for the development of extreme edge and edge cloud. A hard question, that the dynamicity causes here, is that given a workflow to schedule, we do not have any a-priori knowledge on the resources that

are available. To address it, we propose to implement a Publish-Subscribe based mechanism for resource discovery and allocation. The mechanism is implemented in a prior system we developed for mimic desktop grid environments: the RedisDG system.

Recall that the Publish-Subscribe paradigm is an asynchronous mode for communicating between entities. Some users, namely the subscribers or clients or consumers, express and record their interests under the form of subscriptions, and are notified later by another event produced by other users, namely the producers.

This communication mode is multi-point, anonymous and implicit. Thus, it allows spatial decoupling (the interacting entities do not know each other), and time decoupling (the interacting entities do not need to participate at the same time). This total decoupling between the production and the consumption of services increases the scalability by eliminating many sorts of explicit dependencies between participating entities. Eliminating dependencies reduces the coordination needs and consequently the synchronizations between entities. These advantages make the communicating infrastructure well suited to the management of distributed systems and simplify the development of a middleware for the coordination of components in our workflow engine context or for applications running in different domains and communicating through middleware solutions deployed on clouds.

Indeed, we support the thesis that for building systems for heterogeneous and highly dynamic environments we need to be compliant with:

1. a publish-subscribe layer for the orchestration of the components of the system;
2. a set of opportunistic strategies for allocating work/tasks that are also based on the publish-subscribe layer;
3. a small number of software dependencies for the system and the ability to deploy the system and its applications on demand. This point is of particular interest in this paper and we promote the 'easy to use', and systems that can be deployed without a system administrator.

In this paper we propose a solution for item 2 that is implemented into the RedisDG system. We consider a CFD problem and we execute our CFD solution, obtained from transforming a parallel code into a workflow, on top of the RedisDG system. We conduct experiments to validate our approach.

The organisation of the presentation for the workshop is as follows. First we introduce the numerical problem we are faced to. We also summarize some related works. Second we introduce a parallel solution of the problem in the spirit of MPI programming. Third, we explain how to provide with a workflow oriented view for solving the numerical problem. Then, we explain our new strategy for allocating tasks of the workflow and we show experimental results to demonstrate the potential of this strategy. Experiments are conducted on six geographically distributed clusters in the Grid'5000 testbed. Finally, we conclude the presentation.

DESIGN OF CACHE-EFFICIENT MULTITHREADED SPARSE MATRIX FORMAT FOR MODERN ERA

Ivan Kotenkov, Ivan Šimeček, Daniel Langr
Faculty of Information Technology, Czech Technical University in Prague, Czech Republic
koteniv1@fit.cvut.cz, xsimecek@fit.cvut.cz, daniel.langr@fit.cvut.cz

The most common routines in numerical linear algebra are sparse matrix-vector multiplication and transposed sparse matrix-vector multiplication. In the further text, we denote these operations as *sparse multiplication*. These operations are crucial (and the most time-consuming) in iterative solvers of sparse systems of linear equations. In these applications, large number of sparse multiplications with the same matrix A is executed.

Matrices emerging in HPC applications running on these systems need to be mapped to nodes such that each node contains in its memory some portion of matrix nonzero elements. The overall performance and scalability depend strongly on matrix partitioning, matrix-to-node mapping, and the used matrix storage format. Many storage formats for sparse matrices have been developed. Since the commonly used storage formats (like COO or CSR) are not sufficient for high-performance computations (they were introduced more than 45 years ago), extensive research has been conducted about maximal computational efficiency of these routines [1].

We present a new storage format and related algorithms for sparse multiplication that is designed for efficient parallel (multithreaded) execution and has the following features:

- It is a hierarchical format (inspired by ELL idea) that has a good spatial and temporal locality.
- For highly parallel architectures, some of the challenges of the sparse multiplication are thread divergence and load imbalance when operating on matrices with high standard deviation in the number of non-zero elements per row [2]. We address these problems by grouping rows in bins in a modification of the ELL format.
- Using target-specific versions of algorithms we achieve a very good overall performance.

References

[1] I. Šimeček and D. Langr. Space and execution efficient formats for modern processor architectures. In *2015 17th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing (SYNASC)*, 2015, pages 98–105. DOI: 10.1109/SYNASC.2015.24.

[2] A. Ashari, N. Sedaghati, J. Eisenlohr, S. Parthasarath, and P. Sadayappan. Fast Sparse Matrix-Vector Multiplication on GPUs for Graph Applications. *SC14: International Conference for High Performance Computing, Networking, Storage and Analysis*, 2014. DOI: 10.1109/sc.2014.69.

CONVERGENCE AND PARALLELIZATION OF NONNEGATIVE MATRIX FACTORIZATION WITH NEWTON ITERATION

Rade Kutil[1], Markus Flatz[1], Marián Vajteršic[1,2]
Department of Computer Sciences, University of Salzburg, Austria[1]
Department of Informatics, Slovak Academy of Sciences, Bratislava[2]
rkutil@cosy.sbg.ac.at

The goal of Nonnegative Matrix Factorization (NMF) is to represent a large nonnegative matrix A in an approximate way as a product WH of two significantly smaller nonnegative matrices. Applications of the NMF include text mining, document classification, clustering, spectral data analysis, face recognition, and computational biology.

Among several algorithms to calculate the NMF, such as the multiplicative update algorithm (MU) [1], there are Newton-type methods [2]. In [3], we proved that Newton methods can be parallelized very well because Newton iterations can be performed in parallel without exchanging data between processes. Therefore, they have an advantage on parallel architectures over other methods. However, these methods can show problematic convergence behavior, limiting their efficiency.

We present a modified algorithm that has stable convergence. Like all algorithms, it minimizes the approximation error by alternatingly improving W while holding H fixed, and vice versa. Newton iteration is applied in each step. There are several differences to existing algorithms. While [2] uses unconstrained optimization and active set methods, our method uses Karush-Kuhn-Tucker (KKT) conditions and a reflective technique. While [3] uses backtracking line search in order not to violate KKT conditions, we use this technique only to guarantee global convergence, i.e. that the approximation error decreases in each iteration. The KKT conditions are enforced by using a modified target function with the same zeros.

