Space-Time Parallel Methods for Linear Elastodynamics

Stéphanie Chaillat, Marion Darbas, Martin J. Gander, Laurence Halpern

1 Introduction

We are interested in solving numerically in parallel linear elastodynamics problems, i.e., the movement of a linear isotropic elastic solid in $\mathbb{R}^d$, $d = 2$ or 3. Elastodynamics is widely used for example to simulate seismic waves in large scale domains. In contrast to time harmonic approaches [1] that lead to linear systems that are even harder to solve by iterative methods than the Helmholtz equation [3], we want to tackle directly elastodynamics problems in space-time using solvers that are Parallel in Time (PinT). We will consider two major techniques: Schwarz Waveform Relaxation (SWR) with a specific coloring of the subdomains and generous overlap leading to a new Unmapped Tent Pitching algorithm (UTP) for elastodynamics, and ParaDiag II methods, both applied directly to the all at once space-time elastodynamics system.

Our new results for the vector valued elastodynamics are guided by previous analyses for the scalar wave equation, a hyperbolic problem which already has the main features that make the parallel computation by PinT methods challenging; see [6, 5] for classical and optimized Schwarz waveform relaxation, [9] for the parareal algorithm, [13, 7, 8] for diagonalization and the two ParaDiag families of algorithms, and [11, 10], for Mapped Tent Pitching (MTP), and [2] for the new idea of UTP for scalar wave equations in 1D.

We focus here on elastodynamics because of its major role in real life applications, but we emphasize that all the results below (e.g., Theorem 1) are valid for

Stéphanie Chaillat
Laboratoire POEMS, CNRS-INRIA-ENSTA Paris, e-mail: stephanie.chaillat@ensta-paris.fr

Marion Darbas
Université Sorbonne Paris Nord, e-mail: marion.darbas@paris13.fr

Martin J. Gander
Université de Genève, e-mail: martin.gander@unige.ch

Laurence Halpern
Université Sorbonne Paris Nord, e-mail: laurence.halpern@paris13.fr

general symmetric hyperbolic systems, like the Euler system (modeling the motion of a homogeneous compressible fluid supposed to be perfect with constant temperature, neglecting the viscosity effects), the Saint-Venant system (modeling the 3D flow in shallow water), the Maxwell system (relating electric and magnetic fields), or the Dirac equations (modeling the movement of an electron in relativistic quantic mechanics). All these equations have the important property of finite speed of propagation, a property shared by all symmetric hyperbolic systems. This property can be obtained for instance by identifying the cone of propagation through energy estimates.

2 Schwarz Waveform Relaxation and Unmapped Tent Pitching

We start by defining and analyzing the classical Schwarz Waveform Relaxation method applied to linear elastodynamics modeled by the Navier equations, which are

$$\mathcal{L}(\mathbf{u}) := \rho\partial_{tt}\mathbf{u} - \Delta^e\mathbf{u} = \mathbf{f} \quad \text{in } \Omega \times (0,T], \tag{1}$$

where $\mathbf{u} \in \mathbb{R}^d$, $d = 2$ or 3, is the deformation to be determined of the solid domain Ω, $\mathbf{f}$ is the external force applied to the solid, the elasticity operator $\Delta^e\mathbf{u} := \mu\Delta\mathbf{u} + (\lambda+\mu)\nabla(\nabla\cdot\mathbf{u})$, where λ and μ are the Lamé parameters, ρ is the density, and the problem needs to be completed with initial conditions and boundary conditions. A classical overlapping Schwarz Waveform Relaxation (SWR) algorithm is based on a decomposition of the domain Ω into overlapping subdomains Ω_j, $j = 1, 2, \ldots, J$. Starting with an initial guess $\mathbf{u}^0$ in space time satisfying the boundary conditions, $\mathbf{u}^0 = \mathbf{g}$ on $\partial\Omega$, SWR then computes for iteration index $n = 1, 2, \ldots$ the space-time subdomain approximations $\mathbf{u}_j^n$ by solving[1]

$$\begin{aligned} \mathcal{L}\mathbf{u}_j^n &= \mathbf{f} && \text{in } \Omega_j \times (0,T), \\ \mathbf{u}_j^n &= \mathbf{u}^{n-1} && \text{on } \partial\Omega_j \times (0,T), \\ \mathbf{u}_j^n(\cdot,0) &= \mathbf{u}_0 && \text{in } \Omega_j, \\ \partial_t\mathbf{u}_j^n(\cdot,0) &= \tilde{\mathbf{u}}_0 && \text{in } \Omega_j, \end{aligned} \tag{2}$$

from which then a global approximation $\mathbf{u}^n$ in space-time is obtained using a partition of unity χ_j associated with the overlapping subdomains Ω_j,

$$\mathbf{u}^n = \sum_{j=1}^{J} \chi_j \mathbf{u}_j^n. \tag{3}$$

The specificity for elastic waves is that two types of waves propagate: primary (P or compressional) waves propagate with velocity $c_P = \sqrt{\frac{\lambda+2\mu}{\rho}}$, and secondary (S or

