

# Parareal for diffusion problems with space- and time-dependent coefficients

Daniel Ruprecht<sup>1</sup>, Robert Speck<sup>1,2</sup>, and Rolf Krause<sup>1</sup>

## 1 Introduction

The very rapidly increasing number of cores in state-of-the-art supercomputers fuels both the need for and the interest in novel numerical algorithms inherently designed to feature concurrency. In addition to the mature field of space-parallel approaches (e.g. domain decomposition techniques), time-parallel methods that allow concurrency along the temporal dimension are now an increasingly active field of research, although first ideas, like in Nievergelt [1964], go back several decades. A prominent and widely studied algorithm in this area is Parareal, introduced in Lions et al. [2001], which has the advantage that one can couple and reuse classical time-stepping schemes in an iterative fashion to parallelize in time. However, there also exist a number of other approaches, e.g. the "parallel implicit time algorithm" (PITA) from Farhat and Chandesris [2003], the "parallel full approximation scheme in space and time" (PFASST) from Emmett and Minion [2012] or "revisionist integral deferred corrections" (RIDC) from Christlieb et al. [2010] to name a few. Parareal in particular and temporal parallelism in general has been considered early as an addition to spatial parallelism in order to extend strong scaling limits, see Maday and Turinici [2005]. Efficacy of this approach in large-scale parallel simulations on hundreds of thousands of cores has been demonstrated for the PFASST algorithm in Speck et al. [2012].

For Parareal, multiple works exist that demonstrate its efficiency for diffusion problems: Gander and Vandewalle [2007] prove super-linear convergence of Parareal for the standard 1D heat equation. A more general theorem showing super-linear convergence for nonlinear ODEs is proven by Gander and Hairer [2008], while Bal [2005] presents a convergence theorem for linear

---

Institute of Computational Science, Università della Svizzera italiana, Lugano, Switzerland, {daniel.ruprecht,rolf.krause}@usi.ch · Jülich Supercomputing Centre, Forschungszentrum Jülich GmbH, Jülich, Germany, r.speck@fz-juelich.de

parabolic PDEs with constant coefficients. The present paper investigates the effect of space- and time-dependent coefficients in the two-dimensional heat equation on the convergence of Parareal. This is done by means of numerical examples, including one that shows how convergence of Parareal can be estimated by the maximum singular value of a Parareal iteration matrix.

## 2 Parareal

To match the numerical examples in §3, the presentation of Parareal given here starts with an initial value problem

$$My_t(t) = f(y(t), t), \quad y(0) = b \in \mathbb{R}^d, \quad t \in [0, T], \quad (1)$$

with a mass matrix  $M$  and right-hand side  $f$  arising from a finite element discretization of a partial differential equation. Let  $(t_n)_{n=0}^N$  with  $t_0 = 0$  and  $t_N = T$  be a decomposition of  $[0, T]$  into  $N$  so-called time-slices  $[t_n, t_{n+1}]$  which, for the sake of simplicity, are assumed to be of equal length here. Furthermore, let  $y_n$  be an approximation to the solution at  $t_n$ , that is  $y_n \approx y(t_n)$ .

Denote by  $\mathcal{F}$  a "fine", computationally expensive and accurate integration method with a time step  $\delta t$  (e.g. a higher-order Runge-Kutta method) and by  $\mathcal{G}$  a "coarse", computationally cheap and probably inaccurate method with a time step  $\Delta t \gg \delta t$  (e.g. implicit Euler). Assume here that the constant length of the time-slices is a multiple of both  $\delta t$  and  $\Delta t$ , so that the fine as well as the coarse method can integrate over one time-slice using a fixed integer number of time-steps. Denote the result of integrating over the slice  $[t_n, t_{n+1}]$ , starting from an initial value  $y$  at  $t_n$ , using the fine or coarse method as  $\mathcal{F}(y, t_{n+1}, t_n)$  and  $\mathcal{G}(y, t_{n+1}, t_n)$  respectively. Serial integration using the fine method would then correspond to computing

$$y_{n+1} = \mathcal{F}(y_n, t_{n+1}, t_n), \quad n = 0, \dots, N-1, \quad (2)$$

step-by-step with  $y_0 := b$ . Instead, Parareal computes the iteration given by

