Parallel-in-time simulation of an electrical machine using MGRIT

Matthias Bolten¹ · Stephanie Friedhoff¹ · Jens Hahne¹ · Sebastian Schöps²

Received: 6 December 2019 / Accepted: 15 July 2020 / Published online: 8 October 2020
© The Author(s) 2020

Abstract

We apply the multigrid-reduction-in-time (MGRIT) algorithm to an eddy current simulation of a two-dimensional induction machine supplied by a pulse-width-modulation signal. To resolve the fast-switching excitations, small time steps are needed, such that parallelization in time becomes highly relevant for reducing the simulation time. The MGRIT algorithm is an iterative method that allows calculating multiple time steps simultaneously by using a time-grid hierarchy. It is particularly well suited for introducing time parallelism in the simulation of electrical machines using existing application codes, as MGRIT is a non-intrusive approach that essentially uses the same time integrator as a traditional time-stepping algorithm. However, the key difficulty when using time-stepping routines of existing application codes for the MGRIT algorithm is that the cost of the time integrator on coarse time grids must be less expensive than on the fine grid to allow for speedup over sequential time stepping on the fine grid. To overcome this difficulty, we consider reducing the costs of the coarse-level problems by adding spatial coarsening. We investigate effects of spatial coarsening on MGRIT convergence when applied to two numerical models of an induction machine, one with linear material laws and a full nonlinear model. Parallel results demonstrate significant speedup in the simulation time compared to sequential time stepping, even for moderate numbers of processors.

Keywords Parallel-in-time · Multigrid-reduction-in-time (MGRIT) · Spatial coarsening · Electrical machine

1 Introduction

Induction motors are electrical machines in which the magnetic field in the rotor is obtained by an asynchronous motion with respect to the field in the stator, e.g. [1]. In the design process of such machines, motion and electromagnetic fields are numerically determined by space and time discretization of Maxwell's equations. The spatial discretization, typically by finite elements, may lead to large

models, e.g., due to resolving the so-called skin effect [2], but this is often mitigated by considering only two-dimensional models. More computational burden comes from the fact that large time intervals must be considered for simulating the start-up phase, i.e., until the machine reaches its steady state. Moreover, machines are commonly excited by pulse-width modulated (PWM) excitations. PWM uses fast-switching pulses whose width is controlled such that the time-average corresponds to a specific waveform, typically a sine wave such that very small time steps are needed. As a consequence, the computational complexity of transient simulations is extremely high.

Carrying out the simulations on modern parallel compute systems using existing application codes may still result in runtimes that are beyond what is feasible in practice, since parallelism in existing codes is limited to the spatial aspect of the problem. Due to the stagnation of processor clock speeds in recent years, modern compute systems are characterized by growing numbers of processors. Existing application codes cannot take full advantage of these systems, since spatial parallelism limits the amount of processors that can be exploited. A promising approach for reducing simulation times are parallel-in-time methods that introduce parallelism

Communicated by Daniel Ruprecht.

✉ Jens Hahne
jens.hahne@math.uni-wuppertal.de

Matthias Bolten
bolten@math.uni-wuppertal.de

Stephanie Friedhoff
friedhoff@math.uni-wuppertal.de

Sebastian Schöps
schoeps@temf.tu-darmstadt.de

¹ Fakultät für Mathematik und Naturwissenschaften, Bergische Universität Wuppertal, 42097 Wuppertal, Germany

² Centre for Computational Engineering, Technische Universität Darmstadt, 64293 Darmstadt, Germany

in the temporal dimension. Especially for simulations over long time intervals with small time steps as considered for electrical machines, parallelization in time can become even more relevant than in space.

To this end, various approaches have been proposed for speeding up the transient simulation of electrical machines, e.g., the time-periodic explicit error correction method [3,4] or the extraction of a circuit model [5]. Also, two of the most well-known parallel-in-time methods, Parareal [6] and MGRIT [7], have been applied to eddy current problems. In particular, Parareal has been considered for RL-circuit models and an induction machine [8–10], and MGRIT has been applied to a coaxial cable model [11,12]. This paper focuses on MGRIT applied to two numerical models of an induction machine, providing both an extension of ideas of the work on Parareal for problems involving discontinuous or multirate excitations to a multilevel setting as well as broadening the application areas of MGRIT to include more complex eddy current problems.