[1] The boundary condition $\mathbf{g}$ is automatically imposed through the global approximation $\mathbf{u}^{n-1}$ that satisfies it, since $\mathbf{u}^0$ initially satisfies it.

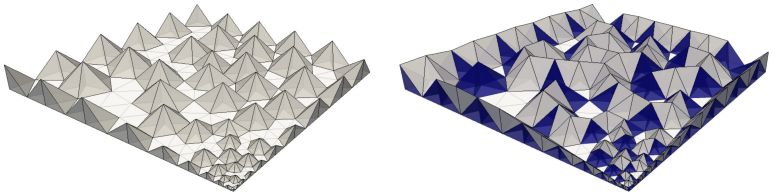

Fig. 1 Mapped Tent Pitching for hyperbolic problems.

shear) waves with velocity $c_S = \sqrt{\frac{\mu}{\rho}}$. Due to the hyperbolic nature of the Navier equations, the solutions of the equation propagate with finite speed bounded by $c_P > c_S$. This implies for overlapping SWR that in each space-time subdomain, the errors, which are only present on the interfaces, since the initial condition is known, will propagate into the interior of the subdomain at finite speed, and thus the solution in a cone in each subdomain is obtained exactly (a key property used later, see the discussion of Figure 2). This solution is then used in the next iteration as boundary data for the neighbors, and the part within the cone is now the exact boundary data. Hence, in each iteration, a bigger and bigger portion in time of exact boundary data is generated, and therefore, overlapping SWR converges in a finite number of steps on bounded time intervals, $T < \infty$, *i.e.*, it is a direct solver. This result, proved for the wave equation in [6, 5], holds also for elastodynamics, and hyperbolic problems in general:

Theorem 1 (Finite Step Convergence) *For* $\mathbf{u}_0 \in H^1(\Omega)$, $\tilde{\mathbf{u}}_0 \in L^2(\Omega)$, $\mathbf{f} \in L^2(0,T;L^2(\Omega))$, $\mathbf{g} \in L^2(0,T;H^{\frac{1}{2}}(\Gamma))$ *and* $\mathbf{u}^0 \in L^2(0,T;H^1(\Omega))$, *overlapping SWR has converged in* $L^2(0,T;H^1(\Omega))$ *as soon as the iteration index n satisfies*

$$n > \frac{T c_P}{\delta}, \quad \delta \text{ smallest overlap,} \tag{4}$$

where $c_P = \sqrt{\frac{\lambda+2\mu}{\rho}}$ *is the speed of the P-waves.*

Proof. The proof follows as in [5] for the wave equation, but now with the fastest propagation speed of elastodynamics. □

A very different technique from SWR in PinT, at first sight, is Mapped Tent Pitching (MTP) [11]. The original Tent Pitching (TP) goes back to space-time finite element methods, see for example [12], where the space-time finite elements need to satisfy slope constraints based on the speed of propagation to get a well posed method. The more recent MTP algorithm [11] for hyperbolic problems uses also the fact that the propagation speed of information is finite in such problems, and is best understood by looking at a picture from [11], see Figure 1. As in TP, the algorithm constructs space-time tents, but these represent now true space-time subdomains not just one finite element each. In each of these tents, the solution can be computed independently of one another, due to the finite speed of propagation, the solution can

be computed in all gray tents in parallel (see Figure 1 left). The tents are then marked in blue to indicate the solution there is known. The algorithm next constructs a new set of gray tents (see Figure 1 right), where the solution can again be computed in parallel since the solution is now known in the blue tents. To compute the solution in the tents, MTP maps the tents into space-time cylinders, solves the problem there by time stepping, and then maps it back, which led to its name. MTP is probably the best PinT Maxwell space-time solver currently available, see [10]. It has however the drawback that the mapping has to be computed, and when using classical time steppers in the mapped tents like for example Runge Kutta methods, order reduction occurs, and specialized structure aware time steppers need to be used [10].

