

10th Workshop on Parallel-in-Time Integration

Parallel-in-Time Scientific Committee:

Stephanie Friedhoff
Laurence Halpern
Benjamin Ong*
Kenji Ono
Debasmita Samaddar
Jacob Schroder*
Jemma Shipton*

August 2 – 6, 2021

*PinT 2021 co-organizers

Welcome Address

The PinT 2021 workshop is being held virtually, August 2–6, 2021, due to COVID-19. We thank the speakers and the participants for their flexibility. Each day will feature a series of talks scheduled within a 2-3 hour block. To accommodate our disparate time zones, the conference will begin earlier on August 5th and 6th.

- Sessions #1–#3, Aug 2 – 4 (3 – 5:30 pm GMT)
- Sessions #4–#5, Aug 5 –6 (9 – 11 am GMT)
- Virtual Social Event, Aug 3 (5 – 6 pm GMT)

Please join the zoom webinar at <https://michigantech.zoom.us/j/89536135023>

Some meeting logistics:

- We recommend that you download and install the latest version of the zoom app, <https://zoom.us/support/download>.
- Only the presenters and host will be able to share video/screens.
- We invite participants to submit questions via the Q & A feature throughout the lectures. The host will address them to the speaker during the Q & A period.

Related PinT news:

- We encourage all participants to join our parallel-in-time mailing list. To join, please send an email to parallelintime+subscribe@googlegroups.com.
- An NSF-sponsored PinT training conference will be held ??-??, 2022. Please visit that conference website at: <http://conferences.math.mtu.edu/cbms2020/>
- PinT 2022 will be held in Marseille, July 11 – 16, 2022. Please visit that conference website at: <https://www.chairejeanmorlet.com/2642.html>
- Keep up with general PinT news at www.parallel-in-time.org

Contents

Welcome Address, Monday August 2, 2021 (2:55pm GMT)	i
Session #1, Monday August 2, 2021 (3 – 6:00pm GMT)	1
Convergence of Parareal with spatial coarsening (<i>Daniel Ruprecht</i>)	1
Multigrid Reduction in Time: Two-Level Convergence Theory with Spatial Coarsening (<i>Jacob Schroder</i>)	2
Parareal with Exponential Integrators (<i>Tommaso Buvoli and Michael Minion</i>) . .	3
Optimized, parallel time integrators for better accuracy with large time steps (<i>Hans Johansen</i>)	4
(Plenary) Time parallel in PDEs using machine learning tools (<i>George Karniadakis</i>)	5
Session #2, Tuesday August 3, 2021 (3 – 5pm GMT)	6
Error bounds for PFASST and related Block Spectral-Deferred-Correction algorithms (<i>Thibaut Lunet</i>)	6
Quantum Algorithms for Solving Ordinary Differential Equations via Classical Integration Methods (<i>Benjamin Zanger</i>)	7
In Search of a Classical Convergence Proof for Waveform Relaxation (<i>Martin Gander</i>)	8
Parallel exponential Runge–Kutta methods (<i>Luan Vu Thai</i>)	9
Virtual Social Hour, Tuesday August 3, 2021 (5 – 6pm GMT)	10
Session #3, Wednesday August 4, 2021 (3 – 5:30pm GMT)	11
A Space-Time DPG Method for the Heat Equation (<i>Johannes Storn</i>)	11
Parallel space-time residual minimization for parabolic PDEs (<i>Raymond van Venetië, Jan Westerdiep</i>)	12
Diagonal SDC Preconditioning via Differentiable Programming and Reinforcement Learning (<i>Jan Ebert, Jenia Jitsev, Stefan Kesselheim, Michael Minion, Ruth Schöbel and Robert Speck</i>)	13
(Plenary) Space-time block preconditioning (<i>Ben S. Southworth, Federico Danieli, and Andrew J. Wathen</i>)	14
Session #4, Thursday August 5, 2021 (9 – 11:30 am GMT)	15

