

Convergence of the multigrid-reduction-in-time algorithm for the linear elasticity equations

A. Hessesenthaler^{1*}, D. Nordsletten², O. Röhrle¹, J. B. Schroder³, R. D. Falgout³

¹*Institute of Applied Mechanics (CE), University of Stuttgart, Pfaffenwaldring 7, 70569 Stuttgart, Germany*

²*Division of Imaging Sciences and Biomedical Engineering, King's College London, 4th Floor, Lambeth Wing St. Thomas Hospital London, SE1 7EH, UK*

³*Center for Applied Scientific Computing, Lawrence Livermore National Laboratory, P.O. Box 808, L-561, Livermore, CA 94551*

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344, LLNL-JRNL-731168.

SUMMARY

This paper presents some recent advances for parallel-in-time methods applied to linear elasticity. With recent computer architecture changes leading to stagnant clock speeds, but ever increasing numbers of cores, future speedups will be available through increased concurrency. Thus, sequential algorithms, such as time stepping, will suffer a bottleneck. This paper explores multigrid reduction in time (MGRIT) for an important application area for many time stepping codes, linear elasticity. Previously, efforts at parallel-in-time for elasticity have experienced difficulties, for example, the beating phenomenon, leading to the current state of no practical parallel-in-time algorithm existing for this application area. This paper proposes some solutions made possible by MGRIT (*e.g.*, slow temporal coarsening and FCF-relaxation) and more importantly, a different formulation of the problem that is more amenable to parallel-in-time methods. Using a recently developed convergence theory for MGRIT and Parareal, we show that the changed formulation of the problem avoids the instability issues and allows reduction of the error using two temporal grids. The paper ends with supporting numerical results showing a practical algorithm. Copyright © 2010 John Wiley & Sons, Ltd.

Received . . .

KEY WORDS: multigrid reduction in time (MGRIT); parallel-in-time; linear elasticity; convergence estimate

1. INTRODUCTION

Clock rates of recent computer architectures have become stagnant, while the available number of parallel processors has increased and continues to increase rapidly. Thus, simulation algorithms need to allow greater concurrency to exploit massively parallel hardware and further reduce wall clock time. One severe sequential bottleneck in many parallel application codes is the use of sequential time integration methods. This sequential bottleneck limits parallelism to the spatial component of an investigated space-time problem.

One such research area where sequential time-stepping is almost omnipresent is fluid-structure interaction (FSI) research. Here, the system under consideration studies the interaction between fluids and (often deformable) solid structures or particles, for example, in biomedical or aerospace

*Correspondence to: A. Hessesenthaler, Institute of Applied Mechanics (CE), University of Stuttgart, Pfaffenwaldring 7, 70569 Stuttgart, Germany. Email: hessesenthaler@mechbau.uni-stuttgart.de

engineering applications [1, 2, 3]. Well-established parallelization techniques, such as spatial domain decomposition methods, provide a straightforward and scalable approach for reducing the wall clock time for many FSI algorithms. However, spatial parallelism saturates when communication tasks become dominant over computation tasks. This often prohibits use of parallel resources beyond an optimal processor count, leaving large numbers of processors unused. On the other hand, parallel-in-time integration methods provide means of introducing an additional layer of parallelism with the potential to greatly enhance efficient and more exhaustive use of parallel resources. Despite its long history [4, 5], parallel-in-time has very little uptake in the large-scale simulation community. Only few works exist that use parallel-in-time ideas in real-world simulations, including reservoir simulation [6], fusion research [7] and numerical weather prediction [8].

In the field of FSI research, limitations to using parallel-in-time methods usually exist because of an instability arising in the parallel-in-time integration of one of its subproblems, that is, dynamic structural mechanics. The instability is known as the beating phenomenon [9, 10]. Although this instability can be overcome by filtering the natural modes of a given structural dynamics problem [10, 11], only small scale parallelism was explored. This can mainly be related to a relatively expensive projection step and the restriction of the method to two time grid levels. Thus, it is of prime importance to obtain a stable and robust parallel-in-time technique that can explore greater parallelism while avoiding the previously mentioned instability. Such a method can then be the foundation of a fast and efficient parallel-in-time method applicable to FSI applications.

The focus of this work is the investigation of convergence of the multigrid reduction in time (MGRIT) method for the second-order partial differential equation governing the dynamic linear-elastic response of an incompressible solid structure. In a special case (two temporal grids with F-relaxation), MGRIT is equivalent [12] with Parareal [13] but has several benefits. For example, MGRIT allows use of FC-relaxation, slow temporal coarsening and is a non-intrusiveness algorithm. Moreover, MGRIT is a true multilevel algorithm. In this paper, we will make use of the recent work of the authors of [12] to use the presented two-grid convergence theory as a design tool for convergent algorithms. The beating phenomenon was described qualitatively in [9, 10]. Here, we will show that it is also founded in the two-grid theory. Furthermore, the same analysis provides a powerful tool for estimating convergence *a priori* that leads to a convergent two-grid MGRIT algorithm for the dynamic second-order elasticity equations and further, to a convergent multigrid MGRIT algorithm.

In Section 2, two different backward Euler time discretization schemes are presented (referred to as Scheme I and Scheme II) for a finite element implementation of the second-order hyperbolic elasticity equations. The two schemes are embedded in the MGRIT algorithm, where we provide a description of the parallel-in-time algorithm and discuss the application of the convergence theory presented in [12]. We discuss the difference between both time discretization schemes in Section 3.2. In Section 3, we relate the theoretical analysis of the two-grid convergence to observations in numerical experiments. Further, we demonstrate how to derive a convergent MGRIT algorithm that does not exhibit the instabilities reported in previous works [9, 10] by using [12] as an algorithm design tool. We then extend the numerical experiments to a multilevel case with FMG-cycles in Section 3.4. The section concludes with preliminary speedup results, emphasizing the potential for parallel speedup when using MGRIT for linear elasticity problems.

2. METHODOLOGY

2.1. Notation

Consider the domain $\Omega = \Omega(t) \subset \mathbb{R}^d \times [0, T]$ with Dirichlet boundary Γ^D in d spatial dimensions. Then, $\mathbf{X} \in \Omega(0)$ and $\mathbf{x} \in \Omega$ denote the reference and current position of a material point and $\nabla_{\mathbf{X}}$ and $\nabla_{\mathbf{x}}$ denote the respective Lagrangian and Eulerian gradient operators. Further, we define the deformation gradient $\mathbf{F} = \nabla_{\mathbf{X}} \mathbf{x} = \nabla_{\mathbf{X}} \mathbf{u} + \mathbf{I}$ where $\mathbf{u} = \mathbf{u}(\mathbf{x}, t) = \mathbf{x} - \mathbf{X}$ is the displacement of

a material point with respect to its position in the reference configuration. The partial derivative operator is denoted as ∂_t and $\partial_{tt} = \partial_t \partial_t$.

2.2. Governing equations

Consider the governing equations for the dynamic and linear elastic response of an incompressible solid structure with given initial data and Dirichlet boundary condition data,

$$\rho \partial_{tt} \mathbf{u} - \nabla_{\mathbf{x}} \cdot \boldsymbol{\sigma} = \mathbf{0} \quad \text{in } \Omega, \quad (1)$$

$$\nabla_{\mathbf{x}} \cdot \partial_t \mathbf{u} = 0 \quad \text{in } \Omega, \quad (2)$$

$$\mathbf{u}(\cdot, t) = \mathbf{0} \quad \text{on } \Gamma^D, \quad (3)$$

$$\mathbf{u}(\cdot, 0) = \mathbf{0}, \quad \mathbf{v}(\cdot, 0) = \hat{\mathbf{v}}_0 \quad \text{in } \Omega(0), \quad (4)$$

with density ρ , Cauchy stress tensor $\boldsymbol{\sigma}(\mathbf{u}, p) = \mu(\mathbf{F} - \mathbf{I}) - p\mathbf{I}$, material stiffness parameter μ , the hydrostatic pressure variable p and initial velocity $\hat{\mathbf{v}}_0$.

Equation 1 can be transformed to a system of first-order equations,

$$\partial_t \mathbf{u} = \mathbf{v} \quad \text{in } \Omega, \quad (5)$$

$$\rho \partial_t \mathbf{v} = \nabla_{\mathbf{x}} \cdot \boldsymbol{\sigma} \quad \text{in } \Omega, \quad (6)$$

$$\nabla_{\mathbf{x}} \cdot \mathbf{v} = 0 \quad \text{in } \Omega, \quad (7)$$

$$\mathbf{u}(\cdot, t) = \mathbf{0} \quad \text{on } \Gamma^D, \quad (8)$$

$$\mathbf{u}(\cdot, 0) = \mathbf{0}, \quad \mathbf{v}(\cdot, 0) = \hat{\mathbf{v}}_0 \quad \text{in } \Omega(0), \quad (9)$$

with velocity \mathbf{v} .

To reduce the complexity of the computational model, Equation 5 is eliminated from the system of equations in the following by including it implicitly. That is, we solve for velocity \mathbf{v} and hydrostatic pressure p and update the displacement variable based on the solution for the velocity variable.

Further, we note that all quantities are computed on the reference domain $\Omega(0)$. That is, linear elastic response is assumed and higher order effects of the deforming domain are neglected.

2.3. Time discretization

We decompose the temporal domain $[0, T]$ by $N_t + 1$ equidistant time points, such that

$$t_i = i \cdot \delta_{N_t}, \quad \text{with } i = 0, \dots, N_t, \quad (10)$$

with time step size $\delta_{N_t} = T/N_t$, initial time $t_0 = 0$ and final time $t_{N_t} = T$.