In order to avoid these drawbacks, in [2] a new, Unmapped Tent Pitching algorithm (UTP) was introduced for a one dimensional wave equation, based on overlapping SWR. We now develop this technique as a new space-time solver for elastodynamics in two spatial dimensions. To explain the algorithm, it is best to consider the specific example of a square part of a domain Ω (see the illustration in Figure 2). In the top row, we present the subdomains for the first space-time parallel solve in red. On the left we show the spatial subdomains, and in the middle the space-time subdomains. On the right the red pyramids indicate where the exact solution is obtained, as a result of the finite speed of propagation when the red space-time subdomain problems shown in the middle are solved, independently of the initial guess used. In the second row, we show the green subdomains that are solved for the second step. On the right, the green tetrahedra represent the exact solution now obtained. In the third row, the blue subdomains (which could be computed simultaneously with the green ones), and in the last row the black subdomains, represent again the space-time region where the exact solution is obtained. The algorithm then continues again with the green and blue subdomains, but further up in time, followed by the red subdomains playing now the role of the black subdomains.

This overlapping SWR algorithm for the specific red, green, blue and black subdomains in space time can be very easily implemented using Restricted Additive Schwarz (RAS) in space-time (see [4] for this general equivalence in space), which leads to the following algorithm applied to the all at once space-time problem $A\mathbf{u} = \mathbf{f}$: one simply computes for an initial guess $\mathbf{u}^0$ in space-time and iteration index $n = 0, 6, 12, \ldots$

$$
\begin{aligned}
\mathbf{u}^{n+1} &= \mathbf{u}^n + \textstyle\sum \tilde{R}^T_{r_{n+1}} A^{-1}_{r_{n+1}} R_{r_{n+1}} (\mathbf{f} - A\mathbf{u}^n),\\
\mathbf{u}^{n+2} &= \mathbf{u}^{n+1} + \textstyle\sum \tilde{R}^T_{g_{n+2}} A^{-1}_{g_{n+2}} R_{g_{n+2}} (\mathbf{f} - A\mathbf{u}^{n+1}),\\
\mathbf{u}^{n+3} &= \mathbf{u}^{n+2} + \textstyle\sum \tilde{R}^T_{b_{n+3}} A^{-1}_{b_{n+3}} R_{b_{n+3}} (\mathbf{f} - A\mathbf{u}^{n+2}),\\
\mathbf{u}^{n+4} &= \mathbf{u}^{n+3} + \textstyle\sum \tilde{R}^T_{k_{n+4}} A^{-1}_{k_{n+4}} R_{k_{n+4}} (\mathbf{f} - A\mathbf{u}^{n+3}),\\
\mathbf{u}^{n+5} &= \mathbf{u}^{n+4} + \textstyle\sum \tilde{R}^T_{g_{n+5}} A^{-1}_{g_{n+5}} R_{g_{n+5}} (\mathbf{f} - A\mathbf{u}^{n+4}),\\
\mathbf{u}^{n+6} &= \mathbf{u}^{n+5} + \textstyle\sum \tilde{R}^T_{b_{n+6}} A^{-1}_{b_{n+6}} R_{b_{n+6}} (\mathbf{f} - A\mathbf{u}^{n+5}),
\end{aligned}
\tag{5}
$$

where the restriction matrices R and $\tilde{R}$ come from RAS for the four colored subdomains. Note that while we presented here the specific case of a square decomposition, algorithm (5) can run with an arbitrary spatial decomposition and will produce good