An Efficient Parallel-in-Time Method for Explicit Time-Marching Schemes (<i>Yen-Chen Chen</i>)	15
Parallel-in-Time Preconditioner for Optimal Control of Wave Equations (<i>Shulin Wu</i>)	16
Asynchronous Truncated Multigrid-reduction-in-time (AT-MGRIT) (<i>Jens Hahne, Ben Southworth, and Stephanie Friedhoff</i>)	17
(Plenary) From Parallel-in-Time to Full Space-Time Parallelization (<i>Ulrich Langer</i>)	18
Session #5, Friday August 6, 2021 (9 – 11am GMT)	19
Spatial re-distribution on the temporal coarse-level for Multigrid Reduction in Time (<i>Ryo Yoda</i>)	19
Parareal for higher index differential algebraic equations (<i>Iryna Kulchytska-Ruchka</i>)	20
Low Rank Parareal (<i>Benjamin Carrel</i>)	21
P-SPIKE Algorithm for solving Elastohydrodynamic Lubrication problem using parallel computers (<i>Peeyush Singh</i>)	22
Author Index	23

Monday August 2, 2021

Convergence of Parareal with spatial coarsening

Daniel Ruprecht

ruprecht@tuhh.de

Hamburg University of Technology

Aug 2
3:00pm
GMT

When solving PDEs, using a coarser spatial mesh in Parareal's coarse propagator can be an effective way to reduce its cost and improve potential speedup. However, the cost reduction has to be balanced against potentially slower convergence. In the talk, we will prove a theoretical best-case bound for the norm of Parareal's iteration matrix when spatial coarsening is used. We will discuss implications of the result and compare it to numerical experiments. One consequence of the bound is that for hyperbolic problems, where Parareal is known to struggle, spatial coarsening will eliminate any chance for a theoretical guarantee of fast convergence.

Multigrid Reduction in Time: Two-Level Convergence Theory with Spatial Coarsening

Aug 2
3:30pm
GMT

Jacob Schroder

jbschroder@unm.edu

University of New Mexico

The need for parallel-in-time is being driven by changes in computer architectures, where future speedups will be available through greater concurrency, but not faster clock speeds, which are stagnant. In this talk, we examine the parallel-in-time method, multigrid reduction in time (MGRIT), which applies multigrid to the time dimension for the (non)linear systems that arise when solving for multiple time steps simultaneously. The result is a flexible and nonintrusive approach that wraps existing time-stepping codes; however, basic MGRIT coarsens only in one dimension, time. For explicit time-stepping schemes (i.e., for stability) and for efficiency reasons in general, coarsening in spatial dimensions is attractive. Thus, we will examine the effects of spatial coarsening on MGRIT with a convergence analysis for 1D model problems using injection restriction and linear interpolation in space. The convergence analysis shows that spatial coarsening can be rather detrimental to convergence; however, the use of so-called FCF-relaxation can ameliorate some of these effects. Supporting numerical results are provided for model 1D advection and heat equations.

This talk is based on joint work with Tz. Kolev and V. Dobrev from Lawrence Livermore National Laboratory. LLNL-PRES-824440

Parareal with Exponential Integrators

Aug 2
4:00pm
GMT

Tommaso Buvoli^a and Michael Minion^b

tbuvo1@ucmerced.edu

^aUniversity of California, Merced

^bBerkeley Lab

In this talk we analyze the stability and convergence properties of Parareal with exponential integrators, and compare our results to those obtained with IMEX-based Parareal. We will show that exponential integrators are overall more robust with regards to parameter selection, especially on equations with no diffusion. Lastly, we also present several simple numerical experiments to demonstrate real-world parallel speedup.

Optimized, parallel time integrators for better accuracy with large time steps

Aug 2
4:30pm
GMT

Hans Johansen

hjohansen@lbl.gov

Lawrence Berkeley National Lab

We present an iterative parallel time integrator for hyperbolic-parabolic systems, which has excellent wave-resolving properties even for large time steps and high CFL's. To achieve this, we exploit several levels of parallelism, including in the spatial discretization, the time integrator, and the solver. We motivate the approach from a linear Von Neumann analysis of the spectral properties of the operators, trying to minimize the error over the parameter space while maximizing allowable time step. At the same time, we have to introduce a profitability model that includes work estimates for specialized instructions on modern hardware, such as tensor core operations on NVIDIA GPUs. After a somewhat frustrating brute-force optimization approach, we trust our instincts and demonstrate reasonably impressive results, knowing that someone else will prove optimality in the paper that gets all the attention.