Then, velocity, pressure and displacement at time point t_i are denoted as $[\mathbf{v}_i, p_i, \mathbf{u}_i]^T = [\mathbf{v}(\cdot, t_i), p(\cdot, t_i), \mathbf{u}(\cdot, t_i)]^T$. We now introduce the two discretization schemes (Scheme I and Scheme II) examined here.

2.3.1. Scheme I

Firstly, we approximate the partial derivative operator ∂_t in Equation 5 using the midpoint rule,

$$\mathbf{u}_i = \mathbf{u}_{i-1} + \delta_{N_t} \frac{\mathbf{v}_i + \mathbf{v}_{i-1}}{2} \quad \text{in } \Omega_0. \quad (11)$$

The partial derivative operator in Equation 6 is discretized using the backward Euler scheme where we substitute Equation 11. Thus, we search (\mathbf{v}_i, p_i) for all $i = 1, \dots, N_t$, such that,

$$\rho \mathbf{v}_i - \frac{\mu \delta_{N_t}^2}{2} \nabla_{\mathbf{X}}^2 \mathbf{v}_i + \delta_{N_t} \nabla_{\mathbf{X}} p_i = \rho \mathbf{v}_{i-1} + \frac{\mu \delta_{N_t}^2}{2} \nabla_{\mathbf{X}}^2 \mathbf{v}_{i-1} + \mu \delta_{N_t} \nabla_{\mathbf{X}}^2 \mathbf{u}_{i-1} \quad \text{in } \Omega_0, \quad (12)$$

$$\nabla_{\mathbf{X}} \cdot \mathbf{v}_i = -\nabla_{\mathbf{X}} \cdot \mathbf{v}_{i-1} \quad \text{in } \Omega_0, \quad (13)$$

and compute the displacement \mathbf{u}^i according to Equation 11 after each solve.

2.3.2. Scheme II

We can obtain a slightly modified scheme by approximating the partial derivative operator in Equation 5 as,

$$\mathbf{u}_i = \mathbf{u}_{i-1} + \delta_{N_t} \mathbf{v}_i \quad \text{in } \Omega_0. \quad (14)$$

Thus, in Scheme II we search the time-discrete (\mathbf{v}_i, p_i) for all $i = 1, \dots, N_t$, such that,

$$\rho \mathbf{v}_i - \mu \delta_{N_t}^2 \nabla_{\mathbf{X}}^2 \mathbf{v}_i + \delta_{N_t} \nabla_{\mathbf{X}} p_i = \rho \mathbf{v}_{i-1} + \mu \delta_{N_t} \nabla_{\mathbf{X}}^2 \mathbf{u}_{i-1} \quad \text{in } \Omega_0, \quad (15)$$

$$\nabla_{\mathbf{X}} \cdot \mathbf{v}_i = 0 \quad \text{in } \Omega_0, \quad (16)$$

The modification in Equation 14 is motivated by the observation, that Scheme I does not yield a convergent MGRIT algorithm and exhibits the same instability as described in previous works [9, 10], see Section 3.3. Scheme II does not suffer from this instability, however, does not reproduce amplitudes of oscillation as well as Scheme I (see Section 3.2).

2.4. Space discretization

The domain Ω_i was discretized using quadrilateral elements, Ω_i^h . Finite element discretizations were constructed using inf-sup stable $\mathbb{Q}^2 - \mathbb{Q}^1$ Taylor-Hood elements for velocity and pressure, \mathbf{v}_i^h and p_i^h , and \mathbb{Q}^2 elements for displacement, \mathbf{u}_i^h . The superscript h denotes the space-discretized version of the domain and state variables. As we do not consider spatial refinement or coarsening, we omit the superscript h for the remainder of this work.

2.4.1. Scheme I

Spatially discretizing Equation 12 and Equation 13 leads to the problem of seeking for each $i = 1, \dots, N_t$ the space-time discrete solution vector $[\mathbf{v}_i, \mathbf{p}_i]^T$, such that,

$$(\rho \mathbf{M} - \frac{\mu \delta_{N_t}^2}{2} \mathbf{K}) \mathbf{v}_i + \delta_{N_t} \mathbf{B}^T \mathbf{p}_i = (\rho \mathbf{M} + \frac{\mu \delta_{N_t}^2}{2} \mathbf{K}) \mathbf{v}_{i-1} + \mu \delta_{N_t} \mathbf{K} \mathbf{u}_{i-1} \quad \text{in } \Omega_0, \quad (17)$$

$$\mathbf{B} \mathbf{v}_i = -\mathbf{B} \mathbf{v}_{i-1} \quad \text{in } \Omega_0, \quad (18)$$

where \mathbf{M} is the mass matrix and \mathbf{K} and \mathbf{B} refer to the discretized weak form Laplacian and divergence operators $\nabla_{\mathbf{X}}^2(\cdot)$ and $\nabla_{\mathbf{X}} \cdot (\cdot)$. Note, after solving for a given $[\mathbf{v}_i, \mathbf{p}_i]^T$, we can update the displacement \mathbf{u}_i from Equation 11. Writing Equation 17 and Equation 18 in matrix form and including the update given in Equation 11 yields the following linear system,

$$\begin{bmatrix} \rho \mathbf{M} - \frac{\mu \delta_{N_t}^2}{2} \mathbf{K} & \delta_{N_t} \mathbf{B}^T & \mathbf{0} \\ \mathbf{B} & \mathbf{0} & \mathbf{0} \\ -\frac{\delta_{N_t}}{2} \mathbf{I} & \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{v}_i \\ \mathbf{p}_i \\ \mathbf{u}_i \end{bmatrix} = \begin{bmatrix} \rho \mathbf{M} + \frac{\mu \delta_{N_t}^2}{2} \mathbf{K} & \mathbf{0} & \mu \delta_{N_t} \mathbf{K} \\ -\mathbf{B} & \mathbf{0} & \mathbf{0} \\ \frac{1}{2} \delta_{N_t} \mathbf{I} & \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{v}_{i-1} \\ \mathbf{p}_{i-1} \\ \mathbf{u}_{i-1} \end{bmatrix}. \quad (19)$$

Letting the linear operators on the left and right hand sides be denoted,

$$\mathbf{D}^I := \begin{bmatrix} \rho \mathbf{M} - \frac{\mu \delta_{N_t}^2}{2} \mathbf{K} & \delta_{N_t} \mathbf{B}^T & \mathbf{0} \\ \mathbf{B} & \mathbf{0} & \mathbf{0} \\ -\frac{\delta_{N_t}}{2} \mathbf{I} & \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad \mathbf{C}^I := \begin{bmatrix} \rho \mathbf{M} + \frac{\mu \delta_{N_t}^2}{2} \mathbf{K} & \mathbf{0} & \mu \delta_{N_t} \mathbf{K} \\ -\mathbf{B} & \mathbf{0} & \mathbf{0} \\ \frac{1}{2} \delta_{N_t} \mathbf{I} & \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad (20)$$

then $\Phi^I := [\mathbf{D}^I]^{-1} \mathbf{C}^I$ and $\mathbf{s}_0^I = \hat{\mathbf{s}}_0 := [\hat{\mathbf{v}}_0, \mathbf{0}, \mathbf{0}]^T$, the state variables $\mathbf{s}_i^I := [\mathbf{v}_i, \mathbf{p}_i, \mathbf{u}_i]^T$ can be computed by the following equation,

$$\mathbf{s}_i^I = \Phi^I \mathbf{s}_{i-1}^I \quad \text{for } i = 1, \dots, N_t. \quad (21)$$

Note, that in the linear elastic incompressible case with constant spatial resolution, the operator Φ^I is only dependent on time step size δ_{N_t} .

2.4.2. Scheme II

In a similar way, Scheme II can be written as: we seek for each $i = 1, \dots, N_t$ the space-time discrete $[\mathbf{v}_i, \mathbf{p}_i]^T$, such that,

$$(\rho\mathbf{M} - \mu\delta_{N_t}^2\mathbf{K})\mathbf{v}_i + \delta_{N_t}\mathbf{B}^T\mathbf{p}_i = \rho\mathbf{M}\mathbf{v}_{i-1} + \mu\delta_{N_t}\mathbf{K}\mathbf{u}_{i-1} \quad \text{on } \Omega_0, \quad (22)$$

$$\mathbf{B}\mathbf{v}_i = 0 \quad \text{on } \Omega_0, \quad (23)$$

and update the displacement \mathbf{u}_i from Equation 14 after each solve. In matrix notation, we can write,

$$\begin{bmatrix} \rho\mathbf{M} - \mu\delta_{N_t}^2\mathbf{K} & \delta_{N_t}\mathbf{B}^T & \mathbf{0} \\ \mathbf{B} & \mathbf{0} & \mathbf{0} \\ -\delta_{N_t}\mathbf{I} & \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{v}_i \\ \mathbf{p}_i \\ \mathbf{u}_i \end{bmatrix} = \begin{bmatrix} \rho\mathbf{M} & \mathbf{0} & \mu\delta_{N_t}\mathbf{K} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix} \begin{bmatrix} \mathbf{v}_{i-1} \\ \mathbf{p}_{i-1} \\ \mathbf{u}_{i-1} \end{bmatrix}. \quad (24)$$

With

$$\mathbf{D}^{II} := \begin{bmatrix} \rho\mathbf{M} - \mu\delta_{N_t}^2\mathbf{K} & \delta_{N_t}\mathbf{B}^T & \mathbf{0} \\ \mathbf{B} & \mathbf{0} & \mathbf{0} \\ -\delta_{N_t}\mathbf{I} & \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad \text{and} \quad \mathbf{C}^{II} := \begin{bmatrix} \rho\mathbf{M} & \mathbf{0} & \mu\delta_{N_t}\mathbf{K} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{I} \end{bmatrix}, \quad (25)$$

equation 24 can be written as,

$$\mathbf{s}_i^{II} = \mathbf{\Phi}^{II}\mathbf{s}_{i-1}^{II} \quad \text{for } i = 1, \dots, N_t, \quad (26)$$

where $\mathbf{\Phi}^{II} := [\mathbf{D}^{II}]^{-1}\mathbf{C}^{II}$ and $\mathbf{s}_i^{II} := [\mathbf{v}_i, \mathbf{p}_i, \mathbf{u}_i]^T$ with $\mathbf{s}_0^{II} = \hat{\mathbf{s}}_0 := [\hat{\mathbf{v}}_0, \mathbf{0}, \mathbf{0}]^T$. Again, the operator $\mathbf{\Phi}^{II}$ only depends on time step size δ_{N_t} .