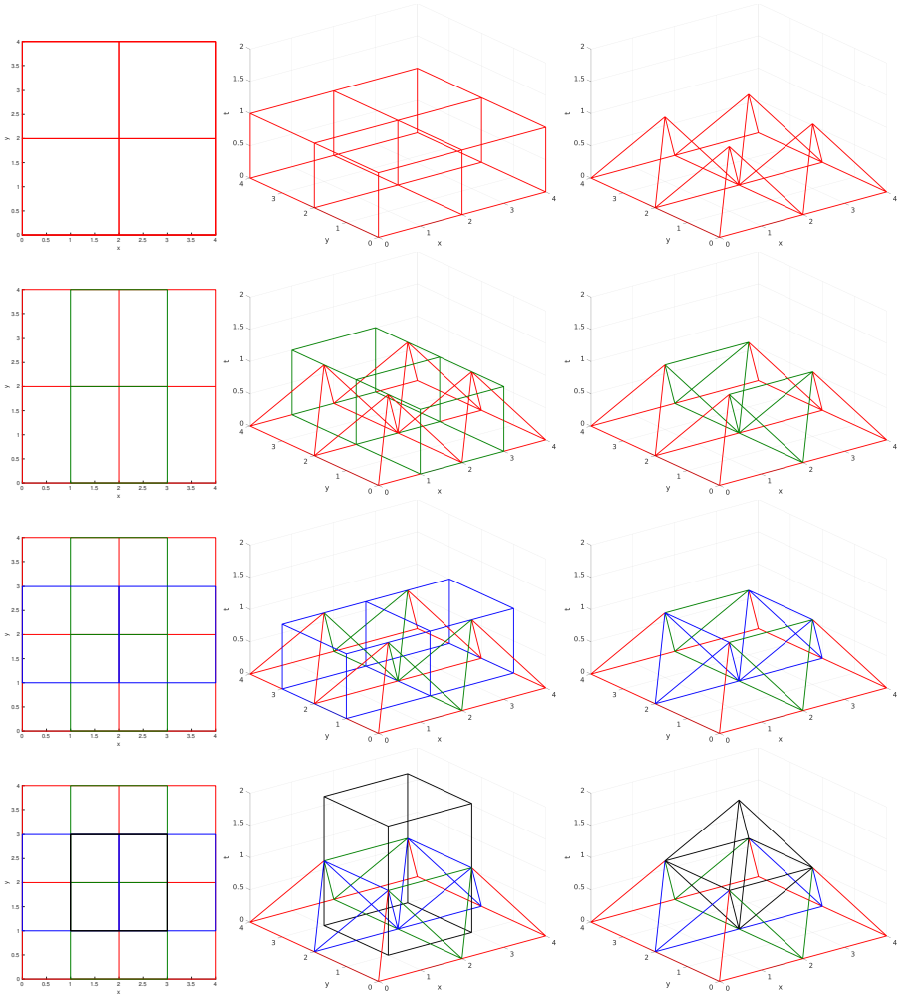

Fig. 2 First four iterations of the new Unmapped Tent Pitching algorithm for elastodynamics in two spatial dimensions, see text for more details.

results with generous overlap, and it is even possible to determine the subdomain height in time adaptively using the residual.

Starting with a random vector valued initial guess $\mathbf{u}^0$ in space and time, we show in Figure 3 for $\rho = \mu = \lambda = 1$ and a standard finite difference discretization of (1) with $\Delta x = \Delta y = \frac{1}{25}$ and $\Delta t = \frac{2}{125}$ snapshots in time of the first iteration of the error in the first solution vector component after the red, green, blue and black solves corresponding to iteration index 1, 2, 3 and 4 in (5). We see in the first column how the error is zero in the geometric structure corresponding to the four red pyramids in Figure 2 (and also close to the boundary since the boundary conditions are known). For the later snapshots in time, the error is still random, since the algorithm has not

