


## MIMS

Innovative space-time  
parallel methods

<a href="#">registration</a>	>
<a href="#">programme</a>	>
<a href="#">list of participants</a>	>
<a href="#">venue</a>	>
<a href="#">accommodation</a>	>
<a href="#">contact details</a>	>

SCHOOL OF  
MATHEMATICS

<a href="#">staff list</a>	>
<a href="#">postgraduate admissions</a>	>

## CONTACT DETAILS

**MIMS**  
The University of Manchester  
School of Mathematics  
Oxford Road  
Manchester  
M13 9PL  
UK  
mims@manchester.ac.uk  
tel: +44 (0)161 275 5812  
fax: +44 (0)161 306 3669

## Innovative space-time-parallel methods: Analysis and Applications June 18-19, 2013

Space-and-time parallel methods are becoming increasingly popular for making better use of available computational resources for the solution of time-dependent partial differential equations. This is achieved by parallelisation not only taking place on the spatial variables (as with classical domain decomposition), but also in the time direction. The aim of this workshop is to bring together researchers working on space-and-time parallel methods to discuss recent advances and future directions.

Invited speakers:

- [Martin J. Gander](#), Université de Genève, Switzerland
- [Yvon Maday](#), Université Pierre et Marie Curie, France
- [Michael Minion](#), Stanford University, USA
- [Benjamin Ong](#), Michigan State University, USA

The workshop is open to everyone without registration fees.

For any questions contact the organizer [Stefan Güttel](#).

Funded by the [London Mathematical Society](#), the [EPSRC Network on Numerical Algorithms and High Performance Computing](#), and the [School of Mathematics at The University of Manchester](#).



## Programme

All talks take place in the **Frank Adams Room** of the Alan Turing Building. Please find below the workshop time table and abstracts. A single PDF file containing all abstracts submitted to the workshop can be found [here](#). The workshop dinner on Tuesday evening is booked out now. It takes place in the Japanese restaurant SAMSI on 36-38 Whitworth Street (walkable from the Alan Turing Building).

### Time table

	Monday, June 17	Tuesday, June 18	Wednesday, June 19
08:30–09:00		Registration and coffee	Coffee
09:00–10:00		Martin J. Gander	Yvon Maday
10:00–10:30		Coffee	Coffee
10:30–12:00		Felix Kwok Roman Andreev Robert Speck	Olga Mula Christian Wieners Stefan Findeisen
12:00–13:00		Group photo and lunch	Lunch
13:00–14:00			Julien Salomon Joerg Wensch Qiqi Wang Debasmita Samaddar
14:00–15:00		Michael Minion	
15:00–16:00		Tobias Weinzierl Uwe Köcher	Closing
16:00–16:30		Coffee	
16:30–17:30		Benjamin Ong	
17:30–18:00		Georges Klein	
18:00	Meeting in town	Workshop dinner (18:30)	

## Abstracts for Tuesday, June 18

---

Chair for morning: **Stefan Güttel** (The University of Manchester)

### 9:00–9:10 Opening remarks

**9:10–10:00 Martin J. Gander** (University of Geneva, Switzerland)

Multiple Shooting, Parareal, Krylov Parareal, ParaExp and Space-Time Multigrid: an Overview on Time Domain Decomposition Methods

**10:30–11:00 Felix Kwok** (University of Geneva, Switzerland)

Coarse Grid Correction for the Neumann-Neumann Waveform Relaxation Method

**11:00–11:30 Roman Andreev** (University of Maryland, USA)

Sparse space-time Petrov-Galerkin discretizations for parabolic evolution equations

**11:30–12:00 Robert Speck** (Jülich Supercomputing Centre, Germany)

From Spectral Deferred Corrections to the Parallel Full Approximation Scheme in Space and Time

*(See also the related abstract by Mathias Winkel et al.:*

*Multi-level Parallel-In-Time Methods for N-body Plasma Physics Applications)*

Chair for afternoon: **Daniel Ruprecht** (University of Lugano, Switzerland)

**14:00–14:50 Michael Minion** (Stanford University, USA)

Time Parallelization for Second-Order Equations

**15:00–15:30 Tobias Weinzierl** (Technische Universität München, Germany)

Parallel Space-time Spacetrees for Simple Parabolic Benchmarks

**15:30–16:00 Uwe Köcher** (Helmut Schmidt University Hamburg, Germany)

Parallel variational space-time methods for the wave equation

**16:30–17:20 Benjamin Ong** (Michigan State University, USA)

RIDC-DD: A parallel space-time algorithm

**17:30–18:00 Georges Klein** (University of Oxford, UK)

Deferred correction from equispaced data based on efficient high-order rational integration

## Abstracts for Wednesday, June 19

---

Chair for morning: **Debasmita Samaddar** (Culham Science Centre, UK)

**9:00–9:50 Yvon Maday** (Université Pierre et Marie Curie, France)

Conservations with parareal in time algorithm

**10:30–11:00 Olga Mula** (Université Pierre et Marie Curie, France)

MINARET: Towards a time-dependent neutron transport parallel solver

**11:00–11:30 Christian Wieners** (Karlsruhe Institute of Technology, Germany)

Space-Time Methods for Wave Equations

**11:30–12:00 Stefan Findeisen** (Karlsruhe Institute of Technology, Germany)

First step towards Parallel and Adaptive Computation of Maxwell's Equations

Chair for afternoon: **Felix Kwok** (University of Geneva, Switzerland)

**13:00–13:30 Julien Salomon** (Universite Paris-Dauphine, France)

An intermediate state method for the time-parallelized solving of optimal control problems

**13:30–14:00 Joerg Wensch** (TU Dresden, Germany)

Multirate time integration of the Euler equations

**14:00–14:30 Qiqi Wang** (Massachusetts Institute of Technology, USA)

Towards Scalable Parallel Long Time Integration of Chaotic Dynamical Systems

**14:30–15:00 Debasmita Samaddar** (Culham Science Centre, UK)

Temporal parallelization of advanced operation scenario simulations of fusion plasma



---

## List of participants

---

Buraq Abdulkareem, School of Electrical and Electronic Engineering, UoM  
Ismail Adeniran, University of Manchester, UK  
Zaid Ait Haddou, Manchester Business School, UK  
Hanan Almeer, School of Electrical and Electronic Engineering, UoM  
Hayder Al-Tameemi, School of Electrical and Electronic Engineering, UoM  
Roman Andreev, University of Maryland, USA  
Bahar Arslan, School of Mathematics, UoM  
Richard Banach, School of Computer Science, UoM  
Chris Boddington, Faculty of Life Sciences, UoM  
Paul Bolton, University of Manchester, UK  
John Brooke, School of Computer Science, UoM  
Timothy Butters, The University of Manchester, UK  
Yu Cao, School of Chemical Engineering and Analytical Science, UoM  
Paul Connolly, University of Manchester, UK  
Simon Cotter, School of Mathematics, UoM  
Stefan Findeisen, Karlsruhe Institute of Technology, Germany  
Joseph Dada, School of Computer Science, UoM  
Oliver Dorn, School of Mathematics, UoM  
Martin J. Gander, University of Geneva, Switzerland  
Richard Gowers, The University of Manchester, UK  
Tuba Gozel, School of Electrical and Electronic Engineering, UoM  
Stefan Güttel, School of Mathematics, UoM  
Xiaowei Guo, School of Chemical Engineering and Analytical Science, UoM  
Dave Harris, School of Mathematics, UoM  
Tadashi Hemmi, Electrical and Electronic Engineering, UoM  
Stephen Hendry, Arup, UK  
Nick Higham, School of Mathematics, UoM  
Tianjiang Hu, University of Manchester, UK  
Kurt Janssens, KU Leuven, Belgium  
Ramaseshan Kannan, School of Mathematics, UoM  
Amal Khabou, School of Mathematics, UoM  
Georges Klein, University of Oxford, UK  
Uwe Koecher, Helmut-Schmidt-University Hamburg, Germany  
Felix Kwok, University of Geneva, Switzerland  
Lijing Lin, School of Mathematics, UoM

Yvon Maday, Université Pierre et Marie Curie, France  
Michael Minion, Stanford University, USA  
Olga Mula, Laboratoire Jacques Louis Lions, France  
Vanni Noferini, School of Mathematics, UoM  
Benjamin Ong, Michigan State University, USA  
Debasish Pradhan, Defence Institute of Advanced Technolgy, India  
Sam Relton, School of Mathematics, UoM  
Daniel Ruprecht, University of Lugano, Switzerland  
Julien Salomon, Université Paris-Dauphine, France  
Debasmita Samaddar, Culham Centre for Fusion Energy, UK  
Ben Saxby, School of Mathematics, UoM  
Meisam Sharify, School of Mathematics, UoM  
Robert Speck, Forschungszentrum Jülich, Germany  
Adam Stanier, University of Manchester, UK  
Vendel Szeremi, School of Mathematics, UoM  
Leo Taslaman, School of Mathematics, UoM  
Kenan Tekbas, Electrical and Electronic Engineering, UoM  
Niels Walet, School of Physics and Astronomy, UoM  
Qiqi Wang, Massachusetts Institute of Technology, USA  
Yanzhen Wang, School of Chemical Engineering and Analytical Science, UoM  
Yi Wang, School of Materials, UoM  
Tobias Weinzierl, Technische Universität München, Germany  
Joerg Wensch, Technische Universität Dresden, Germany  
Christian Wieners, Karlsruhe Institute of Technology, Germany  
Mathias Winkel, Forschungszentrum Jülich, Germany  
Wei Yi, School of Chemical Engineering and Analytical Science, UoM  
Weijian Zhang, School of Mathematics, UoM

# ABSTRACTS

(in the order of presentation)

Martin J. Gander  
University of Geneva

Multiple Shooting, Parareal, Krylov Parareal, ParaExp and Space-Time  
Multigrid: an Overview on Time Domain Decomposition Methods

Many problems in science and engineering are time dependent, and time stepping methods are used to obtain approximate solutions. If the problems are large scale, or solutions are needed in real time, it is necessary to use the computing power of parallel computers. The classical strategy to parallelize time integration is to parallelize the solution at each time step, and to advance sequentially from time step to time step. This approach however neglects an entire dimension, the time dimension, which could also be used for the parallelization. In contrast to the spatial dimensions, the time dimension has however a direction: the solution later in time depends only on the solution earlier in time, and not vice versa. It therefore seems difficult to do useful computations at a future time step, before the present time step results are known.

There are several algorithms which nevertheless try to do useful computations later in time, before fully accurate results at the present time step are available, and one of the more recent ones is based on multiple shooting: the parareal algorithm. This algorithm is using an approximation of the Jacobian on a coarse grid in the Newton iteration classically used for solving the shooting equations. After reviewing a compact convergence result for this algorithm, I will illustrate its numerical performance for several examples of systems of ordinary and partial differential equations. These examples reveal that while the algorithm performs well for diffusive problems, convergence is unsatisfactory for hyperbolic equations. I will then explain as possible remedies for this problem the Krylov parareal algorithm, and also a different time parallel method called ParaExp. I will finally show two further developments, a space-time algorithm based entirely on multigrid techniques, and a space-time algorithm based on space-time domain decomposition.



