Book of abstracts
PinT 2018

VII Workshop on Parallel in Time Methods

May 2-5, 2018. Roscoff Marine Station, France.

Venue to the VII Workshop on Parallel in Time Methods

Conference room:
Roscoff Marine Station (Website, Google Maps)
Place Georges Teissier
29680 Roscoff

Restaurant and hotel:
Hôtel le Gulf Stream (Website, Google Maps)
400 Rue Marquise de Kergariou
29680 Roscoff

Contact: workshop_pint18@ann.jussieu.fr

http://www.math.univ-paris13.fr/~japhet/PINT2018

VII Workshop on Parallel in Time Methods
PinT 2018

1 About PinT 2018

1.1 Foreword, aims and scope

Solving complex models with high accuracy and within a reasonable computing time motivates the development of domain decomposition methods in space, time or other variables that make an efficient use of modern high performance computing (HPC) systems. The current top HPC architectures have already attained million-way concurrency, and current trends suggest that processor counts will continue to grow rapidly. Both from the pertinence of the approach and the challenges it raises stands the decomposition of the time domain because the sequential nature of the evolution makes it difficult to reconcile with efficient parallelization methods.

The workshop on Parallel in Time Methods (PinT) is a series of meetings held every year whose goal is to bring together researchers working on parallel-in-time methods to discuss recent advances and future directions. Starting with the first event in Lugano (2011), successive workshops have been held at various locations: Manchester (2013), Jülich (2014), Dresden (2015), Toulouse (2016), Banff (2016) and Ascona (2017). The meeting of May 2-5 2018 is the seventh of the series and will take place in the marine station of Roscoff (France). To ensure the highest scientific level, six plenary speakers have been invited:

- Jan Hesthaven (EFPL, Switzerland)
- Rolf Krause (Institute of Computational Science, USI, Lugano, Switzerland)
- Debasmita Samaddar (UK Atomic Research Authority (CCFE), Culham, UK)
- Giovanni Samaey (KU Leuven, Belgium)
- Ian Smears (University College London, UK)
- Shu-Lin Wu (Sichuan University of Science and Engineering, China)

Besides the workshop, the location at Roscoff is also ideal for wonderful trips in the surroundings to enjoy the beauty of Brittany, its coast and its hospitality.

The local organizing committee of PinT 2018

1.2 Committees

Organizing Committee:

- Caroline Japhet (Université Paris XIII, France)
- Yvon Maday (UPMC Sorbonne Universités and Brown University, France and USA)
- Olga Mula (Université Paris Dauphine, France)

Scientific Committee:

- Stephanie Friedhoff (Bergische Universität Wuppertal, Germany)
- Laurence Halpern (Université Paris XIII, France)
- Rolf Krause (University of Lugano, Switzerland)
- Daniel Ruprecht (University of Leeds, UK)
- Debasmita Samaddar (Culham Centre for Fusion Energy, Culham, UK)
- Jacob Schroder (Lawrence Livermore National Laboratory, USA)
- Robert Speck (Jülich Supercomputing Centre, Germany)

1.3 Special events

Discussions and round table with industrial partners:
Apart from the classical oral and poster presentations, we plan to have informal discussions on Thursday and Friday after the morning talks. Also, a round table with industrial partners about the industrial perspectives for parallel-in-time algorithms will take place on Wednesday between 11.30 and 12.15.

Best presentation awards:
The best presentations by early career attendees will be awarded with FoMICS/ICS and CCFE Young Scientist prizes. All PhD students and people who have finished their PhD in the last three years are eligible, with a maximum age of 33 years. Theses prizes will be formally presented to the winners at the end of the conference.

2 Support (see the logos at the end)

We thankfully acknowledge the financial support of:

- UPMC Sorbonne Université
- Laboratoire Jacques-Louis Lions (LJLL)
- Université Paris 13
- Laboratoire d'Analyse, Géométrie et Application (LAGA)
- Centre de Recherche en Mathématiques de la Décision (CEREMADE), Université Paris-Dauphine
- Centre National de la Recherche Scientifique (CNRS)
- Agence Nationale de la Recherche (ANR)
- Institute of Computational Science (ICS)
- Swiss Graduate Program FoMICS "Foundations in Mathematics and Informatics for Computer Simulations in Science and Engineering"
- Institut des Sciences du Calcul et des Données (ISCD)
- LRC Manon
- Institut Carnot Smiles
- Lawrence Livermore National Laboratory (LLNL)
- Pôle MathSTIC de l'Université Paris 13
- Jülich Supercomputing Centre (JSC)
- Action MANU (Méthodes mathématiques et numériques) du programme LEFE
- Agence pour les Mathématiques en Intéraction avec l'Entreprise et la Société (AMIES)
- Investissements d'avenir, Ministère de l'Enseignement Supérieur, de la Recherche et de l'Innovation.
- GdR MaNu Mathématiques pour le Nucléaire
- Groupe CALCUL
- Culham Centre of Fusion Energy (CCFE)

3 Full list of attendees

- Sarah Ali Hassan (INRIA Paris)
- Asad Anees (Clausthal University of Technology)
- Mark Asch (TOTAL)
- Katia Aït Ameur (UPMC Sorbonne Universités & CEA)
- Duc Quang Bui (Paris 13 University)
- Andrew Clarke (University of Leeds)
- Camille Coti (Paris 13 University)
- Anissa Dabachi (Sorbonne Universités)
- Nicolò Dal Santo (EPFL)
- Rob Falgout (Lawrence Livermore National Laboratory)
- Stephanie Friedhoff (Bergische Universität Wuppertal)
- Martin Gander (University of Geneva)
- Venkatesh Gopinath (Institut de Physique du Globe de Paris)
- David Guibert (Bull/Atos)
- Ronald Haynes (Memorial University)
- Jan Hesthaven (EPFL)
- Florent Hédin (École des Ponts - ParisTech)
- Mikio Iizuka (Kyushu University)
- Erell Jamelot (CEA Saclay)
- Caroline Japhet (Paris 13 University)
- Rolf Krause (University of Lugano)
- Iryna Kulchytska (Darmstadt University)
- Thibaut Lunet (University of Geneva)
- Yvon Maday (UPMC Sorbonne Universités)
- Michael Minion (Lawrence Berkeley National Laboratory)

- Olga Mula (Paris Dauphine University)
- Charles Murray (Durham University)
- Van-Thanh Nguyen (Inria Paris)
- Pascal Omnes (CEA Saclay)
- Benedict Philippi (Kiel University)
- Sebastián Reyes-Riffo (Paris Dauphine)
- Guillaume Sall (UPMC Sorbonne Universités)
- Julien Salomon (INRIA Paris)
- Debasmita Samaddar (Culham Center for Fusion Energy)
- Giovanni Samaey (KU Leuven)
- Jacob Schroder (Lawrence Livermore National Laboratory)
- Upanshu Sharma (École des Ponts - ParisTech)
- Ian Smears (University College London)
- Marc Tajchman (CEA Saclay)
- Beth Wingate (University of Exeter)
- Shu-Lin Wu (Sichuan University)

4 Program

Meals:

- Lunch is served at 13.00 sharp at Hôtel le Gulf Stream.
- Diner is served at 19.30 sharp at Hôtel le Gulf Stream.
- The conference banquet will take place on Friday.
  Like for the other diners, it will be served at 19.00 sharp at Hôtel le Gulf Stream.