Our method allows for an inexact approach, where only few Newton iterations are performed per outer iteration. Experiments show that this leads to faster convergence. Although, on the other hand, this increases the communication overhead in the parallel implementation, a single Newton iteration is still the best choice and parallel efficiency is satisfactory.

References

[1] D. D. Lee and H. S. Seung. Learning the parts of objects by non-negative matrix factorization. *Nature*, 401:788–791, 1999.

[2] D. Kim, S. Sra, and I. S. Dhillon. Fast Newton-type methods for the least squares nonnegative matrix approximation problem. In *Proc. of the 2007 SIAM Int. Conf. on Data Min. (SDM07)*, 2007, pages 343–354.

[3] M. Flatz and M. Vajteršic. Parallel Nonnegative Matrix Factorization via Newton Iteration. *Parallel Process. Lett.*, 26(03):1650014, 2016. DOI: 10.1142/S0129626416500146.

A FRAMEWORK FOR NONLINEAR FETI-DP AND BDDC METHODS

Axel Klawonn, Martin Lanser, Matthias Uran
Mathematisches Institut
Universität zu Köln
{axel.klawonn,martin.lanser,m.uran}@uni-koeln.de

Oliver Rheinbach
Institut für Numerische Mathematik und Optimierung
Technische Universität Bergakademie Freiberg
oliver.rheinbach@math.tu-freiberg.de

Highly scalable and robust Newton-Krylov domain decomposition approaches are widely used for the solution of nonlinear implicit problems, e.g., in structural mechanics. In general, in those methods, the nonlinear problem is first linearized and afterwards decomposed into subdomains. By changing this order, i.e., by first decomposing the nonlinear problem, new parallel and nonlinear domain decomposition methods can be designed which can reduce communication by increasing local work. These methods show often a higher robustness than classical Newton-Krylov variants and can be interpreted as nonlinear globalization strategies. In this talk, we discuss different Nonlinear-FETI-DP and BDDC approaches [1, 2, 3], which can be formulated in a common framework and also be interpreted as nonlinear right-preconditioners [4]. We also present weak scaling results to mare than 200K BlueGene/Q cores on JUQUEEN at FZ Jülich.

References

[1] A. Klawonn, M. Lanser, and O. Rheinbach. Nonlinear FETI-DP and BDDC Methods. *SIAM J. Sci. Comput.*, 36(2):A737–A765, 2014. DOI: 10.1137/130920563.

[2] A. Klawonn, M. Lanser, and O. Rheinbach. Toward Extremely Scalable Nonlinear Domain Decomposition Methods for Elliptic Partial Differential Equations. *SIAM J. Sci. Comput.*, 37(6):C667–C696, 2015. DOI: 10.1137/140997907.

[3] A. Klawonn, M. Lanser, O. Rheinbach, and M. Uran. New Nonlinear FETI-DP Methods Based on a Partial Nonlinear Elimination of Variables. In *Proceedings of the 23rd International Conference on Domain Decomposition*. Accepted for Publication. Lect. Notes Comput. Sci. Eng., 2016.

[4] A. Klawonn, M. Lanser, O. Rheinbach, and M. Uran. Nonlinear FETI-DP and BDDC Methods: A Unified Framework and Parallel Results. *SIAM J. Sci. Comput.* Submitted for publication, 2016.

EXPLICIT VECTORIZATION AS A DESIGN TOOL FOR PARALLEL ALGORITHMS ON MODERN HARDWARE ARCHITECTURES

Manfred Liebmann
Institute for Mathematics and Scientific Computing, University of Graz, Austria
manfred.liebmann@uni-graz.at

Modern hardware architectures provide a formidable challenge to the design of algorithms with portable performance across different flavors of multicore CPUs, manycore accelerators, and graphics processors. Three different case studies: small eigenvalue problems in magnetic resonance imaging [1], algorithms for semiclassical quantum dynamics [2, 3], and algebraic multigrid methods for uncertainty quantification [4], show the applicability of explicit vectorization techniques a as general design tool for massively parallel software.

References

[1] M. Presenhuber. Numerische Methoden zur Nullstellenbestimmung für Anwendungen in der quantitativen Magnetresonanztomographie. Diplomarbeit. Karl-Franzens-Universität Graz, 2015. URL: http://resolver.obvsg.at/urn:nbn:at:at-ubg:1-86584.

[2] D. Sattlegger M. Liebmann. Parallel algorithms for semiclassical quantum dynamics. Preprint No. IGDK-2015-25. 2015. URL: http://igdk.eu/foswiki/pub/IGDK1754/Preprints/LiebmannSattlegger_2015.pdf.

[3] D. Sattlegger. Efficient Algorithms for Semiclassical Quantum Dynamics. Dissertation. München: Technische Universität München, 2015. URL: http://nbn-resolving.de/urn/resolver.pl?urn:nbn:de:bvb:91-diss-20151221-1277913-1-7.

[4] D. Schaden. Efficient Parallel PDE-Solvers for Uncertainty Quantification. Master Thesis. Technische Universität München, 2016.

PARALLEL MULTI-DENSITY BASED CLUSTERING

Lukáš Csóka, Peter Laurinec, Mária Lucká
Faculty of Informatics and Information Technologies, Slovak University of Technology in Bratislava, Ilkovičova 2, 842 16 Bratislava, Slovak Republic
xcsokal@stuba.sk, peter.laurinec@stuba.sk, maria.lucka@stuba.sk

Data clustering is a process of joining similar objects into groups. Although many clustering algorithms are known, it is still a challenging research area because of increasing amount of data, requesting thus parallel instead of a sequential approach. In this work, we modify the density based algorithm DBSCAN [1] that achieves excellent clustering results for datasets with equal density, but in general very bad results when applied to data with various densities. Another problem of DBSCAN is parallelization because of its strongly sequential character.