# Coarse Grid Correction for the Neumann–Neumann Waveform Relaxation Method

Felix Kwok  
University of Geneva  
felix.kwok@unige.ch

## Abstract

In the recent paper [2], a new variant of the waveform relaxation (WR) method based on Neumann–Neumann iterations has been proposed for the solution of linear parabolic PDEs. Just like for the steady case, one step of the method consists of solving the subdomain problems using Dirichlet traces, followed by a correction step involving Neumann interface conditions. However, each subdomain problem is now in both space and time, and the interface data to be exchanged are also functions of time. One advantage of the WR framework is that it allows the use of different spatial and time discretizations for each subdomain. Moreover, it has been shown in [1] that for finite time intervals, the Neumann–Neumann waveform relaxation (NNWR) method converges superlinearly both in one spatial dimension and for 2D decompositions into strips. Unfortunately, convergence deteriorates significantly as the number of subdomain increases, since the method does not allow communication between far-away subdomains.

The goal of this talk is to introduce a coarse grid component to the NNWR method in order to make it scalable with respect to the number of subdomains. In 1D, an exact coarse grid correction can be calculated by considering shape functions that satisfy the homogeneous PDE; this yields an iteration that converges to the exact solution in two steps. In 2D, however, the exact coarse grid correction yields a matrix problem that is too large and dense to be of practical value. Thus, we must seek corrections in a smaller subspace, consisting of coarse hat functions or discontinuous piecewise linear elements. We analyze the convergence of the method for different choices of coarse spaces and show its scaling behaviour when we vary the number of subdomains.

This is joint work with Martin J. Gander, Sébastien Loisel and Kévin Santugini.

## References

- [1] M. J. Gander, F. Kwok, and B. C. Mandal. Dirichlet–Neumann and Neumann–Neumann waveform relaxation methods for the time-dependent heat equation. *In preparation*.
- [2] F. Kwok. Neumann–Neumann waveform relaxation for the time-dependent heat equation. *Submitted to the Proceedings of the 21st International Conference on Domain Decomposition Methods*, 2012.

## Sparse space-time Petrov-Galerkin discretizations for parabolic evolution equations

**Roman Andreev** (CSCAMM, University of Maryland, andreevr@umd.edu)

In view of applications such as optimal control problems with parabolic PDE constraints and massively parallel computations of time-dependent problems, space-time compressive discretizations of parabolic evolution equations are of significant interest.

We discuss space-time (sparse) tensor product simultaneous Petrov-Galerkin discretizations of parabolic evolution equations, and propose efficient preconditioners for the iterative solution of the resulting single linear system of equations. Therein, space-time stability of the discretization, i.e., the validity of the discrete inf-sup condition with respect to suitable space-time norms uniformly in the discretization parameters, is essential.

Viewing the Crank-Nicolson time-stepping scheme as a space-time Petrov-Galerkin discretization, we can show that it is conditionally space-time stable for those space-time norms. This motivates a general minimal residual Petrov-Galerkin discretization framework along with space-time stable families of trial and test spaces of (sparse) tensor product type, resulting in space-time compressive discretization algorithms.

Additional interesting properties of the proposed algorithm include: very low regularity requirements on the input data; modularity in the spatial discretization; possibility of high-order nonuniform temporal discretization. Several natural questions, such as the validity of the maximum principle, and incorporation of space-time adaptivity while maintaining stability, are open.

### References:

For a condensed description and a basic Matlab implementation of the algorithm see:

*Space-time discretization of the heat equation. A concise Matlab implementation.*

eprint arXiv:1212.6037 (includes code), **2012** → WWW: <http://arxiv.org/abs/1212.6037>

The essentials of the theoretical background and some numerical results are given in:

*Stability of sparse space-time finite element discretizations of linear parabolic evolution equations.*

IMA J Numer Anal (**2013**) 33(1): 242–260 → DOI: 10.1093/imanum/drs014

Further details and extensions are discussed in:


*Stability of space-time Petrov-Galerkin discretizations for parabolic evolution equations.*

PhD thesis, ETH Zurich, ETH Diss No 20842, **2012** → DOI: 10.3929/ethz-a-007563932

and

R.A. and C. Tobler: *Multilevel preconditioning and low rank tensor iteration for space-time simultaneous discretizations of parabolic PDEs.*

SAM report 2012-16. ETH Zürich, **2012** → WWW: [http://www.sam.math.ethz.ch/sam\\_reports](http://www.sam.math.ethz.ch/sam_reports)



### Acknowledgment:

Swiss NSF Grant No. 127034 and ERC AdG No. 247277 awarded to Ch. Schwab, ETH Zürich

# From Spectral Deferred Corrections to the Parallel Full Approximation Scheme in Space and Time

(A summary of recent work)

R. Speck<sup>\*†</sup>, D. Ruprecht<sup>†</sup>, M. Emmett<sup>‡</sup>, M. Minion<sup>§</sup>, M. Bolten<sup>¶</sup>, R. Krause<sup>†</sup>,

<sup>\*</sup>Jülich Supercomputing Centre, Forschungszentrum Jülich, Germany.

<sup>†</sup>Institute of Computational Science, Università della Svizzera italiana, Switzerland.

<sup>‡</sup>Center for Computational Sciences and Engineering, Lawrence Berkeley National Laboratory, USA.

<sup>§</sup>Institute for Computational and Mathematical Engineering, Stanford University, USA.

<sup>¶</sup>Department of Mathematics, Bergische Universität Wuppertal, Germany.

## I. INTRODUCTION

A recently developed time-parallel method is the "parallel full approximation scheme in space and time" (PFASST) introduced in [1], [2]. It is based on spectral deferred correction methods (SDC) [3], a class of methods that iteratively uses low order methods to obtain an overall method of high order of accuracy. By intertwining the SDC iterations with a Parareal-like iteration (see [4] for Parareal), PFASST features an improved bound on parallel efficiency. The efficacy of PFASST in extreme-scale parallel simulations on a BlueGene/P has been demonstrated in [5].

## II. SDC, MULTI-LEVEL SDC AND PFASST

Below, first the time-serial single level spectral deferred corrections method is described briefly. Then, the time-serial, multi-level SDC (MLSDC) approach is discussed. Finally, the time-parallel, multi-level PFASST algorithm is sketched.

### A. Spectral deferred corrections (SDC)

The SDC method introduced in [3] is an iterative approach to compute a solution of a collocation formula. Given a time-step  $[T_n, T_{n+1}]$ , denote by  $T_n \leq t_0 < \dots < t_M \leq T_{n+1}$  a set of intermediate Gauß collocation points. Typically, Gauß-Lobatto nodes are used, so that  $T_n = t_0$  and  $T_{n+1} = t_M$ . Integrating an initial value problem from  $T_n$  to  $T_{n+1}$  is then equivalent to solving the Picard formulation

$$u(t) = u_n + \int_{T_n}^{T_{n+1}} f(u(\tau), \tau) d\tau. \quad (1)$$

Approximating (1) with a quadrature rule with nodes  $t_m$  results in a linear or nonlinear system of equations (depending on the problem) to be solved for the coefficients of the collocation polynomial. Instead of solving the full system directly, SDC proceeds iteratively using so-called "sweeps" of a low order integration method, typically forward or backward Euler. For a backward Euler, the sweeps are of the form

$$U_{m+1}^{k+1} = U_m^{k+1} + \Delta t_m \left( f(U_{m+1}^{k+1}) - f(U_{m+1}^k) \right) + S_m^k. \quad (2)$$

Here, the operator  $S_m^k$  approximates the Picard integral from  $t_m$  to  $t_{m+1}$ , that is

$$S_m^k \approx \int_{t_m}^{t_{m+1}} f(u^k(\tau), \tau) d\tau. \quad (3)$$

If the iteration converges, the term  $f(U_{m+1}^{k+1}) - f(U_{m+1}^k)$  vanishes and (2) for  $m = 0, \dots, M-1$  can be combined into

$$U_M^{k+1} = U_0^{k+1} + \sum_{m=0}^{M-1} S_m^k \quad (4)$$

which is precisely the collocation approximation of (1).

### B. Multi-level spectral deferred corrections

In [6], a multi-level SDC method (MLSDC) is presented, that in a certain sense provides the "missing link" between single-level, time-serial SDC and multi-level, time-parallel PFASST. In contrast to SDC, MLSDC performs sweeps not on a single level but on a hierarchy of levels, where higher levels use fewer collocations nodes and therefore a coarsened temporal discretization. A FAS correction is employed in order to ensure information is properly transferred between levels. It is shown in [6] that MLSDC provides the same accuracy as SDC, minimally improved stability and that it can reduce the number of iterations required for convergence. Also, the incorporation of weighting matrices required e.g. for the use of compact finite difference stencils is explained. Here, the SDC sweep equation as well as the FAS correction need to be modified to achieve higher-order discretizations in space.

In order to reduce the computational cost of MLSDC at coarser levels, multiple strategies are presented to also coarsen the spatial discretization on the higher levels of the hierarchy:

- Reduced spatial resolution
- Reduced order discretization
- Reduced implicit solve
- Reduced physics

The first two strategies are subsequently investigated in detail for a linear advection-diffusion problem, nonlinear viscous Burgers' equation and a shear layer instability described by the Navier-Stokes equations in vorticity-velocity formulation, see Section III for a tentative summary of the last problem.

One key advantage of MLSDC is that it can be parallelized in time, leading to the "parallel full approximation scheme in space and time" (PFASST) described below in Subsection II-C. Besides providing the possibility to parallelize in time, however, MLSDC is also of interest in its own right: It provides a starting point to extend it to a full space-time multi-grid method and thus enables a completely new approach to the field of space-time parallel multi-grid methods as studied e.g. in [7], [8].

### C. PFASST

The PFASST method is pioneered in [1]. In [2], it is introduced in its ultimate form and its performance is studied for viscous Burgers' equation and the Kuramoto-Sivashinsky equation. PFASST basically corresponds to a number of concurrent MLSDC iterations running for multiple time intervals assigned to different processors plus a frequent exchange of updated values. PFASST employs SDC sweeps on multiple levels and uses a FAS correction as well to efficiently transfer information between levels. In particular, the FAS correction also allows for efficient use of different spatial coarsening strategies very similar to the MLSDC approach. Furthermore, to optimize efficiency and reduce overhead from the coarse level sweeps, PFASST features a pipelining strategy, see [1].

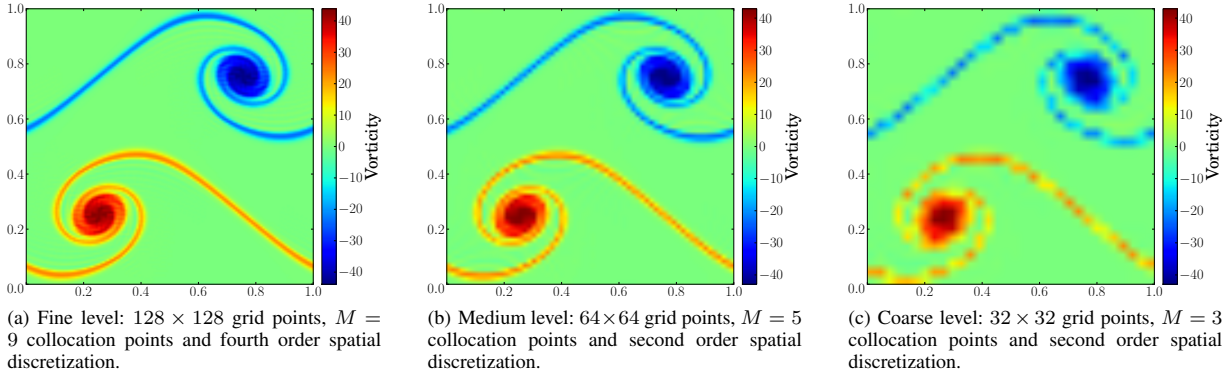


Fig. 1: Vorticity  $\omega$  at  $t = 1.0$  in the shear layer instability example on three levels of a MLSDC hierarchy. The coarse level shows clear signs of underresolution like oscillations in the trailing tails of the vortical structures. MLSDC performs 256 time-steps in this example and requires on average 6.9 iterations on each time-step to converge to a tolerance of  $10^{-12}$ .