(Plenary) Time parallel in PDEs using machine learning tools

Aug 2
5:00pm
GMT

George Karniadakis

George.Karniadakis@brown.edu

Brown University

Tuesday August 3, 2021

Error bounds for PFASST and related Block Spectral-Deferred-Correction algorithms

Thibaut Lunet

thibaut.lunet@tuhh.de

Hamburg University of Technology

Aug 3
3:00pm
GMT

Among the parallel-in-time (PinT) community, PFASST is generally considered to have a good lead in the global competition for the most complex PinT algorithm. Initially developed by M. Minion and M. Emmet, it uses Spectral Deferred Corrections as its base time integrator method, which itself relies on solution updates (sweeps) to increase the accuracy (and order) of the solution on successive time intervals. While the standard form of SDC consists in performing all sweeps interval after interval (Block SDC), PFASST uses different variations of those SDC updates over all time integration intervals, namely Block Gauss-Seidel SDC and Block Jacobi SDC, combined with a two level FAS correction for the Block Gauss-Seidel update.

In this talk, we describe all those different elements of PFASST when applied on a simple linear problem (Dahlquist equation), and show the equivalence of Block Gauss-Seidel with the Parareal algorithm used with specific parts of an SDC integrator. Then we derive convergence bounds for Block SDC, Block Gauss-Seidel SDC and Block Jacobi SDC, using the generating function technique, already used to determine convergence bounds for Parareal. With those bounds, we show the convergence of Block Gauss-Seidel SDC, that can be used as a direct PinT algorithm through pipe-lining of the sweeps. Also, we highlight the particular convergence order for Block Jacobi SDC that explains why order increase in PFASST appears only after a fixed amount of iterations. Finally, we show how the generating function technique can be extended to a Block Gauss-Seidel update with a FAS correction, and ultimately use it to compute a new convergence bound for PFASST.

Quantum Algorithms for Solving Ordinary Differential Equations via Classical Integration Methods

Aug 3
3:30
GMT

Benjamin Zanger

benjamin.zanger@tum.de

Technical University of Munich

Initial value problems are quite challenging to solve using parallel computing, since time integration is a sequential process. In this presentation, I give an introduction to solving initial value problems on quantum hardware. First, I give a short introduction to quantum computing and quantum annealing. Then, I point out the parallel compute capabilities by using high order Runge-Kutta formulations. Afterwards, I present an algorithm of how Runge-Kutta methods can be applied to quantum annealers and discuss a variational approach, which reduces the amount of quantum resources needed. This is followed by results for a linear ordinary differential equation (ODE) obtained on a D-Wave quantum annealer using a 6th order implicit Runge-Kutta method. Furthermore, our results also indicate the feasibility to solve non-linear ODEs on quantum computers. This is followed by a discussion about the time complexity of the algorithm for linear ODEs.

In Search of a Classical Convergence Proof for Waveform Relaxation

Aug 3
4:00pm
GMT

Martin Gander

`martin.gander@unige.ch`

University of Geneva

It is well known in the PinT community that the convergence of classical waveform relaxation is similar to the convergence of the Picard iteration from 1893, with a precise convergence estimate due to Lindelöf in 1894, but it is not easy to find a detailed proof of this in the literature: in [In't Hout, On the convergence of waveform relaxation methods for stiff nonlinear ordinary differential equations, 1995, page 178] a convergence estimate of the form of the classical Lindelöf result is quoted under mild assumptions on the splitting function, with references to [Lindelöf 1894, Nevanlinna 1989, Bellen and Zennaro 1993, Burrage 1995]. The reference [Burrage 1995, equation (3.6)] contains the same result with the comment 'it is easy to prove'; a different result with an extra term is shown in [Bellen and Zennaro 1993, Theorem 2.1], with the comment that 'by some standard analysis, we can easily get the convergence result'. An estimate of the Lindelöf form is given in [Nevanlinna, page 333, equation (2.13)] but for the case of linear problems only, proved using iterated kernel estimates, and this is also indicated as the key reference for the result in the book [Vandewalle 2013, Theorem 2.5.4, see also the page just before], but again without proof. The goal of my presentation is to show how the result announced in [Bellen and Zennaro 1993, Theorem 2.1] with the extra term can be proved.