2.5. Multigrid reduction in time (MGRIT) algorithm

Based on Equation 21 and Equation 26, the global space-time problem can be written in the linear form,

$$\mathbf{A}^S \mathbf{s}^S = \begin{bmatrix} \mathbf{I} & & & & & & \\ -\mathbf{\Phi}^S & \mathbf{I} & & & & & \\ & -\mathbf{\Phi}^S & \mathbf{I} & & & & \\ & & & \ddots & \ddots & & \\ & & & & -\mathbf{\Phi}^S & \mathbf{I} & \end{bmatrix} \begin{bmatrix} \mathbf{s}_0^S \\ \mathbf{s}_1^S \\ \mathbf{s}_2^S \\ \vdots \\ \mathbf{s}_{N_t}^S \end{bmatrix} = \begin{bmatrix} \hat{\mathbf{s}}_0 \\ \mathbf{0} \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} = \hat{\mathbf{s}}, \quad (27)$$

with $S \in \{I, II\}$. A traditional time stepping method would solve Equation 27 in a block-forward fashion, whereas the MGRIT algorithm solves Equation 27 iteratively. Both algorithms are $O(N_t)$, however, the constant is bigger for MGRIT [15, 12]. On the other hand, MGRIT enables parallelism in the temporal domain in contrast to the sequential nature of traditional time stepping. The parallelism is achieved by introducing a time grid hierarchy and applying multigrid techniques in the temporal domain. The coarser time grids provide error corrections to the finest time grid, thus accelerating convergence to the solution; while a relaxation process on each time grid reduces the error that cannot be adequately reduced on coarser grids. These two complementary processes, relaxation and coarse-grid correction, form the core of multigrid methods.

Figure 1 illustrates an example of an m_l -level time grid hierarchy with respective time points, time step sizes and coarsening factors. Here, the maximum number of time grid levels is given as $m_l = 2$ (with levels $m = 0, \dots, m_l - 1$). Each level with $m < m_l - 1$ is composed of F- and C-points. F-points only exist on level m and the C-points compose the coarser grid $m + 1$. The coarsening factor c_{m-1}^m (for $m = 1, \dots, m_l - 1$) defines a coarsening of a given level $m - 1$, such that the time step size on level m is $\delta_{N_t^m} = c_{m-1}^m \cdot \delta_{N_t^{m-1}}$ (for $m = 1, \dots, m_l - 1$), with fine grid time step size $\delta_{N_t^0}$ and the coarsest grid's time step size $\delta_{N_t^{m_l-1}}$. Similarly, $c_0^{m_l-1}$ denotes the total coarsening from level 0 to level m_l . Further, F-relaxation means an update of the F-points based on the previous C-points and C-relaxation refers to an update of a C-point based on the previous F-point. For more details, see [15].




Figure 1. Example of an m_l -level multigrid hierarchy with $m_l = 2$ time grid levels, respective time points t_i , time step sizes $\delta_{N_t^m}$ and coarsening factors c_{m-1}^m .

For the remainder of this work, we distinguish between the numerical space-time solution obtained with Scheme I with sequential time stepping with a subscript s (sequential) and those quantities obtained from using MGRIT with a subscript p (parallel), *e.g.*, the solution vector \mathbf{s}^I obtained from sequential time stepping is denoted as \mathbf{s}_s^I , whereas the solution vector obtained with MGRIT is denoted as \mathbf{s}_p^I . Similarly, we identify the numerical solution obtained with Scheme II by using the notation \mathbf{s}_s^{II} and \mathbf{s}_p^{II} .

Further, we consider two-grid algorithms using V-cycles with various numbers of FC-relaxation steps and multilevel algorithms using full multigrid (FMG) cycles (*i.e.*, F-cycles) as described in the following. Unless noted otherwise, we either use the one-step integrator $\Phi_{\delta_t^0}^I$ on the fine grid and its coarse grid versions $\Phi_{\delta_t^m}^I$ on grid level m or similarly, $\Phi_{\delta_t^0}^{II}$ on the fine grid and its coarse grid versions $\Phi_{\delta_t^m}^{II}$ on grid level m . The subscript δ_t^m denotes $\delta_{N_t^m}$. In Section 3.3, we will also motivate the design of a hybrid algorithm that employs $\Phi_{\delta_t^0}^I$ on the fine grid and $\Phi_{\delta_t^m}^{II}$ on the coarse grid in a two-level time grid hierarchy.

2.5.1. Two-level algorithm using V-cycles

Firstly, a two-grid version of the MGRIT algorithm is considered. In this version, r FC-relaxations ($r = 0, 1, 2$) are applied on the fine grid for Scheme I and Scheme II. For $r = 0$, MGRIT is equivalent to Parareal [13]. To restrict the variables from the fine to the coarse grid, a restriction operator $R(\cdot)$ is chosen that purely injects the respective fine grid quantities to the coarse grid. To transfer the coarse grid quantities to the fine grid, an ideal interpolation is employed [15]. The pseudo-code of the linear two-grid version of the MGRIT algorithm used here is described in Algorithm 1. [†] The coarse-grid operator A_{Δ}^S is equivalent to that defined in (27) except it uses the coarse-grid time-stepper $\Phi_{\delta_t^{m+1}}^S$ and is defined over the smaller coarse time-grid depicted in Figure 1.

The two-grid algorithm with F- and FCF- relaxation will be the basis for predicting and analyzing convergence of the proposed algorithm and for comparing the observed convergence factors with the theoretical upper bounds derived from [12]. The presented analysis will then prove and emphasize the relevance of [12] for the design of convergent parallel-in-time algorithms and motivate the selection of Scheme II to be included in a multilevel hierarchy. In this setting, we also investigate the benefits of switching from V-cycles to F-cycles.

2.5.2. Multilevel algorithm using FMG-cycles

In addition to the two-level algorithms, we consider m_l -grid algorithms ($m_l > 2$) with FMG-cycles. A FMG-cycle is achieved by applying Step 4 of Algorithm 1 recursively and performing one V-cycle as post-relaxation at each level [16].

[†]The nonlinear version of MGRIT is used in practice and is equivalent to the linear version for linear problems [14].

Algorithm 1 Pseudo-code for MGRIT algorithm with V-cycles and $S \in \{I, II\}$, adapted from [12].

- 1: **repeat**
 - 2: Relax on $A^S s^S = \hat{s}$ (r times FC-relaxation, one F-relaxation) using $\Phi_{\delta_t^m}^S$. ▷ In parallel
 - 3: Compute coarse-grid residual $r_\Delta = R(\hat{s} - A^S s^S)$.
 - 4: Solve the coarse grid correction problem $A_\Delta^S e_\Delta^S = r_\Delta$ using $\Phi_{\delta_t^{m+1}}^S$. ▷ Apply recursively
 - 5: Correct the solution at the fine-grid C-points with e_Δ^S . ▷ In parallel
 - 6: **until** norm of residual is small enough

 - 7: Update the solution at the F-points with $\Phi_{\delta_t^0}^S$. ▷ In parallel
-