By intertwining the SDC sweeps on the different levels with the outer iteration, PFASST manages to achieve an improved bound on parallel efficiency compared to plain Parareal. Denote by  $P_T$  the number of processors, by  $K_s$  the number of sweeps of the underlying SDC scheme, by  $\alpha$  the ratio of the execution time of one coarse to one fine sweep and by  $K_p$  the number of PFASST iterations. The speedup provided by PFASST then reads

$$S(P_T, \alpha) = \frac{P_T K_s}{P_T \alpha + K_p (1 + \alpha)} \leq \frac{K_s}{K_p} P_T. \quad (5)$$

Note how strategies to reduce the overhead of coarse level evaluations in MLSDC directly translate into strategies to reduce  $\alpha$  and thus improve the speedup provided by PFASST. The potential of PFASST to extend the strong scaling limit of the N-body tree-code PEPC [9] in extreme-scale parallel simulations has been demonstrated in [5], where timing result from runs on up to 262,144 processors on a BlueGene/P system are presented. There, in order to optimize  $\alpha$ , a tree-code specific coarsening criterion was developed that roughly corresponds to “reduced order in space”. An accuracy study accompanying the performance study and discussing the accuracy of SDC and PFASST in combination with an N-body solver can be found in [10].

### III. EXAMPLE

One example studied in [6] is the performance of MLSDC for a 2D shear layer instability. A tentative summary is given here. The problem is described by the Navier-Stokes equations in vorticity-velocity formulation

$$\omega_t + u \cdot \nabla \omega = \nu \Delta \omega \quad (6)$$

with velocity  $u$  and vorticity  $\omega = \nabla \times u$ . The computational domain  $[0, 1]^2$  is assumed to be periodic in both directions and the initial velocity field is given by

$$u_1^0(x, y) = -1.0 + \tanh(\rho(0.75 - y)) + \tanh(\rho(y - 0.25)) \quad (7)$$

$$u_2^0(x, y) = -\delta \sin(2\pi(x + 0.25)), \quad (8)$$

corresponding to two shear layers at  $y = 0.25$  and  $y = 0.75$  with a thickness parameter  $\rho = 50$  and an initial disturbance in velocity of amplitude  $\delta = 0.05$ . MLSDC uses an IMEX-type sweep, where the advection of vorticity is treated explicitly while the diffusive term is treated implicitly. Three levels are used and both the medium and the coarse level feature a reduced resolution in space and reduced order of the spatial discretization in addition to the reduced number of collocation points, see Figure 1 for the exact values. Despite the fact that clear signs of under-resolution are present on the coarse level, the MLSDC iteration converges quickly and robustly and also conserves

total vorticity. The Poisson problems arising from the implicit part of the IMEX scheme and during the recovery of the vorticity field from a solenoidal stream function are solved using a the multi-grid method PMG [11]. In the current version, the linear problems are always solved to full accuracy and the strategy of using a “reduced implicit solve” to coarsen in space has not yet been investigated.

### IV. OUTLOOK

In the talk, we will provide an overview of the results on MLSDC in [6]. We will also present first results of incorporating the third coarsening strategy, using a “reduced implicit solve” for the implicit part on coarser levels. Further, the connection of MLSDC and PFASST will be illustrated. Studies comparable to the ones conducted for MLSDC will be conducted with PFASST, in order to provide a detailed assessment of the differences between time-serial MLSDC and time-parallel PFASST.

### REFERENCES

- [1] M. L. Minion, “A hybrid parareal spectral deferred corrections method,” *Comm. App. Math. and Comp. Sci.*, vol. 5, no. 2, pp. 265–301, 2010.
- [2] M. Emmett and M. L. Minion, “Toward an efficient parallel in time method for partial differential equations,” *Comm. App. Math. and Comp. Sci.*, vol. 7, pp. 105–132, 2012.
- [3] A. Dutt, L. Greengard, and V. Rokhlin, “Spectral deferred correction methods for ordinary differential equations,” *BIT Numerical Mathematics*, vol. 40, no. 2, pp. 241–266, 2000.
- [4] J.-L. Lions, Y. Maday, and G. Turinici, “A “parareal” in time discretization of PDE’s,” *C. R. Acad. Sci. – Ser. I – Math.*, vol. 332, pp. 661–668, 2001.
- [5] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. Minion, M. Winkel, and P. Gibbon, “A massively space-time parallel N-body solver,” in *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, ser. SC ’12. CA, USA: IEEE Computer Society Press, 2012, pp. 92:1–92:11.
- [6] M. Bolten, M. Emmett, R. Krause, M. Minion, D. Ruprecht, and R. Speck, “A multi-level spectral deferred correction method,” (In preparation).
- [7] W. Hackbusch, “Parabolic multi-grid methods,” *Computing Methods in Applied Sciences and Engineering*, VI, pp. 189–197, 1984.
- [8] G. Horton and S. Vandewalle, “A space-time multigrid method for parabolic partial differential equations,” *SIAM Journal on Scientific Computing*, vol. 16, no. 4, pp. 848–864, 1995.
- [9] M. Winkel, R. Speck, H. Hübner, L. Arnold, R. Krause, and P. Gibbon, “A massively parallel, multi-disciplinary Barnes-Hut tree code for extreme-scale N-body simulations,” *Comp. Phys. Comm.*, vol. 183, no. 4, pp. 880–889, 2012.
- [10] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. Minion, M. Winkel, and P. Gibbon, “Integrating an N-body problem with SDC and PFASST,” in *Proceedings of the 21st International Conference on Domain Decomposition Methods*, ser. Lecture Notes in Computational Science and Engineering, 2012, (Accepted).
- [11] M. Bolten, “Highly scalable multigrid solver for 3-level Toeplitz and circulant matrices,” *Parallel Computing*, (Submitted).

# Multi-level Parallel-In-Time Methods for $N$ -body Plasma Physics Applications

(Prospect on upcoming work)

M. Winkel\*, R. Speck\*<sup>†</sup>, D. Ruprecht<sup>†</sup>, P. Gibbon\*

\*Jülich Supercomputing Centre, Jülich, Germany.

<sup>†</sup>Institute of Computational Science, Università della Svizzera italiana, Lugano, Switzerland.

## I. INTRODUCTION

Parallel-in-time integration methods are a promising approach to extend strong-scaling limits of simulation code at large-scales. With fixed degrees-of-freedom in space and million-way concurrency at hand, classical space-parallel strategies inevitably lead to inefficient use of compute power as strong-scaling saturates. In this contribution, we will describe the ongoing transition from a large-scale, space-parallel simulation code towards a space-time parallel method ready for extreme-scales. To this end, we will join the space-parallel Barnes-Hut tree code PEPC [1] with the time-parallelization library PFAST [2].

## II. LIMITS OF PARALLELIZATION USING MULTIPOLE-BASED FAST SUMMATION SCHEMES

Being dominated by more than 20 Petaflop-class installations, the Top500 list of world's most powerful computers is the evidence of a tremendous upgrade in worldwide computing power. However, this does not solve two very principal problems of all particle-based molecular dynamics simulations. Following Newtonian dynamics, the particle's movement in these simulations is driven by forces from interaction with each of the individual particles. For example, in the electrostatic case, these are mediated through the Coulomb force

$$\vec{F}(\{\vec{r}_i, q_i\}, \{\vec{r}_{1\dots N}, q_{1\dots N}\}) = q_i \sum_{\substack{j=1 \\ j \neq i}}^N q_j \frac{\vec{r}_i - \vec{r}_j}{|\vec{r}_i - \vec{r}_j|^3} + \vec{F}^{(\text{ext})} \quad (1)$$

acting from  $N$  particles  $\vec{r}_{1\dots N} \in \mathbb{R}^3$  with charges  $q_{1\dots N} \in \mathbb{R}$  onto particle  $i$  with charge  $q_i$  at  $\vec{r}_i$ . In addition, usually external forces  $\vec{F}^{(\text{ext})} \equiv \vec{F}^{(\text{ext})}(\vec{r}_i, \vec{v}_i, t)$  are present that – besides the particle's position  $\vec{r}_i$  – can in general also depend on its velocity  $\vec{v}_i = \dot{\vec{r}}_i$  (e.g. Lorentz force on particles in an external magnetic field) and time  $t$  (e.g. in laser-driven systems).

A direct evaluation of (1) for all particles  $i = 1 \dots N$  exhibits a computational complexity of  $\mathcal{O}(N^2)$  that for ever increasing simulation size and thus particle numbers cannot be dealt with by means of simply increasing computation power. Instead, efficient fast summation schemes such as the Barnes-Hut tree code [3] and the Fast Multipole Method [4] have been developed. They rely on a multipole-based approach to replace the interactions between individual particles with interactions between particles and particle groups via their multipole expansion. Thus, the contribution of distant particles is approximated by systematically grouping them into clusters and representing their charge distribution through a series expansion. The coarseness of the spatial approximation can be adjusted through a so-called multipole-acceptance-parameter that determines which sizes of particle groups at which distance are to be accepted as interaction partners. Thus, an overall complexity of  $\mathcal{O}(N \log N)$  or even  $\mathcal{O}(N)$  can be achieved [5].

During the last years, a highly scalable hybrid parallelized implementation of the original Barnes-Hut tree algorithm using MPI and Pthreads has been developed at Jülich Supercomputing Centre under the label *Pretty Efficient Parallel Coulomb solver* (PEPC) [1].

Today, it is not only limited to pure Coulomb interaction but includes numerous interaction-specific backends and respective frontends. Applications are reaching from simulation of interstellar dust dynamics, plasma boundary layer instabilities, laser-driven plasmas and plasma-wall interactions in fusion vessels to hydrodynamics applications with the SPH or the vortex-particle method.<sup>1</sup> PEPC uses a work separation scheme between one communicator thread performing rapid low-latency point-to-point communication for tree node exchange and an adjustable number of worker threads that are responsible for tree structure interpretation and force evaluation on every shared-memory node. For large systems with  $N = 32 \dots 2,048$  million particles, efficient scalability across 294,912 processors of the Blue Gene/P system JUGENE at JSC has been shown recently [1]. However, the inevitably synchronized collective tree construction phase as well as point-to-point communication during tree traversal are becoming dominant as the number of particles  $N$  per processor core  $C$  falls below  $\frac{N}{C} \lesssim 2,000$ . Thus, for smaller problem sizes, the algorithm's efficient operating point is limited to small supercomputer partitions, leading to a strong lower limit for the time-to-solution that is achievable. This strong-scaling limit is the second issue experienced by all parallel particle-based simulation methods. Their parallel scaling is intrinsically constrained and saturates as soon as the number of particles per processor is becoming too low.

## III. CONNECTING PEPC TO PFAST

Instead of decreasing  $\frac{N}{C}$  any further, the additional direction of parallelization offered by parallel-in-time methods provides an appealing possibility to further reduce time-to-solution.

In a previous study, we already applied the *Parallel Full Approximation Scheme in Space and Time* (PFAST) [6], [2] to PEPC, making use of two different levels of spatial discretization [7], [8].

PFAST is an iterative parallel-in-time solver which bases on a coupling of Parareal iterations [9] with the sweeps of a multi-level spectral deferred correction method (MLSDC, see [10], [11]). Similar to Parareal, it profits from decreasing the runtime of the coarse scheme by not only coarsening in time but also using a coarser spatial discretization. PFAST further improves this strategy by using a full approximation scheme (FAS) technique to recycle coarse grid information efficiently and to increase the accuracy of the coarse grid SDC sweeps. While coarsening in time can be easily obtained in PFAST by using less intermediate collocation points, efficient and effective strategies for spatial coarsening are an area of active research.