The idea of MGRIT is to use a time-grid hierarchy, in which each grid level is characterized by a time integrator, with computational costs decreasing from fine to coarse levels. Expensive time-integration routines are applied in parallel on temporal subdomains, and on the coarsest level, time integration is performed over the whole temporal domain. The key difficulty for achieving a significant reduction in simulation time over time-stepping algorithms is the choice of the time integrators such that the cost of the time integrators on coarse time grids is less expensive than on the fine grid. A common approach for defining a time-grid hierarchy is to use different temporal resolutions [7,13,14]. Other techniques such as considering time-integration routines of different orders of accuracy have also been applied, e.g., in power-grid simulations [15,16]. When using time-stepping routines of existing application codes, the combination of various temporal and spatial resolutions is an attractive approach for defining a time-grid hierarchy, since this allows adding time parallelism non-intrusively and reducing costs on coarse levels at the same time.

The use of coarser spatial meshes has been explored for various parallel-in-time methods. For Parareal, this approach has been considered in [17,18]. Furthermore, the combination of temporal coarsening with coarsening in space is natural in space-time multigrid methods, including the methods in [19–22]. The composition of the “parallel full approximation scheme in space and time” (PFASST) [23] with spatial multigrid, including spatial coarsening, has been considered in [24], scaling results can be found in [25,26]. The use of spatial coarsening in MGRIT has been explored for the p -Laplacian [27] and for Burgers equation [28]. In this paper, we demonstrate that spatial coarsening can significantly reduce the runtime of numerical simulations of an induction machine, but it may also degrade MGRIT conver-

gence, as observed for the p -Laplacian in [27]. Our tests show that convergence can be improved by using a stronger relaxation scheme.

The remainder of this paper is structured as follows. In Sect. 2, the MGRIT algorithm is introduced for ordinary differential equations, Sect. 3 discusses the induction machine model and its spatial discretization. Numerical results of the parallel-in-time simulation are presented in Sects. 4 and 5. Finally, conclusions are given in Sect. 6.

2 Multigrid-reduction-in-time

Consider a system of ordinary differential equations (ODEs) of the form

$$\mathbf{u}'(t) = \mathbf{f}(t, \mathbf{u}(t)), \quad \mathbf{u}(t_0) = \mathbf{g}_0, \quad t \in (t_0, t_f], \quad (1)$$

with arbitrary time points t_0 and t_f , where $t_f > t_0$, arising, for example, after the spatial discretization of a space-time partial differential equation (PDE). Discretize the time interval on a uniform distributed temporal mesh $t_i = i \Delta t$, $i = 0, 1, \dots, N_t$, with constant time step $\Delta t = (t_f - t_0)/N_t$, and let $\mathbf{u}_i \approx \mathbf{u}(t_i)$ for $i = 0, \dots, N_t$ with $\mathbf{u}_0 = \mathbf{u}(0)$. Consider a one-step time integration method

$$\mathbf{u}_i = \Phi_i(\mathbf{u}_{i-1}), \quad i = 1, 2, \dots, N_t, \quad (2)$$

with time integrator Φ_i , propagating the solution $\mathbf{u}_{i-1}$ from time point t_{i-1} to time point t_i , including forcing from the right-hand-side of the PDE. Abusing notation for the nonlinear case, Equation (2) can be written in matrix form as

$$\mathbf{A}\mathbf{u} \equiv \begin{bmatrix} I & & & \\ -\Phi_1 & I & & \\ & \ddots & \ddots & \\ & & -\Phi_{N_t} & I \end{bmatrix} \begin{bmatrix} \mathbf{u}_0 \\ \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_{N_t} \end{bmatrix} = \begin{bmatrix} \mathbf{g}_0 \\ \mathbf{0} \\ \vdots \\ \mathbf{0} \end{bmatrix} \equiv \mathbf{g}$$

and solved by a sequential forward block solve. This method is optimal, i.e., $\mathcal{O}(N_t)$, and gives the discrete solution after N_t applications of the time integrator.

In contrast to this sequential-in-time approach, parallel-in-time methods allow for parallelism in the time domain. In the following, we present one of these methods, the multigrid-reduction-in-time (MGRIT) algorithm.

2.1 The MGRIT algorithm

The MGRIT algorithm [7] is an iterative method that applies multigrid reduction (MGR) techniques [29] for solving time-dependent problems like (1). The key practical aspect of MGRIT is its non-intrusiveness, i.e., the algorithm allows reusing existing time propagators and their integration into

Fig. 1 Uniformly spaced fine and coarse temporal meshes. C -points are present on both grids while F -points are only on the fine grid