2.5.3. Theoretical two-grid reduction of the MGRIT residual norm per iteration in the two-grid case

The theoretical two-grid reduction rate estimates of the global space-time error per MGRIT iteration are based on the two-grid cases discussed in [12]. Firstly, we note that $\Phi_{\delta_t^0}^S$ (and the coarse-grid version $\Phi_{\delta_t^1}^S$) have the sparsity pattern,

$$\begin{bmatrix} [\Phi_{\delta_t^m}^S]_{11} & [\Phi_{\delta_t^m}^S]_{12} & [\Phi_{\delta_t^m}^S]_{13} \\ [\Phi_{\delta_t^m}^S]_{21} & [\Phi_{\delta_t^m}^S]_{22} & [\Phi_{\delta_t^m}^S]_{23} \\ [\Phi_{\delta_t^m}^S]_{31} & [\Phi_{\delta_t^m}^S]_{32} & [\Phi_{\delta_t^m}^S]_{33} \end{bmatrix} = \begin{bmatrix} [\Phi_{\delta_t^m}^S]_{11} & \mathbf{0} & [\Phi_{\delta_t^m}^S]_{13} \\ [\Phi_{\delta_t^m}^S]_{21} & \mathbf{0} & [\Phi_{\delta_t^m}^S]_{23} \\ [\Phi_{\delta_t^m}^S]_{31} & \mathbf{0} & [\Phi_{\delta_t^m}^S]_{33} \end{bmatrix}, \quad (28)$$

with $S \in \{I, II\}$ and $m = 0, 1$. This highlights that the current state vector s_i^S is not dependent on the previous pressure value (*i.e.*, the pressure variable and its associated rows and columns in Φ can be ignored without affecting u or v .) Thus, we eliminate the rows and columns related to the pressure variable and proceed with the analysis by simultaneously diagonalizing the time stepping matrix,

$$T^{-1} \begin{bmatrix} [\Phi_{\delta_t^0}^S]_{11} & [\Phi_{\delta_t^0}^S]_{13} \\ [\Phi_{\delta_t^0}^S]_{31} & [\Phi_{\delta_t^0}^S]_{33} \end{bmatrix} T = \text{diag}(\lambda_1, \lambda_2, \dots), \quad (29)$$

$$T^{-1} \begin{bmatrix} [\Phi_{\delta_t^1}^S]_{11} & [\Phi_{\delta_t^1}^S]_{13} \\ [\Phi_{\delta_t^1}^S]_{31} & [\Phi_{\delta_t^1}^S]_{33} \end{bmatrix} T = \text{diag}(\gamma_1, \gamma_2, \dots). \quad (30)$$

Then, the authors of [12] prove that the global space-time error vector at the C-points can be reduced in the mass matrix norm for the two-grid case with F- and FCF-relaxation by a convergence factor of at least,

$$c_f^F = \max_\omega \left\{ |\lambda_\omega^{c_0^1} - \gamma_\omega| \frac{1 - |\gamma_\omega|^{N_t^1}}{1 - |\gamma_\omega|} \right\}, \quad (31)$$

$$c_f^{FCF} = \max_\omega \left\{ |\lambda_\omega^{c_0^1} - \gamma_\omega| \frac{1 - |\gamma_\omega|^{N_t^1 - 1}}{1 - |\gamma_\omega|} |\lambda_\omega|^{c_0^1} \right\}, \quad (32)$$

$$(33)$$

respectively.

2.5.4. Computation of MGRIT residual norm in experiment

The standard Euclidean norm of the MGRIT residual in numerical experiments is computed using the coarse-grid residual r_Δ (see Step 3 of Algorithm 1),

$$\begin{aligned} \|r_\Delta\|_2^2 &= (w_0 \|r_\Delta^v\|_2^2 + w_1 \|r_\Delta^p\|_2^2 + w_2 \|r_\Delta^u\|_2^2)^{1/2} \\ &= (w_0 (r_\Delta^v)^T r_\Delta^v + w_1 (r_\Delta^p)^T r_\Delta^p + w_2 (r_\Delta^u)^T r_\Delta^u)^{1/2} \end{aligned} \quad (34)$$

where \mathbf{r}_Δ^v , \mathbf{r}_Δ^p and \mathbf{r}_Δ^u denote components in the residual vector corresponding to velocity, pressure and displacement. Further, the weights w_0 , w_1 and w_3 are equal to 1 in the standard case.

However, the predicted decrease of the residual is measured in the mass matrix norm

$$\|\mathbf{r}_\Delta\|_{\mathfrak{M}}^2 = \|\mathbf{r}_\Delta^v\|_{\mathfrak{M}^v}^2 + \|\mathbf{r}_\Delta^p\|_{\mathfrak{M}^p}^2 + \|\mathbf{r}_\Delta^u\|_{\mathfrak{M}^u}^2, \quad (35)$$

where \mathfrak{M}^x is a block $N_t^1 \times N_t^1$ diagonal matrix with each block corresponding to a spatial mass matrix for variable x (see Section 2.5.3 and [12]). Thus, a modification of measuring the observed reduction of the residual norm per MGRIT iteration is advisable, and hence, proposed: instead of computing the computationally more demanding mass matrix norm and the less accurate standard Euclidean norm, the following weights to improve the measurement of experimental convergence factors are chosen: $w_0 = 1$, $w_1 = 0$ and $w_2 = 1/c_0^1$. This yields

$$(\mathbf{r}_\Delta^v)^T \mathfrak{M}^v \mathbf{r}_\Delta^v \approx (\mathbf{r}_\Delta^v)^T \mathbf{I} \mathbf{r}_\Delta^v, \quad (36)$$

$$(\mathbf{r}_\Delta^u)^T \mathfrak{M}^u \mathbf{r}_\Delta^u \approx (\mathbf{r}_\Delta^u)^T \left(\frac{1}{c_0^1} \mathbf{I} \right) \mathbf{r}_\Delta^u, \quad (37)$$

Thus, the approximate value of the mass matrix norm of the residual at the C-points is given by,

$$\|\mathbf{r}_\Delta\|_{\mathfrak{M}} \approx \left((\mathbf{r}_\Delta^v)^T \mathbf{r}_\Delta^v + \frac{1}{c_0^1} (\mathbf{r}_\Delta^u)^T \mathbf{r}_\Delta^u \right)^{1/2}. \quad (38)$$

It is important to note that this proposition only changes how the solution progress (*i.e.*, reduction of residual norm) is measured. But it does neither affect the coarse-grid update nor does it change the numerical solution.

2.6. Implementation details

Scheme I and Scheme II were implemented in the finite element software tool CHeart [17]. Wrapper routines were written to incorporate the MGRIT algorithm into CHeart using the open-source library XBraid [18], a non-intrusive implementation of the MGRIT algorithm.

Here, we introduced separate MPI groups and communicators to maintain the capability of CHeart to parallelize in the spatial domain by using domain decomposition methods while enabling independent parallelization in the temporal domain. That is, one can parallelize in space, time or in space-time.

Furthermore, a given relaxation step is performed even if the initial spatial residual satisfies the convergence criteria for the spatial problem. Moreover, the XBraid option to skip work on the first down-cycle is used. Note, due to the linearity of the problem, the matrix $\Phi_{\delta_t^m}^S$ is only computed once for each time step size (*i.e.*, time grid level). This significantly reduces computational work compared to re-computing the matrix for each time step.

3. RESULTS

In the following, the methods described in Section 2 are used for a linear beam problem. For this purpose, consider the domain $\Omega(0) = [0, 8] \times [0, 1]$, final time $T = 1024$ and material parameters $\mu = \rho = 1$. The spatial discretization uses 16×2 quadrilateral elements and initial conditions are given as $\hat{\mathbf{v}}_0 = [-x^2/640, x^2 \cdot (8-x)/640]^T$, see Figure 2. The initial condition is also used as initial guess at all time steps for MGRIT. Further, no displacement boundary conditions are prescribed at $\Gamma^D = \Omega|_{x=0}$.

We consider fine grid time step sizes $\delta_{N_t^0} \in \{1, 0.1, 0.01, 0.001\}$ and coarsening factors $c_0^1 \in \{2, 4, 8, 16, 32\}$ and $c_{m-1}^m = 2$ for $2 \leq m \leq m_l$. The stopping criterion on the residual norm is selected as $\|\mathbf{r}_\Delta\|_{\mathfrak{M}} \leq 10^{-8} / \sqrt{\delta_{N_t^0}}$ (see Equation 38) with the maximum number of MGRIT iterations set to 60 iterations. Unless noted otherwise, reported experimental convergence factors are the global maximum values.




Figure 2. Discretization using 16 x 2 quadrilateral elements and initial velocity.

3.1. Numerical solution of beam oscillation

The initial velocity distribution over the cantilever beam length (see Figure 2) causes the free end to first deflect in negative x - and positive y -direction, see Figure 3. The elastic stresses cause the beam to decelerate and move downward, passing its initial position and deflecting in negative y -direction. The beam deformation follows a repeatable deflection pattern.

3.2. Effect of time step size on amplitude of oscillation

Backward Euler time integration schemes introduce artificial numerical diffusion. Its effect on the sum of kinematic and potential energy of the system over time depends on the time step size where we expect energy conservation in the asymptotic limit $\delta_{N_t^0} \rightarrow 0$. The artificial damping of the system causes the amplitude of oscillation to become smaller over time. The backward Euler time integration scheme exhibits quick damping for $\delta_{N_t^0} = 1$ irrespectively of the chosen scheme (Scheme I and Scheme II). The effect of numerical damping becomes smaller for smaller $\delta_{N_t^0}$, where we note that both schemes reproduce the amplitudes of the beam oscillation with comparable quality for $\delta_{N_t^0} = 0.001$, see Figure 4.

It is important to note, that the varying magnitude of numerical damping is also present with respect to varying time step sizes in an MGRIT time grid hierarchy. Nevertheless, MGRIT converges to the same solution (within the selected solver tolerance) as sequential time stepping on the fine grid. Thus, the converged numerical solution obtained with the MGRIT algorithm suffers from the same amount of numerical damping as the numerical solution from sequential time stepping on the fine grid.

3.3. Convergence in the two-grid case

3.3.1. MGRIT with Scheme I

If Scheme I is employed as a one-step integrator on the fine and coarse grid (that is, $\Phi_{\delta_t^0}^I$ and $\Phi_{\delta_t^1}^I$) in a two-level algorithm, we observe divergence in the numerical experiments for all considered test cases.

