Swiss Numerical Analysis Day 2024

unige.ch/math/snd24
snd24@unige.ch

TUESDAY
SEPTEMBER 10, 2024

INVITED PLENARY SPEAKERS
Prof. Gabriele Ciaramella
(MOX Lab, Politecnico di Milano)
Prof. Annika Lang
(Chalmers/University of Gothenburg)

TOPICS
Scientific computing, numerical analysis, high-performance computing, visualization, numerical analysis in engineering, computational chemistry, biology, etc.

REGISTRATION
unige.ch/math/snd24
Abstract submission deadline: **July 31st, 2024**
Registration deadline: **August 29th, 2024**

VENUE
University of Geneva, Sciences II & III
Lecture Hall A150 (ground floor)
30 Quai Ernest-Ansermet, Geneva

ORGANIZATION
Martin J. Gander, Bart Vandereycken, Gilles Vilmart
(University of Geneva, Section of Mathematics)

SPONSORS
We acknowledge partial support from the Swiss National Science Foundation.

We look forward to seeing you in Geneva!

FACULTY OF SCIENCE
SECTION OF MATHEMATICS

UNIVERSITÉ DE GENÈVE

Explicit stabilized implementation of implicit Runge-Kutta methods

Author and Presenter:
Ibrahim Almuslimani (EPFL, Switzerland)

Co-authors:
Gilles Vilmart (University of Geneva, Switzerland)
Konstantinos Zygalakis (University of Edinburgh, UK)

Abstract: We introduce a new fully explicit implementation of implicit Runge-Kutta methods of any order based on optimization techniques and a stabilized Runge-Kutta-Chebyshev (RKC) method. Inspired by the RKCD optimization algorithm introduced in [2], this is done by reformulating the implicit Runge-Kutta scheme as an optimization problem and iterating on its gradient flow using an explicit stabilized partitioned RKC integrator. Numerical experiments that show promising results by comparing with existing methods are presented.

References

[1] I. Almuslimani, G. Vilmart, and K. C. Zygalakis, Explicit stabilized implementation of implicit Runge-Kutta methods. 2024. (in preparation)

[2] A. Eftekhari, B. Vandereycken, G. Vilmart, and K. C. Zygalakis, Explicit stabilised gradient descent for faster strongly convex optimisation. 2021. *BIT Numerical Mathematics*, 61(1):119–139.

Energy-stable spatial-temporal finite difference discretization of the two-fluid turbulent plasma model

Author and Presenter:
Micol, Bassanini (Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland)

Co-authors:
Simone, Deparis (Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland)
Paolo, Ricci (Ecole Polytechnique Federale de Lausanne (EPFL), Switzerland)

Abstract: In this work, we investigate a reduced two-fluid cold plasma model describing plasma behavior in boundary region of a fusion device. This model captures the dynamics of the electrostatic shear Alfvén waves (SAWs), [3], an anisotropic fast wave present in many plasma systems, alongside the $\mathbf{E} \times \mathbf{B}$-drift dynamics characterized by a slower time scale. We show that the model conserves energy analytically and we design a numerical algorithm to guarantee energy-stable simulations.

We first ensure an energy-stable spatial discretization of the system considering only the fast-oscillating SAWs. In wave propagation problems discretized with finite difference methods, staggered grids are commonly used to avoid checkerboard patterns and to improve accuracy in the approximation of short-wavelength components of the solutions. In this study, we develop a finite difference method on staggered grids for wave-like problems that we prove to be energy-preserving and accurate [1]. This method mimics some characteristics of the summation-by-parts (SBP) operators framework [4], in particular, it preserves the divergence theorem at the discrete level. Its design is intended to be versatile and applicable to a broad spectrum of wave problems characterized by a divergence-free velocity.

As a second step, we apply an implicit-explicit runge-kutta scheme BPR(3,5,3) [2] to the full two-fluid turbulent plasma model. The SAWs are discretized in space with the algorithm described above and are integrated implicitly with an unconditionally stable method, while the other components with slower timescales are integrated explicitly. As a result, the time step of the overall algorithm is no longer dictated by the fast-oscillating components of the equations.