Fig. 2 F - and C -relaxation for a grid with seven points and a temporal coarsening factor of three

a parallel framework. For the construction of a multilevel algorithm, reusing the time integrators Φ_i , we require a definition of a coarse-grid system, a restriction and a prolongation operator between temporal grids and a relaxation scheme. Therefore, given a positive integer coarsening factor $m > 1$, we define a splitting of the fine-grid points into F - and C -points, whereby every m -th point is a C -point and the others are F -points, resulting in a coarse time grid $T_{i_c} = i_c \Delta T$, $i = 0, 1, \dots, N_T$, with $N_T = N_t/m$ and time step $\Delta T = m\Delta t$ as depicted in Fig. 1. For the coarse grid, we define a coarse-grid time integrator Φ_{i_c} resulting from a re-discretization with time step ΔT of the problem; other choices, such as coarsening in the order of the discretization [30,31] are also possible, but are not considered in this paper. We define two types of block relaxation schemes. The so-called F -relaxation performs a relaxation on F -points and propagates the solution from one C -point to all F -points up to the next C -point. Similarly, the C -relaxation propagates the solution from the preceding F -point to a C -point. Both relaxation schemes, depicted in Fig. 2, are highly parallel and can be combined to new relaxation schemes. For example, the FCF -relaxation, an F -relaxation followed by a C - and another F -relaxation, is typically used in the MGRIT algorithm. We define two grid transfer operators, injection as restriction and the “ideal” prolongation, for the transfer between temporal grids. The injection operator takes the values at the corresponding C -points in order to obtain values on the coarse grid. The ideal interpolation is defined as the transpose of the injection at C -points followed by an F -relaxation. That is, to obtain the values on the fine mesh the values at the C -points are obtained from the corresponding points on the coarse grid and these values are then propagated to the F -points in the current interval. The attribute “ideal” originates from the theory of algebraic multigrid as it is presented in [32], as for this choice the Galerkin coarse grid operator is the Schur complement after reordering the system matrix according to the C/F -splitting.

Fig. 3 Structure of V - and F -cycles for four grid levels

Using these components and the full approximation storage (FAS) framework [33], the two-level MGRIT-FAS algorithm [34], as extended by spatial coarsening in [27], can be written as Algorithm 1. Note that all statements of Algorithm 1, except for the solving in step 6, are highly parallel and can be executed simultaneously for several time steps.

Algorithm 1 FAS-MGRIT($\mathcal{A}, \mathbf{u}, \mathbf{g}$)

- 1: Apply F -relaxation to $\mathcal{A}_1(\mathbf{u}^{(1)}) = \mathbf{g}^{(1)}$
 - 2: For 0 to γ :
 - 3: Apply CF -relaxation to $\mathcal{A}_1(\mathbf{u}^{(1)}) = \mathbf{g}^{(1)}$
 - 4: Inject the approximation and its residual to the coarse grid:

$$\mathbf{u}^{(2)} = \mathbf{R}_I(\mathbf{u}^{(1)}),$$

$$\mathbf{g}^{(2)} = \mathbf{R}_I(\mathbf{g}^{(1)} - \mathcal{A}_1 \mathbf{u}^{(1)})$$
 - 5: If spatial coarsening:

$$\mathbf{u}^{(2)} = \mathbf{R}_s(\mathbf{u}^{(2)})$$

$$\mathbf{g}^{(2)} = \mathbf{R}_s(\mathbf{g}^{(2)})$$
 - 6: Solve $\mathcal{A}_2(\mathbf{v}^{(2)}) = \mathcal{A}_2(\mathbf{u}^{(2)}) + \mathbf{g}^{(2)}$
 - 7: Compute the error approximation: $\mathbf{e} = \mathbf{v}^{(2)} - \mathbf{u}^{(2)}$
 - 8: If spatial coarsening:

$$\mathbf{e} = \mathbf{P}_s(\mathbf{e})$$
 - 9: Correct using ideal interpolation: $\mathbf{u}^{(1)} = \mathbf{u}^{(1)} + \mathbf{P}(\mathbf{e})$
-

In Algorithm 1, $\mathcal{A}_l(\mathbf{u}^{(l)}) = \mathbf{g}^{(l)}$ denotes the space-time system of equations on levels $l = 1, 2$, which can be either linear or nonlinear, with each block row corresponding to one time step. The relaxation scheme of the algorithm can be controlled by the parameter γ , and is, in general, for the MGRIT algorithm chosen as $\gamma = 1$, i.e., an FCF -relaxation scheme. But other choices are also possible. Note that the two-level variant with F -relaxation is equivalent to the parareal method [6]. We denote the temporal grid transfer operators by $\mathbf{R}_I$ and $\mathbf{P}$, whereby $\mathbf{P}$ is the ideal prolongation and $\mathbf{R}_I$ the restriction operator based on injection. Additionally, we define $\mathbf{R}_s$ as spatial restriction and $\mathbf{P}_s$ as spatial prolongation (see Sect. 2.2).