To qualitatively investigate what leads to divergence of MGRIT in conjunction with Scheme I, we track the current position of the tip of the cantilever beam (initial coordinate $[8, 0.5]^T$ at $t = 0$) over time with smaller time horizon $t \in [0, 64]$. The time step size is $\delta_{N_t^0} = 1$ with coarsening factor $c_0^1 = 2$. FCF-relaxation is employed. Figure 5 shows the current approximation of the tip's displacement with respect to the initial position over time for a number of algorithmic steps, for example, after FCF-relaxation, after restriction, after the coarse-grid solve, *etc.* (where we extract the values on return from applying $\Phi_{\delta_t^0}^I$ and $\Phi_{\delta_t^1}^I$, respectively). The data in Figure 5 highlight how the current approximation of the cantilever's tip first improves. Though, already during the first MGRIT iteration an instability is introduced by the coarse-grid update which is then amplified in subsequent steps. The observed phenomena are neither physical nor part of the mathematical model, however, are in line with observations in literature, see [9, 10]. Previously this has been a limiting factor for parallel-in-time integration methods of the dynamic elasticity equation.

On the other hand, divergence of the numerical algorithm is reflected in very large theoretical convergence bounds, *i.e.*, $c_f^F, c_f^{FCF} \gg 1$, confirming experimental observations.

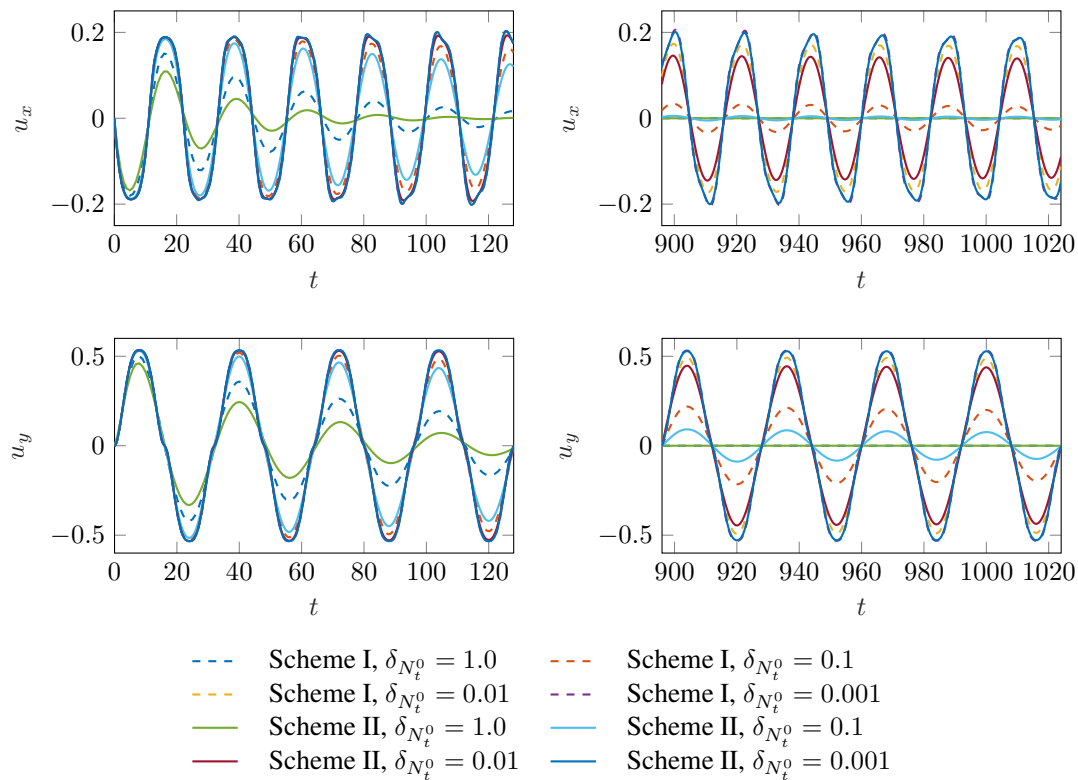


Figure 3. Displacement \mathbf{u} of the tip of the cantilever beam (initial position $[8, 0.5]^T$) in x - and y -direction for $t \in [0, 128]$ and $[896, 1024]$ for Scheme I and Scheme II and time step sizes $\delta_{N_t^0} \in \{1, 0.1, 0.01, 0.001\}$. Note that numerical damping is reduced with $\delta_{N_t^0} \rightarrow 0$.

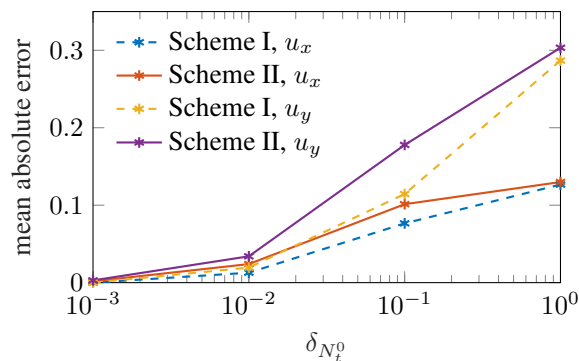


Figure 4. Mean absolute error of the displacement of the beam's tip for $t \in [0, 1024]$. Here, Scheme I with $\delta_{N_t^0} = 0.001$ was used as reference.

3.3.2. MGRIT with Scheme II

Using Scheme II as fine- and coarse-grid integrators (that is, $\Phi_{\delta_t^0}^{II}$ and $\Phi_{\delta_t^1}^{II}$) in a two-level algorithm, we observe good convergence for a range of different coarsening factors c_0^I and for all considered time step sizes, see Figure 6.

We note that for $\delta_{N_t^0} = 1$, both experimental and predicted convergence factors are in excellent agreement and that the predicted values of c_f^F and c_f^{FCF} are a sharp upper bound despite the approximation of the computed residual norm, see Section 2.5.4. Here, the convergence factors

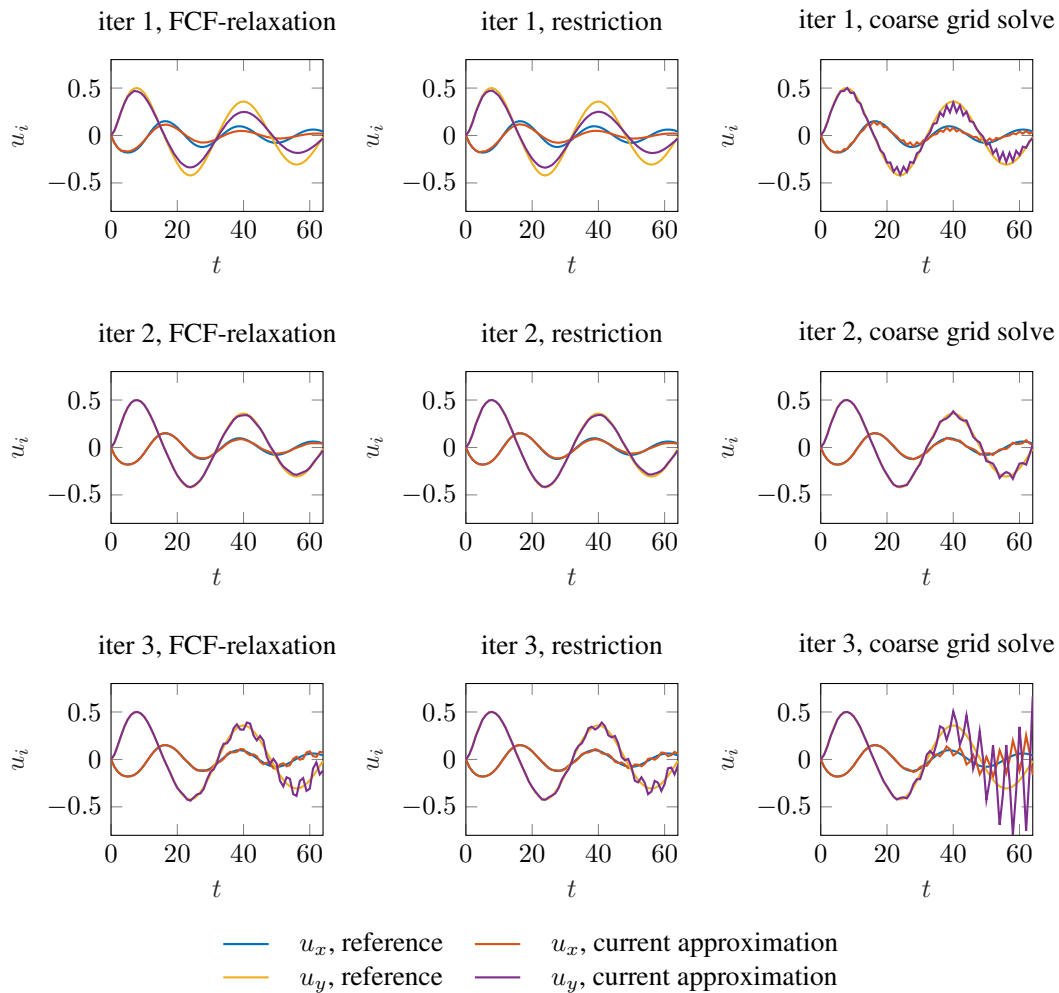


Figure 5. Current approximation of the tip displacement for two-grid MGRIT using Scheme I with coarsening factor $c_0^1 = 2$, fine grid time step size $\delta_{N_t^0} = 1$, $t \in [0, 64]$ and FCF-relaxation compared with a reference solution from sequential time stepping. Note, how the coarse grid update introduces an instability which is amplified in subsequent steps.

first increase with increasing coarsening factor c_0^1 before decreasing due to the small coarse-grid size $N_t^{m_l-1}$.