The proposed spatial-temporal discretization is tested in a 3D Cartesian geometry with a magnetic field configuration resembling that of a tokamak device. Compared to an explicit time integration, we are able to have energy-stable simulation for higher time steps. Ultimately, we show that this IMEX approach reduces the computational cost per unit of time.

References

[1] M. Bassanini, S. Deparis, and P. Ricci, Energy-preserving finite difference discretizations of anisotropic wave motion with applications to plasma physics. 2024. (in preparation)

[2] S. Boscarino, L. Pareschi, and G. Russo, Implicit-explicit runge–kutta schemes for hyperbolic systems and kinetic equations in the diffusion limit. 2013. SIAM Journal on Scientific Computing, 35(1), A22–A51.

[3] S. Jolliet et al., Numerical approach to the parallel gradient operator in tokamak scrape-off layer turbulence simulations and application to the GBS code. 2015. Computer Physics Communications 188.

[4] M. Svärd and J. Nordström, Review of summation-by-parts schemes for initial–boundary-value problems. 2014. Journal of Computational Physics 268.

A Two-Step Method Coupling Eddy Currents and Magneto-Statics

Author and Presenter:
Martina, Busetto (ABB Corporate Research, Switzerland)

Co-authors:
Christoph, Winkelmann (ABB Corporate Research, Switzerland)

Abstract: We present the mathematical theory and its numerical validation of a method tailored to include eddy-current effects only in a part of the domain which does not change over time. This results in a heterogeneous problem combining an eddy-current model in a subset of the computational domain with a magneto-quasistatic model in the remainder of the domain. We adopt a two-domain two-step approach in which the primary variables of the problem are the electric scalar potential and the magnetic vector potential. We discuss technical aspects like boundary conditions and limitations of the approach, and we show numerical results that validate the formulation.

Non-intrusive projected exponential methods for stiff dynamical low-rank approximation

Author and Presenter:
Benjamin Carrel (UNIGE)

Co-authors:
Bart Vandereycken (UNIGE)

Abstract:

The numerical integration of stiff equations is a challenging problem that needs to be approached by specialized numerical methods. Exponential integrators form a popular class of such methods since they are provably robust to stiffness and have been successfully applied to a variety of problems. On the other hand, the dynamical low-rank approximation is a recent technique for solving high-dimensional differential equations by means of low-rank approximations. However, the domain is lacking numerical methods for stiff equations since existing methods are either not robust-to-stiffness or have unreasonably large hidden constants.

In this talk, we focus on solving large-scale stiff matrix differential equations with a Sylvester-like structure,

$$X'(t) = AX(t) + X(t)B + G(t, X(t)), X_0 = X(0),$$

that admit good low-rank approximations. We propose two new methods that have good convergence properties, small memory footprints and that are fast to compute. The theoretical analysis shows that the new methods have order one and two, respectively. We also propose a practical implementation based on Krylov techniques. The approximation error is analyzed, leading to a priori error bounds and, therefore, a mean for choosing the size of the Krylov space. Numerical experiments are performed on several examples, confirming the theory and showing good speedup in comparison to existing techniques.

If time permits, we will also present a new low-rank approximation model that reduces intrusiveness, and therefore allows for fast non-linear evaluations. The model requires the theory of differential inclusions, which is new in the context of low-rank approximations. When the non-linearity is applied componentwise, numerical experiments show huge gains in performance.

