12^{th} Workshop on Parallel-in-Time Integration

July 17th - 21th 2022

Hamburg University of Technology

Plenary Speakers:

Iryna Kulchytska-Ruchka (Robert Bosch GmbH)
Colin Cotter (Imperial College London)
Frédéric Legoll (École Nationale des Ponts & Chaussées, INRIA)
Tommaso Buvoli (Tulane University)

Local Organizers : Daniel Ruprecht, Philipp Neumann, Thibaut Lunet,
Marco Wolkner, Peter Baasch, Felix Hoppe

More information:
`https://www.mat.tuhh.de/veranstaltungen/pint2023/`

We thankfully acknowledge support from our sponsors, who made it possible to host the conference without charging a registration fee.

- Technische Universität Hamburg (TUHH)
- Helmut-Schmidt Universität Hamburg (HSU)
- European High-Performance Computing (EuroHPC) Joint Undertaking (JU) under grant agreement No 955701
- Bundesministerium für Bildung und Forschung (BMBF) unter Förderkennzeichen 16HPC048
- Zentrum für Digitalisierungs- und Technologieforschung der Bundeswehr dtec.bw
- Deutsche Forschungsgemeinschaft (DFG) Projekt 525695584
- Lawrence Livermore National Laboratory (LLNL)

Contents

Programme

Monday, July 17th 2023

12:00 – 13:45	Registration
12:00 – 14:00	Lunch (Mensa)
13:45 – 14:00	**Welcome** (Building H, Room 0.16)

Contributed Talks (Building H, Room 0.16)	
Chair: Thibaut Lunet	
14:00 – 14:30	**Timothy Andrews** *Triadic initial conditions for testing a phase-averaged coarse propagator for Parareal*
14:30 – 15:00	**Benedict Philippi** *The Micro-Macro Parareal Algorithm Applied to FESOM2*
15:00 – 15:30	**Judith Angel** *PDE-constrained Optimisation with Asymptotic Parareal*
15:30 – 16:00	Coffee Break
16:00 – 16:30	**Sriramkrishnan Muralikrishnan** *Parallel-in-time integration with Parareal for kinetic plasma simulations*

16:30 – 18:00	Discussions (anyplace, anywhere)

Tuesday, 18 July 2023

Plenary (Building H, Room 0.16)	
Chair: Sebastian Götschel	
09:00 – 10:00	**Colin Cotter** *Parallel numerics for phase averaging; phase averaging for parallel numerics*
10:00 – 10:30	Coffee Break

Contributed Talks (Building H, Room 0.16)	
Chair: Sebastian Götschel	
10:30 – 11:00	**Ignace Bossuyt** *Micro-macro Parareal algorithms for weak solutions of McKean-Vlasov SDEs: numerical convergence analysis and comparison*
11:00 – 11:30	**Iñigo Jimenez Ciga** *Space-time parallel methods based on parareal algorithm for evolutionary problems*
11:30 – 12:00	**Arne Bouillon** *Multilevel Monte Carlo interacting-particle methods for Bayesian inversion*

12:00 – 14:00	Lunch (Mensa)

Contributed Talks (Building H, Room 0.16)	
Chair: Juliane Rosemeier	
14:00 – 14:30	**Aušra Pogoželskytė** *Optimizing Space-Time Multigrid with Block-Jacobi Smoother*
14:30 – 15:00	**Christoph Lohmann** *On the design of global-in-time Navier-Stokes solvers*
15:00 – 15:30	**Jonas Dünnebacke** *Space-time multigrid methods for stabilized convection-diffusion equations arising from flow problems*
15:30 – 16:00	Coffee Break

15:30 – 16:30	**PinT Workshop Scientific Committee Meeting** Building H, Room 0.10

16:00 – 18:00	**Postersession** Building H, Foyer

Wednesday, 19 July 2023

Attention: Please note the change of venue on Wednesday.

Plenary (Building I, Audimax II)	
Chair: Jemma Shipton	
09:00 – 10:00	**Iryna Kulchytska-Ruchka** *E-machine design: simulation and optimization approaches*
10:00 – 10:30	Coffee Break

Contributed Talks (Building I, Audimax II)	
Chair: Jemma Shipton	
10:30 – 11:00	**Andreas Schafelner** *Multigrid in Time for Time-Periodic Parabolic Evolution Problems*
11:00 – 11:30	**Pedro Peixoto** *Towards Exponential Semi-Lagrangian Parallel-in-Time Methods for the Shallow Water Equations on the Rotating Sphere*
11:30 – 12:00	**Jacob Schröder** *Multigrid Reduction in Time for Chaotic Problems*

12:00 – 14:00	Lunch (Mensa)

Contributed Talks (Building I, Audimax II)	
Chair: Pedro Peixoto	
14:00 – 14:30	**Ryo Yoda** *Multilevel extension of parameterized Runge-Kutta methods for Multigrid Reduction in Time*
14:30 – 15:00	**Wiebke Drews** *Numerical Analysis of a Time-Simultaneous Multigrid Solver for Stabilized Convection-Dominated Transport Problems*
15:00 – 15:30	**Robert Falgout** *Multigrid Reduction in Time for Nonlinear Hyperbolic Problems*
15:30 – 16:00	Coffee Break

16:00 – 18:00	**Training** Building H, Rooms 0.05 and 0.06
16:00 – 18:00	**Time-X Annual Meeting** Building I, WiKomm

Thursday, 20 July 2023

Attention: The bus for for the Airbus tour leaves at 13:30 from TUHH campus.