The two-level MGRIT-FAS algorithm can be extended to a multilevel setting by applying the two-level method recursively to the system in step 6. Depending on the number of recursive calls, different multilevel schemes, e.g., V -, F - or W -cycles as depicted in Fig. 3, can be defined. Note that, in exact arithmetic the MGRIT algorithm with F - or FCF -relaxation solves for the discrete fine-grid solution in N_t/m or $N_t/(2m)$ iterations [7], respectively.

In Algorithm 1, an initial guess $\mathbf{u}$ is needed. In general, if prior knowledge of the solution $\mathbf{u}$ exists, this knowledge

should be used as an initial guess. If no such prior knowledge exist, an improved initial guess can be computed by using the nested iterations [35,36] strategy. The idea of nested iterations is to compute an initial approximation on the coarsest grid and interpolate the approximation to the finer grids, applying one V -cycle per level. This procedure results in the same structure as the F -cycle, with skipping the down-cycle at the beginning.

2.2 MGRIT with spatial coarsening

Spatial coarsening is added to the MGRIT-FAS algorithm by the spatial restriction operator $\mathbf{R}_s$ in line 5 and the spatial prolongation operator $\mathbf{P}_s$ in line 8. Note, that spatial coarsening is an extra option that can be applied after semi-coarsening in time and, therefore, does not touch the non-intrusive character of the algorithm. So both operators have to be specified manually, but are independent of the time-integration scheme. Obviously, the spatial grid transfer operators directly influence the convergence behavior of the algorithm and should be chosen wisely.

In a multilevel setting, multiple coarsening strategies can be chosen depending on the number of available grids in time and space. Both coarsening dimensions are independent of each other and, therefore, can be applied in different ways. Assuming that more temporal grids than spatial grids are available, the most cost efficient approach is to use a “direct” spatial coarsening strategy, starting with spatial coarsening as early as possible. This strategy would reduce the computational costs per iteration for all spatial coarse level problems, but studies [27] have demonstrated that this “direct” strategy can degrade the convergence behavior of the MGRIT algorithm for linear and nonlinear problems. To overcome the degradation in MGRIT convergence, in [27] a “delayed” spatial coarsening strategy is proposed. This strategy applies spatial coarsening as late as possible, i.e., only on the coarsest time grids. Figure 4 illustrates the different coarsening strategies for a five-level V -cycle for the case of three spatial grids, characterized by spatial grid spacings of Δx , $2\Delta x$, and $4\Delta x$, i.e., assuming factor-two coarsening in space. The temporal coarsening factor is chosen to be $m = 4$, i.e., the time step on each temporal grid is given by $4^{l-1}\Delta t$, $l > 0$. The left V -cycle represents the MGRIT algorithm with coarsening only in the time dimension. In the middle, the direct spatial coarsening strategy is shown, where spatial coarsening is applied on the first and second time levels. Finally, the right V -cycle illustrates the delayed strategy, with spatial coarsening starting on the third level to make use of the coarsest space grid on the coarsest level.

3 Simulation of induction machines

The simulation of electromagnetic fields is based on the solution of Maxwell’s equations [2], but, for a variety of applications, it is sufficient to solve a (quasistatic) approximation. In this section, we introduce one of these simplifications, the so-called eddy current problem, that is typically used in the simulation of energy converters. Further, we introduce a numerical model of an induction machine that is used for computations in Sects. 4 and 5.

3.1 Induction machine model

The standard approach in the simulation of electrical machines is to neglect the displacement current density with respect to the source current density. The resulting magnetoquasistatic approximation of Maxwell’s equations is the well-known eddy current problem.

Let Ω denote a 2D cross-section in x, y -coordinates of the 3D geometry (with depth ℓ_z). Then, the problem can be solved in terms of the z -component of the magnetic vector potential $A_z : \Omega \times (t_0, t_f] \rightarrow \mathbb{R}$. We split the spatial domain into stator and rotor $\bar{\Omega} = \bar{\Omega}_1 \cup \bar{\Omega}_2$ and introduce separate unknowns $A_k = A_z|_{\Omega_k}$ on each domain ($k = 1, 2$). The problem reads for the stator

$$-\nabla \cdot (\nu \nabla A_1) - \sum_{s=1}^{n_1} \chi_s i_{1,s} = 0 \quad (3a)$$

$$\frac{d}{dt} \int_{\Omega} \chi_s A_1 \ell_z dS + R_s(\mathbf{i}_1) = v_{1,s}, \quad (3b)$$

where $s = 1, \dots, n_1$, and for the rotor

$$\sigma \partial_t A_2 - \nabla \cdot (\nu \nabla A_2) - \sum_{s=1}^{n_2} \sigma \xi_s v_{2,s} = 0 \quad (3c)$$

$$-\frac{d}{dt} \int_{\Omega} \xi_s \sigma A_2 \ell_z dS + G_s(\mathbf{v}_2, \mathbf{i}_2) = 0 \quad (3d)$$