References

[1] B. Carrel, B. Vandereycken
Projected exponential methods for stiff dynamical low-rank approximation problems
2023. arXiv preprint arXiv:2312.00172

A sweeping domain decomposition method for elliptic problems

Author and Presenter:
Bastien Chaudet-Dumas (HEIG-VD, Switzerland)

Co-authors:
Martin J. Gander (University of Geneva, Switzerland)

Abstract: In the context of partial differential equations, domain decompositions methods constitute an effective way to build efficient iterative solvers. Among other advantages, they are naturally adapted to run in parallel, and their convergence properties do not depend on the chosen discretization method, nor on the mesh size, since they can be applied at the continuous level. However, in order for these iterative methods to be truly efficient, we must be able to have some control over the number of iterations necessary to reach convergence.
In this presentation, we introduce a new domain decompositon method that converges in a finite number of iterations. For a given set of subdomains, the method relies on a corresponding decomposition of the solution into pseudo-even/odd components (one for each dimension variable). Thanks to this decomposition, we are able to compute exactly the solution in boundary subdomains, and then we can propagate the information step by step towards the center of the domain. For a simple elliptic problem, we prove that the method converges in 1D and 2D for generic sets of regular subdomains. Finally, we illustrate our results with numerical experiments.

The Green's function for an acoustic half-space problem with impedance boundary conditions (oral)

Author and Presenter:
Chuhe, Lin (University of Zurich, Switzerland)

Co-authors:
Stefan, Sauter (University of Zurich, Switzerland)
Markus, Melenk (TU Vienna, Austria)

Abstract: In this talk we show that the acoustic Green's function for a half-space impedance problem in general d spatial dimensions can be written as a sum of two terms, each of which is the product of an exponential function with the eikonal in the argument and a slowly varying function. We introduce the notion of families of slowly varying functions and formulate this statement as a theorem along a sketch of its proof. This talk comprises joint work with Stefan Sauter, University of Zurich and Markus Melenk, TU Vienna.

References

[1] C. Lin, J.M. Melenk, S. Sauter. The Green's function for an acoustic half-space problem with impedance boundary conditions. Part I: Representation formula, 2024, in prep.

[2] C. Lin, J.M. Melenk, and S. Sauter. The Green's function for an acoustic, half-space impedance problem. Part II: Analysis of the slowly varying and the plane wave component, 2024, in prep.

On space-time RAS methods for wave-type equations

Author and Presenter:
Gabriele Ciaramella, MOX Lab, Politecnico di Milano, Italy

Abstract: The parallel-in-time integration of wave-type equations is well known to be a difficult task. When applying Schwarz waveform-relaxation methods, one generally experiences rapid error growths before converging in a finite number of iterations. This unfortunate behavior prevents the successful application of these domain decomposition methods for long time intervals and suggests the need for a decomposition in time. Thus, a new space-time restricted additive Schwarz (XT-RAS) framework is presented for the treatment of wave problems defined on long time intervals. Moreover, XT-RAS is strictly related to tent-pitching strategies. This relation will be studied in detail and will allow adapting XT-RAS to the specific wave propagation problem. Numerical experiments support the presented theoretical findings and show the efficiency of the new proposed frameworks.

Parallel Neural Network Training via Nonlinearly Preconditioned Trust-Region Method

Author and Presenter:
Samuel Cruz

Co-authors:
Ken Trotti, Alena Kopanicakova, Rolf Krause

Neural networks and the data sets used to train them continue to grow in size. Traditional training methods like SGD and Adam require extensive hyperparameter tuning, which is increasingly incompatible with today's demands as current applications require methods that ideally do not need hyperparameter tuning. With this in mind, we investigate additive domain decomposition methods for neural network training. Due to their underlying additive decomposition, our methods are parallelizable, and owing to a trust-region strategy, they largely obviate the need for hyperparameter tuning. Our test results suggest that the investigated methods are potential candidates for efficient neural network training.

Fractures and thin heterogeneities as interface conditions

Author and Presenter:
Marco, Favino (Università della Svizzera italiana, Switzerland)

Abstract

The simulation of seismic wave propagation in fluid-filled fractured porous media poses a formidable challenge, requiring the resolution of a hydromechanical coupling problem, described by Biot's equations, in heterogenous media characterized by the presence of complex networks of thin heterogeneities (fractures). Key challenges in this endeavor include the generation of meshes to accurately represent realistic fracture networks, a significant hurdle in itself. Additionally, the low regularity of solutions impacts the convergence of discretization methods and high-performance solution techniques, such as geometric multigrid.