Plenary (Building H, Room 0.16)	
Chair: Jacob Schroder	
09:00 – 10:00	**Frederic Legoll** *Adaptive parareal algorithms: application to molecular dynamics and the simulation of atomistic defects*
10:00 – 10:30	Coffee Break

Contributed Talks (Building H, Room 0.16)	
Chair: Jacob Schroder	
10:30 – 11:00	**Jun Cao** *Parallel Immersed Boundary Lattice Boltzmann Computation for Prediction of Tornado Disaster on Multiple-Building Configuration*
11:00 – 11:30	**Werner Bauer** *Higher order phase averaging for highly oscillatory systems*
11:30 – 12:00	**Martin J. Gander** *Unmapped Tent Pitching for the time parallel solution of hyperbolic problems*

12:00 – 13:30	Lunch (Mensa)

13:30 – 18:00	**Social Event** Airbus factory tour

From 18:30	**Workshop Dinner** at Lieger Caesar

Friday, 21 July 2023

Plenary (Building H, Room 0.16)	
Chair: Abdelouahed Ouardghi	
09:00 – 10:00	**Tommaso Buvoli** *Exponential Integrators, Parareal, and Polynomial Time Integrators*
10:00 – 10:30	Coffee Break

Contributed Talks (Building H, Room 0.16)	
Chair: Abdelouahed Ouardghi	
10:30 – 11:00	**Lisa Wimmer** *Discontinuity Handling for Spectral Deferred Corrections*
11:00 – 11:30	**Josh Hope-Collins** *ParaDiag methods for nonlinear problems in atmospheric modelling*
11:30 – 12:00	**Abdul Qadir Ibrahim** *Parareal with a physics informed neural network as coarse propagator*

12:00 – 12:15	**Closing Remarks** (Building H, Room 0.16)

12:15 – 14:00	Lunch (Mensa)

Plenary Talks

Parallel numerics for phase averaging; phase averaging for parallel numerics

Colin Cotter
Imperial College London
Tuesday, July 18th, 09:00 – 10:00 (Building H, Room 0.16)

In highly oscillatory problems, such as those arising in large scale models of the atmosphere and ocean, solutions often consist of a slowly evolving component ("weather") perturbed with fast but weak amplitude oscillations ("waves"). Phase averaging is an analytical technique for producing approximate equations that aim to predict the evolution of the fast component only. It is built around the use of operator exponentials of the fast linear operator. Whilst this technique has been used for deriving and analysing model equations (in concert with asymptotic limits), Haut and Wingate (2014) proposed for the first time to compute phase averaged equations numerically, distributing the independent terms in the average in parallel to obtain a computationally feasible method. In this talk I will introduce the phase averaging technique via results of experiments using rotating shallow water equations on the sphere. These results demonstrate that the error due to the averaging can be smaller than the error due to time integration, showing that we can interpret phase averaging in its own right. Then I will introduce a new formulation of phase averaging which interlaces parallel computation of the operator exponentials with parallelisation over the average, making use of the "ParaDiag" technique for solving all-at-once systems for linear PDEs with constant coefficients. I will discuss the block diagonal problems that arise from this technique in the case of rotating shallow water equations (serving as a prototype for 3D geophysical fluid models), and will present the latest numerical results from this approach as they become available. Finally I will briefly discuss the computation of higher order corrections to the phase averaging, and their possible potential for use in time parallel methods.

E-machine design: simulation and optimization approaches

Iryna Kulchytska-Ruchka
Robert Bosch GmbH
Wednesday, July 19th, 09:00 – 10:00 (Building I, Audimax II)

This talk describes the modern industrial e-machine design process with focus on numerical simulation and optimization. Based on customer's requirements, it covers different physical domains such as electromagnetics, mechanics, or thermodynamics, and is incorporated in a multi-objective optimization framework. Additionally, behavioral models for e-drive (i.e., e-machine together with the connected power electronics) and several parallel-in-time (PinT) approaches tailored to such problem settings are presented. The talk is summarized with a list of possible applications of the PinT methods to industrial e-machine design.

Adaptive parareal algorithms: application to molecular dynamics and the simulation of atomistic defects

Frederic Legoll
École Nationale des Ponts & Chaussées, INRIA
Thursday, July 20th, 09:00 – 10:00 (Building H, Room 0.16)

In this talk, we consider parareal algorithms in the context of molecular dynamics simulations, that is long-time simulations of Stochastic Differential Equations. Although the parareal algorithm, in its original formulation, always converges, it suffers from various limitations in the context of molecular dynamics. In particular, it is observed that the algorithm does not provide any computational gain (in terms of wall-clock time compared to a standard sequential integration) in the limit of increasingly long time-horizons. This numerical observation is backed up with theoretical discussions. We introduce here a modified version of the parareal algorithm where the algorithm adaptively divides the entire time-horizon into smaller time slabs. We numerically show that the adaptive algorithm overcomes the various limitations of the standard parareal algorithm, thereby allowing for significantly improved computational gains. Several numerical examples (based on the Langevin equation) will be discussed. In particular, we consider the simulation using the LAMMPS software of defects diffusing in a tungsten lattice, a test-case for which key quantities of interest are of statistical nature. This talk is based on joint works with O. Gorynina, T. Lelievre, D. Perez and U. Sharma.

Exponential Integrators, Parareal, and Polynomial Time Integrators

Tommaso Buvoli
Tulane University
Friday, July 21th, 09:00 – 10:00 (Building H, Room 0.16)

Exponential integrators are a class of time integration methods for efficiently solving stiff ordinary differential equations. A key characteristic of these methods is that they treat a linear component exactly while approximating all remaining terms explicitly. In this talk, I will present two families of exponential PinT methods: the first are Parareal methods with exponential Runge-Kutta integrators and the second are parallel, multivalued integrators that are constructed from a recently introduced polynomial time integration framework. The Parareal methods highlight large-scale time parallelism while the polynomial integrators possess more modest parallelism capabilities, albeit with higher parallel efficiency. I will discuss the performance of these methods for solving non-diffusive partial differential equations and broadly highlight the advantages of exponential integration for solving these types of problems in parallel.