$$L_s \left(\mathbf{v}_2, \frac{d}{dt} \mathbf{i}_2 \right) = 0, \quad (3e)$$

where $s = 1, \dots, n_2$, with interface conditions on $\Omega_1 \cap \Omega_2$ within the air-gap for all $t \in (t_0, t_f]$,

$$A_1(\mathbf{x}_1, t) = A_2(\mathbf{x}_2, t) \quad (3f)$$

$$\mathbf{n} \cdot \nabla A_1(\mathbf{x}_1, t) = \mathbf{n} \cdot \nabla A_2(\mathbf{x}_2, t), \quad (3g)$$

where $\mathbf{x}_2 = \mathbf{r}(\mathbf{x}_1, \theta)$ imposes the rotation by angle θ . The problem is completed by the Dirichlet boundary condition $A_z = 0$ on $\partial\Omega$ and the initial value $A_z(\mathbf{x}, t_0) = A_0(\mathbf{x})$. The electrical conductivity $\sigma(\mathbf{x}) \geq 0$ and the (isotropic, nonlinear) magnetic reluctivity $\nu(\mathbf{x}, |\nabla A|) > 0$ encode the geometry. The current and voltage distribution functions

Fig. 4 Five-level MGRIT V-cycle algorithm with three spatial grids and different space-time coarsening strategies. The left figure shows the MGRIT algorithm with coarsening only in time, the middle one illustrates direct spatial coarsening, and at right, the delayed spatial coarsening strategy is applied

$\chi_s, \zeta_s : \Omega_s \rightarrow \mathbb{R}$ model the stator coils by stranded conductors and the rotor bars by interconnected solid conductors [37] in terms of resistive circuits R_s and G_s , inductive circuits L_s , voltages $\mathbf{v}_k = [v_{k,1}, \dots, v_{k,n_k}]^\top$, and currents $\mathbf{i}_k = [i_{k,1}, \dots, i_{k,n_k}]^\top$ for $k = 1, 2$.

A common quantity of interest in the design of electrical machines are the Joule losses

$$P_{\text{loss}} = \int_{\Omega} \partial_t A_z \sigma \partial_t A_z \ell_z d\Omega. \quad (4)$$

To include the rotation of the rotor, the problem above is extended by an additional motion equation. Movements are determined by the mechanical equation

$$\omega(t) = \frac{d\theta(t)}{dt}, \quad t \in (t_0, t_f], \quad (5a)$$

$$I \frac{d^2\theta(t)}{dt^2} + C \frac{d\theta(t)}{dt} = T_{\text{mag}}(A_z) \quad \text{in } (t_0, t_f], \quad (5b)$$

$$\theta(t_0) = \theta_0, \quad (5c)$$

$$\omega(t_0) = \omega_0, \quad (5d)$$

where ω is the angular velocity of the rotor, I denotes the moment of inertia, and C is the friction coefficient, respectively. The torque T_{mag} defines the mechanical excitation of system (5) given by the magnetic field. Note that there are several common special cases, i.e., a fixed rotor ($\theta = \text{const}$) without movement and given rotor speed (synchronous or asynchronous).

Discretizing (3) in space using finite elements with n_a degrees of freedom and using the moving band approach [38] for motion yields a system of equations of the form

$$M \mathbf{u}'(t) + K(\mathbf{u}(t)) \mathbf{u}(t) = \mathbf{f}(t), \quad t \in (t_0, t_f] \quad (6a)$$

$$\mathbf{u}(t_0) = \mathbf{u}_0, \quad (6b)$$

with initial condition $\mathbf{u}_0 \in \mathbb{R}^n$ and unknown $\mathbf{u}^\top = [\mathbf{a}^\top, \mathbf{i}_1^\top, \mathbf{i}_2^\top, \mathbf{v}_2^\top, \theta, \omega] : (t_0, t_f] \rightarrow \mathbb{R}^n$. The solution $\mathbf{u}$ consists of the magnetic vector potential $\mathbf{a}$, currents $\mathbf{i}_1, \mathbf{i}_2$, voltages $\mathbf{v}_2$, the rotor angle θ , and the rotor's angular velocity ω . The stator voltages $v_{1,s}(t)$ drive the problem and, thus, determine the right-hand side $\mathbf{f}(t)$. Integrating the semi-discrete system (6) using the backward Euler method results

in a system of the form

$$\left(\frac{1}{\Delta t} M + K(\mathbf{u}_i) \right) \mathbf{u}_i = \mathbf{f}_i + \frac{1}{\Delta t} M \mathbf{u}_{i-1}, \quad (7a)$$

$$\mathbf{u}_0 = \mathbf{u}(t_0). \quad (7b)$$

This field/circuit coupled problem consists of differential-algebraic equations (DAEs) of index-1 [39,40] depending on the circuit topology. A discussion of the index-1 DAE aspects in a parallel-in-time setting can be found in [8], i.e., if the mass matrix M is constant and the backward Euler method is used then there is no need for explicitly enforcing consistency of the (algebraic) initial values.