Parallel exponential Runge–Kutta methods

Aug 3
4:30pm
GMT

Luan Vu Thai

luan@math.msstate.edu

Mississippi State University

Exponential Runge–Kutta (ExpRK) methods have shown to be competitive for the time integration of stiff semilinear parabolic PDEs. The current construction of ExpRK methods, however, relies on a convergence result that requires weakening many of the order conditions, resulting in schemes whose stages must be implemented in a sequential way. In this talk, after showing a stronger convergence result, we are able to derive two new families of fourth- and fifth-order ExpRK methods, which, in contrast to the existing methods, have multiple stages that are independent of one another and share the same format, thereby allowing them to be implemented in parallel or simultaneously, and making the methods to behave like using with much fewer stages. Moreover, all their stages involve only one linear combination of the product of φ -functions (using the same argument) with vectors. Overall, these features make these new methods to be much more efficient to implement when compared to the existing methods of the same orders. Numerical experiments on a one-dimensional semilinear parabolic problem, a nonlinear Schrödinger equation, and a two-dimensional Gray-Scott model are given to confirm the accuracy and efficiency of the two newly constructed methods.

Virtual Social Hour

Wednesday August 4, 2021

A Space-Time DPG Method for the Heat Equation

Johannes Storn

jstorn@math.uni-bielefeld.de

Bielefeld University

Aug 4
3:00pm
GMT

This talk introduces an ultra-weak space-time DPG method for the heat equation. We prove well-posedness of the variational formulation with broken test functions and verify quasi-optimality of a practical DPG scheme. Numerical experiments visualize beneficial properties of an adaptive and parabolically scaled mesh-refinement driven by the built-in error control of the DPG method.

Parallel space-time residual minimization for parabolic PDEs

Aug 4
3:30pm
GMT

Raymond van Venetië, Jan Westerdiep

janner@gmail.com

Korteweg–de Vries Institute for Mathematics, University of Amsterdam

In 2012, Roman Andreev introduced a minimal residual method for space-time numerical solutions for the heat equation. In a previous work, we showed that this method solves parabolic equations of the form

$$\frac{\partial}{\partial t}u + Au = g, \quad u|_{t=0} = u_0,$$

where g and u_0 are forcing- and initial data, and A is a linear coercive spatial operator. Informally, with X^δ some discrete trial space, the method balances PDE- and initial residuals:

$$u^\delta := \arg \min_{w \in X^\delta} \left[\left\| \frac{\partial}{\partial t}w + Aw - g \right\|^2 + \|w|_{t=0} - u_0\|^2 \right].$$

The error $\|u - u^\delta\|$ in natural norm is *quasi-optimal*: u^δ is (up to a fixed constant) the best approximation to u from the discrete trial space X^δ . Together with a wavelet-in-time multigrid-in-space preconditioner and preconditioned CG, we compute discrete solutions very efficiently.

In this talk, we dive into the specifics of the algorithmic complexity of such a solve. As it turns out, we can compute quasi-optimal approximations to u at *optimal* linear cost (in terms of $\dim X^\delta$) on a single processor. More interestingly though, defining *parallel complexity* as the asymptotic runtime given sufficiently many processors and assuming no communication cost, we can compute such approximations in *polylogarithmic* time, which is on par with the best-known results for elliptic problems. Numerical experiments will show that these abstract results translate to a highly efficient algorithm in practice.

Results in this talk build on our chapter in the proceedings of the previous PinT workshop.