With its multipole-acceptance-parameter, the Barnes-Hut approach as used in PEPC already offers a convenient possibility to choose between fast, inaccurate force computation for the serial initial value propagation and a precise but slow evaluation for the parallel fine propagator. This strategy is similar to the operator-based coarsening strategy described in [11]. Another approach sketched in [11] but not investigated there is the implementation of *different physics* on both levels, which is envisaged in the current project. For the coarse level

<sup>1</sup><http://www.fz-juelich.de/ias/jsc/pepc>

an adiabatic approximation, i.e. decoupling of electronic and ionic motion in plasma simulations, seems natural at least for moderate and weak coupling between the species. Thus, the ions will for example move in front of an electronic background with reduced dynamics. On the fine level, the full electron dynamics will be taken into account. For PEPC, this will go beyond the previously used simple approach and has the potential to further promote parallel molecular dynamics simulations at their strong scaling limit. Thus it will be possible to study previously strictly separated time scales in a consistent model within one simulation.

#### IV. COARSENING VIA REDUCED PHYSICS FOR PLASMA-PHYSICAL APPLICATIONS

The field of plasma physics comprises a broad range of physical configurations as the presence of free charges is common to a multitude of physical regimes. From the cold and ultra-thin interstellar matter to high densities in bulk metals and high temperatures in fusion reactions, densities of  $10^{-5}$  to  $10^{25}$  particles per  $\text{cm}^3$  and temperatures from  $10^{-2}$  to  $10^{10}$  K are covered.

The relevance of mutual interactions in the system is characterized by the coupling parameter  $\Gamma = \frac{E^{(\text{Coul})}}{E^{(\text{therm})}}$ , which is the ratio of (Coulomb) interaction energy  $E^{(\text{Coul})} = \frac{q^2}{4\pi\epsilon_0 d}$  to thermal energy  $E^{(\text{therm})} = k_B T$ . Here,  $d = (\frac{4}{3}\pi n)^{-\frac{1}{3}}$  is the average interparticle distance,  $n$  the particle's number density,  $T$  their temperature,  $\epsilon_0$  the vacuum permittivity, and  $k_B$  Boltzman's constant.

For systems of low density or high temperature, i.e.  $\Gamma \ll 1$ , the system can be treated as an ideal gas. Approaching  $\Gamma \lesssim 1$ , a rigorous perturbation theoretical treatment that includes particle interactions in lower orders has been proven successful, see e.g. [12] for a review. However, for stronger coupling  $\Gamma \gtrsim 1$ , matter is approaching fluid and finally solid state and these methods fail as – with the exception of regular lattices – the emerging correlated structure cannot be expressed analytically.

In this regime, numerical methods such as molecular dynamics simulations are today's methods of choice to proceed towards studying cold and dense plasmas. However, the constituent particles – usually negatively charged electrons and positively charged ions – do not only differ in charge but also carry considerably different masses. This introduces very different dynamic time scales as the mass ratio between electrons and protons  $\frac{m_p}{m_e} \approx 1836$  is a lower bound for  $\frac{m_{\text{ion}}}{m_e}$ . For accurately resolving dynamics on all relevant scales, short physical time steps and thus long total simulation times are equally necessary.

In order to further extend the approach of combining the large-scale spatial parallelization of PEPC with the parallel-in-time method PFASST, we want to make use of this vast dynamic range of time scales for the particle species by decoupling electronic and ionic motion. To this end, the coarse level will fully resolve slow processes while fast dynamics is treated in an averaged or simplified model. The fine level will include full dynamics for slow and rapidly moving particle species.

Among others, the following three applications seem to provide an ideal setup for upcoming studies:

*A. Laser-driven ion acceleration.* In laser-driven ion acceleration, usually charged particles are expelled from a micrometer-sized metallic target by a strong external field. Using an appropriate configuration of laser and target parameters, good directionality of the escaping particles can be achieved. The interaction between travelling ions and electrons is expected to keep the particle bunches compact after leaving the target. Beam propagation outside the target requires long-running open-boundary simulations that have to resolve the dynamics of the electronic and the ionic subsystem, which is currently impossible for the real physical mass ratio given above.

Here, electrons can for example be treated as a single charge cloud on the coarse level.

*B. Instabilities at a plasma-vacuum boundary for magnetized systems.* Due to thermally driven charge separation at a plasma-vacuum interface and the  $\vec{E} \times \vec{B}$  drift resulting from the restoring motion, instabilities can emerge at a plasma-vacuum interface in an external magnetic field. These are currently studied using a two-dimensional simulation with PEPC with the constraint of resolving the gyro-motion of electrons only as the period of the cyclotron motion scales with the particles mass. To sufficiently include ionic gyration that will presumably have considerable influence onto the instability pattern structure, much longer simulation times than currently reachable would be necessary. Here, on the coarse propagator level a guiding-center movement for the rapid electron gyration can be used to allow for larger timesteps and progress of the ionic motion. On the fine level, fully resolved dynamics will include electronic gyration.

*C. Optical and transport properties in strongly correlated media.* Optical properties such as the dielectric function that finally describes propagation of electromagnetic waves inside matter is primarily governed by the rapidly moving quasi-free electrons. However, their trajectories and collision rates strongly depend on the ionic structure and the correlations between electrons and ions. The preparation of strongly correlated, i.e. cold high-density systems in simulations is difficult, though. The formation of long-range structures at phase transitions from gaseous via liquid to bulk state is a slow process as the ions involved are rather heavy. Again, using a parallel-in-time approach for finally resolving both – fast and slow dynamics – will improve our understanding of these phase transitions, the resulting structures and finally optical properties of strongly correlated media.

The projected combination of PFASST and PEPC with spatial coarsening via a *reduced physics approach* will not only boost these plasma physics applications. It will also be an important cornerstone towards efficiently utilizing the ever-increasing computational power of today's supercomputers and their even more powerful successors and will thus pave the ground for even shorter time-to-solution.

#### REFERENCES

- [1] M. Winkel, R. Speck, H. Hübner, L. Arnold, R. Krause, and P. Gibbon, "A massively parallel, multi-disciplinary Barnes-Hut tree code for extreme-scale N-body simulations," *Comput. Phys. Commun.*, vol. 183, no. 4, pp. 880–889, Apr. 2012.
- [2] M. Emmett and M. L. Minion, "Toward an efficient parallel in time method for partial differential equations," *Comm. App. Math. and Comp. Sci.*, vol. 7, pp. 105–132, 2012.
- [3] J. E. Barnes and P. Hut, "A hierarchical  $O(N \log N)$  force-calculation algorithm," *Nature*, vol. 324, no. 6096, pp. 446–449, Dec. 1986.
- [4] L. Greengard and V. Rokhlin, "A fast algorithm for particle simulations," *J. Comput. Phys.*, vol. 73, no. 2, pp. 325–348, Dec. 1987.
- [5] S. Pfalzner and P. Gibbon, *Many-Body Tree Methods in Physics*. Cambridge: Cambridge University Press, 1996.
- [6] M. L. Minion, "A hybrid parareal spectral deferred corrections method," *Comm. App. Math. and Comp. Sci.*, vol. 5, no. 2, pp. 265–301, 2010.
- [7] R. Speck, D. Ruprecht, R. Krause, M. Emmett, M. L. Minion, M. Winkel, and P. Gibbon, "A massively space-time parallel N-body solver," in *Proc. of the Int. Conf. on High Performance Computing, Networking, Storage and Analysis*, ser. SC '12, Salt Lake City, Utah. Los Alamitos, CA, USA: IEEE Computer Society Press, 2012, pp. 92:1–92:11.
- [8] —, "Integrating an N-body problem with SDC and PFASST," in *21st Int. Conf. on Domain Decomposition Methods*, [in press], Rennes, France, 2012.
- [9] J.-L. Lions, Y. Maday, and G. Turinici, "A "parareal" in time discretization of PDE's," *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, vol. 332, pp. 661–668, 2001.
- [10] A. Dutt, L. Greengard, and V. Rokhlin, "Spectral Deferred Correction Methods for Ordinary Differential Equations," *BIT Numerical Mathematics*, vol. 40, no. 2, pp. 241–266, 2000.
- [11] M. Bolten, M. Emmett, R. Krause, M. Minion, D. Ruprecht, and R. Speck, "A multi-level spectral deferred correction method," (In preparation).
- [12] H. Reinholz, "Dielectric and optical properties of dense plasmas," *Annales de Physique*, vol. 30, no. 4-5, pp. 1–187, 2005.

## Time Parallelization for Second-Order Equations

Michael Minion  
Stanford University

This talk will discuss recent work on extending the parallel full approximation scheme in space and time (PFASST) to second-order differential equations, in particular Hamiltonian systems. I will present a new variant of spectral deferred corrections appropriate for second-order equations and discuss options for generating a hierarchy of time-space problems to increase the parallel efficiency of PFASST. I will also discuss some recent progress on applying these ideas to molecular dynamics simulations using the OpenMM framework.

# Parallel Space-time Spacetrees for Simple Parabolic Benchmarks

Tobias Weinzierl

Technische Universität München,  
Boltzmannstrasse 3,  
85748 Garching  
weinzier@in.tum.de

Space-time schemes for parabolic partial differential equations (PDEs) promise advantages compared to traditional time stepping: First guesses of the solution’s evolution in time are delivered quickly (an important ingredient for computational steering or interactive computing), periodic boundary conditions are straightforward to implement (important for steady cycle systems), backward problems interacting with the PDE can be solved directly (important for optimisation and calibration based upon the adjoint), and so forth [4]. The most prominent selling points of space-time however are the increased level of concurrency and superconvergence exploiting solution smoothness in time (see for example [1] for historic remarks). All these properties render space-time a promising candidate for the dawning exascale age.

Three major showstoppers for space-time meshings do exist: memory, efficient solvers, and software complexity. We propose to pick up the concept of spacetimes, a generalisation of the well-known quadtree/octree concept for two and three dimensions, to use three-partitioning, and to apply it to a four-dimensional space-time setting: here, a bounded space-time domain is embedded into a  $4d$ -hypercube. This cube then is cut into  $3^4$  equally-sized subcubes<sup>1</sup>. The process continues recursively and yields a cascade of adaptive Cartesian grids. Such grids can be traversed and serialised along a space-filling curve and thus stored efficiently with basically one bit per vertex [2, 5]. Such a grid cascade delivers a multiscale representation of the computational domain well-suited for geometric multigrid solvers. Such a grid finally has simple tensor-product structure. Existing solvers for Cartesian meshes with standard single step time integration can straightforwardly be rewritten as space-time code.

The present talk studies a plain geometric multigrid solver acting on the space-time domain, and it discusses some of the space-time advantages: We observe that the spacetime’s adaptivity in both space and time can, if multiple snapshots at different time steps have to be held anyway, coarse more aggressively than time stepping and provide a nice framework to realise local time stepping in combination with dynamic adaptivity in space. We observe that the spacetime’s multiscale data structure can, in combination with a full approximation storage scheme, yield first guesses of the solution quickly and, in accordance with textbook knowledge, superconverge. We observe that the additional temporal

---

<sup>1</sup> The factor three results from the use of the PDE framework Peano [3] as software base.



degree of freedom increases the concurrency. Different to traditional approaches, the spacetree allows us to distribute not only time slices or follow a spatial domain decomposition, but it facilitates to deploy whole space-time subdomains to different ranks. However, the load balancing for such a data structure is delicate, the adaptivity in time interplays with the communication and data flow, the  $h$ -adaptivity of the spacetree suffers from a lack of accuracy at the domain boundaries due to the  $\mathcal{O}(h)$  accuracy, and so forth. Some of these issues and potential solutions are addressed and sketched.