$$y_{n+1}^{k+1} = \mathcal{G}(y_n^{k+1}, t_{n+1}, t_n) + \mathcal{F}(y_n^k, t_{n+1}, t_n) - \mathcal{G}(y_n^k, t_{n+1}, t_n) \quad (3)$$

where the evaluation of the fine method over the  $N$  time-slices can be distributed over  $N$  processors (see Lions et al. [2001] for details). The iteration converges to the serial fine solution as  $k \rightarrow N$ . Speedup can be achieved if  $\mathcal{G}$  is sufficiently cheap compared to  $\mathcal{F}$  and if the iteration converges in  $K \ll N$  iterations. Therefore, rapid convergence is critical for Parareal to be efficient. In the examples below, the defect

$$d^k := \max_{i=0, \dots, N} \|y_i - y_i^k\|_\infty \quad (4)$$

between the solution provided by the Parareal iteration (3) after  $k$  iterations and the serial fine solution (2) is used to measure convergence.

### 3 Heat equation with non-constant coefficients

The test problem used here to study the convergence of Parareal for non-constant coefficients is the two-dimensional heat equation

$$u_t(x, y, t) = \nu(t) \nabla \cdot (a(x, y) \nabla u(x, y, t)) \quad (5)$$

on a square  $\Omega = [0, 1]^2$ . The initial values are given by

$$u_0(x, y) = \exp \left[ - \left( (x - 0.5)^2 + (y - 0.5)^2 \right) / \sigma^2 \right], \quad \sigma = 0.35, \quad (6)$$

and the problem is run until  $T = 4.0$ . The interval  $[0, T]$  is divided into  $N = 40$  time-slices and an implicit Euler method with  $\Delta t = 1/100$  is used for  $\mathcal{G}$  and a third order RadauIIA(3) method with  $\delta t = 1/200$  for  $\mathcal{F}$ . The spatial domain  $\Omega$  is divided into three "strips"

$$\Omega_1 = [0, x_0] \times [0, 1], \quad (7)$$

$$\Omega_2 = [x_0, x_0 + w] \times [0, 1], \quad (8)$$

$$\Omega_3 = [x_0 + w, 1] \times [0, 1], \quad (9)$$

and a different constant value for  $a$  is prescribed on every strip, that is

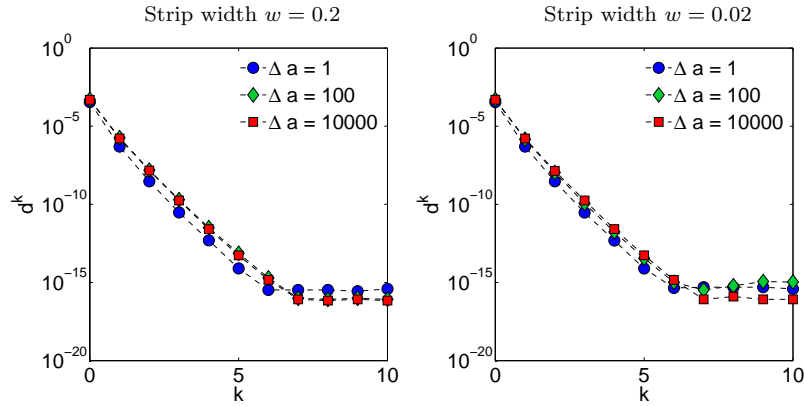
$$a(x, y) = \begin{cases} a_1 : (x, y) \in \Omega_1 \\ a_2 : (x, y) \in \Omega_2 \\ a_3 : (x, y) \in \Omega_3. \end{cases} \quad (10)$$

Furthermore, the effect of varying the width  $w$  of the middle strip  $\Omega_2$  is investigated. Conforming triangle meshes aligned with the strips  $\Omega_i$  are generated for values of  $w \in \{0.2, 0.1, 0.05, 0.02\}$ . Then, for every value of  $w$ , a number of uniform refinement steps is performed in order to produce meshes of comparable mesh width. After refinement, the minimum element sizes for the different values of  $w$  range from  $h_{\min} = 0.01$  to  $h_{\min} = 0.005$  and the maximum element sizes from  $h_{\max} = 0.02$  to  $h_{\max} = 0.035$ , so that the resolutions are comparable. All experiments reported below use linear finite elements, but preliminary tests not documented here suggest that the results are not significantly affected by the use of higher-order FEM. Homogeneous Dirichlet boundary conditions are employed. Simulations are run with  $a_1 = a_3 = 0.01$  fixed and  $a_2 \in \{0.01, 1.0, 100\}$ , resulting in ratios  $\Delta a = a_2/a_3 = a_2/a_1 \in \{1, 100, 10000\}$ .