Diagonal SDC Preconditioning via Differentiable Programming and Reinforcement Learning

Aug 4
4:00pm
GMT

Jan Ebert^a, Jenia Jitsev^a, Stefan Kesselheim^a, Michael Minion^b, Ruth Schöbel^a and Robert Speck^a

^oruthschoebel116@gmail.com

^aJuelich Supercomputing Centre, Forschungszentrum Juelich GmbH, ^bBerkeley Lab

The spectral deferred correction method is an iterative solver for time-dependent partial differential equations. It can be interpreted as a preconditioned Picard iteration for the collocation problem. The key to an efficient method is the choice of the preconditioner: it defines the speed of convergence and the level of parallelism. While the de-facto standard is a fast-converging, serial preconditioner, our goal is to find a fast-converging, diagonal and therefore parallel one. To achieve that, we combine techniques from differentiable programming and reinforcement learning. We look at Dahlquist's equation and train a network to select favorable, diagonal preconditioners depending on the parameter of the equation. The trained network can then be used directly to suggest parallel preconditioners for more complex problems: Using spectral methods, we can apply this to partial differential equations with a linear stiff part that we treat decoupled and implicitly with tailored preconditioners chosen by the network, whereas nonlinear parts are treated explicitly in real space. In this talk we present the overall approach, the training technique and show first results for linear and nonlinear PDEs.

(Plenary) Space-time block preconditioning

Aug 4
4:30pm
GMT

Ben S. Southworth^a, Federico Danieli^b, and Andrew J. Wathen^b

ben.s.southworth@gmail.com

Los Alamos National Laboratory^a, University of Oxford^b

For systems of PDEs, block preconditioning techniques are widely used to solve the linear(ized) system of equations that arise at each (sequential) time step, or for a steady state solution. Convergence of 2x2 block preconditioned fixed-point and Krylov iterations is fully determined by the preconditioning of the underlying Schur complement. In practice, many block preconditioners are developed through some assumptions of approximate commutation of differential operators to form an approximate Schur complement that is relatively easy to invert. This talk introduces the framework of space-time block preconditioning, extending the concepts of approximate commutation used in spatial block preconditioning to the space-time setting. In doing so, we can reduce the iterative solution of space-time systems of PDEs to the solution of a single-variable space-time equation, and corresponding Schur complement approximation. Space-time multigrid and algebraic multigrid methods have already demonstrated significant speedup over sequential time stepping for the former, and, with careful construction, the latter can be straightforward to compute in parallel as well. We demonstrate the principles of space-time block preconditioning applied to the incompressible Navier Stokes equations and the incompressible MHD equations. Theoretical motivation is provided for the new methods, and its robustness and scalability is shown on multiple model problems. More broadly, space-time block preconditioning offers a path to the fast, parallel, space-time solution of systems of PDEs, through the simpler task of solving single-variable space-time equations.

Thursday August 5, 2021

An Efficient Parallel-in-Time Method for Explicit Time-Marching Schemes

Yen-Chen Chen

chen-yenchen842@g.ecc.u-tokyo.ac.jp

The University of Tokyo

Aug 5
9:00am
GMT

Nowadays, conventional High-Performance Computing (HPC) methods for PDE equations numerical solvers are reaching their limits. As we move toward the exaflops supercomputing era, numerical PDE solvers parallelization reaches its saturation point in the spatial dimension and is restricted by the time domain instead. This leads to the development of parallel-in-time methods. Various parallel-in-time methods have been studied since the year 1964. Well-known methods such as parareal and multigrid reduction in time (MGRIT) have been shown to provide reasonable acceleration to partial differential equations (PDE) implicit schemes. However, only a few have applied pure explicit schemes. This research introduces a parallel-in-time method optimized for explicit schemes. The proposed method constructs a multiple coarsening layer structure similar to MGRIT and solves the parareal algorithm through coarse to fine layers. Furthermore, the relaxation method is defined to solve across the whole time segment divided by the number of processors. This way improves the efficiency of parallel-in-time solvers with a limited amount of processors. This research conducts numerical experiments for a one-dimensional advection equation and a 2-dimensional simulation of compressible viscous flow around a circular cylinder, using explicit time-marching schemes as relaxation methods. The research result shows that the proposed parallel-in-time method could improve the computation efficiency of explicit solvers compared to pure spatial parallelization.