In [1], we proposed an automatic strategy to automatically generated meshes for fractured media. The idea is to start from a uniform mesh and a fracture distribution and refine the elements which have a non-empty overlap with at least one fracture. Iterating this process, this strategy produces a mesh which does not resolve the boundary of the fractures but it allows to approximate it with a desired accuracy. However, the resulting mesh, and hence, the linear system associated to the discretization of these meshes is extremely large.

Another common approach to model fractures involves using hybrid-dimensional models, where fractures are downscaled to objects of a lower dimension and separate problems are solved in the background and fracture domains with suitable coupling conditions [2, 3].

In this talk, we propose a novel formulation for modeling fractures, which considers fractures as cuts in the domain and models their effects using Wentzell interface conditions [4]. The resulting model is elegant and straightforward, consisting in a primal formulation where the only unknown is the pressure in the background domain. The model can naturally handle fractures that cross each other.

We analyze the discretization error for different orders of finite element and material properties, and we compare the arising conditions with standard analytical models used for fractured porous and poroelastic media [5].

References

[1] Favino, Marco and Hunziker, Jürg and Caspari, Eva and Quintal, Beatriz and Holliger, Klaus and Krause, Rolf, *Fully-automated adaptive mesh refinement for media embedding complex heterogeneities: application to poroelastic fluid pressure diffusion*, Computational Geosciences (2020)

[2] Flemisch, B., et al. *Benchmarks for single-phase flow in fractured porous media*, Advances in Water Resources, 111, 239-258 (2018).

[3] Berre, Inga, et al. *Verification benchmarks for single-phase flow in three-dimensional fractured porous media*, Advances in Water Resources 147 (2021): 103759.

[4] P. G. Ciarlet (1978). Kashiwabara, T., Colciago, C. M., Dedè, L., Quarteroni, A. (2015). *Well-posedness, regularity, and convergence analysis of the finite element approximation of a generalized Robin boundary value problem.* SIAM Journal on Numerical Analysis, 53(1), 105-126.

[5] Nakagawa, Seiji, and Michael A. Schoenberg. *Poroelastic modeling of seismic boundary conditions across a fracture*, The Journal of the Acoustical Society of America 122.2 (2007): 831-847.

An adaptive finite elements - neural network method applied to parametric PDEs

Author and Presenter :
Maude Girardin (EPFL, Switzerland)

Co-authors :
Alexandre Caboussat (University of Applied Sciences and Arts Western Switzerland, Switzerland)
Marco Picasso (EPFL, Switzerland)

Abtract : We consider a generic parametric PDE

$$\mathcal{F}(u(x;\mu);\mu) = 0, \quad x \in \Omega, \quad \mu \in \mathcal{P},$$

where Ω and $\mathcal{P}$ denote respectively the physical and the parameter spaces. Given a parameter μ, classical numerical methods, such as finite elements, can be used to find an approximated solution $u_h(\cdot;\mu)$. Nevertheless, these methods may be highly time consuming, depending on the PDE of interest. This becomes an issue if the solution corresponding to a new parameter has to be obtained in real time, or if the problem has to be solved several times in a row, for example to solve an inverse problem. In order to overcome this issue, we aim to train a neural network surrogate that approximates the map $(x;\mu) \mapsto u(x;\mu)$, using data that come from finite element simulations.

Given a norm $|||\cdot|||$ on $\Omega \times \mathcal{P}$, the error between the network approximation $u_{\mathcal{N}}$ and the analytical solution u can then be splitted as:

$$|||u - u_{\mathcal{N}}||| \leq |||u - u_h||| + |||u_h - u_{\mathcal{N}}|||.$$

The first term corresponds to the error of the finite element method, that is embedded in the training set, and the second term is the error due to the neural network approximation. We discuss a stategy to estimate and control these two terms, in order to ensure that the overall error of the method is below some preset tolerance [1,2].

Numerical results will be presented for different test cases.

References :

[1] A. Caboussat, M. Girardin and M. Picasso. Error assessment for a finite elements-neural networks approach applied to parametric PDEs. The transport equation. Proceedings of the 11th edition of the International Conference on Adaptive Modeling and Simulation (ADMOS 2023).