Schedule: See next page.

Excursions:

- Thursday: *Excursion to the Laber Bay*.
  - Departure from the Hôtel Gulf Stream at 14.30.
  - You may want to take shoes to walk in the bay where there will be a bit of water and on the rocks.
- Saturday: *Excursion to the Batz Island*.
  - Departure from the conference hall at 12.00 right after the talks and without going back to the hotel.
  - We will store the luggages in cars.
  - We shall have lunch on the island (sandwich or restaurant depending on the weather)
  - Return around 16.15.

Schedule

	TUESDAY	WEDNESDAY	THURSDAY	FRIDAY	SATURDAY
8.30 - 08.50		*Registration*			
8.50 - 09.00		*Welcome*			
9.00 - 10.00		Krause	Smears	Hesthaven	Samaey
10.00 - 10.30		*Coffe Break*	*Coffe Break*	*Coffe Break*	*Coffe Break*
10.30 - 11.00		Falgout	Ali Hassan	Gander	Hédin
11.00 - 11.30		Schroder	Aït Ameur	Lunet	Minion
11.30 - 12.00		*Round table*	Mula	Bui	Murray
12.00 - 12.30		*Round table*	*Discussion*	Salomon	*Excursion Batz Island*
12.30 - 13.00					
13.00 - 14.00		*Lunch*	*Lunch*	*Lunch*	
14.00 - 14.30			*Excursion Laber Bay*		
14.30 - 15.30		Samaddar		Wu	
15.30 - 16.00		Sall		*Discussion*	
16.00 - 16.30		*Coffe Break*		*Coffe Break*	
16.30 - 17.00		Coti		Kulchytska	
17.00 - 17.30		Gopinath		Anees	
17.30 - 18.00		Clarke		Iizuka	
18.00 - 19.30					
19.30 - 20.30	*Diner*	*Diner*	*Diner*	*Banquet at 19.00*	

5 List of contributions

Parareal for Transport-Dominated Problems

J. S. Hesthaven[1] **and A. Nielsen**[1]**,**

[1] *EPFL-SB-MATH-MCSS, Ecole Polytechnique Fédérale de Lausanne, Lausanne, CH* jan.hesthaven@epfl.ch

In an attempt to overcome limits on strong scaling and take full advantage of large computing platforms, there has been substantial research devoted to the development of parallel-in-time methods during the last decade, Among several techniques, the Parareal method has been the topic of substantial attention and has been successfully applied to a range of different problems, including molecular dynamics and diffusion dominated problems.

However, the advances for transport dominated problems has been considerably slower with classic methods being plagued by stability problems or slow convergence.

We shall begin by discussing several past attempts that seek to enable the use of the Parareal method to accelerate the parallel solution of transport problems by overcoming what appears to be a stability problem. This leads to the initial conclusion that the quality of the coarse grid solver is particularly important for hyperbolic problems and must be developed with care. As we shall demonstrate, such strategies can either be based on replacing the coarse solver or by carefully managing the work flow of the Parareal method in a general parallel environment to ensure a sufficiently accurate coarse solver while maintaining the potential for parallel speedup.

A second look at the Parareal algorithm reveals that the scheme primarily corrects for amplitude errors while phase errors, often dominating the solution of transport problems, remain essentially uncorrected. This supports the initial observation that an accurate coarse solver allows for rapid convergence as that introduces the required phase error corrections. The challenge remains that a phase accurate coarse solver is not likely to be achievable at low computational cost.

Based on this insight, we proposed a phase-corrected Parareal method and demonstrate that this converges substantially faster for purely hyperbolic problems and transport dominated problems without requiring the use of complex coarse solvers. It also offers a different and more general approach to formulate parallel-in-time methods based on assumptions on the error behavior.

We shall discuss the performance of these new approaches through a number of examples, including the demonstration of time-parallel efficiency exceeding 30% for the fully parallel solution of nonlinear hyperbolic problems.

Multigrid Approaches and Space-Time Discretizations for Non-linear Problems

R. Krause[1], **P. Benedusi**[1], **P. Zulian**[1],

[1] *Institute of Computational Science*
Università della Svizzera italiana, Lugano, Switzerland

We present a parallel and efficient multilevel solution strategy for solving non-linear time-dependent problems. We consider in particular the mono-domain model, a non-linear reaction-diffusion equation arising from a problem in electrophysiology: the electrical activation in the human heart. Different strategies for the space-time discretization and solution of the mono-domain equation are discussed, which are based on domain decomposition and multi-level methods. For the latter, we propose a semi-geometric multigrid method, for which the coarse level approximation spaces are created using arbitrary hierarchies of non-nested meshes. Interpolation and restriction in the multilevel context is then realized by means of a discrete L^2-projection between the non-matching meshes. This approach allows for creating the coarser levels of a multigrid hierarchy, even if only a single "fine" mesh is available. Hence, multigrid hierarchies can be created for arbitrary geometries in any dimension. We discuss how this approach can be applied to the monodomain equation discretised with space-time finite elements.

While we use continuous finite elements in space, for stability reasons we adopt discontinuous elements in time. We discuss shortly the properties of this time discretization scheme.

We investigate how different block smoothers, coarsening strategies and ordering of the space-time variables effect the overall convergence and robustness of the solver.

Furthermore, we comment on local time-stepping for space-time discretizations.

Finally, we investigate numerically the scalability and the convergence of our multilevel and domain decomposition solution strategies.

Applications of the Parareal Algorithm in strongly coupled non-linear problems in fusion plasma – the story so far

D. Samaddar[1], D. P. Coster[2], X. Bonnin[3], W. R. Elwasif[4], D. B. Batchelor[4], L. A. Berry[4], E. Havlickova[1], T. Casper[3], S. H. Kim[3], D. E. Newman[5], R. Sanchez[6],

[1] *Culham Centre for Fusion Energy, Culham Science Centre, Abingdon, OX14 3DB, UK* Debasmita.Samaddar@ukaea.uk
[2] *Max-Planck-Institut für Plasmaphysik, Germany*
[3] *ITER Organization, Route de Vinon-sur-Verdon, CS 90 046, 13067 St. Paul Lez Durance, Cedex, France*
[4] *Oak Ridge National Laboratory, Oak Ridge, TN, USA*
[5] *University of Alaska Fairbanks, AK, USA*
[6] *University Carlos III de Madrid, Spain*

Achieving long time scale simulations of magnetically confined fusion plasma arguably poses one of the grand challenges of the 21st century for the HPC community. While parallel architectures on modern supercomputers play a crucial role in these computations, a saturation in terms of improving computational gain with respect to increasing computing cores is encountered on existing Petascale machines. Temporal parallelization provides a crucial step in exploiting the computing power of modern supercomputers. The Parareal algorithm [1] has so far been the most widely studied method in fusion plasma, reducing wall clock times by factors of 10 and 20 in complex non-linear problems. This talk will discuss a series of applications in fusion plasma, ranging from fluid turbulence at the plasma core to multi-fluid scrape off layer simulations at the plasma edge where plasma-wall interactions play a dominant role. Applications in plasma scenario modelling will also illustrate some key issues. Since the Parareal algorithm relies on a predictor-corrector approach, choosing a suitable coarse predictor often becomes the focal point of research in complex non-linear problems such as fusion. The talk will discuss both achievements and roadblocks encountered along the way. The views and opinions expressed herein do not necessarily reflect those of the European Commission or the ITER Organization.