Contributed Talks

Triadic initial conditions for testing a phase-averaged coarse propagator for Parareal

Timothy Andrews
University of Exeter
Monday, July 17th, 14:00 – 14:30 (Building H, Room 0.16)

When developing the next generation of weather and climate models, it is important to ensure that these meet the operational wall-clock time limits. One method to achieve this is through a parallel-in-time implementation. However, the geophysical fluid equations are challenging to solve accurately, as they contain both fast linear oscillations and slower nonlinear interactions. A set of 'triadic' test cases were recently developed to highlight this feature in the Rotating Shallow Water Equations (RSWEs) – a simpler version of the full model which still retains the complexity of multiple timescale dynamics. These initial conditions excite specific fast and slow linear waves to generate nonlinear triadic interactions. It was found that large time-steps with standard serial methods, used to reduce the wall-clock time, often miss the slower dynamics created by these triads. This talk will discuss the construction of a coarse propagator that uses phase averaging in this RSWE system. This is one approach to ensure convergence for Parareal, but is dependent on the averaging window length – there exists an optimal window to balance time-stepping and averaging errors. Using this triadic test case provides intuition around the window length selection for the coarse propagator. It also tests how different averaging windows affect the replication of the fast and slow components of the multiscale dynamics.

The Micro-Macro Parareal Algorithm Applied to FESOM2

Benedict Philippi
CAU Kiel
Monday, July 17th, 14:30 – 15:00 (Building H, Room 0.16)

The ocean circulation model FESOM2 represents a state-of-the-art solver in climate research. The talk is intended to give an overview of the technical challenges that can be expected when attempting to apply a parallel-in-time algorithm. With the presentation of the results the talk concludes with a discussion of whether and how an efficient application of Parareal could be achieved.

PDE-constrained Optimisation with Asymptotic Parareal

Judith Angel
Hamburg University of Technology
Monday, July 17th, 15:00 – 15:30 (Building H, Room 0.16)

Inverse problems, for example reconstructing the bottom topography of a river from measurements of the water height at the surface, can be approached as optimisation problems with partial differential equations (PDEs) as constraints. These kind of problems can be computationally expensive to solve since they require a large number of PDE solutions. In the "First optimise, then discretise"-approach we have two PDEs to solve at each step of our optimisation algorithm, the forward problem and the adjoint equation. In order to accelerate the optimization procedure, we propose to use Parareal to solve the forward and adjoint problem faster. We will show and discuss a range of results for the bathymetry reconstruction case where the forward problem is given by the shallow water equations. Additionally, we will discuss the use of asymptotic Parareal to improve performance of Parareal for the forward model, since vanilla Parareal has well documented problems for hyperbolic PDEs. Finally, we want to take advantage of the linearity of the adjoint equations, where an exponential propagation could be a useful approach. The talk will examine the benefits and performance of the combination of gradient descent and different variants of Parareal and show numerical results.

Parallel-in-time integration with Parareal for kinetic plasma simulations

Sriramkrishnan Muralikrishnan
Jülich Supercomputing Centre
Monday, July 17th, 16:00 – 16:30 (Building H, Room 0.16)

Particle-In-Fourier (PIF) schemes are attractive for long-time integration kinetic plasma simulations as they conserve charge, momentum and energy, exhibit a variational structure, do not have aliasing and have excellent stability properties. However, they are typically more expensive than the commonly used Particle-In-Cell (PIC) schemes due to the requirement of non-uniform DFTs or FFTs. In this talk, we propose a Parareal-based parallel-in-time integration method for PIF schemes with PIF as the fine propagator and standard PIC scheme as the coarse propagator towards the goal of performing long-time integration simulations with PIF schemes. The resulting scheme is implemented on the performance portable library IPPL and we numerically investigate the convergence of it with respect to the discretization parameters of PIF and PIC. We present space-time parallel simulations with the proposed scheme for a variety of benchmarks such as the Landau damping, the two-stream instability and the Penning trap on A100 multi-GPUs on the JUWELS Booster supercomputer.

Micro-macro Parareal algorithms for weak solutions of McKean-Vlasov SDEs: numerical convergence analysis and comparison

Ignace Bossuyt
KU Leuven
Tuesday, July 18th, 10:30 – 11:00 (Building H, Room 0.16)

Time-parallel methods can reduce the wall clock time required for the numerical solution of a McKean-Vlasov Stochastic Differential Equation (SDE) by parallelizing across the time-dimension. In this talk, we study the convergence behavior a of a micro-macro Parareal method for McKean-Vlasov SDEs: the fine propagator of the SDE is based on a high-dimensional microscopic model; the coarse propagator uses a deterministic moment ODE (and thereby possesses a model error with respect to the underlying stochastic dynamics). We then compare the convergence of the 2-level method with the convergence of a 3-level micro-macro Parareal method, where the coarse propagators are derived from a hierarchy of reduced models that capture the low-dimensional effective dynamics such as the moment ODE. We investigate numerically how the model error from the approximate models influences the convergence of this algorithm.

Space-time parallel methods based on parareal algorithm for evolutionary problems

Iñigo Jimenez Ciga
Public University of Navarre (UPNA)
Tuesday, July 18th, 11:00 – 11:30 (Building H, Room 0.16)

In a context of increasingly complex computer architectures, parallel programming has become a remarkable tool for the solution of problems of different nature. Regarding differential equations, the concept of parallelization can be implemented for integrating evolutionary problems, in which space and time variables are involved. As a first approximation, parallelization in space can be suitable for certain problems, but, if a large number of calculation cores is available, time parallel time integrators are also required. In this framework, we present new space-time parallel methods based on the parareal algorithm. This time parallel time integrator considers two propagators, one of them cheap and fast and the other one expensive but more accurate that is used in a parallel way over the time windows. The algorithm is combined with time-splitting schemes. These integrators can be parallelized in space, thus yielding combined methods which are parallel in space and time. The resulting algorithms permit us to integrate parabolic problems significantly faster by optimizing the use of available connected CPUs. We show stability and convergence properties for the methods, remarking the importance of choosing L-stable methods in order to obtain stable space-time parallel schemes with faster convergence. Along with theoretical results, we propose some numerical experiments to confirm the power of the designed integrators, illustrating their potential when very fine meshes in time and space are needed.