### 3.1 Space-dependent coefficients

First, set  $\nu \equiv 1$  in order to study only the effect of spatially varying coefficients. Figure 1 shows the resulting convergence of Parareal for the different values of  $\Delta a$  and  $w = 0.2$  (left) and  $w = 0.02$  (right). Convergence in the cases with jumping coefficients is slightly slower, but the effect is very small. Also, the reduction in convergence speed seems to be rather independent of the magnitude of the jump in the diffusion coefficient: In both plots, the lines for  $\Delta a = 100$  and  $\Delta a = 10000$  are more or less indistinguishable.

Convergence of Parareal is utterly oblivious to the width  $w$  of the middle strip  $\Omega_2$ : When plotting the defects for fixed  $\Delta a$  and different values of  $w$ , the resulting data points all essentially coincide so that the corresponding plots are rather uninteresting and are therefore omitted.



**Fig. 1** Defect  $d^k$  between Parareal and the serial fine solution versus the iteration number  $k$  depending on the magnitude of the jump in the diffusion coefficient from  $\Omega_1, \Omega_3$  to  $\Omega_2$ .

### 3.2 Space- and time-dependent coefficients

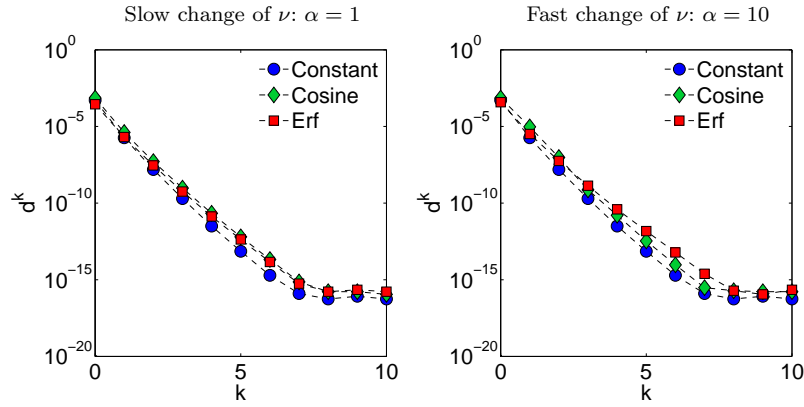
To investigate the effect of a time-dependent diffusion coefficient on the convergence of Parareal, fix the strip width to  $w = 0.2$  and the coefficient jump to  $\Delta a = 100$ . Furthermore, use the following three different profiles for the time-dependent diffusion coefficient  $\nu$  :

$$\nu(t) = 1 \quad (\text{"constant"}), \quad (11)$$

$$\nu(t) = \frac{1}{2} \left( 1 + \cos \left( \alpha \frac{\pi}{2} t \right) \right) \quad (\text{"cosine"}), \quad (12)$$

$$\nu(t) = \frac{1}{2} (1 + \operatorname{erf}(\alpha(t - 2))) \quad (\text{"erf"}). \quad (13)$$

Initial value and boundary conditions are set as described above. Two sets of simulations are performed, one with  $\alpha = 1$  corresponding to a very slowly changing  $\nu$  and one with  $\alpha = 10$  corresponding to a more rapid change. The resulting convergence of Parareal is shown in Figure 2. In both cases, the slow as well as the fast varying one, Parareal's convergence is only marginally affected by the space- and time-dependent diffusion coefficients. The resulting defects are slightly larger than for the reference case and the difference is a little more pronounced for  $\alpha = 10$ , but the overall effect is not drastic: In the fast varying case with the error function profile (13), Parareal requires only a single additional iteration compared to the constant reference in order to reach the same defect level.



**Fig. 2** Defect  $d^k$  of Parareal versus the iteration number  $k$  for different time-dependent  $\nu$ -profiles with  $\Delta a = 100$ ,  $\alpha = 1$  (left) and  $\alpha = 10$  (right).