References

[1] J. Lions, Y. Maday, and G. Turinici. A "parareal" in time discretization of PDE's. *Comptes Rendus de l'Académie des Sciences - Series I - Mathematics*, 332(7):661–668, 2001.

Micro-macro parareal methods: low-dimensional, effective coarse propagators and matching of high-dimensional fine states

Giovanni Samaey[1]

[1] *NUMA, Department of Computer Science, KU Leuven (Belgium)*
giovanni.samaey@kuleuven.be

Many systems of current interest exhibit behavior on a wide range of time scales, which cannot be simulated directly on long (macroscopic) time intervals. In such a setting, efficient time-parallel methods can be constructed that iteratively improve an approximate coarse-grained model by time-parallel fine-scale corrections. We present and analyse a micro-macro parareal method that couples the microscopic model, of which we have full knowledge, to a macroscopic level, of which we assume only limited information. The macroscopic model is described by a finite set of macroscopic state variables — averages over the microscopic distribution. A crucial part of such an algorithm is the appropriate inference of a fine-scale state that is consistent with an imposed macroscopic state. We call this step matching.

In this talk, we will present and discuss micro-macro parareal methods based on the above idea for a set of model problems of increasing complexity: a singularly perturbed ODE, a one-dimensional climate model, and a stochastic model representing molecular dynamics. In each case, we discuss the nature of the matching operator and its convergence properties.

This presentation is based on joint work with Frédéric Legoll, Tony Lelièvre, Keith Myerscough, Thomas Slawig, and Przemyslaw Zielinski.

References

[1] Frédéric Legoll, Tony Lelievre, and Giovanni Samaey. A micro-macro parareal algorithm: application to singularly perturbed ordinary differential equations. *SIAM Journal on Scientific Computing*, 35(4):A1951–A1986, 2013.

[2] Kristian Debrabant, Giovanni Samaey, and Przemysław Zieli?ski. A micro-macro acceleration method for the monte carlo simulation of stochastic differential equations. *SIAM Journal on Numerical Analysis*, 55(6):2745–2786, 2017.

[3] Tony Lelièvre, Giovanni Samaey, and Przemysław Zieliński. Analysis of a micro-macro acceleration method with minimum relative entropy moment matching. *arXiv preprint arXiv:1801.01740*, 2018.

Time-parallel iterative solvers for parabolic evolution equations: an inf-sup theoretic approach

M. Neumüller[1], **I. Smears**[2]

[1] *Institute of Computational Mathematics, Johannes Kepler University Linz, 4040 Linz, Austria,*
[2] *Department of Mathematics, University College London, 25 Gordon Street, London WC1E 6BT, United Kingdom,* i.smears@ucl.ac.uk

Many parallel-in-time methods can be viewed as iterative solvers for a large time-global nonsymmetric linear system arising from the discretization of a time-dependent equation. The nonsymmetry of these systems represents a key challenge in the analysis, since the available theory for iterative methods for nonsymmetric systems is much more limited than for their symmetric counterparts.

In this talk [1], we will show how the underlying inf-sup theory of continuous and discretized parabolic problems provides an effective approach to the construction and rigorous analysis of parallel-in-time solvers. In particular, we consider the implicit Euler discretization of a general linear parabolic evolution equation with time-dependent self-adjoint spatial operators. We first show that the discrete system admits a similar inf-sup condition as for the underlying continuous operator. We use this to show that the standard nonsymmetric time-global system can be equivalently reformulated as an original symmetric saddle-point system that remains inf-sup stable in the same norms. The essential idea is that the mapping from trial functions to their optimal test functions in the inf-sup condition defines a left-preconditioner that symmetrizes the system in a stable way.

We then propose and analyse an inexact Uzawa method for the saddle-point reformulation based on an efficient parallel-in-time preconditioner. The preconditioners is non-intrusive and easy to implement in practice, since it simply combines existing spatial preconditioners with parallel Fast Fourier Transforms (FFT) in time. We prove robust spectral bounds, leading to convergence rates that are independent of the number of time-steps, final time, or spatial mesh sizes. The theoretical parallel complexity of the method then grows only logarithmically with respect to the number of time-steps, owing to the parallel FFT. Numerical experiments of large-scale parallel computations, with up to 131 072 processors and more than 2 billion unknowns, show the effectiveness of the method, along with its good weak and strong scaling properties.

References

[1] M. Neumüller and I. Smears. Time-parallel iterative solvers for parabolic evolution equations. *ArXiv e-print 1802.08126*, 2018.

A Parallel-In-Time Algorithm For Time-Periodic PDE-constrained Optimal Control Problems

Shu-Lin Wu

Sichuan University of Science and Engineering, Zigong 643000, Sichuan Province, P.R.China
wushulin84@hotmail.com

We are interested in designing and analyzing parallel-in-time computation, via the *diagonalization* technique [1,3], for time-periodic PDE-constrained optimization problems [4]. (This is a point of departure and our ultimate goal is to handle parabolic PDE-constrained optimization problems with initial-value condition.) For a class of model problems, following an *optimize-then-discretize* strategy, at the time discrete level we have to solve the steady-state PDE system:

$$\left(\underbrace{\begin{bmatrix} B & \mu^{-1}I_t \\ -I_t & B^\top \end{bmatrix}}_{:=\mathbf{B}} \otimes \mathcal{I}_x\right)\begin{bmatrix}\mathcal{Y}(x)\\ \mathcal{P}(x)\end{bmatrix} - \left(\begin{bmatrix} I_t & \\ & I_t\end{bmatrix}\otimes\Delta\right)\begin{bmatrix}\mathcal{Y}(x)\\ \mathcal{P}(x)\end{bmatrix} = \begin{bmatrix}\mathcal{Y}_d(x)\\ 0\end{bmatrix}, \quad (1)$$

where $B \in \mathbb{R}^{n_t\times n_t}$ denotes the time discretization, $I_t \in \mathbb{R}^{n_t\times n_t}$ is the identity matrix, $\mathcal{I}_x$ is the spatial identity operator, Δ is the Laplacian and $\mu > 0$ is the regularization parameter. By looking into the special structure of B, we show that the 2×2 block matrix $\mathbf{B}$ can be diagonalized as $\mathbf{B} = \mathbf{V}\mathbf{D}\mathbf{V}^{-1}$. Then, following [1, 3] we can solve (1) for all the n_t time points in a direct parallel pattern. We get explicit formulas for $\mathbf{V}$, $\mathbf{V}^{-1}$ and $\mathbf{D}$ and we prove that the condition number of $\mathbf{V}$ is a moderate number and is independent of Δt and μ (this implies that the roundoff error arising from the diagonalization procedure can be well controlled).