[2] A. Caboussat, M. Girardin and M. Picasso. Error assessment of an adaptive finite elements neural networks method for an elliptic parametric PDE. Computer Method in Applied Mechanics and Engineering, 2024, vol. 421.

Subspace Acceleration for a Sequence of Linear Systems and Application to Plasma Simulation

Author and Presenter:
Margherita Guido, École Polytechnique Fédérale de Lausanne (EPFL), Institute of Mathematics - Swiss Plasma Center, Switzerland

Co-authors:
Daniel Kressner, École Polytechnique Fédérale de Lausanne (EPFL), Institute of Mathematics, Switzerland
Paolo Ricci, École Polytechnique Fédérale de Lausanne (EPFL), Swiss Plasma Center, Switzerland

Abstract: We present an acceleration method for sequences of large-scale linear systems, such as the ones arising from the numerical solution of time-dependent partial differential equations coupled with algebraic constraints. We leverage the subspace containing the history of solutions computed at previous time steps in order to generate a good initial guess for the iterative solver. In particular, we propose a novel combination of reduced-order projection with randomized linear algebra techniques, which drastically reduces the number of iterations needed for convergence. The accuracy of the initial guess produced by the reduced-order projection is analyzed when the coefficients of the linear system depend analytically on time, showing that it improves rapidly as the size of the history increases. The developed method is applied to the simulation of plasma turbulence in the boundary of a fusion device, showing that the time needed for solving the linear systems is significantly reduced.

References

[1] M. Guido, D. Kressner, and P. Ricci, Subspace acceleration for a sequence of linear systems and application to plasma simulation *Journal of Scientific Computing*, 99(3):68, 2024.

Stochastic partial differential equations on surfaces and evolving random surfaces: a computational approach

Author and Presenter:
Annika Lang, Chalmers/University of Gothenburg, Sweden

Abstract: Looking around us, many surfaces including the Earth are no plain Euclidean domains but special cases of Riemannian manifolds. One way of describing uncertain physical phenomena on these surfaces is via stochastic partial differential equations. In this talk, I will introduce how to compute approximations of solutions to such equations and give convergence results to characterize the quality of the approximations. Furthermore, I will show how these solutions on surfaces are a first step towards the computation of time-evolving stochastic manifolds.

Time domain decomposition and application to PDE-constrained optimization problems

Author and Presenter:
Liu-Di, Lu (University of Geneva, Switzerland)

Co-authors:
Martin Jakob, Gander (University of Geneva, Switzerland)

Abstract: In this talk, we will explore domain decomposition methods and their application to parabolic PDE-constrained optimization problems. We will first compare the difference between decomposing in space and in time. Then, we are going to discuss some properties of time domain decomposition methods based on the forward-backward structure of the optimality system. Variants can be identified, some of these are only good smoothers, while others could lead to efficient solvers.

A spline-based hexahedral mesh generator for coronary arteries

Author and Presenter:
Fabio, Marcinnó[1]

Co-authors:
Jochen, Hinz[1]
Annalisa, Buffa[1,2]
Simone, Deparis[1]

[1] Institute of Mathematics, École Polytechnique Fédérale de Lausanne (EPFL), Lausanne, Switzerland
[2] Instituto di Matematica Applicata e Tecnologie Informatiche 'E. Magenes', Pavia, Italy

Abstract: Coronary arteries are remarkably subject to the formation of atherosclerotic plaques (stenoses), therefore increasing the risk of myocardial infarction (MI) [1]. Several indices have been proposed in order to quantify the severity of stenotic formations, such as the fractional flow reserve (FFR) and descriptors based on the wall shear stress (WSS) [2]. This work is the result of a collaboration with the CHUV (Centre hospitalier universitaire vaudois) which provided the data on the coronary arteries. The purpose of this study is to develop a computational framework for hexahedral grid generation and numerical simulation of hæmodynamics in patient-specific coronary artery geometries. Regarding the meshing, we present a fast and robust hexahedral mesh generator which is capable of adapting to the bifurcations and to a possible integration of the cardiac motion. On the fluid dynamics side, we take into account the first bifurcation of the coronary tree (LAD-CX) and the presence of stenosis. We show numerical results on a select number of geometries. The FFR results are furthermore compared against clinical measures [3].