### 3.3 Error bound from singular values

Parareal can be considered as a fixed point iteration, see e.g. Amodio and Brugnano [2009] or Friedhoff et al. [2013] for more detailed explanations. For  $\nu \equiv 1$  and the linear problem considered here, the action of the propagators  $\mathcal{F}$  and  $\mathcal{G}$  can be expressed as multiplication by matrices  $G$  or  $F$ . Then, running the fine or coarse method over all  $N$  time-slices can be expressed as inversion

of size  $Nd \times Nd$  matrices

$$\mathbf{M}_f = \begin{bmatrix} I & \dots & & & \\ -F & I & & & \\ & & \ddots & \ddots & \\ & & & & -F & I \end{bmatrix}, \quad \mathbf{M}_g = \begin{bmatrix} I & \dots & & & \\ -G & I & & & \\ & & \ddots & \ddots & \\ & & & & -G & I \end{bmatrix}, \quad (14)$$

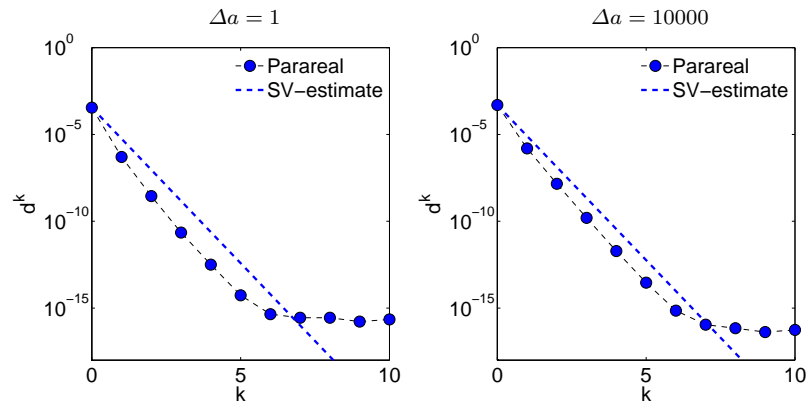
so that computing the fine solution through (2) corresponds to a block-wise solution of  $\mathbf{M}_f \mathbf{y} = \mathbf{b}$  with  $\mathbf{y} = (y_0, \dots, y_N)^\top$  and  $\mathbf{b} = (b, 0, \dots, 0)^\top$ . The Parareal iteration (3) can then be written as the preconditioned fixed point iteration

$$\mathbf{M}_g \mathbf{y}^{k+1} = (\mathbf{M}_g - \mathbf{M}_f) \mathbf{y}^k + \mathbf{b}, \quad (15)$$

where inverting  $\mathbf{M}_g$  corresponds to running the coarse method. A straightforward computation shows that the iteration matrix  $\mathbf{I} - \mathbf{M}_g^{-1} \mathbf{M}_f$  is nilpotent and thus that its spectral radius is zero, corresponding to the well-known fact that Parareal always converges to the fine solution after  $N$  iterations, see e.g. Gander and Vandewalle [2007] (although Parareal won't provide any speedup in this case). A bound for the convergence rate can be obtained by computing the maximum singular value instead. In order to keep the size of the iteration matrix manageable, the example studied here uses only the  $w = 0.2$  geometry with a coarser grid with  $h_{\min} = 0.04$ ,  $h_{\max} = 0.068$  and only  $N = 20$  time-slices. The maximum singular values of the iteration matrix are computed with Matlab's `svds` function and are  $\sigma_{\max} \approx 0.162$  for  $\Delta a = 1$  and  $\sigma_{\max} \approx 0.163$  for  $\Delta a = 10000$ : The minimal difference gives an additional indication that the coefficient jump should not influence Parareal's convergence. Figure 3 shows the convergence rates of Parareal for this example for  $\Delta a = 1$ , i.e. with a constant coefficient (left) and with  $\Delta a = 10000$  (right), as well as the estimate  $d^0 \times (\sigma_{\max})^k$  resulting from the maximum singular value. In both cases, actual convergence is a little better than expected but  $\sigma_{\max}$  gives a reasonable estimate. Again, the jumping coefficients affect Parareal's convergence only marginally. Note that interpreting variants like the "Krylov-subspace-enhanced Parareal", introduced in Gander and Petcu [2008] and studied further in Ruprecht and Krause [2012], as a non-stationary fixed point iteration could be an interesting approach for a mathematical analysis.