The modification DBSCAN-DLP [2] solves the problem of various densities, but it is still strongly sequential and unusable for larger dataset clustering. A successful parallel version of DBSCAN based on the disjoint-set data structure is presented in [3].

In our work, we combine and modify these two approaches for clustering large data sets with various densities in parallel. Proposed algorithm aims to find multiple Density Level Sets (DLS) based on a statistical analysis of data. To each DLS few data points from the dataset are assigned, but not all points belong to any DLS in general. The standard parameter of DBSCAN ϵ, characterizing the distance of neighborhood points, is then computed for every DLS. The algorithm continues with DBSCAN clustering on the data points that were already assigned to any DLS. After, an expansion of clusters to not assigned points is performed.

The method is compared with the well-known K-Means, standard DBSCAN and DBSCAN-DLP methods on artificial datasets with various densities, where it achieves better results concerning quality of classification and performance. On datasets with equal densities, it behaves similar to DBSCAN, but with better identification of outliers. We have used both OpenMP and Message Passing Interface (MPI) approaches and showed that solving this big data problem is without proper parallelization approach almost impossible.

References

[1] M. Ester, H. Kriegel, J. Sander, and X. Xu. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Second Int. Conf. on Knowledge Discovery and Data Mining*, 1996, pages 226–231.

[2] Z. Xiong, R. Chen, Y. Zhang, and X. Zhang. Multi-density dbscan algorithm based on density levels partitioning. *Int. J. Comput. Inf. Sci.*, 9(10):2739–2749, 2012.

[3] M. M. A. Patwary, D. Palsetia, A. Agrawal, W. K. Liao, F. Manne, and A. Choudhary. A new scalable parallel DBSCAN algorithm using the disjoint-set data structure. In *High Perform. Computing, Networking, Storage and Analysis (SC), 2012 Int. Conf. for*, 2012, pages 1–11.

ENERGY AWARE COMPUTATIONS ON MANYCORE SYSTEMS

Alban Lumi and Gundolf Haase
Institute for Mathematics and Scientific Computing, University of Graz, Austria
alban.lumi@uni-graz.at, gundolf.haase@uni-graz.at

Besides accuracy of the results, the overall solution time is the main quantity programmers focussing on. On the other hand the compute nodes transform electrical energy to heat which has to be dissipated afterwards. Extrapolating the recent hardware developments by ARM and Intel as well a by NVIDIA and AMD we have to scope with many cores on one chip that all have to transfer data to/from the main memory transpassing a hierarchy of caches and/or faster memory.

We will present available tools [1] to determine the power consumption when executing various application codes on different hardware as a conventional CPU-Cluster, the Intel's Knights Landing and the ThunderX by ARM. The application codes range from the simple Jacobi iteration to fully coupled cardiovascular simulations.

Choosing the Eikonal solver [2] as one special application we will demonstrate how algorithmic changes and different of memory access patterns will reduce the overall energy consumption although the overall runtime might not be reduced.

Supported by the Horizon 2020 project MontBlanc3.

References

[1] V. M. Weaver, M. Johnson, K. Kasichayanula, J. Ralph, P. Luszczek, D. Terpstra, and S. Moore. Measuring Energy and Power with PAPI. In. PASA Workshop, 2012.

[2] D. Ganellari and G. Haase. Fast many-core solvers for the Eikonal equations in cardiovascular simulations. In *2016 International Conference on High Performance Computing Simulation (HPCS)*. peer-reviewed. IEEE, 2016, pages 278–285. DOI: 10.1109/HPCSim.2016.7568347.

COMMUNITY SOFTWARE ECOSYSTEMS FOR HIGH-PERFORMANCE COMPUTATIONAL SCIENCE: OPPORTUNITIES AND CHALLENGES

Lois Curfman McInnes
Mathematics and Computer Science Division, Argonne National Laboratory, USA
mcinnes@mcs.anl.gov

Numerical libraries have proven effective in providing widely reusable software that is robust, efficient, and scalable—delivering advanced algorithms and data structures that enable scientific discovery for a broad range of applications. However, as we exploit emerging extreme-scale architectures to address more advanced modeling, simulation, and analysis, daunting challenges arise in software productivity and sustainability. Difficulties include increasing complexity of algorithms and computer science techniques required in multiscale and multiphysics applications, the imperative of portable performance in the midst of dramatic and disruptive architectural changes, the realities of large legacy code bases, and human factors arising in distributed multidisciplinary research teams. New architectures require fundamental algorithm and software refactoring, while at the same time the demand is increasing for greater reproducibility of simulation and analysis results for predictive science. This situation brings with it the unique opportunity to fundamentally change how scientific software is designed, developed, and sustained.

This presentation will introduce the Extreme-scale Scientific Software Development Kit (xSDK) [1], which defines community policies (`https://xsdk.info/policies`) to improve code quality and compatibility across independently developed packages. The vision of the xSDK is to provide infrastructure for and interoperability of a collection of related and complementary software elements—developed by diverse, independent teams throughout the community—that provide the building blocks, tools, models, processes, and related artifacts for rapid and efficient development of high-quality extreme-scale applications. The xSDK currently includes four major open-source numerical libraries (hypre, SuperLU, PETSc, and Trilinos) and two domain components (Alquimia and PFLOTRAN). The xSDK approach provides turnkey installation of member software packages and seamless combination of aggregate capabilities. We will discuss experiences in creating xSDK foundations—first steps toward realizing an extreme-scale scientific software ecosystem. We welcome contributions to the xSDK, feedback on draft xSDK community policies, and dialogue about work toward broader community ecosystems for scientific software.