Multilevel Monte Carlo interacting-particle methods for Bayesian inversion

Arne Bouillon
KU Leuven
Tuesday, July 18th, 11:30 – 12:00 (Building H, Room 0.16)

The solution of Bayesian inverse problems is classically a process for which sequential algorithms such as Markov chain Monte Carlo are used. Recently, there has been a growing interest in methods that use ensembles of interacting particles, moving around in parameter space. These particles follow dynamics carefully designed to move towards an approximation of the posterior distribution, or of the maximum a posteriori (MAP) estimate. These interacting-particle methods can be gradient-free and are trivially parallelisable. For the case where the forward model is expensive and a hierarchy of cheaper approximations is available, we present an adaptation of these methods based on the multilevel Monte Carlo (MLMC) idea. Inspired by the multilevel Kalman filter, our algorithm can yield faster convergence than the single-level variant. We discuss the conditions on the interacting-particle method that are needed for good multilevel performance.

Optimizing Space-Time Multigrid with Block-Jacobi Smoother

Aušra Pogoželskytė
University of Geneva
Tuesday, July 18th, 14:00 – 14:30 (Building H, Room 0.16)

For time-dependent problems, Parallel-in-Time (PinT) algorithms allow us to parallelize problems in the time dimension when space parallelization alone creates communication bottlenecks. Parareal and Multigrid Reduction-in-Time (MGRIT) are two examples of such PinT algorithms based on multigrid techniques, but they are not truly scalable since they coarsen the problem only in the time dimension. We will focus on a more intrusive method: the Space-Time Multigrid algorithm with block-Jacobi relaxation introduced by Gander and Neumüller. This algorithm provides excellent scalability for parabolic problems up to millions of cores, while still being equally fast as forward substitution on one core only. We will show that the performance of this algorithm can be further improved by the optimization of the smoothing parameters. This will allow the algorithm to be up to twice as fast as the original one. We will also discuss coarsening strategies that would allow us to always coarsen in space and in time. Joint work with Bastien Chaudet-Dumas and Martin Gander.

On the design of global-in-time Navier-Stokes solvers

Christoph Lohmann
TU Dortmund
Tuesday, July 18th, 14:30 – 15:00 (Building H, Room 0.16)

The work to be presented in this talk focuses on the design of a new global-in-time multigrid solution strategy for incompressible flow problems, which highly exploits the Pressure Schur complement (PSC) approach and provides the possibility to use massively parallelizable solution components [1]. For linear problems like the incompressible Stokes equations discretized in space using an inf-sup-stable finite element pair, the fundamental idea is to block the linear systems of equations associated with individual time steps into a single all-at-once saddle point problem for all velocity and pressure unknowns. Then the Pressure Schur complement can be used to eliminate the velocity fields and set up a linear system for all pressure variables only. This algebraic manipulation allows the construction of parallel-in-time preconditioners for the corresponding all-at-once Picard iteration by extending frequently used sequential PSC preconditioners in a straightforward manner (cf. [2]). We show that those preconditioners can be applied very efficiently on modern high performance computing facilities and are asymptotically exact in the limit of vanishing time increments. To accelerate the convergence of the proposed fixed-point iteration, this iterative solver is embedded as a smoother into a space-time multigrid algorithm, where the computational complexity of the coarse grid problem highly depends on the coarsening strategy in space and/or time. While coarsening in space using commonly used FE intergrid transfer operators significantly reduces the size of the space-time problem, most promising results for convection-dominated problems could be obtained by only coarsening in time using tailor-made prolongation and restriction operators. This procedure even allows the efficient solution of the incompressible Navier-Stokes equations using a nonlinear viscosity model for many time steps by employing Newton's method for linearization. At the end, the presented multigrid solution strategy only requires the solution of time-dependent linear convection-diffusion-reaction equations and many Poisson problems, which both can be performed efficiently by using algorithms that exploit massive parallelism. The potential of this approach for CFD simulations with large time intervals is illustrated in numerical examples.

[1] Ch. Lohmann and St. Turek, On the design of global-in-time Newton-Multigrid-Pressure Schur complement solvers for incompressible flow problems. Journal of Mathematical Fluid Mechanics (accepted).

[2] F. Danieli, B. S. Southworth, and A. J. Wathen, Space-Time Block Preconditioning for Incompressible Flow. SIAM Journal on Scientific Computing, 44.1, 2022.

Space-time multigrid methods for stabilized convection-diffusion equations arising from flow problems

Jonas Dünnebacke
TU Dortmund
Tuesday, July 18th, 15:00 – 15:30 (Building H, Room 0.16)

Usually time dependent evolution equations are solved time step by time step where in each step a system of equations corresponding to the spatial discretization has to be solved. Such methods can only be parallelized in space, but the size of the spatial problem limits the strong scaling behavior. Using multigrid methods that treat multiple time steps in an all-at-once system the parallel scaling can be improved significantly in comparison to geometric multigrid solvers in a time stepping application [2, 4]. Due to the improved communication pattern between the parallel processes, this holds true even if a time-simultaneous multigrid method without temporal parallelization is used [1]. Here, we numerically analyze how such multigrid methods behave for convection-diffusion(-reaction) equations found in flow problems, when the problems are increasingly transport dominated. Then, we show how stabilization techniques, e.g. the variational multiscale method [3], can remedy the emerging problems depending on the discretization of the problem. Furthermore, we show the behavior of such methods for problems with space and time dependent diffusion and convection parameters, that arise in global-in-time solution strategies of the Navier-Stokes equations.

[1] J. Dünnebacke, S. Turek, C. Lohmann, A. Sokolov, P. Zajac (2021). Increased space-parallelism via time-simultaneous newton-multigrid methods for nonstationary nonlinear pde problems. Int. J. High Perform. Comput. Appl., 35(3), 211–225.

[2] M.J. Gander, M. Neumüller (2014). Analysis of a New Space-Time Parallel Multigrid Algorithm for Parabolic Problems SIAM J. Sci. Comput., 38(4), A2173 - A2208.

[3] V. John, S. Kaya and W. Layton (2006). A two-level variational multiscale method for convection- dominated convection-diffusion equations. Comput. Methods Appl. Mech. Eng., 195(33-36), 4594–4603.

[4] C. Lubich, A. Ostermann (1987). Multi-grid dynamic iteration for parabolic equations. BIT, 27(2), 216–234.