3.2 Pulse-width-modulation

Electrical machines are commonly driven by PWM excitations since they have technical advantages for the power electronics that is implemented to control the machine. While multiple forms of PWM exists [41], in this paper, we consider a PWM signal that is produced by comparing a reference wave with a triangular wave. Idealizing the circuitry to generate the PWM, i.e. diodes are replaced by switches, results in the following excitation

$$b(t) = \text{sgn}[r(t) - c(t)], \quad t \in \mathbb{R}, \quad (8)$$

with reference signal $r(t)$ and where $c(t)$ denotes the carrier signal. Furthermore, the modulation factor is defined as the ratio of the amplitude of the reference signal and the amplitude of the carrier signal. Figure 5 shows an example of a PWM signal. The red line illustrates a sine reference signal of 50 Hz with an amplitude of 0.8 and the carrier signal, modeled by a sawtooth carrier [42] of 500 Hz and with an amplitude of one, is shown as a blue line. The resulting PWM signal with switching frequency of 500 Hz is presented in the lower subplot of Fig. 5.

3.3 Numerical model

We model the semi-discrete eddy-current problem (6), supplied by a three-phase PWM voltage source, using the multi-slice FE model “im_3_kw” of an electrical machine

Fig. 5 Illustration of a PWM signal (black), generated by a sine wave reference signal (red) and a sawtooth carrier signal (blue)

first presented in [43] and modified in [9] for the usage of a PWM voltage source. In more detail, the two-dimensional model “im_3_kw”, depicted in Fig. 6, of a four-pole squirrel cage induction machine with no-load operation condition is used for numerical computations. Further, as typical for this kind of symmetric models, we consider only a quarter of the machine with periodic boundaries to reduce simulation costs. The machine is supplied by a three-phase PWM voltage source

$$v_s^p(t) = \text{sign}[r_s(t) - c_p(t)], \quad t \in (t_0, t_f] \quad (9)$$

with reference signals $r_s(t)$, $s = 1, 2, 3$, and carrier signal $c_p(t)$ with p pulses. The reference signal for the three phases is given by

$$r_s(t) = \sin\left(\frac{2\pi}{T}t + \varphi_s\right), \quad (10)$$

with $\varphi_1 = 0$, $\varphi_2 = 2/3\pi$, $\varphi_3 = -4/3\pi$ and electric period T . The carrier signal, given by a bipolar trailing-edge modulation using a sawtooth carrier [42], is defined by

$$c_p(t) = 2\left(\frac{p}{T}t - \lfloor \frac{p}{T}t \rfloor\right) - 1. \quad (11)$$

For our simulation we consider a frequency of 50 Hz, corresponding to an electric period of $T = 0.02$ s, and choose $p = 400$ pulses for one period, resulting in a PWM voltage source of 20 kHz. The model provides a linear, i.e., v does not depend on $\mathbf{A}$, and a nonlinear version of the induction machine. In the nonlinear setting, the voltage source v_s^p , defined in (9), is multiplied with the time-dependent factor

$$\alpha(t) = 0.5(1 - \cos(\pi t/2)) \quad (12)$$

Fig. 6 Mesh view of the four-pole squirrel cage machine model “im_3_kw” of [43]

Fig. 7 Three spatial discretizations of one quarter of the machine generated by Gmsh. Finer meshes Ω_2 (middle) and Ω_1 (right) are created by refinements of the left mesh Ω_3

in the first two periods, i.e., for $t \in [0, 2T]$, which reduces the transient behavior of the motor and allows to obtain the steady-state more quickly [43].

The model computes the value of Joule losses P_{loss} for each time point, i.e., the loss in the transfer of electrical energy. The size of Joule losses indicates a significant quality in the prototyping of designs, and, therefore, it is one possible stopping criterion for iterative methods.

3.4 Spatial discretization

Gmsh [44,45] is used to generate mesh representations of the model. We consider a hierarchy of three nested spatial grids, illustrated in Fig. 7. The coarsest model, defined on the mesh Ω_3 , has $n_a = 4449$ degrees of freedom. Two more accurate discretizations are defined on grids Ω_2 and Ω_1 that are constructed by refinements of the coarsest mesh, resulting in grids with $n_a = 17,496$ and $n_a = 69,384$ degrees of freedom, respectively.