References

[1] R. Bartlett, I. Demeshko, T. Gamblin, G. Hammond, M. Heroux, J. Johnson, A. Klinvex, X. Li, L. Curfman McInnes, J. D. Moulton, D. Osei-Kuffuor, J. Sarich, B. Smith, J. Willenbring, and U. Meier Yang. xSDK Foundations: Toward an Extreme-scale Scientific Software Development Kit. available via `https://arxiv.org/abs/1702.08425`, to appear in *Supercomputing Frontiers and Innovations*. 2017.

A HIGHLY SCALABLE MPI PARALLELIZATION OF THE FAST MULTIPOLE METHOD

Michael Obersteiner, Nikola Tchipev, Philipp Neumann, Hans-Joachim Bungartz

Chair of Scientific Computing, Technical University of Munich, Germany

oberstei@in.tum.de, tchipev@in.tum.de, neumanph@in.tum.de bungartz@in.tum.de

In this talk an MPI Parallelization strategy of the Fast Multipole Method (FMM) is discussed that scales up to 32k processors [1]. The implementation is based on a well-known parallelization scheme of the list-based FMM [2], which splits the octree in a local and a global part. This scheme uses local operations of type reduce to obtain the results in the global tree part and therefore avoids reduce operations involving all processors which is shown to be critical for large scale simulations. By utilizing an adaptation of the Neutral Territory Method [3] only 6 processors in the local tree and 31 in the global tree in comparison to, respectively, 26 and 189 for a full-shell approach, are involved in the send as well as receive operations for each level. Furthermore, import loads are reduced significantly for the global tree part and for up to three levels in the local tree part. Additional optimizations are an auto-tuning scheme and a method for reducing the communication by fusing neighboring domains in the global tree which can in some cases improve the performance. In this way, relative speedups of 2.3 for a small scenario with 64 local cells on the finest level and 5.6 for a larger scenario with 512 local cells on the finest level were obtained in the range of 4096 to 32768 processors on the Shaheen cluster of the King Abdullah University of Science and Technology [4].

References

[1] M. Obersteiner. Parallel Implementation of the Fast Multipole Method. Master's thesis. Technical University of Munich, 2016.

[2] R. Yokota, G. Turkiyyah, and D. Keyes. Communication Complexity of the Fast Multipole Method and its Algebraic Variants. *ArXiv e-prints*, 2014. arXiv: 1406.1974 [cs.DC].

[3] D. Shaw. A fast, scalable method for the parallel evaluation of distance-limited pairwise particle interactions. *J. Comput. Chem.*, 26(13):1318–1328, 2005.

[4] Shaheen II. URL: https://www.hpc.kaust.edu.sa/content/shaheen-ii.

CONVERGENCE OF THE PARALLEL BLOCK-JACOBI EVD ALGORITHM FOR HERMITIAN MATRICES

Gabriel Okša
Institute of Mathematics, Slovak Academy of Sciences, Bratislava, Slovakia
Gabriel.Oksa@savba.sk

Yusaku Yamamoto
The University of Electro-Communications, Tokyo, Japan
yusaku.yamamoto@uec.ac.jp

Martin Bečka
Institute of Mathematics, Slovak Academy of Sciences, Bratislava, Slovakia
Martin.Becka@savba.sk

Marián Vajteršic
University of Salzburg, Salzburg, Austria, and
Institute of Mathematics, Slovak Academy of Sciences, Bratislava, Slovakia
marian.vajtersic@sbg.ac.at

Let a Hermitian matrix A of order n be divided into a $w \times w$ block structure with $w = 2p$, where p is the number of processors (cores). The aim is to compute the eigenvalue decomposition (EVD) of A in parallel using the two-sided block-Jacobi method with the dynamic ordering defined as follows. At parallel iteration step k, $2p$ off-diagonal blocks of $A^{(k)}$ with block indices (X_{k_1}, Y_{k_1}), (Y_{k_1}, X_{k_1}), ..., (X_{k_p}, Y_{k_p}), (Y_{k_p}, X_{k_p}), $X_{k_i} < Y_{k_i}$ for all i, are eliminated using the *greedy implementation of parallel dynamic ordering* (GIPDO):

1. At iteration step k, all pairs of the off-diagonal blocks are ordered decreasingly with respect to their weights
$$w_{IJ}^{(k)} = \|A_{IJ}^{(k)}\|_F^2 + \|A_{JI}^{(k)}\|_F^2, \ I \neq J.$$
2. After choosing the first pair, $\|A_{X_{k_1}Y_{k_1}}^{(k)}\|_F^2 = \|A_{Y_{k_1}X_{k_1}}^{(k)}\|_F^2 = \max_{I \neq J} \|A_{IJ}^{(k)}\|_F^2$, additional $p-1$ pairs are chosen for annihilation with a decreasing weight and each new pair must have its block-row and block-column indices different from all already chosen blocks.

Processor i, $1 \leq i \leq p$, solves the 2×2-block EVD subproblem:

$$\begin{pmatrix} P_{X_{k_i}X_{k_i}}^{(k)} & P_{X_{k_i}Y_{k_i}}^{(k)} \\ P_{Y_{k_i}X_{k_i}}^{(k)} & P_{Y_{k_i}Y_{k_i}}^{(k)} \end{pmatrix}^H \begin{pmatrix} A_{X_{k_i}X_{k_i}}^{(k)} & A_{X_{k_i}Y_{k_i}}^{(k)} \\ A_{Y_{k_i}X_{k_i}}^{(k)} & A_{Y_{k_i}Y_{k_i}}^{(k)} \end{pmatrix} \begin{pmatrix} P_{X_{k_i}X_{k_i}}^{(k)} & P_{X_{k_i}Y_{k_i}}^{(k)} \\ P_{Y_{k_i}X_{k_i}}^{(k)} & P_{Y_{k_i}Y_{k_i}}^{(k)} \end{pmatrix} = \begin{pmatrix} \hat{A}_{X_{k_i}X_{k_i}}^{(k+1)} & 0 \\ 0 & \hat{A}_{Y_{k_i}Y_{k_i}}^{(k+1)} \end{pmatrix},$$

where the diagonal blocks $\hat{A}_{X_{k_i}X_{k_i}}^{(k+1)}$ and $\hat{A}_{Y_{k_i}Y_{k_i}}^{(k+1)}$ are square, diagonal matrices of order $\ell = n/w$. Subsequently, the orthogonal matrix of local eigenvectors is used in the update of block columns and rows (X_{k_i}, Y_{k_i}) in parallel.