For the case that there is a bound constraint for the control variable: $\underline{u}(x,t) \le u(x,t) \le \overline{u}(x,t)$, following the active-set based semismooth Newton method and the Moreau-Yosida approach [2,4], we need to solve a linear system (1) with $\mu^{-1}I_t$ being replaced by $\mu^{-1}I_t + \frac{\chi_k(x)}{\varepsilon}$ at each Newton iteration, where k is the iteration index, $\varepsilon > 0$ is the Moreau-Yosida regularization parameter and $\chi_k(x) \in \mathbb{R}^{n_t\times n_t}$ is a diagonal matrix with diagonal elements being 0 or 1. To solve such a large scale linear system, we propose two preconditioners and analyze the Rayleigh quotient for each of the two preconditioned matrices. The inversion of these two preconditioners can be implemented parallel-in-time via the diagonalization technique.

The matrix B denotes several time discretizations, e.g., the linear θ-method with $\theta \in [\frac{1}{2}, 1]$ and the second-order Lobatto IIIC method.

1. M. J. Gander, L. Halpern, J. Ryan, and T. T. B. Tran. *A direct solver for time parallelization*. LNCSE **104** (2016): 491-499.
2. R. Herzog and E. Sachs. *Preconditioned conjugate gradient method for optimal control problems with control and state constraints*. SIAM J. Matrix Anal. Appl. **31** (2010): 2291-2317.
3. Y. Maday and E. M. Rønquist. *Parallelization in time through tensor product space-time solvers*. CRM **346** (2008): 113-118.
4. M. Stoll. *One-shot solution of a time-dependent time-periodic PDE-constrained optimization problem*. IMA J. Numer. Anal. **34** (2014): 1554-1577.

A posteriori stopping criteria for space-time domain decomposition for porous media flow

E. Ahmed[1,2,3], S. Ali Hassan[1], C. Japhet[2], M. Kern[1], M. Vohralík[1],

[1] *SERENA, Inria Paris & CERMICS (ENPC)* sarah.ali-hassan@inria.fr
[2] *LAGA, Université Paris 13, Sorbonne Paris Cité*
[3] *Department of Mathematics, University of Bergen*

In this work we develop a posteriori error estimates and stopping criteria for global-in-time domain decomposition methods with optimized Robin or Ventcell transmission conditions on the interface between different subdomains. The heat equation using the mixed finite element method in space and the discontinuous Galerkin scheme in time is first analysed. Then we consider a two-phase flow problem in a porous medium composed of two different rock types, so that the capillary pressure field is discontinuous at the interface between the rocks. Numerical approximation is achieved by a cell centered finite volume scheme, using the Matlab Reservoir Simulation Toolbox. We bound the error between the exact solution of the partial differential equation and the approximate numerical solution at each iteration of the space-time domain decomposition algorithm. Different error components (domain decomposition, space discretization, time discretization, Newton linearization) are distinguished, which allows us to define efficient stopping criteria for the DD algorithm. The estimates are based on a reconstruction technique for pressure and flux respectively in the spaces H^1 and $\mathbf{H}(\mathrm{div})$ [1, 2, 3]. For the flux, local Neumann problems in small bands arround the interfaces extracted from the subdomains are solved. Numerical experiments illustrate the theoretical findings.

References

[1] Sarah Ali Hassan, Caroline Japhet, Michel Kern, and Martin Vohralík. A posteriori stopping criteria for optimized Schwarz domain decomposition algorithms in mixed formulations. *Comput. Methods Appl. Math*, 2018, to appear.

[2] Sarah Ali Hassan, Caroline Japhet, and Martin Vohralík. A posteriori stopping criteria for space-time domain decomposition for the heat equation in mixed formulations. HAL Preprint 01586862, submitted for publication, 2017.

[3] Elyes Ahmed, Sarah Ali Hassan, Caroline Japhet, Michel Kern, and Martin Vohralík. A posteriori error estimates and stopping criteria for space-time domain decomposition for two-phase flow between different rock types. HAL Preprint 01540956, submitted for publication, 2017.

Time Domain Finite Element Methods for Maxwell's Equations

A. Anees[1], **L. Angermann**[1],

[1] *Institute of Mathematics, Clausthal University of Technology, Clausthal-Zellerfeld, D-38678, Germany*

asad.anees@tu-clausthal.de

In this presentation, we discuss a time-domain finite element methods to the solution of linear and nonlinear Maxwell's equations. A weak formulation is derived for the electric and magnetic field with appropriate initial and boundary conditions, and the problem is discretised both in space and time. Nédeléc curl conforming and Raviart Thomas divergence conforming finite elements are used to discretized the electric and magnetic fields respectively in space. In time domain, problems are discretized by symplatic and backward Euler methods, and provide some numerical results to validate our simulation for fully discretized problems. The proposed methods are accurate in time up to order 4, in case of symplectic time integration and they are conditionally stable. Our methods also allow to treat a complex geometries of various physical system coupled to electromagnetics fields. Electric and magnetic fields are also visualized at each time step on tetrahedron meshes. Our proposed methods are also parallel, but time integration is not parallel. We are working to add parallel time integrations in our proposed mixed finite element methods for electromagnetic using XBraid.

References

[1] A. Anees and L. Angermann, *Time Domain Finite Element Methods for Maxwell's Equations in Three Dimensions*, IEEE and International Applied Computational Electromagnetics Society (ACES) Symposium, March 24-29, 2018, Denver, Colorado, U.S.A,

[2] A. Anees and L. Angermann, *Mixed Finite Element Methods for the Maxwell's Equations with Matrix Parameters*, IEEE and International Applied Computational Electromagnetics Society (ACES) Symposium, March 24-29, 2018, Denver, Colorado, U.S.A,

[3] G. Rodrigue and D. White, *A vector finite element Time-Domain method for solving Maxwell's equations on unstructured Hexeahedral grids*, Society for Industrial and Applied Mathematics, 23, pp. 683-706, 2001.

Parareal algorithm for two phase flows simulation

K. Aït Ameur[1,2], **Y. Maday**[1] **and M. Tajchman**[2],

[1] *Sorbonne Université, UPMC*
[2] *CEA Saclay* aitameur.katia@gmail.com

In the nuclear energy domain, system codes are dedicated to the thermal-hydraulics analysis of nuclear reactors for safety studies. We are here interested in the Cathare code developed by CEA, EDF, AREVA-NP and IRSN.

Cathare solves the 6-equations/two-fluid model by considering a set of conservation laws (mass, momentum and energy) for each phase liquid and vapor. The discretization is based on a finite volumes method on staggered grid in space and on a fully implicit time integration method. Typical cases involve up to a million of numerical time iterations, computing the approximate solution during long physical simulation times. On the other hand, the discretization level is kept intentionally at a coarse level to be able to handle whole systems simulations.

Cathare is used in a simulator of a reactor submitted to accidental events. This platform is dedicated to train operators and prepare crisis management exercises thus requiring real-time response of the code. To optimize the response time, we consider a strategy of time domain decomposition, complementing the current space domain decomposition.

This strategy is based on the parareal method, introduced in [1], that provides a strategy for "parallel-in-time" computations and offers the potential for an increased level of parallelism.

Here we apply the parareal algorithm to the resolution of an oscillating manometer. This test case is proposed in [2] for system codes to test the ability of each numerical scheme to preserve system mass and to retain the gas-liquid interface.

References

[1] G. Turinici J.-L. Lions, Y. Maday. Résolution par un schéma en temps "pararéel". 2001.

[2] N. Zuber G.F. Hewitt, J.M. Delhaye. *Multiphase science and technology*, volume 6. 1991.