## 4 Conclusions

The paper presents a numerical study of the convergence behavior of the time-parallel Parareal method for the heat equation with space- and time-dependent coefficients. It demonstrates that the good convergence of Parareal for diffusive problems is only marginally affected by both jumps in the dif-



**Fig. 3** Convergence of Parareal and error estimate from the largest singular value of the Parareal iteration matrix (dashed line).

fusion coefficients and a diffusion coefficient that changes in time. For linear problems, Parareal can be interpreted as a preconditioned fixed point iteration and, at least for small enough problems, the iteration matrix and its maximum singular value can be computed numerically. An example is shown that demonstrates that the largest singular value gives a reasonable estimate for the convergence of Parareal. Extending the analysis presented here to more complicated cases e.g. in three dimensions with complicated geometries, with coefficient jumps not aligned with the mesh or cases that also include advection would be an interesting direction of future research.

### Acknowledgments.

This work was supported by the Swiss National Science Foundation (SNSF) under the lead agency agreement through the project "ExaSolvers" within the Priority Programme 1648 "Software for Exascale Computing" (SPPEXA) of the Deutsche Forschungsgemeinschaft (DFG). The authors thankfully acknowledge support from Achim Schädle, who provided parts of the used code.

### References

- P. Amodio and L. Brugnano. Parallel solution in time of odes: some achievements and perspectives. *Applied Numerical Mathematics*, 59(3-4):424 – 435, 2009.
- G. Bal. On the convergence and the stability of the parareal algorithm to solve partial differential equations. In Ralf Kornhuber and et al., editors,

- Domain Decomposition Methods in Science and Engineering*, volume 40 of *Lecture Notes in Computational Science and Engineering*, pages 426–432, Berlin, 2005. Springer.
- Andrew J Christlieb, Colin B Macdonald, and Benjamin W Ong. Parallel high-order integrators. *SIAM Journal on Scientific Computing*, 32(2):818–835, 2010.
- M. Emmett and M. L. Minion. Toward an efficient parallel in time method for partial differential equations. *Communications in Applied Mathematics and Computational Science*, 7:105–132, 2012.
- C. Farhat and M. Chandesris. Time-decomposed parallel time-integrators: theory and feasibility studies for fluid, structure, and fluid-structure applications. *International Journal for Numerical Methods in Engineering*, 58(9):1397–1434, 2003.
- S Friedhoff, R D Falgout, T V Kolev, S MacLachlan, and J B Schroder. A multigrid-in-time algorithm for solving evolution equations in parallel. In *Presented at: Sixteenth Copper Mountain Conference on Multigrid Methods, Copper Mountain, CO, United States, Mar 17 - Mar 22, 2013*, 2013.
- M. Gander and E. Hairer. Nonlinear convergence analysis for the parareal algorithm. In U. Langer, O. Widlund, and D. Keyes, editors, *Domain Decomposition Methods in Science and Engineering*, volume 60 of *Lecture Notes in Computational Science and Engineering*, pages 45–56. Springer, 2008.
- M. Gander and M. Petcu. Analysis of a Krylov subspace enhanced parareal algorithm for linear problems. *ESAIM: Proc.*, 25:114–129, 2008.
- M. J. Gander and S. Vandewalle. Analysis of the parareal time-parallel time-integration method. *SIAM Journal on Scientific Computing*, 29(2):556–578, 2007.
- 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*, 332:661–668, 2001.
- Y. Maday and G. Turinici. The parareal in time iterative solver: A further direction to parallel implementation. In Ralf Kornhuber and et al., editors, *Domain Decomposition Methods in Science and Engineering*, volume 40 of *Lecture Notes in Computational Science and Engineering*, pages 441–448, Berlin, 2005. Springer.
- J. Nievergelt. Parallel methods for integrating ordinary differential equations. *Commun. ACM*, 7(12):731–733, 1964.
- D. Ruprecht and R. Krause. Explicit parallel-in-time integration of a linear acoustic-advection system. *Computers & Fluids*, 59(0):72 – 83, 2012.
- 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, SC '12*, pages 92:1–92:11, Los Alamitos, CA, USA, 2012. IEEE Computer Society Press.