## References

1. M. Gander and S. Vandewalle. Analysis of the parareal time-parallel time-integration method. *SIAM Journal on Scientific Computing*, 29(2):556–578, 2007.
2. T. Weinzierl. *A Framework for Parallel PDE Solvers on Multiscale Adaptive Cartesian Grids*. Verlag Dr. Hut, 2009.
3. T. Weinzierl et al. Peano—a Framework for PDE Solvers on Spacetree Grids, 2012. [www.peano-framework.org](http://www.peano-framework.org).
4. T. Weinzierl and T. Köppl. A geometric space-time multigrid algorithm for the heat equation. *Numerical Mathematics: Theory, Methods and Applications*, 5(1):110–130, February 2012.
5. T. Weinzierl and M. Mehl. Peano – a traversal and storage scheme for octree-like adaptive cartesian multiscale grids. *SIAM Journal on Scientific Computing*, 33(5):2732–2760, October 2011.

# Parallel variational space–time methods for the wave equation

Uwe Köcher<sup>†</sup>, Markus Bause<sup>†</sup>

<sup>†</sup>Helmut Schmidt University, University of the Federal Armed Forces Hamburg  
Department of Mechanical Engineering · Numerical Methods in Computational Engineering  
Holstenhofweg 85, 22043 Hamburg, Germany  
{koecher, bause}@hsu-hamburg.de

**Keywords:** parallel variational space–time methods, higher order space–time finite element methods.

## Abstract

The accurate and reliable numerical approximation of the hyperbolic wave equation is of fundamental importance to the simulation of acoustic, electromagnetic and elastic wave propagation phenomena. Here, we present families of variational space–time discretisation methods of higher order for the acoustic wave equation as a prototype model for the three-dimensional elastic wave equation. In temporal domain we present two families of continuous and discontinuous discretisation schemes. For the spatial discretisation a symmetric interior penalty discontinuous Galerkin method is used. From these classes of uniform Galerkin discretisations in space and time an approach of fourth-order accuracy is analysed carefully. The efficient solution of the resulting block-matrix system and inherently parallel numerical simulation through domain decomposition is addressed. The performance properties of the schemes are illustrated by sophisticated and challenging numerical experiments with complex wave propagation phenomena in heterogeneous media.

## 1 Motivation

Our interest in developing numerical approximation schemes of higher order accuracy for the wave equation comes from mechanical engineering. Multiple layer fibre reinforced composites have become one of the most promising materials to build light-weight structures for several fields of application, for instance in aerospace and automotive fields. These composites are able to combine high strength and rigidity of the reinforced fiber with excellent properties of synthetic resins in best possible way. Nondestructive material inspection by piezoelectric induced ultrasonic waves is a relatively new and challenging technique to monitor the healthiness of such components. Several material damages (delamination of layers, matrix cracks, fibre breaks) may occur and have to be detected by some autonomous structural health monitoring system. For the design of structural health monitoring systems it

is strictly necessary to understand phenomenologically and quantitatively wave propagation in layered fibre reinforced composites and the influence of the geometrical and mechanical properties of the system structure. Numerical simulation is a promising way to achieve this goal; cf. [2] and Fig. 1. Therefore, the ability to solve numerically the wave equation in three space dimensions is particularly important from the point of view of physical realism. However, this is still a challenging task and an active area of research.




Figure 1: Ultrasonic waves in carbon fibre composite.

## 2 Introduction to variational space–time methods

In the field of numerical wave propagation the spatial discretisation by some discontinuous Galerkin finite element method (DGM) has attracted the interest of researchers; cf., e.g. [4]. Advantages of the FEM are the flexibility with which it can accommodate discontinuities in the model, material parameters and boundary conditions and the ability to approximate the wavefield with high degree polynomials. The spatial DGM has the further advantage that it can accommodate discontinuities also in the wavefield, it can be energy conservative, and it is suitable for inherently parallel simulations. The mass matrix of the spatial DGM is block-diagonal, where each block size coincides with the degrees of freedom of the associated element, such that its inverse is available at very low computational cost. Recently, variational space–time discretisation schemes were proposed and studied for the parabolic heat equation and for systems of ordinary differential equations [3]. In this contribution we will focus on the presentation of continuous and discontinuous variational temporal discretisation schemes from the variational

space–time approach for the hyperbolic wave equation. For the spatial discretisation a symmetric interior penalty discontinuous Galerkin method is used; cf. [4].

From these families of uniform variational discretisations a scheme of fourth order accuracy with respect to the temporal and spatial variables is studied carefully. It will be shown that the block-diagonal structure of the spatial mass matrix, resulting from the discontinuous Galerkin approach, can be used to decouple efficiently the arising temporal block linear system. The performance properties and computational cost of the numerical scheme are illustrated by some numerical convergence studies. Moreover, the schemes are applied to wave propagation phenomena in heterogeneous media admitting multiple sharp wave fronts; cf. [1] and Fig. 2.




Figure 2: Inherently parallel acoustic wave simulation.

### 3 Temporal discretisation schemes

Exemplarily, our family of continuous-in-time variational discretisation schemes for the wave equation

$$\begin{aligned} \rho(\mathbf{x}) \partial_t v(\mathbf{x}, t) + a(u(\mathbf{x}, t)) &= f(\mathbf{x}, t), \\ \partial_t u(\mathbf{x}, t) - v(\mathbf{x}, t) &= 0, \end{aligned} \quad (1)$$

written as velocity-displacement formulation and equipped with initial conditions  $u(0) = u_0$ ,  $v(0) = v_0$  and homogeneous Dirichlet boundary conditions, is presented here. For the discontinuous-in-time counterpart of this approach we refer to [1].

We decompose  $I = (0, T]$  into  $N$  subintervals  $I_n = (t_{n-1}, t_n]$ . For some Hilbert space  $\mathcal{H}$ , let

$$\begin{aligned} \mathcal{X}_C^r(\mathcal{H}) &= \{u \in \mathcal{C}(I, \mathcal{H}) : u|_{I_n} \in \mathbb{P}_r(I_n, \mathcal{H})\}, \\ \mathcal{Y}^r(\mathcal{H}) &= \{w \in L^2(I, \mathcal{H}) : w|_{I_n} \in \mathbb{P}_{r-1}(I_n, \mathcal{H})\}, \\ \mathbb{P}_r(I_n, \mathcal{H}) &= \left\{ u : I_n \rightarrow \mathcal{H} : u = \sum_{j=0}^r \xi_n^j t^j, \xi_n^j \in \mathcal{H} \right\}. \end{aligned}$$

Our continuous-in-time variational approximation of (1) then reads as: Find  $u_\tau \in \mathcal{X}_C^r(I, H_0^1(\Omega))$ ,  $v_\tau \in \mathcal{X}_C^r(I, L^2(\Omega))$ , such that  $u_\tau(0) = u_0$ ,  $v_\tau(0) = v_0$  and

$$\begin{aligned} \int_0^T \left\{ (\partial_t v_\tau, \hat{w}_\tau)_\Omega + a(u_\tau, \hat{w}_\tau) \right\} dt &= \int_0^T (f, \hat{w}_\tau)_\Omega dt \\ \int_0^T \left\{ (\partial_t u_\tau, \tilde{w}_\tau)_\Omega - (v_\tau, \tilde{w}_\tau)_\Omega \right\} dt &= 0 \end{aligned}$$

for all  $\hat{w}_\tau \in \mathcal{Y}^r(H_0^1(\Omega))$  and  $\tilde{w}_\tau \in \mathcal{Y}^r(L^2(\Omega))$ . Here,  $(\cdot, \cdot)_\Omega$  denotes the  $L^2(\Omega)$  inner product. Precisely, we have a temporal Galerkin-Petrov method, since the trial and test spaces do not coincide. Since the test space imposes no continuity constraints between elements, we can rewrite the problem as time marching scheme. Finally, we apply an interior penalty discontinuous Galerkin method in spatial domain and call this approach a cGP(r)-dG(p) method. We obtain the Crank-Nicolson scheme for  $r = 1$  and for  $r \geq 2$  we observe superconvergence in the integration points, cf. [1, 3].

### 4 Future Prospects

By using variational space–time methods for the discretisation of the wave equation we have a uniform variational approach in space and time which may be advantageous for the future analysis of the fully discrete problem and the construction of simultaneous space–time adaptive methods. Further, it is very natural to construct temporal methods of even higher order than presented here. The well-known finite element stability concepts of the temporal Galerkin-Petrov or discontinuous Galerkin methods can be used to obtain at least A-stable methods. For future developments, the well-known adaptive finite element techniques can be applied for changing the polynomial degree and the length of the time steps, cf. [3, 6, 5].

### References

- [1] U. Köcher, M. Bause: Variational space–time methods for the wave equation, in preparation, 2013.
- [2] M. Bause, U. Köcher: Parallel numerical simulation of elastic wave propagation in fiber reinforced material, submitted, 2012.
- [3] G. Matthies, F. Schieweck: Higher order variational time discretizations for nonlinear systems of ordinary differential equations, Preprint No. 23/2011, OvG-Universität Magdeburg, 2011.
- [4] M. J. Grote, D. Schötzau: Optimal error estimates for the fully discrete interior penalty DG method for the wave equation, J. Sci. Comput., 40:257–272, 2009.
- [5] J. Diaz, M. J. Grote: Energy conserving explicit local time stepping for second-order wave equations, J. Sci. Comput., 31(3):1985–2014, 2009.
- [6] W. Bangerth, M. Geiger, R. Rannacher: Adaptive Galerkin finite element methods for the wave equation, Comput. Meth. Appl. Math., 10(1):3–48, 2010.

# RIDC-DD: A parallel space-time algorithm

Ronald Haynes <sup>α</sup> and Benjamin Ong <sup>β</sup>

<sup>α</sup> Memorial University of Newfoundland, Saint Johns, NL, Canada

<sup>β</sup> Michigan State University, East Lansing, MI, USA

We present recent efforts to develop fully parallel space-time algorithms, which will improve the fidelity and scalability of numerical simulations across many areas in computational science. Our approach combines Revisionist Integral Defect Correction (RIDC) methods [2, 3], a family of time-parallel integrators, with Domain Decomposition (DD) methods [6, 5], an approach to split boundary value problems into smaller boundary value problems on sub-domains and iterating to coordinate the solution between adjacent sub-domains. The basic idea is as follows: while a set of  $N$  processing cores are computing a low-order approximation to the domain-decomposed solution at time  $t^n$ , additional sets of  $N$  cores are simultaneously computing corrections to the low(er)-order approximations at time  $t^m$ , where  $m < n$ .

Consider the general problem of interest:

$$u_t = \mathcal{L}(t, u), \quad x \in \Omega \times [0, T] \quad (1a)$$

$$\mathcal{B}(t, u) = 0, \quad x \in \partial\Omega \times [0, T] \quad (1b)$$

$$u(0, x) = g(x), \quad x \in \Omega. \quad (1c)$$

Denote an approximate solution to (1) as  $\eta(t, x)$ , and the residual as  $\epsilon(t, x) = \eta_t - \mathcal{L}(t, \eta)$ . Then, the following correction PDE can be derived,

$$\left[ e + \int_0^t \epsilon(\tau, x) d\tau \right]_t = \mathcal{L}(t, \eta + e) - \mathcal{L}(t, \eta), \quad x \in \Omega \times [0, T] \quad (2a)$$

$$e(t, x) = 0, \quad x \in \partial\Omega \times [0, T] \quad (2b)$$

$$e(0, x) = 0, \quad x \in \Omega, \quad (2c)$$

where  $e$  is the correction to an approximate solution  $\eta$ . *Pipeline* parallelism can be employed by staggering solutions to the PDE of interest (1) and multiple correction PDEs (2) (after initial start-up costs).