We consider the standard finite element interpolation or nodal interpolation as the grid transfer operator between the spatial meshes. To avoid numerical instabilities at the boundaries between the stator and the rotor, we apply the standard finite element interpolation separately for both regions. Recall from Sect. 3.1 that in our model problem, only the magnetic vector potential is a function that depends on space, and, therefore, only the magnetic vector potential is discretized on the different spatial meshes, whereas currents,

rotor angle, and the rotor's angular velocity are functions that are independent of space. Taking into account that System (6) consists of grid-dependent and grid-independent unknowns, we define the spatial interpolation and restriction operators as block operators with two blocks as follows: For the interpolation of grid-dependent unknowns, we use the standard finite element interpolation, whereas the grid-independent unknowns are injected.

3.5 Implementation details

The MGRIT algorithm is implemented in PyMGRIT [46], a Python package using the Message Passing Interface (MPI). The implementation conforms to the non-intrusive character of the MGRIT method and requires an implementation of the time stepper routine by the user. For the machine model, the time stepper routine calls the GetDP [47–49] library for performing the simulation of the electrical machine which implements the implicit Euler method. As we aim at demonstrating the benefits of parallelization in time, for the implementation on a distributed memory computer, we consider a domain-decomposition approach only in time, i.e., only the time interval is distributed across processors. All tests were performed on an Intel Xeon Phi Cluster consisting of 272 1.4 GHz Intel Xeon Phi processors.

3.6 Storage requirements

The MGRIT algorithm requires storing solution values at all C -points. On coarse levels, two additional vectors are stored: one for a restricted copy of the fine-grid solution and one for the FAS right-hand-side. Denoting the number of MGRIT levels by L and the storage cost of a vector on level l by s_l , $l = 0, \dots, L - 1$, the total storage requirement of MGRIT per temporal processor is given by [27]

$$S_{\text{MGRIT}} \approx \sum_{l=0}^{L-1} \left\lceil \frac{N_t}{m^{l+1}p} \right\rceil s_l + \sum_{l=1}^{L-1} 2 \left\lceil \frac{N_t}{m^l p} \right\rceil s_l,$$

where p is the number of processors in time. The first term corresponds to the storage requirements of solution values, and the second term is that of the additional vectors on the coarse levels. Note that spatial coarsening reduces the size of all vectors on coarse grids.

In comparison, the sequential time-stepping algorithm requires storing only one solution value, corresponding to a total storage requirement of s_0 . Thus, the total storage cost of the MGRIT algorithm is larger than that of sequential time-stepping by a factor of S_{MGRIT}/s_0 . However, note that this factor decreases with increasing the number of processors and in the case of spatial coarsening.

4 Linear material model

In this section, we investigate the behavior of the MGRIT algorithm for the eddy current problem of a two-dimensional induction machine with linear material characteristics. Especially, we are interested in the behavior of adding spatial coarsening to the MGRIT algorithm for system (6), containing scalar-valued as well as grid-dependent unknowns. For this study, we choose a linear model of the induction machine, also available in the model “im_3_kw”, to reduce computational costs. In Sect. 5, we transfer our results to the nonlinear model of the machine.

4.1 Algorithmic parameters

The MGRIT-FAS algorithm has a large number of parameters and variations, considering V -, F - and W -cycles, the number of temporal grid levels, coarsening strategies in space and time, number and type of relaxation schemes, and so forth. Based on results presented in the literature as well as on personal experiences, we consider a two- and a five-level MGRIT variant in this study. Further, we choose V - and F -cycles. V -cycles offer better parallel scaling than F -cycles, but F -cycles have better convergence behavior due to additional work on coarse temporal grids. We consider FCF -relaxation for all variants, and, additionally, F -relaxation is used in the two-level setting as well as in MGRIT F -cycles.

The main goal of the MGRIT algorithm is to reduce the runtime of the simulation by making use of parallel computer architectures. To achieve the best runtime results for our model problem, the temporal coarsening strategy of the algorithm has to be adapted to the number of available processors. Therefore, we choose a non-uniform temporal coarsening strategy directly connected to the number of points on the finest time grid and the number of processors. On the one hand, a large coarsening factor reduces the number of points for all following time grids and, thereby, the costs. On the other hand, to make use of all available processors, the coarsening factor cannot be chosen too large. As a compromise and based on experience with MGRIT for eddy current problems in a parallel setting [11, 12], we choose a large coarsening factor on the first level, such that the number of points on the second level corresponds to the number of processors available. On coarse levels, a small coarsening factor is used. For example, consider a fine temporal grid with 129 points, 16 processors available and a five-level MGRIT method. The first temporal coarsening factor is chosen to be eight, such that the work of the first grid is distributed evenly among all processors; here, eight time points per processor. For the coarser levels, a coarsening factor of two is used to reduce the number of time points per grid, such that the coarsest grid has only three points. With this strategy, the first level is perfectly parallelized and the amount of work on coarser levels

that cannot be parallelized is minimized. Note that this is a heuristic choice purely based on distributing computational work evenly among a given number of processors. Due to the high computational cost of the external solver and to the use of moderate numbers of processors, it is reasonable to neglect communication costs for this problem. Optimization of the temporal coarsening strategy that takes additional parameters that affect parallel performance, such as communication costs and MGRIT convergence behavior into account, is a topic of future work.