On the other hand, for $\delta_{N_t^0} \in \{0.1, 0.01, 0.001\}$ we observe an increase in predicted and observed convergence factors with increasing coarsening factor c_0^1 . Again, predicted and observed convergence factors are in excellent agreement for almost all considered cases. Only for $\delta_{N_t^0} = 0.01$ the maximum observed convergence factor is smaller for the F-relaxation case and approximately 1% larger for (i) F-relaxation and $c_0^1 = 32$ and (ii) FCF-relaxation and $c_0^1 \in \{16, 32\}$. Again, this is likely due to the approximate residual norm computation.

Further, we note that additional relaxation steps can be beneficial for small fine grid time step sizes, whereas the effect is negligible for $\delta_{N_t^0} = 0.001$, thus suggesting that relaxation can be omitted for small fine grid time step sizes to reduce computational work without sacrificing convergence.

Finally, we note that in all considered cases we do not observe any instability as described in Section 3.3.1 and previous works [9, 10]. For example, Figure 7 illustrates the position of the tip of the cantilever after the first three MGRIT iteration with $t \in [0, 64]$, $\delta_{N_t^0} = 0.1$ and $c_0^1 = 2$. Here, no artificial amplification of the amplitude of oscillation is observed as is the case for MGRIT with Scheme I, see Figure 5.

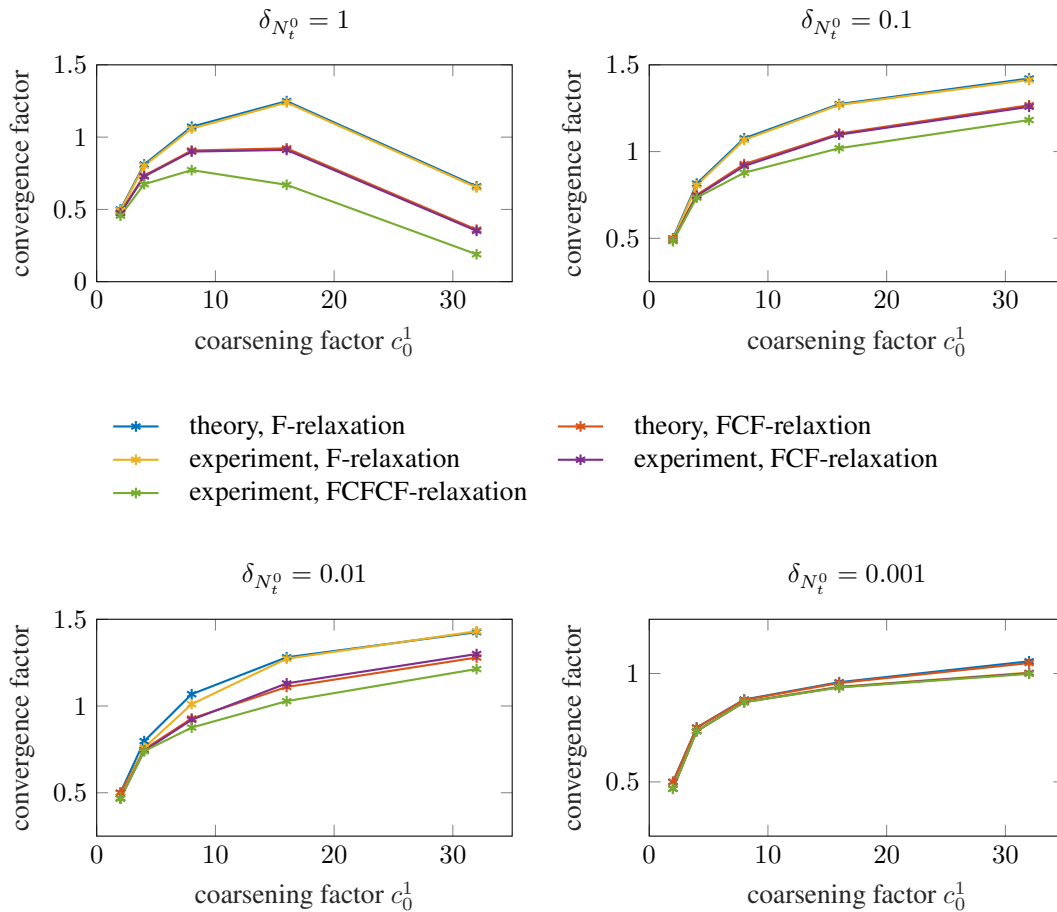


Figure 6. Predicted and measured convergence factor for a two-grid algorithm with various fine grid time step sizes $\delta_{N_t^0} \in \{1.0, 0.1, 0.01, 0.001\}$ and coarsening factors $c_0^1 \in \{2, 4, 8, 16, 32\}$.

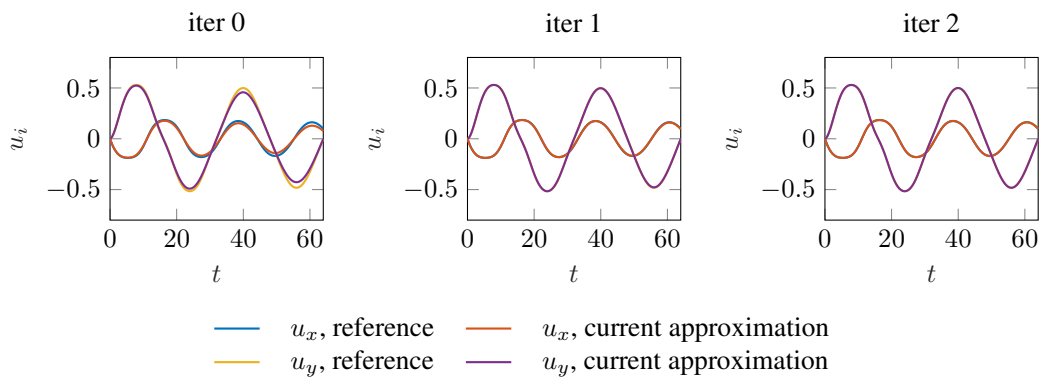


Figure 7. Current approximation of the tip displacement for two-grid MGRIT using Scheme II with coarsening factor $c_0^1 = 2$, fine grid time step size $\delta_{N_t^0} = 0.1$, $t \in [0, 64]$ and FCF-relaxation compared with reference solution from sequential time stepping. No instability is observed with Scheme II in contrast to using Scheme I, see Figure 5.

The results in this section highlight the benefit of using theoretical upper bounds as given in [12] as tool to estimate experimental convergence *a priori* and to design convergent MGRIT algorithms with guaranteed worst case convergence factors.

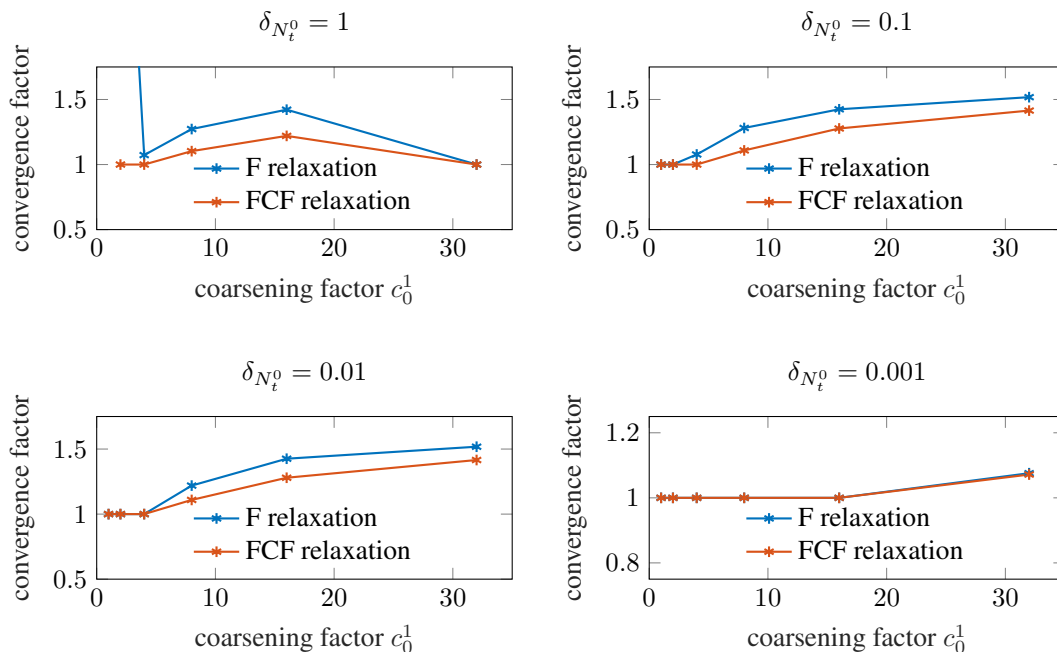


Figure 8. Predicted worst-case convergence for MGRIT with the hybrid approach. The hybrid approach avoids the instability observed for MGRIT with Scheme I, however, cannot yield guaranteed convergence in the two-grid case.