*Data* parallelism is layered on top of the time (pipeline) parallelism through DD. As a concrete example, discretize (1) and (2) using a first-order backwards Euler integrator. After algebraic manipulations, one recovers a model boundary value problem from both (1) and (2),

$$(1 - \alpha\mathcal{L})u = f(x), \quad x \in \Omega \quad (3a)$$

$$\mathcal{C}(u) = 0, \quad x \in \partial\Omega. \quad (3b)$$

Decomposing the computational domain  $\Omega$  into  $N$  non-overlapping domains,  $\Omega = \bigcup_{i=1}^N \Omega_i$ , one reformulates (3) into the coupled system of boundary value problems. Also known as a Schwarz-type approach, the reformulated problem is to find  $(u_i)_{1 \leq i \leq N}$  such that

$$(1 - \alpha\mathcal{L})u_i = f(x), \quad x \in \Omega_i \quad (4a)$$

$$\mathcal{C}(u_i) = 0, \quad x \in \partial\Omega_i \cap \partial\Omega, \quad (4b)$$

$$\mathcal{T}_{ij}(u_i) = \mathcal{T}_{ij}(u_j), \quad x \in \partial\Omega_i \cap \partial\Omega_j, \quad (4c)$$

where  $(\mathcal{T}_{ij})_{1 \leq i, j \leq N}$  are transmission conditions on the interfaces between the sub-domains. The coupled system (5) is solved iteratively using a Jacobi algorithm,

$$(1 - \alpha\mathcal{L})u_i^{k+1} = f(x), \quad x \in \Omega_i \quad (5a)$$

$$\mathcal{C}(u_i^{k+1}) = 0, \quad x \in \partial\Omega_i \cap \partial\Omega, \quad (5b)$$

$$\mathcal{T}_{ij}(u_i^{k+1}) = \mathcal{T}_{ij}(u_j^k), \quad x \in \partial\Omega_i \cap \partial\Omega_j. \quad (5c)$$

Since this formulation of a parallel space–time algorithm involves pipeline and data parallelism, our implementation uses multiple modes of programming. OpenMP is used to handle the pipeline (time) parallelism; MPI is used to handle the data (DD) parallelism. The hybrid OpenMPI–MPI framework used to is discussed in [1, 4]. The C++ software was compiled and bench-marked on a general purpose x86 cluster. Each node, consisting of dual-socket Intel Sandy Bridge processors, are connected using FDR infiniband. The linear heat equation in 2D with optimized transmission conditions was used as a test problem.




Figure 1: The linear heat equation is solved in  $\mathbb{R}^2$ . The domain is subdivided into a  $10 \times 10$  grid of non-overlapping subdomains. A second order finite discretization is used spatially, and an eighth-order RIDC method is used for the time discretization. The number of concurrent threads that the RIDC integrator is allowed to use is varied. This result shows that we can take a spatially parallel code, and layer on time-parallelism with relatively high efficiency.

## References

- [1] A. Christlieb, R. Haynes, and B. Ong. A parallel space-time algorithm. *SIAM J. Sci. Comput.*, 34(5):233–248, 2012.
- [2] A. Christlieb, C. Macdonald, and B. Ong. Parallel high-order integrators. *SIAM J. Sci. Comput.*, 32(2):818–835, 2010.
- [3] A. Christlieb and B. Ong. Implicit parallel time integrators. *J. Sci. Comput.*, 49(2):167–179, 2011.
- [4] R. Haynes and B. Ong. A hybrid MPI-OpenMP algorithm for the parallel space-time solution of time dependent PDEs. In *Domain Decomposition Methods in Science and Engineering XXI*, 2013. to appear.
- [5] T. Mathew. *Domain decomposition methods for the numerical solution of partial differential equations*, volume 61 of *Lecture Notes in Computational Science and Engineering*. Springer-Verlag, Berlin, 2008.
- [6] A. Toselli and O. Widlund. *Domain decomposition methods—algorithms and theory*, volume 34 of *Springer Series in Computational Mathematics*. Springer-Verlag, Berlin, 2005.

# Deferred correction from equispaced data based on efficient high-order rational integration

Georges Klein\* with Stefan Güttel†

In this talk, we present *rational deferred correction* (RDC) methods for the solution of initial value problems. Inspired by *spectral deferred correction* (SDC) methods from [1] by Dutt, Greengard & Rokhlin, we demonstrate that similar accuracy and stability can be achieved with equispaced points instead of Gauss–Legendre points if one resorts to the linear barycentric rational interpolants [2] constructed by Floater and Hormann.

To be more specific, we are interested in solving initial-value problems for a function  $u : [0, T] \rightarrow \mathbb{C}^N$ ,

$$u'(t) = f(t, u(t)), \quad u(0) = u_0 \in \mathbb{C}^N. \quad (1)$$

Assume we have already computed a discrete low order approximation  $\tilde{u} \approx u$ , e.g., with a forward or backward Euler scheme. In order to achieve higher precision,  $\tilde{u}$  is iteratively corrected as follows.

The problem (1) is reformulated as a Picard integral to avoid numerical differentiation,

$$u(t) = u(0) + \int_0^t f(\tau, u(\tau)) \, d\tau, \quad (2)$$

or equivalently, with  $e = u - \tilde{u}$  the approximation error,

$$\tilde{u}(t) + e(t) = u(0) + \int_0^t f(\tau, \tilde{u}(\tau) + e(\tau)) \, d\tau. \quad (3)$$

Using (2) to define the residual

$$r(t) = u(0) + \int_0^t f(\tau, \tilde{u}(\tau)) \, d\tau - \tilde{u}(t), \quad (4)$$

we immediately find from (3)

$$e(t) = r(t) + \int_0^t f(\tau, \tilde{u}(\tau) + e(\tau)) - f(\tau, \tilde{u}(\tau)) \, d\tau, \quad (5)$$

which is a Picard-type formulation for the error  $e$ . Equation (5) can then be solved with the same time-stepping method which was used to obtain the initial approximation  $\tilde{u}$ . With this estimation of the error  $\tilde{u}$  is corrected via  $\tilde{u}_{\text{new}} = \tilde{u} + e$ , so as to conclude one deferred correction sweep. The procedure can be iterated until stagnation occurs, typically when the precision of the collocation solution of (1) is reached.

Since the approximation  $\tilde{u}$  is available only as  $\tilde{u}_j$  at discrete values of the time variable once (2) is solved, for (4) to make sense, a continuous approximation must be constructed from the  $\tilde{u}_j$ , e.g., via interpolation. The procedure proposed in [1] involves polynomial interpolants of degree  $n$ , where  $n$  can be quite large to achieve sufficient accuracy. To prevent this polynomial interpolation from being unstable, it was advocated in [1] to interpolate the integrand in (4) at Gauss–Legendre points

---

\*Mathematical Institute, University of Oxford, 24-29 St Giles', Oxford OX1 3LB, United Kingdom. Supported by the Swiss National Science Foundation under Grant No. PBFPR2-142826

†School of Mathematics, University of Manchester, Oxford Road, Manchester M13 9PL, United Kingdom

or Chebyshev points. Because such an interpolation process will achieve spectral accuracy in time, the resulting method is called spectral deferred correction (SDC).

Recent investigations have revealed that linear rational interpolation [2] with equispaced points can be stable and achieve high accuracy even with large numbers of points. We define rational deferred correction (RDC) as the analogue of the above described SDC with linear *rational* interpolation and *equispaced* points.

Before we present the RDC integrator in more detail and compare its performance with SDC, we review the construction and properties of the linear rational interpolation scheme. Assume we are given  $n + 1$  points  $t_0 < t_1 < \dots < t_n$  in a closed interval  $[a, b]$  and corresponding values of a function  $f_0, f_1, \dots, f_n$ . Each choice of the nonnegative parameter  $d \leq n$  defines a member of the family of rational interpolants,

$$r_n(t) = \frac{\sum_{i=0}^{n-d} \lambda_i(t) p_i(t)}{\sum_{i=0}^{n-d} \lambda_i(t)}, \quad \lambda_i(t) = \frac{(-1)^i}{(t - t_i) \cdots (t - t_{i+d})},$$

as a blend of polynomial interpolants of degree  $d$ . With equispaced points in particular and an adequate choice of  $d$ , which we will explain, these interpolants are well-conditioned and lead to approximations of order  $d + 1$  of sufficiently smooth functions. Moreover the linearity in the data conveniently leads to quadrature rules of order  $d + 2$ , and makes them very appealing for the solution of initial value problems as described above.

We review theoretical results, present numerical examples and compare some properties of RDC with those of SDC. Our investigations are initial studies of integration methods based on linear rational interpolation and can be extended in numerous ways.

## References

- [1] A. Dutt, L. Greengard, and V. Rokhlin. Spectral deferred correction methods for ordinary differential equations. *BIT*, 40:241–266, 2000.
- [2] M. S. Floater and K. Hormann. Barycentric rational interpolation with no poles and high rates of approximation. *Numer. Math.*, 107:315–331, 2007.

Yvon Maday  
Laboratoire Jacques-Louis Lions  
Universit Pierre et Marie Curie

Conservations with parareal in time algorithm

In this presentation we shall cover the basics of the parareal in time algorithm, we shall recall the theoretical reasons why it works for dissipative problems. We shall present some ingredients that needs to be added in order to tackle problems for which conservation is important.



# MINARET: Towards a time-dependent neutron transport parallel solver

Baudron, A.-M.<sup>d,e</sup>, Lautard, J.J.<sup>d,e</sup>, Maday, Y.<sup>a,b,c,e</sup> and Mula, O.<sup>a,d,e</sup>

<sup>a</sup> UPMC Univ Paris 06, UMR 7598, Laboratoire Jacques-Louis Lions, F-75005, Paris, France.

<sup>b</sup> Institut Universitaire de France

<sup>c</sup> Division of Applied Mathematics, Brown University, Providence RI, USA.

<sup>d</sup> CEA Saclay, DEN/DANS/DM2S/SERMA/LLPR, 91191 Gif-Sur-Yvette CEDEX - France

<sup>e</sup> LRC MANON, Laboratoire de Recherche Conventionnée – CEA/DEN/DANS/DM2S and UPMC-CNRS/LJLL.

## Abstract

The advances of computer power in the last decades have today resulted in an increased memory storage per processor and also in a raising number of available processors to compute a given task. From the numerical simulation point of view, this context has allowed the implementation of more and more accurate — and therefore more computationally expensive and time-consuming — solvers that can be run in a reasonable time if advantage is taken of launching concurrently several processors. This speed-up in the computing time requires, however, the investigation of innovative acceleration techniques such as domain decomposition methods.

In the field of nuclear core calculations and, more particularly, regarding the neutron transport equation, this context has led to significant advances from the accuracy point of view: recent developments of time-dependent neutron transport codes such as DORT-TD [1] or TORT-TD [2] have indeed overcome the traditional diffusion, improved quasi-static [3] or point kinetics traditional approximations.

However, although the problem of merging a neutron transport accurate solver with parallel computations has already been addressed before for the angular [4] and energy [5] variables in the stationary case, no study has been done so far in order to specifically speed-up time-dependent transport calculations and this work is a step in this direction. On this purpose, a multigroup 3D kinetic transport  $S_n$  code has been implemented (in a solver called MINARET [6]) and the parallelization of the angular and time variables are currently being explored. In particular, the time variable is parallelized by a domain decomposition technique: the parareal time algorithm (see [7], [8], [9]).