For such an algorithm and under reasonable assumptions, we prove its asymptotic quadratic convergence (AQC) to a diagonal matrix for all possible distributions of eigenvalues (simple, multiple, clusters). Numerical examples confirm the developed theory.

EFFICIENT TRANSFORMATION OF THE GENERAL EIGENPROBLEM WITH SYMMETRIC BANDED MATRICES TO A BANDED STANDARD EIGENPROBLEM

Michael Rippl, Thomas Huckle
Chair of Scientific Computing, Technical University of Munich, Germany
ripplm@in.tum.de

Bruno Lang
Applied Computer Science Group, Bergische Universität Wuppertal, Germany

The solution of symmetric eigenproblems plays a key role in many computational simulations. Generalized eigenproblems are transformed to a standard problem and solved with a common approach for this problem. This transformation has the drawback that for banded matrices in the generalized eigenproblem the banded structure is not preserved. The matrix of the standard eigenproblem will generally be a full matrix.
We followed the ideas of the group of Lang (University of Wuppertal) who modified Crawford's algorithm [1]. Crawford's algorithm proposes a way to immediatelly remove the fill-in when applying the factorization of B to A. This algorithm requires for both matrices a common bandwith. The new approach only requires that the bandwith of matrix A is not bigger than the bandwith of matrix B.
We implemented this procedure to the ELPA project [2]. ELPA offers a two step approach for solving the standard eigenvalue problem. The first step transfers the matrix of the standard problem to a banded matrix and the second step transfers the banded matrix to tridiagonalized form where the eigenvalues can be determined easily.
By using Lang's Twisted-Crawford algorithm the transformation to banded form and also the corresponding back transformation of the eigenvectors can be skipped. Furthermore it provides some interesting blocking and parallelization posibilities, which allow to achieve a good speedup compared to the Crawford's method or Cholesky factorization.

References

[1] C. R. Crawford. Reduction of a Band-Symmetric Generalized Eigenvalue Problem Comm. *Comm. ACM*, (16):41–44, 1973.

[2] T. Auckenthaler. Highly scalable eigensolvers for petaflop applications. PhD thesis. Universität München, 2012.

OPENACC PARALLELIZATION FOR THE SOLUTION OF THE BIDOMAIN EQUATIONS

Stefan Rosenberger
Institute for Mathematics and Scientific Computing, University of Graz, Austria
SFB Research Center, Mathematical Optimization and Applications in Biomedical Sciences
stefan.rosenberger@uni-graz.at

Cardiovascular simulations include coupled PDEs (partial differential equations) for electrical potentials, non-linear deformations and systems of ODEs (ordinary differential equations) all of them are contained in the simulation software CARP (Cardiac Arrhythmia Research Package). We focus in this talk on the solvers for the elliptical part of the bidomain equations describing the electric stimulation of the heart for an anisotropic tissue. The existing conjugate gradient solver with an algebraic multigrid preconditioner is already parallelized by MPI+OpenMP/CUDA.

We investigate the OpenACC parallelization of this solver on one GPU especially its competitiveness with respect to the highly optimized CUDA implementation on recent GPUs. The OpenACC performance seems to be quite close the CUDA performance when typical traps are avoided. We will present additionally first results of these solver parts on Intel's KNL (Knights Landing).

THE FACTORS IN THE SR DECOMPOSITION AND THEIR CONDITIONING

Miroslav Rozložník
Institute of Mathematics, Czech Academy of Sciences, Žitná 25, CZ-115 67 Prague 1, Czech Republic.
miro@cs.cas.cz

Heike Faßbender
Institut Computational Mathematics, AG Numerik, Technische Universität Braunschweig, Pockelsstr. 14, 38106 Braunschweig, Germany.
h.fassbender@tu-braunschweig.de

Almost every nonsingular matrix $A \in R^{2m,2m}$ can be decomposed into the product of a symplectic matrix S and an upper J-triangular matrix R. This decomposition is not unique. In this contribution we analyze the freedom of choice in the symplectic and the upper J-triangular factors and review several existing suggestions on how to choose the free parameters in the SR decomposition. In particular we consider two choices leading to the minimization of the condition number of the diagonal blocks in the upper J-triangular factor and to the minimization of the conditioning of the corresponding blocks in the symplectic factor. We develop bounds for the extremal singular values of the whole upper J-triangular factor and the whole symplectic factor in terms of the spectral properties of even-dimensioned principal submatrices of the skew-symmetric matrix associated with the SR decomposition. The theoretical results are illustrated on two small examples.

This research is supported by the Grant Agency of the Czech Republic under the project GA17-12925S.

PARTIAL INVERSES OF BLOCK TRIDIAGONAL NON-HERMITIAN MATRICES

Louise Spellacy
School of Mathematics, Trinity College, Dublin
spellal@tcd.ie

Darach Golden
Research I.T., Trinity College, Dublin
dgolden@tcd.ie

The SMEAGOL electronic code uses a combination of density function theory (DFT) and Non-Equilibrium Green's Functions (NEGF) to study nanoscale electronic transport under the effect of an applied bias potential [1]. Inversion of a block tridiagonal non-Hermitian matrix is required to obtain the Green's function used by the SMEAGOL code. In many cases, only the block tridiagonal part of the inverse is needed. Currently the SMEAGOL code is limited by single node, multicore matrix inverses. The addition of parallel sparse matrix inverse functionality will allow significantly larger systems to be addressed.