Coupled Parareal-Optimized Schwarz Waveform relaxation method for advection-diffusion equation

D.Q. Bui[1], **C. Japhet**[1], **Y. Maday**[2], **P. Omnes**[1,3],

[1] *LAGA, Université Paris 13, Sorbonne Paris Cité,* bui@math.univ-paris13.fr
[2] *LJLL, UPMC Sorbonne Université*
[3] *CEA, DEN, DM2S-SFME*

Parareal method [1, 2] is a numerical method to solve time - evolutional problems in parallel, which uses two propagators: the coarse - fast and inaccurate - and the fine - slow but more accurate. Instead of running the fine propagator on the whole time interval, we divide the time space into small time intervals, where we can run the fine propagator in parallel to obtain the desired solution, with the help of the coarse propagator and through parareal steps. Furthermore, each local subproblem can be solved by an iterative method, and instead of doing this local iterative method until convergence, one may perform only a few iterations of it, during parareal iterations. Propagators then become much cheaper but sharply lose their accuracy, and we hope that the convergence will be achieved across parareal iterations.
In this contribution, we propose to couple Parareal with a well-known iterative method - Optimized Schwarz Waveform Relaxation (OSWR) [3] - with only few OSWR iterations in the fine propagator and with a simple coarse propagator deduced from Backward Euler method. We present the analysis of this coupled method for 1-dimensional advection reaction diffusion equation, for this case the convergence is almost linear. We also give some numerical illustrations for 1D and 2D equations, which shows that the convergence is much faster in practice.

References

[1] Jacques-Louis Lions, Yvon Maday, and Gabriel Turinici. A parareal in time discretization of pde's. *C.R. Acad. Sci. Paris, Serie I*, 332:661–668, 2001.

[2] Martin J. Gander and Stefan Vandewalle. Analysis of the parareal time-parallel time-integration method. *SIAM J. Sci. Comp.*, 29(2):556–578, 2007.

[3] Martin J. Gander and Laurence Halpern. Optimized Schwarz waveform relaxation methods for advection reaction diffusion problems. *SIAM J. Numer. Anal.*, 45(2):666–697, 2007.

Parallel in time integration of the Kinematic Dynamo

A. Clarke[1], **D. Ruprecht**[1], **S. Tobias**[1], **C. Davies**[1]

[1] *University of Leeds* scatc@leeds.ac.uk

The precise mechanisms responsible for the natural dynamos in the Earth and Sun are still not fully understood. Numerical simulations of natural dynamos are extremely computationally intensive, and are carried out in parameter spaces many orders of magnitude away from real conditions.

Parallelization in space is a common strategy to speed up simulations on high performance computers, but eventually hits a scaling limit. Additional directions of parallelization are desirable to utilise the high number of processor cores now available. Parallel-in-time methods can deliver speed-up in addition to that offered by spatial partitioning.

This talk will investigate Parareal's [1] ability to speed up simulations of the kinematic dynamo, where the velocity field is a prescribed field in the induction equation. We will describe an implementation of Parareal into the open-source Python software Dedalus [2], which implements spectral methods with implicit-explicit (IMEX) time-stepping. Good performance in Parareal depends on an efficient coarse propagator, which should be faster than the fine propagator, whilst still giving adequate accuracy for good convergence. In this case, reduced spatial resolution is used for the coarse solver. This also allows larger time step sizes, further increasing the speed of the coarse propagator.

Convergence properties, speed-up, and efficiency of the Parareal algorithm for the Roberts [3] flow and other types of dynamos will be presented.

References

[1] Jacques-Louis Lions, Yvon Maday, and Gabriel Turinici. Résolution d'edp par un schéma en temps «pararéel». *Comptes Rendus de l'Académie des Sciences-Series I-Mathematics*, 332(7):661–668, 2001.

[2] Keaton J Burns, Geoffrey M Vasil, Jeffrey S Oishi, Daniel Lecoanet, Benjamin Brown, and Eliot Quataert. Dedalus: Flexible framework for spectrally solving differential equations. (In preparation).

[3] Gareth O Roberts. Dynamo action of fluid motions with two-dimensional periodicity. *Phil. Trans. R. Soc. Lond. A*, 271(1216):411–454, 1972.

Scalable, Fault-Tolerant Parallel Matrix Factorization

Camille Coti†

† *LIPN, CNRS UMR 7030, Université Paris 13, Sorbonne-Paris-Cité*
camille.coti@univ-paris13.fr

Large-scale parallel system are now more usual and accessible to scientists. The latest release of the Top 500 (www.top500.org), in November 2017, show that, in the world, 181 systems achieve a performance of 1 Petaflops and more; 58 systems feature more than 100 000 cores and 499 feature more than 10 000 cores.

Besides, at large-scale, failures are statistically frequent, and need to be taken into account by the computation. As a matter of fact, at large-scale, failures are so common that it is expected that some nodes will fail during the computation [1]. As a consequence, the computation needs to expect these failures and adapt in order to proceed with the computation in spite of these failures.

I will present a recent model for fault tolerance in parallel applications extending the MPI standard and called *user-level failure mitigation*[2]. Failure detection and resolution are local, and depend on communications: only the processes that try to communicate with a failed process are aware that it has failed. The rest of the system continue its execution unknowingly. This model is particularly interesting at large scale, since it involves limited synchronizations between processes.

In this model, I will present two fault-tolerant, direct algorithms for QR and LU dense matrix factorizations. These algorithms are based on communication-avoiding algorithms [3], that have shown their good scalability on a wide range of parallel architectures. I will present how some algorithmic and algebraic properties can be taken advantage of in order to insert or exploit partial result redundancy and obtain algorithms that are intrinsically fault-tolerant [4].

References

[1] Daniel A. Reed, Charng da Lu, and Celso L. Mendes. Reliability challenges in large systems. *Future Generation Computer Systems*, 22(3):293 – 302, 2006.

[2] Wesley Bland, Aurelien Bouteiller, Thomas Hérault, Joshua Hursey, George Bosilca, and Jack J. Dongarra. An evaluation of user-level failure mitigation support in MPI. *Computing*, 95(12):1171–1184, 2013.

[3] James Demmel, Laura Grigori, Mark Hoemmen, and Julien Langou. Communication-optimal parallel and sequential QR and LU factorizations. *SIAM Journal on Scientific Computing*, 34(1):A206–A239, 2012.

[4] Camille Coti. Scalable, robust, fault-tolerant parallel qr factorization. In *Distributed Computing and Applications to Business, Engineering and Science (DCABES), 2016 15th International Symposium on*. IEEE, 2016.

Recent advances with multigrid reduction in time (MGRIT) for hyperbolic problems

R. Falgout[1], **H. De Sterck**[2], **A. Howse**[3], **S. MacLachlan**[4], **J. Schroder**[1]

[1] *Lawrence Livermore National Laboratory, USA* rfalgout@llnl.gov
[2] *Monash University, Australia*
[3] *Red Deer College, Canada*
[4] *Memorial University of Newfoundland, Canada*

The focus of this talk is on the application of the multigrid reduction in time (MGRIT) method to hyperbolic partial differential equations. We first consider the MGRIT approach developed in [1] that coarsens in both time and space. In the case of explicit time-stepping, spatial coarsening is needed to ensure stability on all levels, but it is also useful for implicit time-stepping by producing cheaper multigrid cycles. Unfortunately, uniform spatial coarsening results in extremely slow convergence when the wave speed is near zero, even if only locally. An adaptive spatial coarsening strategy addresses this issue for the variable coefficient linear advection equation and the inviscid Burgers equation using first-order explicit or implicit time-stepping methods. Numerical results show that this offers significant improvements over classical coarsening, and parallel scaling tests indicate that speedups over serial time-stepping strategies are possible. We also discuss progress on other recent ideas for handling hyperbolic problems, one based on more accurately approximating the Petrov-Galerkin coarse-grid operator, and another based on adding higher-order dispersion to minimize diffusiveness while improving MGRIT convergence.