3.3.3. MGRIT with hybrid scheme

Section 3.2 highlights that for large time step sizes, Scheme I can better reproduce the amplitude of oscillation compared to Scheme II. Though, Scheme I with MGRIT does not yield a convergent parallel-in-time algorithm. On the other hand, Scheme II with MGRIT does not exhibit the same instability as Scheme I with MGRIT, see Section 3.3.2. Additionally, MGRIT yields the same solution as sequential time stepping on the fine grid within the selected solver tolerance, as was noted earlier. These aspects motivate the design of a hybrid algorithm with the aim to preserve amplitudes as good as Scheme I but with the MGRIT convergence properties of Scheme II. Thus, the hybrid scheme employs $\Phi_{\delta_t^0}^I$ and $\Phi_{\delta_t^1}^{II}$ on the fine and coarse grid, respectively.

Although MGRIT with the hybrid scheme does not suffer from the same instability as MGRIT with Scheme I, convergence deteriorates. The predicted upper bound for the convergence factor either indicates divergence with $c_f^F, c_f^{FCF} > 1$, in the worst case, or a stalling solver with $c_f^F, c_f^{FCF} = 1$. This is illustrated in Figure 8, where one can observe that convergence cannot be guaranteed for all fine grid time step sizes. Thus, MGRIT with the hybrid scheme has no significance for designing robust parallel-in-time solvers in the two-grid case.

3.4. Convergence in the multigrid case

Since a convergent algorithm could only be achieved for MGRIT with Scheme II in the two-grid case, only this case is considered in the following to investigate convergence in the multigrid case.

3.4.1. MGRIT with Scheme II

To explore convergence in the multigrid case, we employ multilevel hierarchies with $m_l \in \{3, 4, 5\}$ for $\delta_{N_t^0} \in \{1, 0.1, 0.01\}$ and $m_l \in \{3, 4, 5, 6\}$ for $\delta_{N_t^0} = 0.001$. Further, we consider combined coarsening factors of $c_0^{m_l-1} \in \{4, 8, 16, 32, 64\}$, where we only vary c_0^1 but select $c_{m-1}^m = 2$ for $m = 2, \dots, m_l - 1$. The MGRIT algorithm is started with a forward solve on the coarsest time grid (*i.e.* use of the XBraid skip-first-down option), performs an initial V-cycle and full multigrid (FMG)

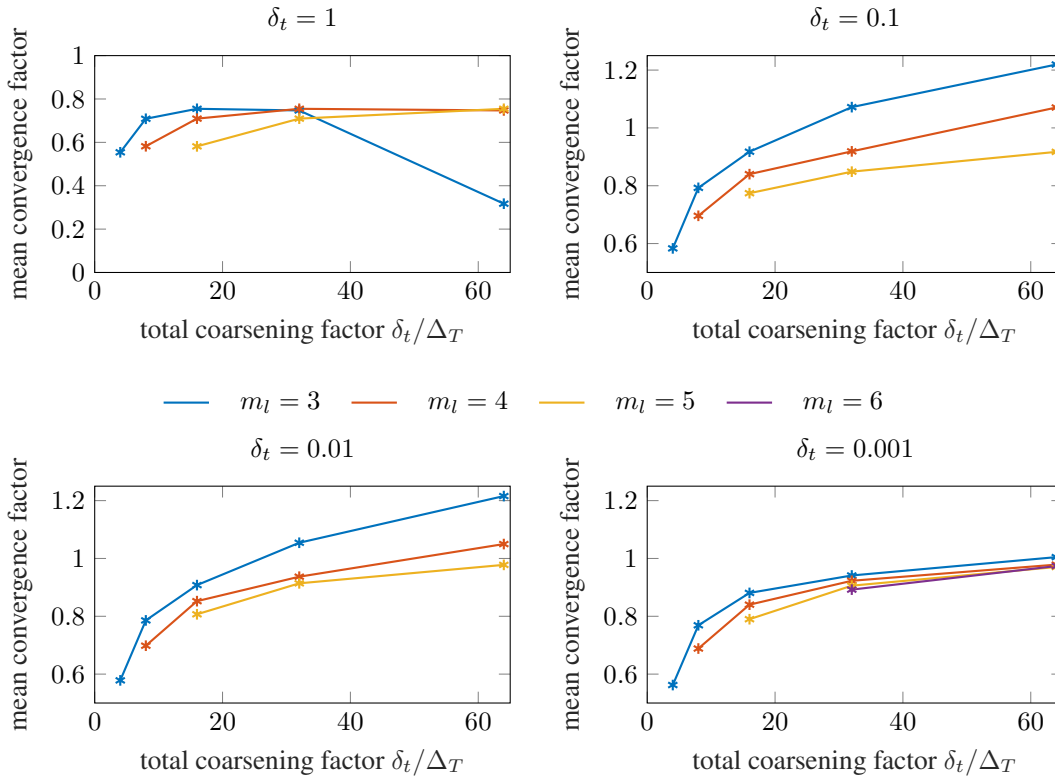


Figure 9. Measured mean convergence factor for m_l -grid MGRIT algorithm with F-cycles and Scheme II for fine grid time step sizes $\delta_{N_t^0} \in \{1, 0.1, 0.01, 0.001\}$. Note, that the x -axis corresponds to the combined coarsening factor $c_0^{m_l-1}$.

cycles for all following MGRIT iterations. Scheme II is employed as one-step integrator on all grid levels and one V-cycle is performed as post-relaxation at each FMG level. To provide a better indicator for overall performance of the MGRIT algorithm, we report the mean of the experimental convergence factor over all MGRIT iterations.

Figure 9 reports mean experimental convergence factors over the combined coarsening factor $c_0^{m_l-1} = \prod_{m=0}^{m_l-2} c_m^{m+1}$. As illustrated, the best convergence factor available for a particular combined coarsening factor is available through the use of FMG cycles with more levels and slower coarsening, between levels, as opposed to FMG cycles with more aggressive coarsening between levels. Observed convergence factors are generally significantly smaller than in the two-grid case. Thus, the use of FMG-cycles allows more aggressive coarsening and yields more potential for parallelism in the temporal domain.

If $\delta_{N_t^0}$ and $c_0^{m_l-1}$ are kept fixed, the mean convergence factor improves with increasing m_l . For example, for $\delta_{N_t^0} = 0.1$ and $c_0^{m_l-1} = 64$ the mean convergence factor is approximately 0.92 for $m_l = 5$ but 1.07 for $m_l = 4$ and 1.22 for $m_l = 3$. Thus, one can obtain a moderately convergent instead of a slowly divergent algorithm simply by introducing an additional intermediate time grid level but with the same fine and coarsest grid size.

3.4.2. Timing results

In this section, we present preliminary speedup results for an MGRIT algorithm that employs Scheme II and $m_l = 5$ time grid levels with a coarsening of two between all time grid levels. Similar to Section 3.4.1, an initial V-cycle is performed before switching to FMG-cycles with one V-cycle as post-relaxation at each FMG-level. The time step size is selected as $\delta_{N_t^0} = 0.001$. The standard Euclidean norm is used to measure solution progress. To investigate the dependency of the wall

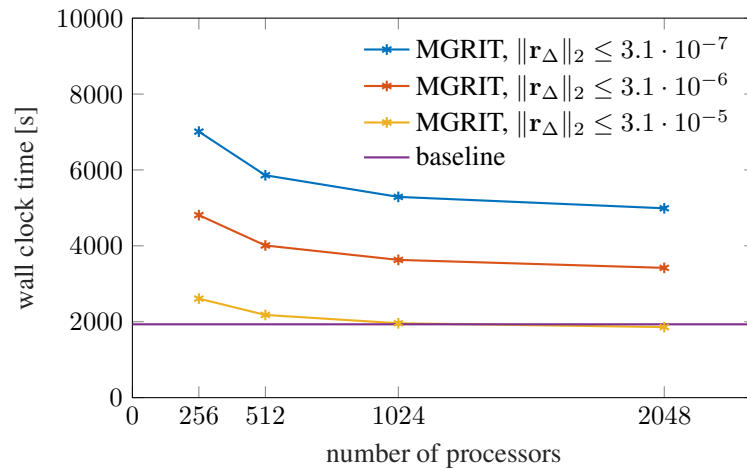


Figure 10. Wall clock time for MGRIT with Scheme II and FMG-cycles with $m_l = 5$ and $c_{m-1}^m = 2$ for all m . Note, that even for small processor counts a marginal speedup can be achieved with relaxed convergence criteria besides the relatively large coarse grid size with $c_0^{m_l-1} = 16$.

clock time for MGRIT on the convergence criterion, we vary the tolerance between $3.1 \cdot 10^{-7}$ and $3.1 \cdot 10^{-5}$.

All reported timing results were obtained on CAB (Intel Xeon E5-2670 architecture, 16 cores 2.6GHz and 32GB RAM per node) at Lawrence Livermore National Laboratory, CA, USA. Data export was switched off and status messages were restricted to a bare minimum to ensure that reported wall clock times are almost entirely dominated by computation and communication costs but not data I/O tasks.

The elapsed wall clock time for the sequential time stepping algorithm (using Scheme II) was 1932 seconds, which is established as baseline in Figure 10. As illustrated in Figure 10, the elapsed wall clock time for the five-grid MGRIT algorithm drops from approximately 7010 seconds to 4990 seconds when using 2048 processors instead of 256 processors for parallelizing in the temporal domain. By relaxing the convergence criterion for the time-parallel algorithm, the wall clock time of the time-parallel algorithm drops below the wall clock time for the sequential algorithm with a marginal speedup. Here, it is important to note that although the achieved speedup is almost negligible, one can expect to improve the parallel speedup, for example, by selecting a more efficient cycling and relaxation strategy (*e.g.*, cheaper MGRIT iterations and improved reduction of residual norm per MGRIT iteration) and a larger coarsening factor $c_0^{m_l-1}$ in combination with growing processor counts.