The algorithm presented here is an extension of a previous work in [2] and [3], where a method for parallel inversion of Hermitian block tridiagonal matrices is detailed. This method extends [2] and [3] to the non-Hermitian case and allows for the case of varying block sizes. The tridiagonal blocks of the matrix are evenly distributed across p processes. The local blocks are used to form a "super matrix" on each process. These matrices are inverted locally and the local inverses are combined in a pairwise manner. There are $\log(p)$ combination steps. At each combination step, the updates to the global inverse are represented by updating "matrix maps" on each process. The matrix maps are finally applied to the original local blocks to retrieve the block tridiagonal elements of the inverse. This extended algorithm requires the computation and communication of a greater number of matrix maps than the algorithm detailed in [3]. This "pairwise" algorithm has been implemented as a standalone program, written in Fortran and MPI. It has been tested on local clusters in the Trinity Centre for High Performance Computing. The algorithm is discussed in detail in the presentation. Inverses calculated using the "pairwise" implementation are compared with inverses calculated using well known parallel matrix libraries such as ScaLAPACK and MUMPS. Results are obtained for random test matrices and for matrices arising from DFT calculations.

References

[1] SMEAGOL: Non-equilibrium Electronic Transport. www.smeagol.tcd.ie.

[2] S. Cauley et al. A Scalable Distributed Method for Quantum-Scale Device Simulation. *J. Appl. Phys.*, 101(12):123715, 2007. DOI: 10.1063/1.2748621.

[3] S. Cauley et al. Distributed Non-Equilibrium Green's Function Algorithms for the Simulation of Nanoelectronic Devices with Scattering. *J. Appl. Phys.*, 110(4):043713, 2011. DOI: 10.1063/1.3624612.

WORKFLOW FOR PARALLEL PROCESSING OF BIOMEDICAL IMAGES

Robert Spir, Karol Mikula
Faculty of Civil Engineering, Slovak University of Technology in Bratislava, Slovakia
spir@math.sk, mikula@math.sk

We present an integrated workflow for processing of the biomedical images of early stages of embryo development of various organisms obtained from the two-photon microscopy. We first start with the geodesic mean curvature flow filtering of the raw data to remove the noise and to improve the image quality [1], then we continue using the level-set center detection to obtain the cell identifiers [2], after that we proceed with the segmentation of cells, membranes or the whole embryo using the generalized subjective surface method [3] and finally we can do an automated cell tracking and cell lineage tree reconstruction by the extraction of the cell trajectories forming the lineage tree from the potential field calculated from the combination of distance functions computed inside the 4D segmentations of the processed data [4]. Each step of our processing workflow is parallelized using various techniques such as MPI for distributed computing, OpenMP for local parallelization, GNU Parallel script for launching parallel tasks and Task Parallel Library for parallelization of .net applications. In addition to the parallelization, our workflow is optimized to run on the computer clusters with NUMA architecture.

References

[1] Z. Kriva, K. Mikula, N. Peyrieras, B. Rizzi, A. Sarti, and O. Stasova. 3D Early Embryogenesis Image Filtering by Nonlinear Partial Differential Equations. *Med. Image Anal.*, 14(4):510–526, 2010. DOI: 10.1016/j.media.2010.03.003.

[2] P. Frolkovic, K. Mikula, N. Peyrieras, and A. Sarti. A counting number of cells and cell segmentation using advection-diffusion equations. *Kybernetika*, 43(6):817–829, 2007.

[3] K. Mikula, N. Peyrieras, M. Remesikova, and A.Sarti. 3D embryogenesis image segmentation by the generalized subjective surface method using the finite volume technique. *Finite Volumes for Complex Applications V, Problems & Perspectives*:585–592, 2008.

[4] K. Mikula, R. Spir, M. Smisek, E. Faure, and N. Peyrieras. Nonlinear PDE based numerical methods for cell tracking in zebrafish embryogenesis. *Appl. Numer. Math.*, 95:250–266, 2015. DOI: 10.1016/j.apnum.2014.09.002.

DOMAIN DECOMPOSITION APPLIED TO THE THIN-PLATE SPLINE SADDLE POINT PROBLEM

Linda Stals
Department of Mathematics, Australian National University. Canberra, Australia
linda.stals@anu.edu.au

Data fitting is an integral part of a number of applications including data mining, 3D reconstruction of geometric models, image warping and medical image analysis. A commonly used method for fitting functions to data is the thin-plate spline method. This method is popular because it is not sensitive to noise in the data.

We have developed a discrete thin-plate spline approximation technique that uses local basis functions [1]. With this approach the system of equations is sparse and its size depends only on the number of points in the discrete grid, not the number of data points. Nevertheless the resulting system is a saddle point problem that can be ill-conditioned for certain choices of parameters. In this talk I will present a domain decomposition based preconditioner that results in a well conditioned system for a wider choice of parameters.

References

[1] L. Stals. Efficient Solution Techniques for a Finite Element Thin Plate Spline Formulation. *J. Sci. Comput.*, 63(2):374–409, 2015. DOI: 10.1007/s10915-014-9898-x.

ON THE NUMERICAL STABILITY ANALYSIS OF PIPELINED KRYLOV SUBSPACE METHODS

Erin C. Carson
Courant Institute of Mathematical Sciences, New York University, New York
rinc@cims.nyu.edu

Miroslav Rozložník
Institute of Computer Science and Institute of Mathematics, Czech Academy of Sciences, Prague, Czech Republic
miro@cs.cas.cz

Zdeněk Strakoš, Petr Tichý, Miroslav Tůma
Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic
strakos@karlin.mff.cuni.cz, ptichy@karlin.mff.cuni.cz, mirektuma@karlin.mff.cuni.cz

Inexact computations in Krylov subspace methods, either due to floating point roundoff error or intentional action motivated by savings in computing time or energy consumption, have two basic effects, namely, slowing down convergence and limiting attainable accuracy. Although the methodologies for their investigation are different, these phenomena are closely related and cannot be separated from one another.