Multigrid in Time for Time-Periodic Parabolic Evolution Problems

Andreas Schafelner
Johannes Kepler University Linz (Building I, Audimax II)
Wednesday, July 19th, 10:30 – 11:00

We present a parallel-in-time multigrid method for the solution of time-periodic parabolic evolution problems, with possibly nonlinear coefficients. After discretizing with finite elements in space and with backward Euler in time, we solve the resulting linear block-system of equations by means of a multigrid in time method. As a smoother we use a block-Jacobi method, which can be applied in parallel. We present numerical experiments based on applications in magneto-quasistatics.

Towards Exponential Semi-Lagrangian Parallel-in-Time Methods for the Shallow Water Equations on the Rotating Sphere

Pedro Peixoto
University of São Paulo
Wednesday, July 19th, 11:00 – 11:30 (Building I, Audimax II)

The hyperbolic nature and the development of small-scale spatial features in mathematical models for atmospheric circulation constitute major challenges for parallel-in-time (PinT) numerical simulations in climate modeling and numerical weather prediction. Therefore, it is of great importance to develop coarse time stepping schemes that are able to improve the stability and convergence of PinT methods. In this work, we explore the application of semi-Lagrangian exponential methods for the two- and multilevel parallel-in-time simulation of the shallow water equations (SWE) on the rotating sphere, using Parareal and MGRIT. As a main feature, exponential methods (e.g. ETDRK) solve exactly the linear terms of the governing equations, and their recently proposed semi-Lagrangian version (SL-ETDRK) was shown to be more stable in the serial integration of the SWE on the plane. We then consider these methods on the sphere and in the PinT context. We perform stability studies on a linearized ODE which indicate that Parareal and MGRIT using ETDRK and SL-ETDRK as coarse schemes have improved stability properties compared to the use of the well-known SL-SI-SETTLS method, which turns out to be highly unstable in the PinT framework despite its successful application in serial simulations of atmospheric circulation. Numerical tests of standard benchmarks confirm that exponential schemes, notably in the semi-Lagrangian framework, are promising choices for parallel-in-time integration.

Multigrid Reduction in Time for Chaotic Problems

Jacob Schröder
University of New Mexico
Wednesday, July 19th, 11:30 – 12:00 (Building I, Audimax II)

In this talk, we consider the parallel-in-time method, multigrid-reduction-in-time (MGRIT). MGRIT adds parallelism to the time dimension by applying multigrid to the (non)linear systems that arise when solving for multiple time steps simultaneously. A hierarchy of successively coarser time-grids are used to provide error corrections on the finest level. MGRIT, like many parallel-in-time methods, is known to be effective for parabolic problems, but to struggle for some other problem types, including chaotic problems. This is because chaotic initial value problems are inherently ill-conditioned, which is problematic for MGRIT and other multilevel methods, because small and subtle inaccuracies on the coarser time-grids interact with the ill-conditioning and result in poor coarse-grid corrections. Here, we propose a modified MGRIT coarse time-grid equation that allows for better agreement between fine and coarse levels for chaotic modes. We additionally propose a new time-coarsening scheme which more accurately captures long-term chaotic behavior on coarse time-grids. Supporting numerical results with parallel speedup for the Lorenz system, the Kuramoto–Sivashinsky equation, and computational fluids are provided. Lastly, we detail the library implementation of this new feature in the MGRIT package, XBraid.

Multilevel extension of parameterized Runge-Kutta methods for Multigrid Reduction in Time

Ryo Yoda
University of Wuppertal
Wednesday, July 19th, 14:00 – 14:30 (Building I, Audimax II)

This work discusses the construction method of coarse-grid time integrators for the multigrid reduction in time (MGRIT), one of the multigrid-based parallel-in-time methods. Conventional construction methods are based on the rediscretization approach, but these are known to have poor convergence for hyperbolic problems, and new construction methods must be considered. As one of the directions, we deal with the parameterized Runge-Kutta method, a time integrator construction method based on the Runge-Kutta method using the coefficients of the Butcher table as parameters. A previous study has proposed a method dedicated to MGRIT convergence using the two-level convergence rate based on the reduction analysis. This work considers a multilevel extension of the parameterized Runge-Kutta method using the convergence rate of the multilevel version. Although the number of parameters to be determined at a time increases compared to the method that applies the two-level method recursively, this method is expected to provide good convergence.

Numerical Analysis of a Time-Simultaneous Multigrid Solver for Stabilized Convection-Dominated Transport Problems

Wiebke Drews
TU Dortmund
Wednesday, July 19th, 14:30 – 15:00 (Building I, Audimax II)

The work to be presented in this talk focuses on the one-dimensional convection-diffusion equation, especially in the regime of small diffusion coefficients, which is solved using a time-simultaneous multigrid algorithm closely related to multigrid waveform relaxation [2]. For spatial discretization we use linear finite elements, while the time integrator is given by e.g. the Crank-Nicolson scheme. Blocking all time steps into a global linear system of equations and rearranging the degrees of freedom leads to a space-only problem with vector-valued unknowns for each spatial node. Then, common iterative solution techniques, such as the block Jacobi method or the preconditioned GMRES method, can be used for the numerical solution of the (spatial) problem and allow parallelization in space. We consider a time-simultaneous multigrid algorithm, which exploits space-only coarsening and the solution techniques mentioned above for smoothing purposes. By treating more time steps simultaneously, the dimension of the system of equations increases significantly and, hence, results in a larger number of degrees of freedom per spatial unknown. This can be used to employ parallel processes more efficiently [1]. In numerical studies, the iterative multigrid solution of a problem with up to thousands of blocked time steps is analyzed. For the special case of the heat equation, it is well known that the number of iterations is bounded from above independently of of the number of blocked time steps, the time step size, and the spatial resolution. Unfortunately, stability problems arise for the standard Galerkin method if the diffusion coefficient is small compared to the grid size and the magnitude of the velocity field. Therefore, the influence of VMS-type stabilization techniques [3] is discussed, which remove artificial oscillations in the solution and aim to improve the convergence behavior of the iterative solution algorithm.