4. DISCUSSION

In the previous sections, two different backward Euler time discretization schemes were presented and investigated. It was shown, that MGRIT with Scheme I exhibits strong instabilities parallel-in-time. The observations are in line with previous works in this field [9, 10]. Although MGRIT with Scheme II uses only a slightly modified time discretization, a convergent scheme can be derived for a range of coarsening factors, both for two-grid V-cycles and multigrid FMG-cycles.

Intuitively, the modification in Scheme II changes how the discretized version of the stress tensor in Equation 6 is evaluated (compare Equation 12 and Equation 15), such that it is evaluated at given discrete time points across all levels in the time grid hierarchy in a consistent manner, see Section 2.5. Here, we note that both time discretizations are consistent and converge to the same

numerical solution for decreasing time step sizes. By rewriting the time-discrete Equation 14,

$$\mathbf{u}_i = \mathbf{u}_{i-1} + \delta_{N_t} \mathbf{v}_i \quad (39)$$

$$= \mathbf{u}_{i-1} + \frac{\delta_{N_t}}{2} (\mathbf{v}_i + \mathbf{v}_{i-1}) + \frac{\delta_{N_t}^2}{2} \left(\frac{\mathbf{v}_i - \mathbf{v}_{i-1}}{\delta_{N_t}} \right), \quad (40)$$

one may argue that, for δ_{N_t} , Equation 40 becomes increasingly similar to the time-discrete Equation 11. Further, rewriting Equation 11,

$$\mathbf{u}_i = \mathbf{u}_{i-1} + \frac{\delta_{N_t}}{2} (\mathbf{v}_i + \mathbf{v}_{i-1}) = \mathbf{u}_{i-1} + \delta_{N_t} \mathbf{v}_{i-1} - \frac{\delta_{N_t}}{2} \mathbf{v}_{i-1} + \frac{\delta_{N_t}}{2} \mathbf{v}_i \quad (41)$$

$$= \mathbf{u}_{i-1} + \delta_{N_t} \mathbf{v}_{i-1} + \frac{\delta_{N_t}^2}{2} \left(\frac{\mathbf{v}_i - \mathbf{v}_{i-1}}{\delta_{N_t}} \right), \quad (42)$$

and assuming that the last term in Equation 42 can be neglected for small δ_{N_t} , one would effectively approximately discretize Equation 5 by an explicit step. This might be a hint for explaining the observed instability of MGRIT with Scheme I. On the other hand, the two-grid theory presented in [12], when used as an *a priori* tool for estimating convergence of MGRIT with Scheme I, provides a more mathematical and clear pathway for relating observed divergence in numerical experiments with the employed time-discretization.

In the case of using MGRIT with Scheme II, convergence was predicted for two-grid algorithms with F- and FCF-relaxation and a range of coarsening factors. The predictions matched quite closely with observations in the numerical experiments, supporting the use of the analysis presented in [12] as a powerful tool to design convergent algorithms *a priori*.

Two-grid results also motivated the design of a hybrid algorithm with the use of Scheme I on the fine grid and Scheme II on the coarse grid. Although, we were not able to derive a convergent hybrid scheme in this work, we showed that the use of Scheme II on the coarse grid removes the instability observed for MGRIT with Scheme I on both time grid levels. Further, Figure 8 illustrates that the predicted worst-case convergence factors are close to one, whereas for MGRIT with Scheme I on all levels yielded immediate and strong divergence with estimates $c_f^F, c_f^{FCF} \gg 1$. Thus, hybrid schemes might be a promising approach to combine the advantages of different time discretization schemes.

Although the analysis presented in [12] is restricted to the two-grid case, the multilevel algorithm using MGRIT with Scheme II (presented in Section 3.4.1) was a direct result from the design of the convergent two-grid algorithm using MGRIT with Scheme II. Here, the use of FMG-cycles was beneficial to accelerate convergence in the true multilevel case, enabling larger combined coarsening factors $c_0^{m_i-1}$ compared to the two-grid case.

Further, preliminary timing results for the time-parallel algorithm were presented in Section 3.4.2. While only a marginal speedup could be achieved in the presented case for MGRIT with Scheme II and a five-grid hierarchy, results are promising. That is, reported wall clock times are almost entirely dominated by computation and communication. However, it is important to note that data export, for example, is a completely serial process for sequential time-stepping while it is parallel for the MGRIT algorithm. Thus, for practical applications that include data I/O tasks, larger speedups can be expected. Furthermore, the reported timing results were also restricted to a relatively small combined coarsening factor $c_0^4 = 16$. Thus, an MGRIT algorithm with improved convergence rates and larger coarsening factors (*i.e.*, smaller coarse-grid size) will exhibit much better speedup over sequential time-stepping.

5. CONCLUSION

In this work, we provided an analysis of the convergence of the multigrid reduction in time algorithm for the dynamic linear elasticity equations. It was shown that the two-grid convergence theory presented in [12] provides a mathematical explanation for the instability observed in previous

works [9, 10]. Using the two-grid convergence theory as a design tool to estimate convergence of a modified backward Euler scheme (MGRIT with Scheme II) *a priori*, we were able to obtain a convergent parallel-in-time algorithm for a range of coarsening factors, advancing the application of parallel-in-time methods for second-order hyperbolic equations. In this study, the predicted theoretical convergence bounds were in excellent agreement with the worst-case convergence factors in numerical experiments. We then extended the algorithm to the multilevel case, noting that FMG-cycles and slow temporal coarsening can improve convergence compared to the two-grid algorithm.

In future work, we will investigate other cycling strategies (*e.g.* switching from FMG- to V-cycles after a number of initial iterates) to improve and accelerate convergence. Moreover, spatial coarsening will be investigated. Further, we aim to generalize the framework to the efficiently solve nonlinear hyperelastic stress-strain relationships with MGRIT to tackle fluid-structure interaction applications, particularly in the field of biomedical engineering.

ACKNOWLEDGEMENT

The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007–2013)/ERC Grant Agreement No. 306757 (LEAD).

REFERENCES

1. McCormick M, Nordsletten DA, Kay D, Smith NP. Simulating left ventricular fluid–solid mechanics through the cardiac cycle under LVAD support. *J Comput Phys* 2013; 244:80–96.
2. Hessenthaler A, Röhrle O, Nordsletten D. Validation of a non-conforming monolithic fluid-structure interaction method using phase-contrast MRI. *Int J Numer Meth Bio* 2016.
3. Bai Y, Sun D, Lin J, Kennedy D, Williams F. Numerical aerodynamic simulations of a NACA airfoil using CFD with block-iterative coupling and turbulence modelling. *Int J Comput Fluid D* 2012; 26(2):119–132.
4. Nievergelt J. Parallel methods for integrating ordinary differential equations. *Commun ACM* 1964/ 7(12):731–733.
5. Gander M. 50 years of time parallel time integration (In: *Multiple Shooting and Time Domain Decomposition Methods*). Springer 2015.
6. Garrido I, Espedal MS, Fladmark GE. A convergent algorithm for time parallelization applied to reservoir simulation (In: *Domain Decomposition Methods in Science and Engineering*). Springer 2005.
7. Samaddar D, Casper TA, Kim SH, Berry LA, Elwasif WR, Batchelor DB, Houlberg WA. Time parallelization of advanced operation scenario simulations of ITER plasma. *J Phys Conf Ser* 2013; 410(1):012032.
8. Schreiber M, Peixoto PS, Haut T, Wingate B. Beyond spatial scalability limitations with a massively parallel method for linear oscillatory problems. *Int J High Perform C* 2017.
9. Farhat C, Chandesris M. Time-decomposed parallel time-integrators: theory and feasibility studies for fluid, structure, and fluid-structure applications. *Int J Numer Meth Eng* 2003; 58(9):1397–1434.
10. Farhat C, Cortial CJ, Dastillung C, Bavestrello H. Time-parallel implicit integrators for the near-real-time prediction of linear structural dynamic responses. *Int J Numer Meth Eng* 2006; 67(5):697–724.
11. Cortial J, Farhat C. A time-parallel implicit method for accelerating the solution of non-linear structural dynamics problems. *Int J Numer Meth Eng* 2009; 77(4):451–470.
12. Dobrev VA, Kolev TZ, Petersson NA, Schroder JB. Two-level convergence theory for multigrid reduction in time (MGRIT). *SIAM J Sci Comput* 2016 (submitted).
13. Gander MJ, Vandewalle S. Analysis of the parareal time-parallel time-integration method. *SIAM J Sci Comput* 2007; 29(2):556–578.
14. Falgout RD, Manteuffel TA, O'Neill B, Schroder JB. Multigrid Reduction in Time for Nonlinear Parabolic Problems: A Case Study. *SIAM J Sci Comput* 2016 (to appear).
15. Falgout RD, Friedhoff S, Kolev TZ, MacLachlan SP, Schroder JB. Parallel time integration with multigrid. *SIAM J Sci Comput* 2014; 36(6):C635–C661.
16. Trottenberg U, Oosterlee C, Schüller A. *Multigrid*. Academic Press, London, UK, 2001.
17. Lee J, Cookson A, Roy I, Kerfoot E, Asner L, Viguera G, Sochi T, Deparis S, Michler C, Smith NP, Nordsletten DA. Multiphysics Computational Modeling in CHeart. *SIAM J Sci Comput* 2016; 38(3):C150–C178.
18. XBraid: Parallel multigrid in time. <http://llnl.gov/casc/xbraid>.