In this talk, after recalling the time-depend neutron transport equation, we will explain its discretization and implementation in MINARET. This will allow to illustrate the importance of the study of speed-up methods for the solver in order to run calculations in a reasonable computing time. After a brief review of the traditional sequential accelerations used in the field of neutron transport, our talk will specifically focus on the analysis of acceleration techniques that involve the parallelization of the variables of the equation. In particular, results concerning the parallelization of the angular and the time variables in MINARET will be exposed.

## Acknowledgements

This work was supported in part by the joint research program MANON between CEA-Saclay and University Pierre et Marie Curie-Paris 6.

## References

- [1] A. Pautz and A. Birkhofer, “DORT-TD: A transient neutron transport code with fully implicit time integration,” *Nuclear Science and Engineering*, vol. 145, pp. 299–319, 2003.

- [2] A. Saubert, A. Sureda, J. Bader, J. Lapins, M. BUCK, and E. Laurien, “The 3-D time-dependent transport code TORT-TD and its coupling with the 3-D thermal-hydraulic code ATTICA3D for HTGR applications,” *Nuclear Engineering and Design*, vol. 251, pp. 173–180, 2012.
- [3] E. Girardi, P. Guérin, S. Dulla, and P. Ravetto, “Comparison of direct and quasi-static methods for neutron kinetic calculations with the EDF R&D COCAGNE code,” in *PHYSOR, Advances in Reactor Physics*, 2012.
- [4] E. Jamelot, J. Dubois, J.-J. Lautard, C. Calvin, and A.-M. Baudron, “High performance 3D neutron transport on petascale and hybrid architectures within APOLLO3 code,” in *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, 2011.
- [5] R. Slaybaugh, T. Evans, G. Davidson, and P. Wilson, “Rayleigh quotient iteration in 3D, deterministic neutron transport,” in *PHYSOR, Advances in Reactor Physics*, 2012.
- [6] J.-Y. Moller and J.-J. Lautard, “MINARET, a deterministic neutron transport solver for nuclear core calculations,” in *International Conference on Mathematics and Computational Methods Applied to Nuclear Science and Engineering*, 2011.
- [7] J. Lions, Y. Maday, and G. Turinici, “Résolution d’EDP par un schéma en temps pararéel,” *C. R. Acad. Sci. Paris*, 2001. t. 332, Série I, p. 661-668.
- [8] L. Baffico, S. Bernard, Y. Maday, G. Turinici, and G. Zérah, “Parallel in time molecular dynamics simulations,” *Phys. Rev. E*, vol. 66, p. 057701, 2002.
- [9] G. Bal and Y. Maday, “A parareal time discretization for nonlinear PDE’s with application to the pricing of an American put,” in *Recent developments in domain decomposition methods, Lect. Notes Comput. Sci. Eng.*, vol. 23, pp. L.F. Pavarino and A. Toselli eds., 189–202, Springer Verlag, 2002.

# Space-Time Methods for Wave Equations

## Discretizations and Convergence Analysis

Christian Wieners

Karlsruhe Institute of Technology  
Department of Mathematics  
Kaiserstr. 89-93, 76128 Karlsruhe, Germany  
christian.wieners@kit.edu

We discuss three full space-time discretizations for linear wave equations: an adaptive discontinuous Galerkin method in space and time (dG), a discontinuous Petrov-Galerkin method (DPG), and a new hybrid variant (hDPG).

**The general setting** Let  $\Omega \subset \mathbb{R}^D$  be a bounded Lipschitz domain, and let  $V \subset L_2(\Omega)^J$  be a Hilbert space with weighted inner product  $(\mathbf{v}, \mathbf{w})_V = (M\mathbf{v}, \mathbf{w})_{0,\Omega}$ , where  $M \in L_\infty(\Omega)^{J \times J}$  is uniformly positive and symmetric. We study the evolution equation

$$M\partial_t \mathbf{u}(t) + A\mathbf{u}(t) = \mathbf{f} \quad t \in [0, T], \quad \mathbf{u}(0) = \mathbf{u}_0, \quad (1)$$

where  $A$  is a linear operator in  $V$  with domain  $\mathcal{D}(A) \subset V$  corresponding to a hyperbolic linear system, i.e.  $(A\mathbf{v}, \mathbf{v})_{0,\Omega} = 0$  for  $\mathbf{v} \in \mathcal{D}(A)$ . For simplicity, we consider only homogeneous boundary conditions on  $\partial\Omega$  which are included in the domain of the operator.

For the existence and uniqueness of the solution we require that  $V(A) \subset V$  is the closure of  $\mathcal{D}(A)$  with respect to the topology in  $V$ , that  $A$  is a closed operator in  $V(A)$  and that the operator  $M^{-1}A$  generates a semigroup in  $V(A)$ . This setting applies, e.g., to acoustic waves with  $M(\mathbf{q}, p) = (\mathbf{q}, \rho p)$  and  $A(\mathbf{q}, p) = (\nabla p, \operatorname{div} \mathbf{q})$ , to elastic waves with  $M(\boldsymbol{\sigma}, \mathbf{v}) = (\mathbf{C}^{-1}\boldsymbol{\sigma}, \rho \mathbf{v})$  and  $A(\boldsymbol{\sigma}, \mathbf{v}) = (-\boldsymbol{\varepsilon}(\mathbf{v}), -\operatorname{div} \boldsymbol{\sigma})$ , and to electro-magnetic waves with  $M(\mathbf{H}, \mathbf{E}) = (\mu \mathbf{H}, \varepsilon \mathbf{E})$ , where the operator  $A(\mathbf{H}, \mathbf{E}) = (\operatorname{curl} \mathbf{E}, -\operatorname{curl} \mathbf{H})$  is defined, e.g., in  $\mathcal{D}(A) = \{(\mathbf{H}, \mathbf{E}) \in H(\operatorname{curl}, \Omega) \times H_0(\operatorname{curl}, \Omega) : \operatorname{div}(\mu \mathbf{H}) = 0, \operatorname{div}(\varepsilon \mathbf{E}) = 0\}$  for a perfect conducting boundary.

Let  $Q = \Omega \times (0, T)$  be the space-time cylinder. We define  $H = L_2(0, T; V(A)) \subset L_2(Q)$ , and we consider the space-time operator  $L = M\partial_t + A$  defined on  $U = \mathcal{D}(L)$ , where we assume that  $U$  is closed with respect to the graph norm of  $L$ . Our analysis is based on the a-priori bound

$$\|\mathbf{u}\|_H \leq 2T \|L\mathbf{u}\|_H,$$

which shows that solution of the equation  $L\mathbf{u} = \mathbf{f}$  is unique. Note that  $L$  is skew-adjoint, i.e.,  $(L\mathbf{u}, \mathbf{v})_Q = -(\mathbf{u}, L\mathbf{v})_Q$ .

For a given decomposition of the space-time cylinder  $\bar{Q} = \bigcup \bar{\tau}$  into space-time cells  $\tau = K \times I$  we discuss three discretizations.

**A discontinuous Galerkin method (dG)** Let  $V_h \subset L_2(\Omega)^J$  be a discontinuous space of polynomials with variable degree, and let  $A_h$  be the discontinuous Galerkin operator with full upwind flux, see [1]. In every time slice  $\Omega \times (t_{n-1}, t_n)$  we use the implicit mid-point rule; note that this is unconditionally stable. This allows for an adaptive local refinement of the space-time discretization and for a parallel multigrid preconditioner on the refinement hierarchy, where the full problem on the coarsest level is solved with a parallel direct solver [2].

**A discontinuous Petrov-Galerkin method (DPG)** On every space-time cell  $\tau$ , integration by parts yields

$$(L\mathbf{u}, \mathbf{v})_\tau = -(\mathbf{u}, L\mathbf{v})_\tau + \langle \gamma_\tau \mathbf{u}, \gamma_\tau^* \mathbf{v} \rangle_\tau,$$

where  $\gamma_\tau$  is the trace operator of  $U|_\tau$  to values on  $\partial\tau$ , and  $\gamma_\tau^*$  is its adjoint. Let  $\hat{U} = \prod \gamma_\tau(U|_\tau)$  be the global trace space. The DPG method computes an approximation  $(\hat{\mathbf{u}}_\tau) \in \hat{U}$  of the trace and local discontinuous approximations  $\mathbf{u}_\tau \in L_2(\tau)^J$  such that

$$-(\mathbf{u}_\tau, L\mathbf{v}_\tau)_\tau + \langle \hat{\mathbf{u}}_\tau, \gamma_\tau^* \mathbf{v}_\tau \rangle_\tau = (\mathbf{f}, \mathbf{v}_\tau)_\tau$$

for all test functions  $\mathbf{v}_\tau$  in an optimal test space. In this setting, optimal space-time a priori estimates exist [4], and the trace approximation is determined by a symmetric positive definite Schur complement problem which can be preconditioned efficiently by multigrid methods [3].

**A hybrid discontinuous Petrov-Galerkin method (hDPG)** A hybrid variant of the DPG method is formally obtained by choosing a nonconforming approximation of the trace space. This makes the data structure and a parallel communication more flexible. Again this allows for a symmetric positive definite Schur complement reduction.

## References

- [1] M. Hochbruck, T. Pazur, A. Schulz, E. Thawinan, and C. Wieners. Efficient time integration for discontinuous Galerkin approximations of linear wave equations. 2013. In preparation.
- [2] D. Maurer and C. Wieners. A parallel block LU decomposition method for distributed finite element matrices. *Parallel Computing*, 37:742–758, 2011.
- [3] C. Wieners. A geometric data structure for parallel finite elements and the application to multigrid methods with block smoothing. *Computing and Visualization in Science* 13:161-175, 2010.
- [4] C. Wieners and B. Wohlmuth. Robust operator estimates and the application to substructuring methods for first-order systems. 2013. Submitted.

# First step towards Parallel and Adaptive Computation of Maxwell's Equations

Stefan Findeisen

Karlsruhe Institut of Technolgy  
Departement of Mathematics  
Kaiserstr. 12, 76128 Karlsruhe, Germany  
stefan.findeisen@kit.edu

Maxwell's equations describe the behavior of an electromagnetic wave in three spacial dimensions over a certain period  $[0, T]$ . The wave consists of two fields, the electric  $\mathbf{E}$  and the magnetic  $\mathbf{H}$ , respectively. They can be computed by the linear first-order Maxwell system

$$\mu \partial_t \mathbf{H} + \nabla \times \mathbf{E} = 0, \quad \varepsilon \partial_t \mathbf{E} - \nabla \times \mathbf{H} = 0, \quad \nabla \cdot (\mu \mathbf{H}) = 0, \quad \nabla \cdot (\varepsilon \mathbf{E}) = 0 \quad (1)$$

with permeability  $\mu$  and permittivity  $\varepsilon$ . Each field consists of three components. Hence one has to compute six components of the fields, which depend on space and time. For a given problem this can lead to huge linear systems. This is why fast space-time codes are needed to solve the problem in a reasonable computation time.