This work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

References

[1] Hans De Sterck, Robert D. Falgout, Alexander J.M. Howse, Scott P. MacLachlan, and Jacob B.Schroder. Parallel-in-time multigrid with adaptive spatial coarsening for the linear advection and inviscid burgers equations. *SIAM J. Sci. Comput.*, submitted. LLNL-JRNL-737050.

Why it is difficult to solve hyperbolic problems with parareal type algorithms

Martin J. Gander

[1] *University of Geneva* martin.gander@unige.ch

The convergence of the parareal algorithm [1] for the time parallel solution of evolution problems is well understood, see [2] for precise convergence estimates for parabolic and hyperbolic partial differential equations, and [3] for non-linear problems. In the hyperbolic case, the parareal algorithm is missing one of the essential convergence mechanisms that make it work well in the parabolic case. I will try to explain this in my presentation using both a precise convergence estimate, and also the method of characteristics from [4]. The difficulty is similar to the difficulty when solving time harmonic problems with iterative methods [5].

References

[1] Jacques-Louis Lions, Yvon Maday, and Gabriel Turinici. A "parareal" in time discretization of PDE's. *Comptes Rendus de l'Académie des Sciences-Series I-Mathematics*, 332(7):661–668, 2001.

[2] Martin J. Gander and Stefan Vandewalle. Analysis of the parareal time-parallel time-integration method. *SIAM Journal on Scientific Computing*, 29(2):556–578, 2007.

[3] Martin J. Gander and Ernst Hairer. Nonlinear convergence analysis for the parareal algorithm. *Lecture Notes in Computational Science and Engineering*, 60:45–56, 2008.

[4] Martin J. Gander. Analysis of the parareal algorithm applied to hyperbolic problems using characteristics. *Boletin de la Sociedad Espanola de Matemática Aplicada*, 42:21–35, 2008.

[5] Oliver G. Ernst and Martin J. Gander. Why it is difficult to solve Helmholtz problems with classical iterative methods. In *Numerical analysis of multiscale problems*, pages 325–363. Springer, 2012.

Time domain integration schemes for the simulation of Earth and planetary cores

V. Gopinath[1], **A. Fournier**[1], **T. Gastine**[1]

[1] *Institut de Physique du Globe de Paris (IPGP)* gopinath@ipgp.fr

The magnetic field of the Earth is understood to be generated by the motion of an electrically conducting fluid inside the outer core. This field protects the Earth from charged particles in the form of cosmic rays or solar winds. Numerical simulations of outer core of the earth have been an essential tool in understanding the evolution of such magnetic fields. In order to maintain such magnetic fields over vast time scales, a convective motion inside the outer core is predicted to happen. Such convective motions can either be driven by thermal convection or by compositional convection.

The simulation of such coupled systems often involves High Performance Computing (HPC). However, even with advanced HPC resources, the simulations are not yet close to operating at the actual parameters and regimes of the Earth. Thus there is the need for efficient numerical strategies to tackle such problems. In this work, we pursue to develop an advanced time integration scheme which would help us explore extreme simulation regimes. A two-dimensional annulus geometry which mimics an equatorial plane of the outer core is modeled. A thermal convection is set up between the outer and inner walls, the strength of which is measured by the Rayleigh number (Ra). Currently, a code for simulating this problem is implemented. It is structured in such a way that we are able to test new schemes in an efficient manner. Various new time integration schemes are implemented and are being currently tested for efficiency. These new schemes will open a new window to simulate certain parameter regimes which were previously considered costly. From this study, we hope to select a coarse and a fine integration scheme which best suits the parallel in time approach, which we will pursue for this problem.

A new implementation of the Generalized Parallel Replica dynamics for the long time simulation of metastable biochemical systems

Florent Hédin[1], **Tony Lelièvre**[1]

CERMICS, École des Ponts–ParisTech, 6 et 8 avenue Blaise Pascal, Cité Descartes – Champs sur Marne, F-77455 Marne-la-Vallée Cedex 2, France

florent.hedin@enpc.fr

Molecular dynamics (MD) simulations are nowadays of a common use for simulating large and complex biological or chemical systems. However, a commonly encountered obstacle is the timescale separation between the fastest conformational changes — occurring at the femtoseconds (fs) level — and the slowest one, occurring from the nanosecond (ns) to second (or more) timescale. The existence of *metastable* regions in the phase space is the main origin of this separation, and the simulation time required for observing state-to-state transition quickly becomes intractable.

Among the methods developed to address this challenge, the "Parallel Replica"(ParRep)[1] method maps the original dynamics to a state-to-state dynamics, relying on quasi-stationary distributions (QSDs).[2] The later developed "Generalized Parallel Replica"[3] method uses advanced techniques in order to estimate *on the fly* if convergence to the QSD is obtained, making the method easier to use when investigating biochemical systems.

We will present a new implementation of the Generalized parallel replica method,[4] targeting frequently encountered metastable biochemical systems, such as conformational equilibria or dissociation of protein–ligand complexes. It will be shown that the method is mathematically accurate, that it can be used for studying a wide variety of processes, and that the *time parallelization* strategy is scalable, making the application ready for High Performance Computing (HPC) machines.

References

[1] A. F. Voter. *Phys. Rev. B*, 57:R13985–R13988, Jun 1998.

[2] C. Le Bris, T. Lelièvre, M. Luskin, and D. Perez. *Monte Carlo Methods and Applications*, 18(2):119–146, Jan 2012. arXiv: 1105.4636.

[3] A. Binder, T. Lelièvre, and G. Simpson. *J. Comput. Phys.*, 284:595–616, Mar 2015. arXiv: 1404.6191.

[4] F. Hédin and T. Lelièvre. *The Journal of Chemical Physics*, to be published, 2018.

Investigation of Convergence of Parareal Method for Advection Equation using Accurate Phase Calculation Method

M. Iizuka[1], **K. Ono**[1],

[1] *Research Institute for Information Technology, Kyushu University, 744 Motooka, Nishi-ku, Fukuoka, 819-0395 Japan*

iizuka.mikio.903@m.kyushu-u.ac.jp

Applying the parareal method to hyperbolic PDEs, such as the advection equation, is strongly desired. However, it is still difficult to solve it even now. Reference [1] reported that the reason of the difficulty is identified as the beating between waves, which is introduced by the discrepancy of the phase between the fine and the coarse solvers, and leads the low convergence rate. So far, huge efforts have been performed to improve the phase accuracy of the hyperbolic PDEs. In this study, we introduce CIP (Constrained Interpolation Profile Scheme) scheme [2], and STRS-CIP scheme, which is the combination of STRS (Space-Time Reversal Symmetry Scheme) [3] and CIP, and try to improve the convergence rate of the parareal method for the advection equation. In this presentation, we show the numerical results of the convergence of the parareal method by the CIP-3rd, CIP-5th order schemes, and the STRS-CIP scheme. Then, we discuss the convergence of the parareal method with highly accurate phase treatment for the advection equation.