As the name suggests, Krylov subspace methods can be viewed as a sequence of projections onto nested subspaces of increasing dimension. They are therefore by their nature implemented as synchronized recurrences. This is the fundamental obstacle to efficient parallel implementation. Standard approaches to overcoming this obstacle described in the literature involve reducing the number of global synchronization points and increasing parallelism in performing arithmetic operations within individual iterations. One such approach, employed by the so-called pipelined Krylov subspace methods, involves overlapping the global communication needed for computing inner products with local arithmetic computations.

Recently, the issues of attainable accuracy and delayed convergence caused by inexact computations became of interest in relation to pipelined Krylov subspace methods. In this contribution based on [1] we recall the related early results and developments in synchronization-reducing Krylov subspace methods, identify the main factors determining possible numerical instabilities, and outline approaches needed for the analysis and understanding of pipelined Krylov subspace methods. We demonstrate the discussed issues numerically using several algorithmic variants of the conjugate gradient method. We conclude with a brief perspective on Krylov subspace methods in the forthcoming exascale era.

References

[1] E. Carson, M. Rozložník, Z. Strakoš, P. Tichý, and M. Tůma. On the numerical stability analysis of pipelined Krylov subspace methods. *Submitted*, 2016.

EMPLOYING HPC FOR ANALYZING NONLINEAR PDE SYSTEMS BEYOND SIMULATION

Jonas Thies, Weiyan Song
German Aerospace Center (DLR)
Simulation and Software Technology
Linder Höhe, 51147 Cologne, Germany
Jonas.Thies@DLR.de, Weiyan.Song@DLR.de

Fred W. Wubs, Sven Baars
University of Groningen
Institute for Mathematics and Computer Science
Nijenborgh 9, 9747 AG Groningen, the Netherlands
F.W.Wubs@RuG.nl, S.Baars@RuG.nl

We review techniques of numerical bifurcation and stability analysis with examples from computational fluid dynamics and biology. The methodology allows insight into the complete dynamics of nonlinear PDE systems, where standard simulation tool chains leave the question of existence, proximity and stability of multiple solutions open.

The main bottleneck in the method are large and sparse linear systems of equations and eigenvalue problems arising from the discretized steady-state PDE. The use of HPC is therefore attractive to increase the achievable resolution, but remains challenging because nonsymmetric and indefinite systems need to be solved. The 'hybrid multi-level solver' HYMLS [1, 2] is a robust multi-level incomplete factorization technique that was designed for this particular class of problems. HYMLS has an intuitive geometric interpretation and good parallelization properties. We present some performance results of a prototypical implementation based on MPI and the Trilinos software. The eigenvalue problems that arise are solved using the Jacobi-Davidson method as implemented in the SPPEXA ESSEX [3] project's phist library [4].

References

[1] F. W. Wubs and J. Thies. A Robust Two-Level Incomplete Factorization for (Navier-)Stokes Saddle Point Matrices. *SIAM J. Matrix Anal. Appl.*, 32(4):1475–1499, 2011. DOI: 10.1137/100789439.

[2] HYMLS, a HYbrid Multi-Level Solver. URL: https://bitbucket.org/hymls/hymls/.

[3] The ESSEX project website. URL: http://blogs.fau.de/essex/.

[4] PHIST, a Pipelined Hybrid-parallel Iterative Solver Toolkit. URL: https://bitbucket.org/essex/phist/.

IMPACT OF INTERCONNECTION NETWORK TOPOLOGY ON PARALLEL PERFORMANCE - A SURVEY

Roman Trobec and Janez Ugovšek
Department of Communication Systems, Jožef Stefan Institue, Ljubljana, Slovenia
roman.trobec@ijs.si, janez@ugovsek.info

Interconnection networks (ICNs) have an important role on the execution time of cooperating computers. In parallel systems with a high number of cooperating computers the performance of data communication systems is becoming more important than the performance of the processors. The technological barrier posed on the further increasing of processor speed is evident in the contemporary high performance computers (HPC) as an ever increasing number of cooperating processors [1]. However, the exchange of temporary data between processors can disturb the balance between calculation and communication time. The ICN importantly determines the efficiency and scalability of a HPC on most real-world parallel applications. It can shorten the execution by more efficiently exploited computers, even if their number grows.

The performance of an ICN depends on technological and topological factors, e.g., network topology, message routing, and flow-control algorithms. The routing and flow-control algorithms have advanced to a state where efficient techniques are known and used. However, further sophistication is possible in the development of network topologies [2], which is the main focus of our work. We present the state-of-the-art technology and topology of several common ICNs used in petascale computers. Their analysis indicates that ICNs with higher performance are needed for future exascale computers [3]. They should be based on high-radix topologies with optical connections [4] for longer links. It is could be also expected that future ICNs will be able to adapt dynamically to the current application in some optimal way.

References

[1] R. Trobec, R. Vasiljević, M. Tomašević, V. Milutinović, and M. Beivide R. Valero. Interconnection Networks in Petascale Computer Systems: A Survey. *ACM Comput. Surv.*, 49(3):44:1–44:24, 2016. DOI: 10.1145/2983387.

[2] W. J. Dally and B. P. Towles. *Principles and Practices of Interconnection Networks.* Morgan Kaufmann Publishers Inc., 2004.

[3] Coteus, P. W. and Knickerbocker, J. U. and Lam, C. H. and Vlasov, Y. A. Technologies for Exascale Systems. *IBM J. Res. Dev.*, 55(5):581–592, 2011. DOI: 10.1147/JRD.2011.2163967.

[4] A. F. Benner, M. Ignatowski, J. A. Kash, D. M. Kuchta, and M. B. Ritter. Exploitation of Optical Interconnects in Future Server Architectures. *IBM J. Res. Dev.*, 49(4/5):755–775, 2005. URL: http://dl.acm.org/citation.cfm?id=1148882.1148902.

MIXED SPARSE-DENSE LINEAR LEAST SQUARES AND PRECONDITIONED ITERATIVE METHODS

Jennifer Scott
STFC Rutherford Appleton Laboratory, Harwell Campus, Didcot, Oxfordshire, OX11 0QX, UK
jennifer.scott@stfc.ac.uk.