[1] J. Dünnebacke, S. Turek, C. Lohmann, A. Sokolov, P. Zajac (2021). Increased space-parallelism via time-simultaneous Newton-multigrid methods for nonstationary nonlinear PDE problems. The International Journal of High Performance Computing Applications 35.3, pp. 211-225.

[2] J. Janssen, S. Vandewalle (1996). Multigrid Waveform Relaxation on Spatial Finite Element Meshes: The Continuous-Time Case. SIAM Journal on Numerical Analysis 33.2, pp. 456-474.

[3] V. John, S. Kaya, W. Layton (2006). A two-level variational multiscale method for convection-dominated convection-diffusion equations. Computer Methods in Applied Mechanics and Engineering 195.33, pp. 4594-4603.

Multigrid Reduction in Time for Nonlinear Hyperbolic Problems

Robert Falgout
Lawrence Livermore National Laboratory
Wednesday, July 19th, 15:00 – 15:30 (Building I, Audimax II)

Hans De Sterck, Robert D. Falgout, Oliver A. Krzysik, and Jacob B. Schroder

The multigrid reduction in time (MGRIT) method is a parallel multigrid-in-time solver designed to be as non-intrusive as possible and take advantage of existing simulation codes and techniques. This has worked well for parabolic equations, but parallel-in-time methods for advection-dominated hyperbolic problems have proven difficult to develop. In previous work, we demonstrated the effectiveness of a modified semi-Lagrangian coarse-grid operator for speeding up the parallel solution of high-order discretizations of variable-wave-speed linear advection problems in both 1D and 2D. In this talk, we will present progress extending this technique for solving nonlinear hyperbolic conservation laws such as the inviscid Burgers and Buckley-Leverett equations.

Parallel Immersed Boundary Lattice Boltzmann Computation for Prediction of Tornado Disaster on Multiple-Building Configuration

Jun Cao
Toronto Metropolitan University
Thursday, July 20th, 10:30 – 11:00 (Building H, Room 0.16)

Tornadoes are disastrous, naturally recurrent atmospheric phenomena. The speed and direction of a tornado may change rapidly. The rotational and translational velocity combination makes up the tornadoes more dreadful. In numerical simulation of tornadoes, the velocity at the outer boundary of the computational domain needs to be unceasingly updated since this type of combined flow appears time-dependent everywhere. In order to get rid of the computationally painful boundary-condition-updating process, the renowned Rankine Combined Vortex Model (RCVM) is re-tailored, in the present study, by using the "relative motion" principle, such that the tornado-building interaction is viewed as the buildings "virtually moving" towards a "center-virtually-pinned" tornado. As a result of this re-interpretation, the flow located sufficiently distant from the "locked" tornado center can be considered purely rotational, which is time-independent, and the corresponding velocity boundary condition no longer needs updating. Accordingly, the immersed boundary (IB) approach is utilized to model the virtual movement of the buildings, and the lattice Boltzmann method (LBM) is employed to find the fluid properties in the computational domain. The present RCVM based IB-LBM framework is applied to investigate a tornadic flow over a complex configuration consisting of five cylinders. The wind loadings such as the force coefficients in the x- and y-directions along with the moment coefficient are examined to identify the maximum wind loadings and, hence, the wind-resistant capability standard for the buildings under investigation.

Higher order phase averaging for highly oscillatory systems

Werner Bauer
University of Surrey, Department of Mathematics
Thursday, July 20th, 11:00 – 11:30 (Building H, Room 0.16)

We introduce a higher order phase averaging method for nonlinear oscillatory systems. Phase averaging is a technique to filter fast motions from the dynamics whilst still accounting for their effect on the slow dynamics. Phase averaging is useful for deriving reduced models that can be solved numerically with more efficiency, since larger timesteps can be taken. Haut and Wingate (2014) introduced the idea of computing finite window numerical phase averages in parallel as the basis for a coarse propagator for a parallel-in-time algorithm. Here, we provide a framework for higher order phase averages that aims to better approximate the unaveraged system whilst still filtering fast motions. Whilst the basic phase average assumes that the solution is independent of changes of phase, the higher order method expands the phase dependency in a basis which the equations are projected onto. We illustrate the properties of this method on an ODE that describes the dynamics of a swinging spring due to Lynch (2002) and explore different choices of a basis. Although idealized, this model shows an interesting analogy to geophysical flows as it exhibits a slow dynamics that arises through the resonance between fast oscillations. On this example, we show convergence to the non-averaged (exact) solution with increasing approximation order also for finite averaging windows. At zeroth order, our method coincides with a standard phase average, but at higher order it is more accurate in the sense that solutions of the phase averaged model track the solutions of the unaveraged equations more accurately. This is a collaboration with Colin Cotter (Imperial College London, UK) and Beth Wingate (University of Exeter, UK).

Unmapped Tent Pitching for the time parallel solution of hyperbolic problems

Martin J. Gander
University of Geneva
Thursday, July 20th, 11:30 – 12:00 (Building H, Room 0.16)

Mapped Tent Pitching is an advanced domain decomposition method for the time parallel solution of general hyperbolic problems, introduced in 2017 by Jay Gopalakrishnan, Joachim Schöberl and Christoph Wintersteiger. Mapped Tent Pitching computes the solution of the hyperbolic problem by constructing new polygonal space-time subdomains, called tents, in a way that the hyperbolic problem can be solved exactly within them, due to the finite speed of propagation of hyperbolic problems. Using this specific property, it is currently among the very best time parallel solvers for hyperbolic problems, and spectacular performance was demonstrated in 2020 for three dimensional Maxwell equations in space and time: 49.2 seconds for solving a problem with 7.632e9 unknowns in a time slab on 64 cores using a method of order 3. Due to the polygonal space-time structure of the subdomains, the numerical solution is obtained by a process that maps the tents into space-time cylinders, computes the solution in the transformed subdomains, and maps it back into the original tents, which led to the name Mapped Tent Pitching. The tent mapping leads however to singularities, and special time integrators are needed to mitigate order reduction. I will show in my presentation how this mapping can be avoided using waveform relaxation techniques, leading to a new, Unmapped Tent Pitching algorithm which computes the same solution as the Mapped Tent Pitching algorithm at the continuous level. Unmapped Tent Pitching is much easier to implement than Mapped Tent Pitching, one can use all standard time integrators since there is no order reduction, and the computational cost is reduced since there is no mapping to be computed, while the volume of computation on the time cylinders remains the same. The trick is to use redundancy in computations. This presentation is joint work with Gabriele Ciaramella and Illario Mazzieri.