References

[1] M. Gander and M. Petcu. Analysis of a krylov subspace enhanced parareal algorithm for linear problems. volume 25, pages 114–129. ESAIM Proc, 2008.

[2] H. Takewaki and T. Yabe. The cubic-interpolated pseudo particle (cip) method: Application to nonlinear and multi-dimensional hyperbolic equations. *J. Comput. Phys.*, 70:353372, 1987.

[3] K. Watanabe. a novel framework to construct amplitude preserving wave propagation schemes. pages 123–124, 2010. Japan Society for Industrial and Applied Mathematics, Annual meeting.

A New Parareal Algorithm for Problems with Discontinuous Sources

I. Kulchytska-Ruchka[1], **M. J. Gander**[2], **S. Schöps**[1], **I. Niyonzima**[3]

[1] *Technical University of Darmstadt* kulchytska@gsc.tu-darmstadt.de
[2] *University of Geneva*
[3] *Columbia University*

The Parareal algorithm [1] allows to solve time-dependent problems in a time-parallel manner. Its convergence rate for nonlinear problems was derived in [2] under the assumption of regular (smooth) inputs. In the present contribution we perform a convergence analysis of a modified Parareal method for systems, which involve discontinuous right-hand sides. Such situations occur, e.g., in power engineering when electric devices are supplied with a pulse-width-modulated signal. In order to develop the theory for such problems we propose to use a low-frequency smooth input for the coarse propagator, e.g., from Fourier analysis. The influence of the input reduction on the overall convergence rate of the algorithm is illustrated by the derived error estimate. Numerical verification complements the theoretical findings and supports the obtained error bounds. Furthermore, similarly to the recent results of [3] on Parareal for the eddy current problem, we applied our modified approach with reduced coarse dynamics to the simulation of an induction machine and compared its performance to that of the standard Parareal method.

Acknowledgments The work of I. Kulchytska-Ruchka is supported by the Excellence Initiative of the German Federal and State Governments and the Graduate School of Computational Engineering at Technische Universität Darmstadt, as well as by the BMBF (grant No. 05M2018RDA) in the framework of project PASIROM.

References

[1] Jacques-Louis Lions, Yvon Maday, and Gabriel Turinici. A parareal in time discretization of PDEs. *Comptes Rendus de l'Académie des Sciences – Series I – Mathematics*, 332(7):661–668, 2001.

[2] Martin J. Gander and Ernst Hairer. *Nonlinear Convergence Analysis for the Parareal Algorithm*, pages 45–56. Springer Berlin Heidelberg, Berlin, Heidelberg, 2008.

[3] Sebastian Schöps, Innocent Niyonzima, and Markus Clemens. Parallel-in-time simulation of eddy current problems using parareal. 54(3):1–4, 2018.

Stable time-parallel integration of advection dominated problems using Parareal with space coarsening.

J. Bodart[1], **S. Gratton**[2], **X. Vasseur**[1], **T. Lunet**[3]

[1] *ISAE-SUPAERO*
[2] *Toulouse-INP-IRIT*
[3] *University of Geneva* thibaut.lunet@unige.ch

A common idea in the PinT community is that Parareal, one of the most popular time-parallel algorithm, is often numerically unstable when applied to hyperbolic problems, such as the advection equation. This has been both numerically observed and studied theoretically in the case of an implicit time integrator as a coarse solver (see, *e.g*, [1]). Such results could discourage the application of Parareal to Computational Fluid Dynamics (CFD) problems, especially for high Reynolds numbers (see, *e.g*, [2]). On the contrary, a recent application of Parareal with spatial coarsening to an explicit CFD solver [3] has shown that not only a stable numerical integration was obtained, but also that the Reynolds number played a minor role in the convergence behaviour, compared to other parameters of the parallel-in-time algorithm.

Hence, in this talk, we present numerical experiments related to the application of Parareal with spatial coarsening to the one-dimensional advection-diffusion problem. We investigate the influence of several parameters on the convergence (importance of the diffusion term, spatial resolution, order of interpolation, regularity of the initial solution, time-slice length, nonlinearity,...). We advocate that "a high Reynolds number" is not a good enough reason for *not using* Parareal, and that a stable and efficient parallel in time integration can be made possible, even for highly advective problems, provided that important algorithmic components are carefully chosen.

References

[1] Daniel Ruprecht. Wave propagation characteristics of Parareal. *arXiv preprint arXiv:1701.01359*, 2017.

[2] J. Steiner, D. Ruprecht, R. Speck, and R. Krause. Convergence of Parareal for the Navier-Stokes equations depending on the Reynolds number. In *Numerical Mathematics and Advanced Applications - ENUMATH 2013*, volume 103, pages 195–202. Springer, 2015.

[3] T. Lunet, J. Bodart, S. Gratton, and X. Vasseur. Time-parallel simulation of the decay of homogeneous turbulence using Parareal with spatial coarsening. *Comput. Vis. Sci.*, (Accepted), 2017.

The Performance of the PFASST Method on a Suite of Shallow Water Test Equations on the Rotating Sphere

M. Minion[1], **F. Hamon**[1], **M. Schreiber**[2]

[1] *Lawrence Berkeley National Lab* mlminion@lbl.gov
[2] *Exeter University*

This talk will report on recent work to assess the parallel in time performance of the PFASST algorithm applied to a popular suite of test equations from the atmospheric modeling community. For this study, the PFASST algorithm is paired with a spatial discretization based on spherical harmonics and applied to popular test cases governed by the shallow water equations on a rotating sphere. The implementation uses the *libpfasst* library and the SWEET (Shallow Water Equation Environment for Tests) package. High-order semi-implicit or IMEX parallel in time schemes are tested with a careful consideration of the effects of viscosity and resolution on the parallel performance.

A scalable adaptive parareal algorithm with online stopping criterion

Y. Maday[1,2,3], **O. Mula**[4]

[1] *Sorbonne Universités, UPMC Univ Paris 06 and CNRS UMR 7598, Laboratoire Jacques-Louis Lions, F-75005, Paris, France.*
[2] *Institut Universitaire de France.*
[3] *Division of Applied Mathematics, Brown University, Providence RI, USA*
[4] *Université Paris-Dauphine, PSL Research University, CNRS, UMR 7534, CEREMADE, 75016 Paris, France* mula@ceremade.dauphine.com

In this talk, we present a new vision of the parareal in time algorithm which carries potential to overcome the classical scalability limitations of domain decomposition schemes for the time domain. For this, we reformulate the algorithm in an infinite dimensional functional space where the realization of each iteration involves approximations of increasing accuracy of time dependent subproblems formulated on small subdomains of the original time interval. We rigorously prove that the resulting parareal algorithm presents a near-optimal convergence rate and has a parallel efficiency which is significantly superior to the traditional version. In some cases, it can even be close to the ideal value of one. We illustrate our findings for two different problems where, in addition, we explore the potential benefits of reusing information from previous parareal iterations.