References

[1] A. Candreva et al., "Current and future applications of computational fluid dynamics in coronary artery disease," Reviews in Cardiovascular Medicine, vol. 23, no. 11, p. 377, 2022.

[2] F. Gijsen et al., "Expert recommendations on the assessment of wall shear stress in human coronary arteries: existing methodologies, technical considerations, and clinical applications," European heart journal, vol. 40, no. 41, pp. 3421–3433, 2019.

[3] F. Marcinnó et al., "A spline-based hexahedral mesh generator for coronary arteries," in preparation.

Analysis of Parareal with Spatial Coarsening

Author and Presenter:
Aušra, Pogoželskytė (Université de Genève, Switzerland)

Co-authors:
Martin J., Gander (Université de Genève, Switzerland)

Abstract: In this talk, we will discuss Parareal, a black-box algorithm introduced by Lions, Maday and Turninici in 2001 to solve ODEs and PDEs by parallelizing the time dimension [1]. Its limitations will be highlighted; one being that it can be unstable when using explicit solvers inside the algorithm. A simple solution to that problem is to use spatial coarsening in conjunction with Parareal [2]. We will explore the implications of such a technique on its convergence [3].

References

[1] J.-L. Lions, Y. Maday, and G. Turninici, Résolution d'EDP par un schéma en temps "pararéel", 2001

[2] P. Fischer, F. Hecht, and Y. Maday, A Parareal in Time Semi-implicit Approximation of the Navier-Stokes Equations, 2005

[3] M. J. Gander, A. Pogoželskytė, Convergence Bounds for Parareal with Spatial Coarsening, 2022

A Posteriori Error Estimates for the Wave Equation with Mesh Change in the Leapfrog Method

Author and Presenter:
Carina, Santos (University Basel, Switzerland)

Co-authors:
Marcus, Grote (University Basel, Switzerland)
Omar, Lakkis (University of Sussex, UK)

Abstract: Adaptive mesh refinement is key for the efficient simulation of wave phenomena in complex geometry. A posteriori based adaptivity for finite element methods (FEMs) is well-developed for elliptic and parabolic problems. For wave equations in contrast, a posteriori error control is less developed as residual-based error control mainly applies to implicit first-order in time DG or RK schemes. Georgoulis et al. [3] derived an a posteriori error estimate for the leapfrog method, arguably the most popular two-step explicit time integration method for the wave equation. Local mesh refinement, however, severely constrains the time-step of any explicit time-stepping method due to the CFL stability condition governed by the smallest elements in the mesh. Leapfrog based local time-stepping (LTS) methods [1, 2] overcome that bottleneck without sacrificing explicitness by taking smaller time-steps only where needed. Here we derive fully discrete a posteriori error estimates for the wave equation when integrated with a leapfrog-based LTS method, incorporating thus adaptivity into explicit time integration with mesh change in time while retaining efficiency.

References

[1] J. Diaz and M. J. Grote, Energy conserving explicit local time stepping for second-order wave equations, *SIAM J. Sci. Comput.* **31** (2009), pp. 1985–2014.

[2] M. J. Grote, S. Michel and S. A. Sauter, Stabilized leapfrog based local time-stepping method for the wave equation, *Math. Comp.* **90** (2021), pp. 2603–2643.

[3] E. H. Georgoulis, O. Lakkis, C. G. Makridakis and J. M. Virtanen, A posteriori error estimates for leap-frog and cosine methods for second order evolution problems, *SIAM J. Numer. Anal.* **54** (2016), pp. 120–136.