Discontinuity Handling for Spectral Deferred Corrections

Lisa Wimmer
Bergische Universität Wuppertal
Friday, July 21th, 10:30 – 11:00 (Building H, Room 0.16)

Differential equations arise naturally in many fields to describe physical systems or processes. Discontinuities are also belong to such a system. In simulation, such discrete events can make the computation of a solution to a very difficult task. Even when the occurence of a discontinuity depends on the system dynamic itself, numerical methods may have problems calculating an accurate solution. Often they occur in fields such as science and engineering. In engineering, converters use high-frequency switching to compare after fixed number of time steps if the actual voltage is equal to a target voltage to control the output. In science, a gas-liquid model describes the output of a tube. The liquid comes out of the tube if the level of the liquid exceeds the dip tube, otherwise gas will leak out. In this presentation, the switch estimator applied to the method of spectral deferred corrections (SDC) is presented together with examples to show its performance. After few iterations the switch estimator predicts the time point of the discrete event using interpolation and root finding techniques, and the time step is adapted. Restarting SDC with the new time step the method is able to resolve the discontinuity more accurate.

ParaDiag methods for nonlinear problems in atmospheric modelling

Josh Hope-Collins
Imperial College London
Friday, July 21th, 11:00 – 11:30 (Building H, Room 0.16)

Weather and climate simulations use large amounts of compute power to solve nonlinear oscillatory models of atmospheric flow. The high number of timesteps required make them ideal candidates for parallel-in-time (PinT) methods. ParaDiag is a PinT algorithm that solves multiple timesteps simultaneously by preconditioning the system Jacobian with a block diagonalisable matrix, which can be solved efficiently in parallel. One advantage of ParaDiag over other PinT algorithms is its ability to handle oscillatory problems. However, for nonlinear problems the preconditioner must be linearised around a constant reference state, which is problematic when the solution varies significantly. Previous approaches have used the time average state and quasi-Newton iterations using the preconditioner. However, using preconditioned (quasi-)Newton-Krylov iterations gives greater freedom in constructing the method. We present a variety of options for the reference state and for which the terms are included in the Jacobian and preconditioner. These are implemented in our open-source library asQ using Firedrake, an automatic code generation library for finite elements. We compare their performance, including the number of Newton iterations, the Krylov iterations per linear solve, and the total performance of the nonlinear solve.

Parareal with a physics informed neural network as coarse propagator

Ibrahim Abdul Qadir
Hamburg University of Technology
Friday, July 21th, 11:30 – 12:00 (Building H, Room 0.16)

Parallel-in-time algorithms provide an additional layer of concurrency for the numerical integration of models based on time-dependent differential equations. Methods like Parareal, which parallelize across multiple time steps, rely on a computationally cheap and coarse integrator to propagate information forward in time, while a parallelizable expensive fine propagator provides accuracy. The talk will explore the use of machine learning-based coarse propagators in Parareal. We consider, as an example, the nonlinear Black-Scholes equation, which may be used to value financial options and to calculate implied volatilities. We will show that Parareal with a PINN coarse propagator provides comparable convergence as a numerical coarse propagator but is significantly faster to execute. This reduces the serial bottleneck posed by the coarse method and improves speedup obtained from Parareal. In single-node benchmarks, a Parareal with a physics-informed neural network (PINN) as coarse propagator produced more than double the speedup than Parareal with a numerical coarse method. We also show that the PINN coarse propagator can be moved to a GPU attached to the node to further reduce its runtime and improve Parareal speedup further.

Scientific Posters

ParaDiag meets PDE-constrained optimization

Gayatri Caklovic
Karlsruhe University of Technology
Tuesday, July 18th, 16:00 – 18:00 (Building H, Foyer)

Large-scale optimization problems governed by time-dependent partial differential equations (PDEs) occur in many applications, for example, inverse problems for non-destructive testing of materials and structures or optimal control problems related to individualized medicine. Algorithms for the numerical solution of such PDE-constrained optimization problems are computationally highly demanding, requiring multiple PDE solves during the iterative optimization process. With today's modern computers, time-to-solution can be decreased through massive parallelization, which is already traditionally done in the spatial dimensions. In addition, time-parallel methods have received increasing interest in recent years to overcome scaling limits. However, coarsening-based methods tend to give poor speedups for hyperbolic problems. ParaDiag is a diagonalization-based iterative time-parallel integration method that does not rely on coarsening, making it suitable for hyperbolic problems and a promising way to extend scaling. This poster presents ideas on how to utilize the method for PDE-constrained optimization problems.

The Multi-level Parareal Method with Averaging for Geophysical Problems

Juliane Rosemeier
University of Exeter
Tuesday, July 18th, 16:00 – 18:00 (Building H, Foyer)

The Multi-level Parareal method with Averaging shall be used to solve time evolution problems, which are characterised by a skew-hermitian linear operator and a bi-linear term. One example is the Rotating Shallow Water Equations. To reduce the oscillatory stiffness a transformation and an averaging technique are applied. With these analytical modifications it is possible to apply the Parareal method with big time-steps for the coarse propagator. Especially, a multi-level version of the Parareal method is considered in the presentation, as it can be used to solve multi-scale problems.

MaMiCo: A PinT Implementation for Molecular-Continuum Flow Simulation

Piet Jarmatz
Helmut-Schmidt Universität Hamburg
Tuesday, July 18th, 16:00 – 18:00 (Building H, Foyer)

On the PinT2022 workshop we have presented a concept that applies a modified Parareal variant to coupled molecular-continuum flow, using a continuum solver as a hydrodynamic predictor to supervise the microscopic system. In this poster, we present our implementation of the approach in the open-source C++ framework MaMiCo and show performance and scalability results for parallel-in-time coupled 3D flow simulations.