Parallel-in-Time Preconditioner for Optimal Control of Wave Equations

Aug 5
9:30am
GMT

Shulin Wu

wushulin84@hotmail.com

Sichuan University of Science & Engineering

In this talk, we discuss a new preconditioner for iteratively solving the large-scale indefinite saddle-point sparse linear system, which arises from discretizing the optimality system in optimal control problems of wave equations with a one-shot second-order finite difference scheme in both space and time. The proposed preconditioner can be implemented in a parallel-in-time (PinT) manner via a carefully designed unitary diagonalization decomposition. We also present the eigenvalue bounds of the preconditioned system, which are shown to be highly clustered around one. Moreover, a simple splitting algorithm that alternates between a linear complementarity problem (LCP) and a quasi-Newton iteration is discussed for handling the case with control constraints. Within the quasi-Newton iteration, the proposed PinT preconditioner can be directly used in preconditioning the Jacobian system of the same structure. Both 1D and 2D numerical examples are given to illustrate the promising convergence performance of our proposed PinT preconditioner in comparison with a recently proposed matching Schur complement (MSC) preconditioner.

Asynchronous Truncated Multigrid-reduction-in-time (AT-MGRIT)

Jens Hahne^a, Ben Southworth^b, and Stephanie Friedhoff^a

jehahne@uni-wuppertal.de

Bergische Universität Wuppertal^a and Los Alamos National Laboratory^b

Aug 6
10:00am
GMT

In this talk, we present the new "asynchronous truncated multigrid-reduction-in-time" (AT-MGRIT) algorithm for introducing time parallelism to the solution of discretized time-dependent problems. The new algorithm is based on the multigrid-reduction-in-time (MGRIT) approach, which, in certain settings, is equivalent to another common multilevel parallel-in-time method, Parareal. In contrast to Parareal and MGRIT that both consider a global temporal grid over the entire time interval on the coarsest level, the AT-MGRIT algorithm uses truncated local time grids on the coarsest level, each grid covering certain temporal subintervals. These local grids can be solved completely independent from each other, which reduces the sequential part of the algorithm and, thus, increases parallelism in the method. Here, we study the effect of using truncated local coarse grids on the convergence of the algorithm, and show, using challenging nonlinear problems, that the new algorithm consistently outperforms classical Parareal/MGRIT in terms of time to solution.

(Plenary) From Parallel-in-Time to Full Space-Time Parallelization

Aug 5
10:30am
GMT

Ulrich Langer

ulanger@numa.uni-linz.ac.at

Institute for Computational Mathematics, Johannes Kepler University Linz,
Johann Radon Institute for Computational and Applied Mathematics, and the
Austrian Academy of Sciences

The traditional approaches to the numerical solution of initial-boundary value problems for parabolic Partial Differential Equations (PDEs) are based on the separation of the discretization in time and space leading to time-stepping methods. This separation of time and space discretizations comes along with some disadvantages with respect to parallelization and adaptivity. Parallel-in-Time methods try to avoid the curse of sequentiality of time-stepping procedures. To overcome the disadvantages connecting with the separation of time and space discretizations, we consider completely unstructured finite element (fe) discretizations of the space-time cylinder into simplicial finite elements, and derive stable fe schemes. Unstructured space-time fe discretizations considerably facilitate the parallelization and simultaneous space-time adaptivity. Moving spatial domains or interfaces can easily be treated since they are fixed in the space-time cylinder. Beside initial-boundary value problems for parabolic PDEs, we will also consider optimal control problems constrained by linear or non-linear parabolic PDEs. Here, unstructured space-time methods are especially suited since the reduced optimality system couples two parabolic equations for the state and adjoint state that are forward and backward in time, respectively. In contrast to time-stepping methods, one has to solve one big linear or non-linear system of algebraic equations. Thus, the memory requirement is an issue. In this connection, adaptivity, parallelization, and matrix-free implementations are very important techniques to overcome this bottleneck. Fast parallel solvers are the most important ingredient of efficient space-time methods.

The talk is based on joint works with C. Hofer, M. Neumüller, A. Schafelner, R. Schneckleitner, O. Steinbach, I. Touloupoulos, F. Tröltzsch, H. Yang, and M. Zank. This research was supported by the Austrian Science Fund (FWF) through the DK W1214-04. This support is gratefully acknowledged.