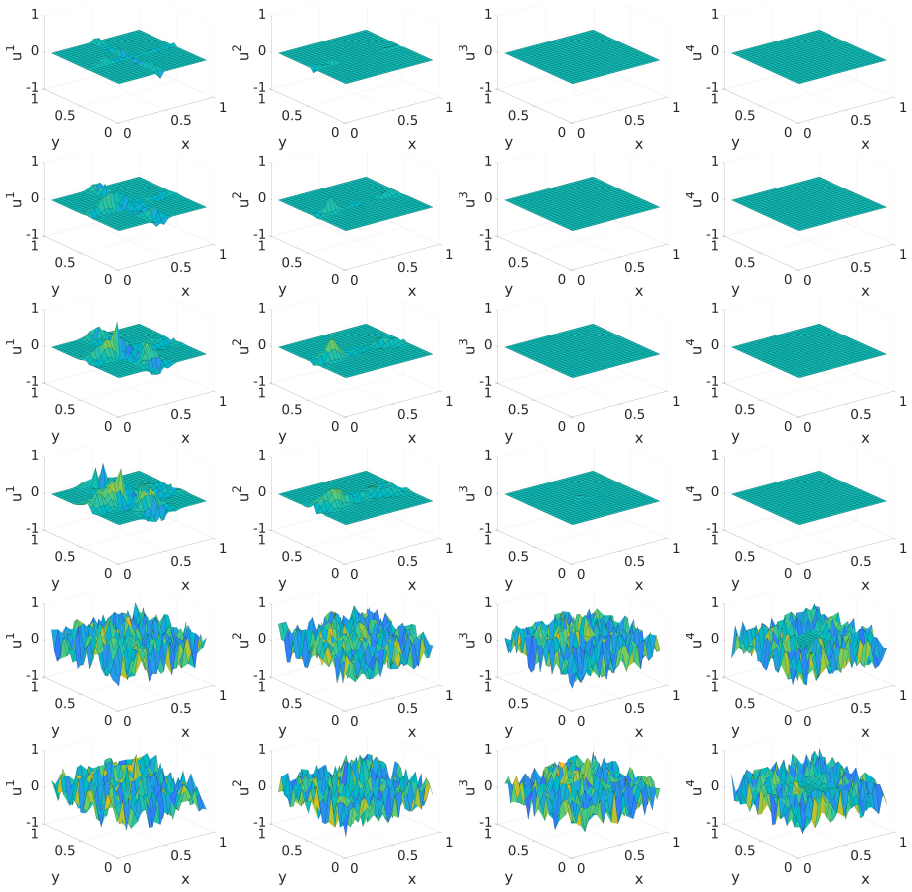

Fig. 3 Snapshots in time for $t = 2\Delta t, 4\Delta t, 6\Delta t, 8\Delta t, 10\Delta t, 12\Delta t$ of the overlapping SWR algorithm applied to elastodynamics on the colored space-time subdomains: first column after the red solve, second after the green solve, third after the blue solve and last after the black solve.

performed any work there yet. In the second column the error is now also zero in the green tetrahedra from Figure 2; in the third column also in the blue tetrahedra, and finally in the last column also in the double black pyramid shown also in Figure 2 (to see the tip of the pyramid, look at the last two plots in the right column which have now zero error in the middle). This shows that overlapping SWR with these colored fully overlapping subdomains in space automatically computes solutions in the tents of the corresponding MTP algorithm, but without the mapping. This algorithm has multiple advantages: it is much easier to implement using the RAS formalism applied to the all at once space-time system than MTP, and it is also cheaper, because first, no mapping has to be computed, and second, the computational volume is the same, since after the mapping, MTP computes in the same space-time volume as UTP. Furthermore, UTP does not require special time integrators to avoid order reduction.

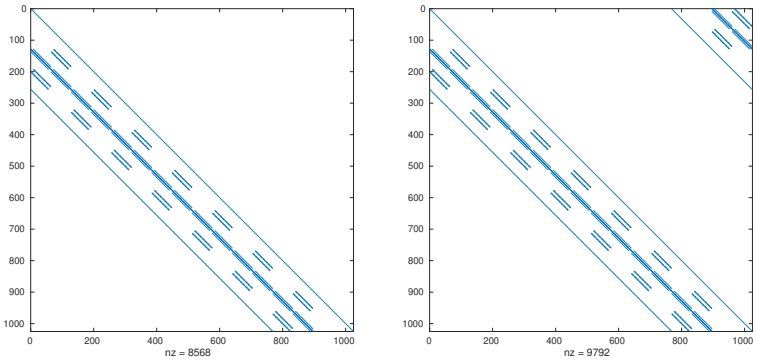

Fig. 4 Sparsity of the all at once space-time system for 2D elastodynamics (left), and structure of the corresponding circulant (or α-circulant) preconditioner.

One might argue that UTP is not truly a PinT method, since it does not solve over very long time in a parallel fashion, it still looks like time marching, albeit over much longer time intervals in parallel, corresponding to the tent heights. We next show for the first time for elastodynamics a very different PinT technique, namely ParaDiag II, which allows one to truly compute parallel in time over long time intervals.

3 A ParaDiag II Algorithm for Elastodynamics

ParaDiag methods are based on the diagonalization of the time stepping matrix [13]. There are two ParaDiag variants currently available. On the one hand, ParaDiag I requires different time steps to be used, see [7] for an application to elastodynamics, but is then a direct PinT solver. On the other hand, ParaDiag II methods modify the time stepping matrix to make it either circulant (periodic) or α-circulant (α-periodic), and then one needs to use an iteration, see [8] for a review. We show in Figure 4 on the left the sparsity structure of the all at once space-time system matrix A for the 2D elastodynamics problem, and on the right the sparsity of the α-circulant preconditioner $\tilde{A}$. Due to the α-circulant structure, the time stepping component in $\tilde{A}$ can be diagonalized using Fast Fourier Transforms (FFT), and then all time steps can be solved parallel in time for the spatial components of the elastodynamics problem. We can then solve the elastodynamics problem iteratively parallel in time performing either the simple residual correction iteration

$$\mathbf{u}^{n+1} = \mathbf{u}^n + \tilde{A}^{-1}(\mathbf{f} - A\mathbf{u}^n), \tag{6}$$

or one can also use $\tilde{A}$ as preconditioner for a Krylov method like GMRES. We show in Figure 5 for the same model parameters and discretization as in Section 2 for Figure 3 the convergence of the residual correction iteration for different values of the parameter α, on the left for the stationary iteration (6) and on the right with