A Generic Tool for Preliminary Performance Analysis of Iterative Parallel-in-Time Methods

Thibaut Lunet
Hamburg University of Technology
Tuesday, July 18th, 16:00 – 18:00 (Building H, Foyer)

As the title says ... if that triggers your curiosity, come see the poster !

Project Posters

SWEET × LibPFASST

Keerthi Gaddameedi
TU München
Tuesday, July 18th, 16:00 – 18:00 (Building H, Foyer)

SWEET is a software that provides implementation of various PDEs and time integration methods. LibPFASST has recently been integrated into SWEET. The idea is to analyse results and convergence of PFASST algorithm for solving shallow water equations on rotating sphere and to further explore addition of hyperviscosity to the SWEs.

ExaOcean: Improving Performance of the ICON-O Oceanmodel on heterogeneous Exascale-Supercomputers with Machine Learning

Philip Freese
Hamburg University of Technology
Tuesday, July 18th, 16:00 – 18:00 (Building H, Foyer)

Ocean models are a key component of any weather or climate model. However, despite the computing power of modern supercomputers, important dynamical features like sub-mesoscale eddies can so far only be resolved in simulations over a few weeks. Simulations of climatologically relevant time scales are not yet feasible. ExaOcean will provide modern mathematical algorithms and integrate them into the ICON-O ocean model to achieve better parallel scaling and faster runtimes in highly resolved simulations on new supercomputers. We will integrate techniques from machine learning into mesh based algorithms, using data from high-resolution short-term simulations to train a correction term for long-term simulations with lower mesh resolution. This will enable "effectively sub-mesoscale resolving simulations" where the effect of the sub-mesoscale vortices on the larger scale dynamics is represented via the ML correction term, even though the mesh resolution is too coarse to represent the vortices explicitly. To maintain stability and mass conservation, the correction will be integrated into a new numerical time stepping method tailored to improve performance on many-core CPUs. Together with further improvements, such as online learning, we expect a significant reduction of runtimes for high-resolution ICON-O simulations. The ultimate goal of ExaOcean is to deliver simulations that correctly account for the effects of sub-mesoscale vortices of a climatologically relevant 30-year time span in less than 30 days of runtime.

HPC.BW Project

Philipp Neumann
Helmut-Schmidt Universität Hamburg
Tuesday, July 18th, 16:00 – 18:00 (Building H, Foyer)

This poster demonstrates our current progress of our project hpc.bw. The aim of our project is to identify HPC issues in used compute intensive software, establish HPC awareness, encourage interdisciplinary exchange and sustainably propagate HPC knowledge to the users. In the last year, we established our container-based HPC center, including the PetaFLOP-capable HPC cluster HSUper. Besides, we worked on the first set of performance engineering projects, that were granted to researchers at the universities of the federal armed forces (UniBws) to speed up various applications. In the first of these projects, that we could already finalize, we achieved a 2x speedup for a three-body-potential-based Monte Carlo simulation. We further established networking on the topic of HPC and related topics within and beyond the UniBws by, amongst others, seminar series, newsletters and hosting/co-organizing meetings such as the PinT conference 2023 or the latest EuroTUG meeting. Collaborations on performance engineering and optimizations have also started with external companies, with a pilot project in progress on optimization of logistics solvers. Our next steps are, amongst others, to establish an interactive scientific computing cloud for the users of the UniBws, to host a second EuroTUG meeting in 2024 and, finally, to intensify our work on the web-based HPC competence platform and related materials.

pySDC: Prototyping Spectral Deferred Corrections

Thomas Baumann
Forschungszentrum Jülich
Tuesday, July 18th, 16:00 – 18:00 (Building H, Foyer)

Spectral deferred corrections (SDC) is a very malleable time-stepping scheme that is especially well-suited for stiff problems. A wide range of parallel-in-time (PinT) extensions exist, from small-scale parallelism across the method to the large-scale PinT algorithm PFASST. SDC iteratively solves fully-implicit Runge-Kutta (RK) schemes with low-order methods. pySDC is a Python library that allows interested users to test how well SDC works for their specific problem or algorithm with very little effort. It includes implementations of PinT algorithms in both simulated parallel versions as well as MPI-based parallel versions. In addition, implementations of many popular serial RK methods are part of pySDC to enable a quick and fair comparison in an internally consistent framework. The code is very actively developed by collaborators in Juelich, Hamburg and Wuppertal. pySDC is publicly hosted on GitHub and extensively tested with continuous integration to ensure reproducibility and maintainability of the code and the results obtained with it. Contribution guidelines and tutorials exist to invite and guide new users. The modular nature allows developers to work on specific features without awareness of all details in the code, which encourages also smaller projects, such as master's theses.

Social Event & Conference Dinner

The social event will be a guided tour through the Airbus aircraft factory. A bus will collect workshop participants that registered for the event at TUHH and drop them off close to the restaurant for the conference dinner. The bus leaves at 13.30 from TUHH Campus. Please make sure you are on time, we cannot delay the departure of the bus!

The tour starts at 14.30 from Airbus Werksführung, Kreetslag 7, 21129 Hamburg.

On the way back, the bus leaves at 17.00 from Airbus. It will drop everyone off at Magellan-Terrassen / Großer Grasbrook, about 300m from the restaurant for the conference dinner. Depending on traffic, we will arrive between 17.30 and 18.00. Feel free to take a walk around the iconic "Speicherstadt" area to pass the time until the start of the dinner at 18.30.

Time: 18:30, Thursday, July 20th

Venue: Lieger Caesar Hamburg Hafencity

Location: Sandtorhafen/Traditionsschiffhafen, Runter Richtung Wasser, Gelbes Hausboot, 20457 Hamburg

Conference Venue

The conference will take place on the campus of TU Hamburg. TU Hamburg is located in Harburg, which is the part of Hamburg south of River Elbe.

Presentations will be given in Building H on Monday, Tuesday, Thursday and Friday and in Building I on Wednesday. See below for a map of the campus :