GPU-enabled distributed dense eigensolver in DLA-Future

Author and Presenter:
Raffaele Solcà (ETH Zürich, CSCS, Switzerland)

Co-authors:
John Biddiscombe (ETH Zürich, CSCS, Switzerland)
Alberto Invernizzi (ETH Zürich, CSCS, Switzerland)
Rocco Meli (ETH Zürich, CSCS, Switzerland)
Auriane Reverdell (ETH Zürich, CSCS, Switzerland)
Mikael Simberg (ETH Zürich, CSCS, Switzerland)

Abstract: DLA-Future [1, 3] is a linear algebra library providing GPU-enabled Hermitian or symmetric eigenvalue solvers for standard problems $A\mathbf{x} = \lambda\mathbf{x}$ and generalized problems $A\mathbf{x} = \lambda B\mathbf{x}$. It is designed to replace ScaLAPACK, the well-known state-of-the-art implementation which employs fork-join parallelism. This method was suitable for older system architectures, but can heavily underperform on modern heterogeneous machines.

The DLA-Future eigensolver employs direct methods to compute the eigenspectrum of a matrix. To reduce the number of memory-bound operations it uses a two-stage tridiagonalization approach. In the first stage the eigenproblem is transformed into a eigenvalue problem of a band matrix, using block application of Householder reflectors. In the second stage the problem is further reduced to a tridiagonal eigenvalue problem using a bulge chasing algorithm. The reduction of memory-bound operations comes at the cost of an increase in compute-bound operations (mainly due to the need of an extra eigenvector back-transformation). The solution of the tridiagonal eigensolver is found using the Cuppen's Divide and Conquer method.

DLA-Future is implemented in `C++` using `std::execution`, a concurrency API which has been accepted for standardization in `C++26`. As `std::execution` is not yet available in standard library implementations, we based the implementations on pika [4], which provides the proposed `std::execution` API implemented on top of a thread-pool runtime. pika also integrates asynchronous MPI and GPU kernel launches with the `std::execution` API.

DLA-Future employs task-based variants of the aforementioned algorithms. Significant benefits arise applying the task-based approach to linear algebra. Reduction of synchronization points compared to the fork-join method, automatic overlap of communication and computation, and algorithm co-scheduling when enough resources are available noticeably reduce the time to solution.

We present the results on different types of GPU systems starting from Piz Daint, moving to more recent architectures available in ALPS, the flagship supercomputer at CSCS (Nvidia A100 GPUs, AMD MI250x GPUs, Nvidia GH200 modules). First we present a comparison of DLA-Future against widely used eigensolver implementations. A second part of results presents the performance of the eigensolver in real applications. We present the case of CP2K which uses

the `Fortran` interface available in DLA-Future-Fortran [2] and Sirius which uses the `C`-API provided in DLA-Future.

References

[1] Raffaele Solcà, Alberto Invernizzi, Auriane Reverdell, John Biddiscombe, Mikael Simberg, Rocco Meli, et al. DLA-Future 0.6.0, June 2024.
`https://doi.org/10.5281/zenodo.12519197`

[2] Rocco Meli, John Biddiscombe, Alberto Invernizzi, Auriane Reverdell, Mikael Simberg, Raffaele Solcà, and Mathieu Taillefumier. DLA-Future-Fortran 0.2.0, July 2024.
`https://doi.org/10.5281/zenodo.12665896`

[3] Raffaele Solcà, Mikael Simberg, Rocco Meli, Alberto Invernizzi, Auriane Reverdell, and John Biddiscombe. DLA-Future: A Task-Based Linear Algebra Library Which Provides a GPU-Enabled Distributed Eigensolver. In Patrick Diehl, Joseph Schuchart, Pedro Valero-Lara, and George Bosilca, editors, *Asynchronous Many-Task Systems and Applications*, pages 135–141, Cham, 2024. Springer Nature Switzerland.
`https://doi.org/10.1007/978-3-031-61763-8_13`

[4] Mikael Simberg, Auriane Reverdell, John Biddiscombe, Hartmut Kaiser, Akhil Nair, Bhumit Attarde, Severin Strobl, Hannes Vogt, Dimitra Karatza, Rocco Meli, et al. pika-org/pika: pika 0.26.0, July 2024.
`https://doi.org/10.5281/zenodo.12687187`