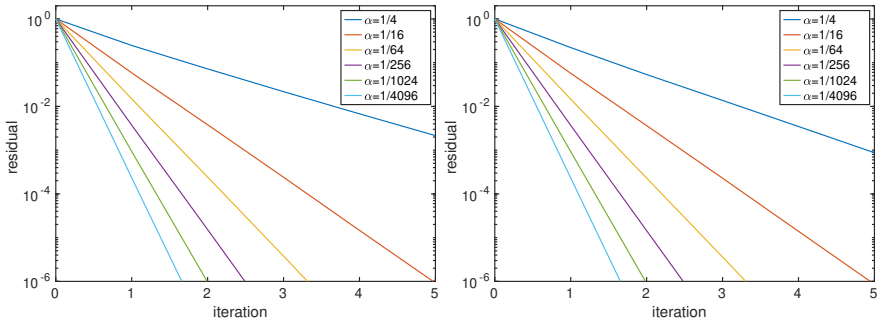

Fig. 5 Preconditioned normalized residual of ParaDiag II applied to the elastodynamics problem for different values of α. Left: stationary iteration (6). Right: with GMRES acceleration.

GMRES acceleration. We see that making α small leads to very rapid convergence, and the ParaDiagII preconditioner is so good that GMRES acceleration does not bring any further gain when α is small, only for $\alpha = 1/4$ there is some gain!

References

1. Brunet, R., Dolean, V., Gander, M.J.: Natural domain decomposition algorithms for the solution of time-harmonic elastic waves. SIAM Journal on Scientific Computing **42**(5), A3313–A3339 (2020)
2. Ciaramella, G., Gander, M.J., Mazzieri, I.: Unmapped tent pitching schemes by waveform relaxation. In: International Conference on Domain Decomposition Methods, pp. 455–462. Springer (2022)
3. Ernst, O.G., Gander, M.J.: Why it is difficult to solve Helmholtz problems with classical iterative methods. Numerical analysis of multiscale problems pp. 325–363 (2011)
4. Gander, M.J.: Schwarz methods over the course of time. Electron. Trans. Numer. Anal **31**(5), 228–255 (2008)
5. Gander, M.J., Halpern, L.: Absorbing boundary conditions for the wave equation and parallel computing. Math. Comp. **74**(249), 153–176 (2005)
6. Gander, M.J., Halpern, L., Nataf, F.: Optimal Schwarz waveform relaxation for the one dimensional wave equation. SIAM J. Numer. Anal. **41**(5), 1643–1681 (2003)
7. Gander, M.J., Halpern, L., Rannou, J., Ryan, J.: A direct time parallel solver by diagonalization for the wave equation. SIAM Journal on Scientific Computing **41**(1), A220–A245 (2019)
8. Gander, M.J., Liu, J., Wu, S.L., Yue, X., Zhou, T.: Paradiag: Parallel-in-time algorithms based on the diagonalization technique. arXiv preprint arXiv:2005.09158 (2020)
9. Gander, M.J., Vandewalle, S.: Analysis of the parareal time-parallel time-integration method. SIAM Journal on Scientific Computing **29**(2), 556–578 (2007)
10. Gopalakrishnan, J., Hochsteger, M., Schöberl, J., Wintersteiger, C.: An explicit mapped tent pitching scheme for Maxwell equations. In: Spectral and high order methods for partial differential equations—ICOSAHOM 2018, pp. 359–369 (2020)
11. Gopalakrishnan, J., Schöberl, J., Wintersteiger, C.: Mapped tent pitching schemes for hyperbolic systems. SIAM Journal on Scientific Computing **39**(6), B1043–B1063 (2017)
12. Lowrie, R.B., Roe, P.L., Van Leer, B.: Space-time methods for hyperbolic conservation laws. In: Barriers and Challenges in Computational Fluid Dynamics, pp. 79–98. Springer (1998)
13. Maday, Y., Rønquist, E.M.: Parallelization in time through tensor-product space–time solvers. Comptes Rendus Mathematique **346**(1), 113–118 (2008)