Friday August 6, 2021

Spatial re-distribution on the temporal coarse-level for Multigrid Reduction in Time

Ryo Yoda

ryo_yoda@mist.i.u-tokyo.ac.jp

The University of Tokyo

Aug 6
9:00am
GMT

Multigrid Reduction in Time (MGRIT) is a parallel-in-time method that extracts temporal parallelism by applying the reduction-based multigrid method for the temporal domain. In order to take advantage of this high parallelism, some cores need to be idle on a temporal coarse-level of MGRIT. In this presentation, we consider an optimization of the implementation that utilizes the temporal idle cores by re-distributing the spatial domain on coarse levels. It reduces the overhead of MGRIT and switches to a parallel-in-time algorithm at a smaller degree of parallelism. Although the spatial re-distribution is contrary to the motivation of parallel-in-time methods, it is expected to be effective in a parallelism-limited environment. This is due to the fact that parallelization-in-time is usually switched on already when the parallelization in space starts to saturate, so one can still benefit from additional parallelism.

Parareal for higher index differential algebraic equations

Aug 6
9:30am
GMT

Iryna Kulchytska-Ruchka

kulchytska@temf.tu-darmstadt.de

Technical University of Darmstadt

Electromagnetic devices such as electric circuits or electric motors are often modeled by systems of differential algebraic equations (DAEs). Application of the parallel-in-time (PinT) method Parareal to such systems requires a special treatment, e.g., of initial conditions. In this contribution we propose an adjustment of the Parareal algorithm which computes consistent initial values and enables the PinT solution of higher index DAE systems.

Low Rank Parareal

Aug 6
10:00am
GMT

Benjamin Carrel

benjamin.carrel@unige.ch

University of Geneva

The Parareal algorithm is a well known time parallel algorithm for evolution problems. It is based on a Newton like iteration, with cheap coarse corrections performed sequentially, and expensive fine solves performed in parallel. We are interested here in evolution problems that admit good low-rank approximations, and where the dynamical low-rank approximation (DLRA), proposed by Koch and Lubich, can be used as effective solvers. Many variants of DLRA have recently been proposed, like the projector-splitting methods and the projected Runge-Kutta methods. The cost and accuracy in DLRA are mostly governed by the rank chosen for the approximation. We want to use these properties in a new method we call Low Rank Parareal, in order to obtain time parallel DLRA solvers for evolution problems. We first describe an idealized variant of Low Rank Parareal, give a detailed convergence analysis, and illustrate our theoretical results numerically. We then describe a more practical variant of Low Rank Parareal for which we do not have a convergence analysis yet, and we illustrate its performance as well with numerical experiments.

P-SPIKE Algorithm for solving Elastohydrodynamic Lubrication problem using parallel computers

Aug 6
10:30am
GMT

Peeyush Singh

peeyush.singh@vitap.ac.in

VIT-AP University

We discuss a novel parallel approach for solving Elasto-hydrodynamic lubrication problems for steady state as well time dependent case.

Author Index

Buvoli
Tommaso, 3

Carrel
Benjamin, 21

Chen
Yen-Chen, 15

Danieli
Federico, 14

Ebert
Jan, 13

Friedhoff
Stephanie, 17

Gander
Martin, 8

Hahne
Jens, 17

Jitsev
Jenia , 13

Johansen
Hans, 4

Karniadakis
George, 5

Kesselheim
Stefan, 13

Kulchytska-Ruchka
Iryna, 20

Langer
Ulrich, 18

Lunet
Thibaut, 6

Minion
Michael, 3, 13

Ruprecht
Daniel, 1

Schöbel
Ruth, 13

Schroder
Jacob, 2

Singh
Peeyush, 22

Southworth
Ben, 14, 17

Speck
Robert , 13

Storn
Johannes, 11

van Venetië
Raymond, 12

Vu Thai
Luan, 9

Wathen
Andrew J., 14

Westerdiep
Jan, 12

Wu
Shulin, 16

Yoda
Ryo, 19

Zanger
Benjamin, 7