Together with an initial condition  $\mathbf{u}_0$ , (1) can be written as

$$M \partial_t \mathbf{u}(t) + A \mathbf{u}(t) = \mathbf{0} \text{ for } t \in [0, T], \quad \mathbf{u}(0) = \mathbf{u}_0,$$

on a bounded Lipschitz domain  $\Omega \subset \mathbb{R}^3$ , where  $M$ ,  $A$ ,  $u$  are given by

$$M := \begin{bmatrix} \mu & 0 \\ 0 & \varepsilon \end{bmatrix}, \quad A := \begin{bmatrix} 0 & \nabla \times \\ -\nabla \times & 0 \end{bmatrix}, \quad \mathbf{u} := \begin{bmatrix} \mathbf{H} \\ \mathbf{E} \end{bmatrix}.$$

Using a discontinuous Galerkin discretization with upwind flux for the spacial discretization, we receive the semi-discrete system of equation

$$M_h \partial_t \mathbf{u}_h(t) + A_h \mathbf{u}_h(t) = 0 \text{ for } t \in [0, T], \quad \mathbf{u}_h(0) = \mathbf{u}_{h,0},$$

where  $M_h$  and  $A_h$  are the corresponding mass and flux matrices, respectively. In order to avoid a CFL (Courant-Friedrichs-Lewy) condition we use the implicit midpoint rule for time integration, which is of order two. Hence we have to solve a linear system  $(M_h + \frac{\tau}{2} A_h) v = b$  in every time step. However an implicit method allows a larger step size  $\tau$  than an explicit scheme.

In the following we present an adaptive programming model which is able to solve the 2D reduction of Maxwell's equations, e.g.,  $\mathbf{u} = (\mathbf{H}_1, \mathbf{H}_2, \mathbf{E}_3)$  and  $\mathbf{H}_3 \equiv \mathbf{E}_1 \equiv \mathbf{E}_2 \equiv 0$  and  $\Omega \subset \mathbb{R}^2$ . The discretization of

our space–time domain  $Q := \Omega \times [0, T]$  is done as follows. First  $\Omega$  is decomposed into a finite number of open elements (e.g. triangles)  $K \subset \Omega$  such that  $\bar{\Omega} = \bigcup \bar{K}$ . Analogously  $[0, T]$  is decomposed into a finite number of  $N$  open intervals  $I_n = (t_n, t_{n+1}) \subset [0, T]$  for  $n = 0, \dots, N - 1$  such that  $[0, T] = \bigcup_{n=0}^{N-1} \bar{I}_n$ . Now we are able to define space–time cells  $K_I := K \times I$  which consists of a spacial element  $K$  and a time interval  $I$ . Hence  $Q$  can be decomposed into a finite number of open space–time elements  $K_I \subset Q$  such that  $\bar{Q} = \bigcup \bar{K}_I$ . The data structure is organized as hash maps. That means that every space–time cell and its components (such as faces, edges, vertices) is identified by its geometric midpoint and stored in a hash map (where the midpoints are used as hash keys). By doing so, it is easy to distribute the space–time cells among the different processes and solve the problem parallel.

So far, our code is  $p$ -adaptive and uses different time steps. That means that the order of a polynomial on a space–time cell is adapted and the space–time cells can be refined in time. In a first step we compute our solution  $\mathbf{u}_h$  on a slice  $S_n := \Omega \times [t_n, t_{n+1}] \subset Q$  with fixed polynomial degree  $p = 0$ . Then we take the flux over the cell faces as an indicator where the polynomial degree should be increased and cells should be refined in time. To be more precise, the indicator  $n_{K_I}$  of a cell  $K_I$  is given by the sum of the face indicators  $\eta_f$ :

$$\eta_{K_I}^2 := \sum_{f \in \mathcal{F}_{K_I}} \eta_f^2,$$

$$\eta_f^2 := h_f \|(\mathbf{F}^*(\mathbf{u}_h) - \mathbf{F}(\mathbf{u}_h)) \cdot \mathbf{n}_{K_I}\|_{L_2}^2,$$

where  $F$  and  $F^*$  are the flux and the numerical flux and  $h_f$ ,  $\mathbf{n}_f$  are the area and the outer normal vector of the face  $f$ . The polynomial degree is increased, if the indicator of the cell fulfills the inequality

$$\eta_{K_I} > (1 - \theta) \max_{K_I \in S_n} \eta_{K_I}$$

for a given parameter  $\theta$  (e.g.  $\theta = 0.9$ ). After that we recompute the solution with the new distribution of the polynomial degrees and time refined cells. Since our code provides polynomials up to order  $p = 4$ , four  $p$ -adaptive refinements are possible. Although the computation of  $\eta_{K_I}$  is heuristic, it leads to good results in a sense that high polynomial degrees are used in areas where a single wavefront is located. In areas with absence of a wave, lowest polynomial degrees are used.

In the future the described programming model will be the basis for a space–time adaptive code to solve Maxwell’s equation in three spacial dimensions. Thus our next steps will be  $h$ -adaptivity in space.

Julien Salomon  
Universite Paris-Dauphine

-----  
An intermediate state method for the time-parallelized solving of  
optimal control problems.

Time parallelization is an efficient way to accelerate the numerical  
solving of optimality systems associated to control problems. In  
this talk, we present a general approach to parallelize optimal  
control solvers. This method is based on the introduction of  
intermediate states and enables to decompose the original optimality  
system into similar sub-systems. these ones can then be treated  
independently using standard solvers. We present a recent improving  
on the method that makes it fully efficient and discuss the role of  
the solver used in parallel.

-----

# Multirate time integration of the Euler equations

Joerg Wensch (TU Dresden)

The simulation of atmospheric dynamics relies on the numerical solution of the Euler equations. These equations exhibit phenomena on different temporal scales. Sound waves propagate approximately ten times faster than the advective waves. An approach to overcome the CFL restriction caused by sound waves are split-explicit methods. By multirate techniques the terms relevant for sound waves are integrated by small time steps with a cheap time integration procedure, whereas the slow processes are solved by an underlying Runge-Kutta method using a larger macro step size. We construct such methods based on TVD-RK schemes and discuss order and stability properties.



Qiqi Wang  
Department of Aeronautics and Astronautics  
Massachusetts Institute of Technology

## Towards Scalable Parallel Long Time Integration of Chaotic Dynamical Systems

Simulations of chaotic dynamical systems, e.g., turbulent fluid flows, often require hundreds of thousands of time steps in order to obtain converged statistics. For applications that requires fast turnaround time, scalable parallel time integration is needed to break the bottle neck of spatial-only parallelization in current-generation simulations. This talk first summarizes existing parallel time integration methods, and analyze the scalability problem encountered in chaotic dynamical systems. We then outlines a new method that can potentially achieve perfect scalability. This new method is based on a least squares problem of the governing equation, instead of an initial value problem. In contrast to many existing time decomposition methods, the number of iterations required by our method is insensitive to the length of the time integration, making our method scalable.

## **Temporal parallelization of advanced operation scenario simulations of fusion plasma**

D. Samaddar<sup>1</sup>, T.A. Casper<sup>2</sup>, S.H. Kim<sup>2</sup>, L.A. Berry<sup>3</sup>, W.R. Elwasif<sup>3</sup>,  
D. Batchelor<sup>3</sup>, W.A. Houlberg<sup>3</sup>

1. Culham Science Centre, Abingdon, Oxon OX14 3DB, UK

2. ITER Organization, Route de Vinon sur Verdon, 13115 Saint Paul Lez Durance, France

3. Oak Ridge National Laboratory, Oak Ridge, USA

This work explores the application of the Parareal algorithm[1] to advanced operation scenario codes in tokamak plasmas, using the CORSICA[2, 3] code as a test bed. CORSICA is an advanced free-boundary equilibrium and transport simulation code used to study plasma scenarios in burning plasma experiments and is of particular importance for making predictions for the ITER experiment under construction.

The “Parareal Framework” developed at the Oak Ridge National Laboratory as part of the SWIM IPS project[4] has been used for implementing the algorithm. This framework allows efficient use of processors using the “moving window scheme”[5] along with multiple levels of concurrency.

In the past the Parareal algorithm has been successfully applied to multiple problems including fully developed plasma turbulence simulations which are high dimensionally chaotic initial value problems[6]. However, temporal parallelization of CORSICA introduces new challenges compared to previous applications.

Since the Parareal algorithm involves a predictor-corrector technique, applying the algorithm in this case requires a new approach to the coarse solver necessary for the algorithm. The integration of the Parareal algorithm into the parallelization of CORSICA also allows multiple levels of concurrency. With temporal parallelization being the highest level of parallelization in this case, different levels of simplification are possible when selecting the coarse solver. For example, the source terms and transport properties may not be updated in the coarse solution.

Successfully implementing the Parareal algorithm to codes like CORSICA generates the possibility of time efficient simulations of ITER-like plasmas.


### References:

- 1) J. Lions, Y. Maday and G. Turinici, CR Acad. Sci. I – Math. 332 (7) (2001), pp. 661–668.
- 2) T A Casper et al, 23rd Int. Conf. on Fusion Energy Conference (Daejeon, Korea) ITR/P1-19, accepted for publication in Nuclear Fusion (2010).
- 3) J A Crotinger et al, LLNL Report UCRL-ID-126284 (1997); NTIS #PB2005-102154.
- 4) L.A. Berry et al, Journal of Computational Physics (2011) (In Press).
- 5) W.R Elwasif et al, 4th IEEE Workshop on Many-Task Computing on Grids and Supercomputers, MTAGS 2011, (2011) (Accepted).
- 6) D. Samaddar, D. E. Newman and R. Sanchez, Journal of Computational Physics 229 (18) (2010), pp. 6558-6573.



## Numerical Analysis Group

The Numerical Analysis group is led by Professor Nick Higham, FRS, Professor David Silvester and Professor Françoise Tisseur and comprises eight permanent members of staff, plus research associates and research fellows. Professor Jack Dongarra holds a part-time position jointly with the School of Computer Science. The group carries out research in various areas of numerical analysis and provides teaching and supervision at the undergraduate (in all four years) and postgraduate (M.Sc. and Ph.D.) levels. It forms part of the Applied Mathematics Unit.



Professor James H. Wilkinson

In 2012 the group was strengthened by the appointment of three new Lecturers: Simon Cotter, Stefan Guettel and Martin Lotz.

Activities of the group include organizing international conferences and a Numerical Analysis and Scientific Computing seminar series; writing textbooks and research monographs; membership of editorial boards of international journals and book series; and contributing software to the NAG and LAPACK libraries and MATLAB. The group has a large grant portfolio supported by bodies such as EPSRC, The European Research Council, The Royal Society, The British Council, and in the USA (though Dongarra) the NSF and DOE. Large grants include an EPSRC Leadership Fellowship (Tisseur) and an ERC Advanced grant (Higham). In addition, Higham and Silvester lead the EPSRC Network on Numerical Algorithms and High Performance Computing (2011-2014).

The strength of numerical analysis in Manchester is indicated by the fact that no less than three of our researchers, Dongarra, Hammarling and Higham, have been elected SIAM Fellows and designated Highly Cited Researchers at ISIHighlyCited.com. Members of the group also feature strongly on the list of Manchester mathematicians on Google Scholar Citations.

Numerical analysis has a long history going back at least to Newton and Gauss, whose names adorn some of today's most-used numerical methods. The UK has a strong tradition of research excellence in the subject, as exemplified by James H. Wilkinson (1919-1986), FRS, who was at the forefront of developments in numerical linear algebra from the days of the first digital computers. In Manchester, numerical analysis has long been an area of strength, and M.Sc. programmes in numerical analysis have run continuously since 1959. See the short historical essay Numerical Analysis at the Victoria University of Manchester, 1957-1979, by Joan Walsh.

## Workshop dinner

For participants registered on the workshop dinner: it will take place on Tuesday evening in the Japanese restaurant SAMSI on 36–38 Whitworth Street. See below for walkable route from the Alan Turing Building (A) to SAMSI (B).