Miroslav Tůma
Department of Numerical Mathematics, Faculty of Mathematics and Physics, Charles University, Prague, Czech Republic
mirektuma@karlin.mff.cuni.cz

The efficient solution of large linear least squares problems in which the system matrix A contains rows with very different densities is challenging. There have been many classical contributions to solving this problem that focus on direct methods; they can be found in the monograph [1]. Such solvers typically perform a splitting of the rows of A into two row blocks, A_s and A_d. The block A_s is such that the sparse factorization of the normal matrix $A_s^T A_s$ is feasible, while the rows in the block A_d have a relatively large number of nonzero entries. These dense rows are initially ignored, a factorization of the sparse part is computed using a sparse direct solver and then the solution updated to take account of the omitted dense rows.

There are two potential weaknesses of this approach. First, in practical applications the number of rows that contain a significant number of entries may not be small. Processing some of the denser rows separately may improve performance. Furthermore, large-scale problems require the use of preconditioned iterative solvers. A straightforward proposal to precondition the iterative solver using only an incomplete factorization of the sparse block while discarding the dense block may not lead to any success. In this presentation, we propose processing A_s separately within a conjugate gradient method using an incomplete factorization preconditioner combined with the factorization of a dense matrix of size equal to the number of rows in A_d. Problems arising from practical applications are used to demonstrate the potential of the new approach; see also [2].

References

[1] A. Björck. *Numerical methods for least squares problems*. Society for Industrial and Applied Mathematics (SIAM), 1996, pages xviii+408. DOI: 10.1137/1.9781611971484.

[2] J. A. Scott and M. Tůma. Solving mixed sparse-dense linear least squares by preconditioned iterative methods. Technical Report RAL-P-2017-001. RAL, 2017.

EFFICIENT GPU-BASED SMOOTHED PARTICLE HYDRODYNAMICS

Matthias Korch
Department of Computer Science
University of Bayreuth
korch@uni-bayreuth.de

Tim Werner
Department of Computer Science
University of Bayreuth
werner@uni-bayreuth.de

Smoothed particle hydrodynamics (SPH) is a numerical method, which simulates a fluid by dividing it into particles interacting with each other. For reducing the computational complexity, SPH simulations typically limit the interactions between particles to a short range. Still, computing those short-ranged interactions is the most computationally intensive task in SPH simulations, typically requiring more than 90 % of the total runtime. Because these interactions can also be computed in parallel, SPH is well suited for parallel processors such as GPUs. However, the performance can be enhanced further by using GPU specific optimization techniques.

In this paper, we investigate how to efficiently compute those short-ranged interactions of particles in SPH. For this purpose, starting from a basic linked cell approach, we iteratively evaluate several different optimization techniques for the kernels, namely removal of the x-loop, decreasing the cell size, simplification of the cell-sphere interaction test, fast forwarding of particles, and temporary Verlet lists. Our main goals are the improvement of the data parallelism and the reduction of the overhead of the grid traversal. The final implementation achieves both goals and yields a significant speedup compared to the basic approach.

IS GOSSIP-INSPIRED REDUCTION COMPETITIVE IN HIGH PERFORMANCE COMPUTING?

Elias Wimmer
TU Wien, Faculty of Informatics, Research Group Parallel Computing, Austria
ewimmer@par.tuwien.ac.at

Marc Casas
Barcelona Supercomputing Center (BSC), Spain
marc.casas@bsc.es

Wilfried N. Gansterer
University of Vienna, Faculty of Computer Science, Austria
wilfried.gansterer@univie.ac.at

The utility of gossip-based reduction algorithms in a High Performance Computing (HPC) context is investigated. They are compared to state-of-the-art deterministic parallel reduction algorithms in terms of fault tolerance and resilience against silent data corruption (SDC) as well as in terms of runtime performance and scalability. New gossip-based reduction algorithms are proposed which significantly improve the state-of-the-art in terms of resilience against SDC. A new gossip-inspired reduction algorithm is proposed which promises a more competitive runtime performance for low accuracy in an HPC context than gossip-based algorithms. It is shown that for very large systems the new gossip-inspired reduction algorithm has the potential to outperform classical reduction algorithm for low accuracy problems.

ROUNDOFF ERROR ANALYSIS OF THE CHOLESKYQR2 AND RELATED ALGORITHMS

Yusaku Yamamoto
Department of Communication Engineering and Informatics,
The University of Electro-Communications, Tokyo, Japan
yusaku.yamamoto@uec.ac.jp

Cholesky QR is an ideal QR factorization algorithm from the viewpoint of high performance computing [1], but it has rarely been used in practice due to numerical instability. Recently, we showed that by repeating Cholesky QR twice, we can greatly improve the stability [2]. In this talk, we present a detailed error analysis of the algorithm, which we call CholeskyQR2. Numerical stability of related algorithms, such as the CholeskyQR2 algorithm in an oblique inner product [3], is also discussed.

References

[1] T. Fukaya, Y. Nakatsukasa, Y Yanagisawa, and Y. Yamamoto. CholeskyQR2: A Simple and Communication-avoiding Algorithm for Computing a Tall-skinny QR Factorization on a Large-scale Parallel System. In *ScalA'14 Proceedings of the 5th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems*. IEEE, 2014, pages 31–38. DOI: 10.1109/ScalA.2014.11.

[2] Y. Yamamoto, Y. Nakatsukasa, Y. Yanagisawa, and T. Fukaya. Roundoff Error Analysis of the CholeskyQR2 Algorithm. *Electron. Trans. Numer. Anal.*, 44:306–326, 2015.

[3] Y. Yamamoto, Y. Nakatsukasa, Y. Yanagisawa, and T. Fukaya. Roundoff Error Analysis of the CholeskyQR2 Algorithm in an Oblique Inner Product. *JSIAM Lett.*, 8:5–8, 2015.