Full space-time multigrid using time stepping specific multigrid restriction and prolongation

C. Murray[1], **T. Weinzierl**[1],

[1] *Durham University* c.d.murray@durham.ac.uk

Classic multigrid methods operate on a grid discretised in space and use a complementing method (time stepping) to advance a solution in time. This synchronises/sequentialises the time steps. A space-time grid can be used to discretise both space and time requiring a stencil extended in the time dimension. Traditional multigrid methods on space-time grids use multigrid exclusively in space and couple space and time on the finest level, while traditional space parallelisation couples whole time slices with each other.

We propose (i) to use multigid on the global space-time grid which is well-known to comprise both classic spatial multigrid and parallel-in-time methods. We propose (ii) to use classic Galerkin operator construction to come up with coarse time stepping rules. The time stepping is read as limit of a convection operator. We propose (ii) to extend the algebraic multigrid technique of Black Box Multigrid (BoxMG) to compute operator- and time stepping-specific multigrid restriction and prolongation. Such an approach is doomed to fail if the proposed realisation is not fast and scaling. Our approach relies on three pillars. First, we apply dynamically adaptive meshes in space in time and we use full multigrid cycles in space and time, this is combined it with Full Approximation Storage (FAS). The latter allows for a straightforward realisation of arbitrary dynamic mesh refinement (AMR), combined with the space-time solver this gives dynamic time-stepping effectively for free. Second, we observe that classic multigrid methods require all stencils, i.e. matrix entries, to be computed prior to the solve, at the same time, the first few iterations determine approximations of limited accuracy only. To be able to exploit massively parallel computers without a ramp up (assembly) phase we propose to kick off multigrid with very rough stencil approximations but to use idle cores to determine better and better stencil integrations on-the-fly. Finally, we note that multigrid suffers from limited arithmetic intensity and concurrency, while the gap between compute power and memory bandwidth is widening. It is thus important to read each unknown as rarely as possible. Additive multigrid variants are well-known to have advantageous scaling properties but tend to be less stable and efficient than their multiplicative counterparts. Asynchronous Fast Adaptive Composite (AFAC) candidates to stabilise additive multigrid as each correction term is damped by an additional multilevel coarse grid solve. We propose a pipelined, single-touch realisation of AFAC in space and time where each degree of freedom is loaded only once per multilevel smoothing step. It is memory-access optimal.

The Parareal Algorithm for American Options

Gilles Pagès[1], **Olivier Pironneau**[2], **Guillaume Sall**[1]

[1] *Laboratoire de Probabilités, Statistique et Modélisation, UMR 8001, Sorbonne Université, Boîte courrier 158, 4 pl. de Jussieu, F-75252 Paris Cedex 5, France*
[2] *Laboratoire Jacques Louis Lions, UMR 7598, Case 187, 4 pl. de Jussieu, F-75252 Paris Cedex 5, France,*

guillaume.sall@sorbonne-universite.fr

With parallelism in mind we investigate the parareal method for American contracts both theoretically and numerically. Least-Square Monte-Carlo (LSMC, see [1, 2, 3]) and parareal time decomposition (see [4, 5]) with two or more levels are used, leading to an efficient parallel implementation which scales linearly with the number of processors and is appropriate to any multiprocessor-memory architecture in its multilevel version. We prove L^2 superlinear convergence for an LSMC backward in time computation of American contracts, when the conditional expectations are known, i.e. before Monte-Carlo discretization. The argument provides also a tool to analyze the multi-level parareal algorithm; in all cases the computing time is increased only by a constant factor, compared to the sequential algorithm on the finest grid, and speed-up is guaranteed when the number of processors is larger than that constant. A numerical implementation confirms the theoretical error estimates

References

[1] J.N. Tsitsiklis and B. van Roy. Regression methods for pricing complex american-style options. *IEEE Trans. Neural Netw*, 12(4):694–703, 2001.

[2] J. Tsitsiklis and B. van Roy. Regression methods for pricing complex american-style options. *IEEE TRANSACTIONS ON NEURAL NETWORKS, VOL. 12, NO. 4, JULY 2001*, 12(4):694–703, 2001.

[3] F. A. Longstaff and E. S. Schwartz. Valuing American options by simulation: a simple least squares approach. *Review of Financial Studies*, 14:113–148, 2001.

[4] G. Bal and Y. Maday. A "parareal" time discretization for non-linear pde's with application to the pricing of an american put. *Recent developments in domain decomposition methods (Zürich, 2001)*, 23:189–202, 2002.

[5] J.-L. Lions, Y. Maday, and G. Turinici. Résolution d'edp par un schéma en temps "pararéel". *C.R. Acad. Sci. Paris sér. I Math.*, 332(7):661–668, 2001.

Time parallelisation for optimal control and data assimilation

M. Gander[1], F. Kwok[2], S. Reyes-Riffo[3], J. Salomon[4]

[1] *Université de Genève,*
[2] *Hong Kong Baptist University,*
[3] *Ceremade, Université Paris-Dauphine,*
[4] *INRIA-Paris & Laboratoire J-L. Lions, Sorbonne Université*
julien.salomo@inria.fr

In this talk, we introduce a method to solve approximately optimality systems associated with optimal control problems. Our approach is based on parareal algorithm idea, i.e., it consists in using a coarse solver to build an approximation of the Jacobian matrix involved in a Newton Loop. This method does not work on an bounded intervals, e.g., when the system needs to be controlled for all $t \in [0, +\infty)$. To deal with such a case, we focus on a data assimilation problem and present an adjoint-free version of the method.

Parallel-in-Time Optimization with the General-Purpose XBraid Package

S. Günther[2], J. B. Schroder[1], N. Gauger[2], R. D. Falgout[1]

[1] *Lawrence Livermore National Laboratory, USA* schroder2@llnl.gov
[2] *TU Kaiserslautern, Germany*

Parallel-in-time methods provide a solution to the serial time-integration bottleneck. This bottleneck is a result of long-term computer architecture trends, where current and future speedups will be available through greater concurrency, not faster clock speeds, which are stagnant. One area ripe for development by parallel-in-time is numerical optimization with unsteady PDEs. Here, the common adjoint-based approach suffers from a large serial bottleneck because it couples multiple sweeps forwards and then backwards in time. The backwards sweeps use the adjoint to compute sensitivities, which are then used to optimize the objective function. The multiple sweeps forwards and backwards in time create a large un-parallelized sequential process. In this talk, we examine the optimal-scaling multigrid reduction in time (MGRIT) method, which is able to parallelize over such nonlinear time-stepping processes. Initial work by Günther et al. [1, 2] showed promise by using MGRIT to compute both sweeps parallel-in-time, in the context of a traditional adjoint-based approach and also a one-shot optimization framework. In this talk, we describe the implementation of this work into the open source XBraid (MGRIT) software package. The goal is to provide the broader community a general-purpose parallel-in-time optimization approach.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52–07NA27344, LLNL-ABS-748131

References

[1] S. Günther, N. R. Gauger, and J. B. Schroder. A non-intrusive parallel-in-time adjoint solver with the xbraid library. *Computing and Visualization in Science*, 2017. (accepted).

[2] S. Günther, N. R. Gauger, and J. B. Schroder. A non-intrusive parallel-in-time approach for simultaneous optimization with unsteady pdes. *Optimization Methods and Software*, 2018. (submitted).

Organizing and sponsoring institutions