In MGRIT multilevel variants, we apply the direct and delayed spatial coarsening strategies described in Sect. 2.2, with the interpolation and restriction operators $\mathbf{P}_s$ and $\mathbf{R}_s$ defined in Sect. 3.4. Note that these operators have to be computed only once and can be used for all MGRIT variants. Furthermore, nested-iterations, described in Sect. 2.1, is used to generate an improved initial guess. For this preprocessing step, we use the reference signal (10) as a voltage source based on the idea in [9]. For the iterations of the MGRIT algorithm, the discontinuous signal (9) is used on every grid level. Since we are interested in the algorithmic behavior of different MGRIT variants, in the following, we report on tests with MGRIT convergence measured based on the absolute space-time residual norm instead of the Joule losses (4). More precisely, the convergence tolerance is chosen to be 10^{-4} , measured in the discrete L_2 -norm. We have checked that for all algorithms, this fixed space-time residual norm is sufficient to achieve discretization error for our linear model problem if the algorithms converge. That is, algorithms either converge to the time-stepping solution of the linear model of the four-pole squirrel cage induction machine with a PWM voltage source or convergence is only observed after N_t/m or $N_t/(2m)$ iterations for F - or FCF -relaxation, respectively.

4.2 Effects of spatial coarsening

We apply several MGRIT variants to the linear machine model problem on the space-time domain $\Omega \times [0, 0.03125]$ s. The time interval, corresponding to about one and a half periods, consists of $2^{15} + 1$ time steps of size $\Delta t = 2^{-20}$ that are distributed evenly among 256 processors. Note that this time-step size resolves physical processes as well as the pulses of the PWM voltage source. As we study the initial behavior of the machine, also note that the final time is not sufficient to reach the steady state. With this choice of number of processors and time points, the temporal coarsening factor is chosen to be $m = 256$ on the first level, resulting in a second-level time grid with 256 points. The five-level algorithms use a coarsening factor of four on all coarse levels.

Table 1 shows the number of iterations, setup, solve, and total time in seconds for the different MGRIT variants without spatial coarsening. The setup time consists of the setup of the algorithm and the generation of an improved initial guess

using nested iterations. Note that for nested iterations, the same relaxation scheme as the algorithm itself is used, e.g., the setup for the five-level F -cycle with F -relaxation is faster than that for the five-level F -cycle with FCF -relaxation. The two-level algorithm with F -relaxation reaches the desired tolerance after 12 MGRIT iterations. While the use of FCF -relaxation is beneficial in the two-level setting, reducing both the number of iterations and the runtime, for F -cycles, F -relaxation performs better than FCF -relaxation. Comparing the total runtimes, the multilevel variants are faster than the two-level methods, with the five-level V -cycle being the fastest. The problem on the second level is still very expensive, and, therefore, using the two-level method recursively for this problem allows for better performance.

Now, we consider the same MGRIT variants with spatial coarsening. More precisely, we add spatial coarsening by using the grids Ω_2 and Ω_3 , following either the direct or delayed coarsening strategy. Note that in the five-level MGRIT variants, direct spatial coarsening refers to using Ω_1 on the finest grid, Ω_2 on the second level, and Ω_3 on all remaining coarse levels. Delayed spatial coarsening, in contrast, uses Ω_3 only on the coarsest time grid, Ω_2 on the second coarsest level, and Ω_1 on the three finest levels. In the two-level methods, Ω_2 or Ω_3 is used on the coarse grid. Table 2 demonstrates the effects of adding spatial coarsening to the MGRIT variants considered in Table 1. All variants with FCF -relaxation converge to the solution in the same number of iterations as the variants without spatial coarsening. Especially for this problem, the direct spatial coarsening strategy does not degrade the convergence behavior. The situation is different for all variants with F -relaxation. Neither the two-level method with spatial coarsening converges to the desired tolerance in significantly fewer than N_t/m iterations, nor does the five-level F -cycle.

To better understand the degradation in convergence when using F -relaxation, Fig. 8 details the convergence behavior for all F -cycle variants considered in Table 2. Shown are the space-time residual norms over the first ten MGRIT iterations for all F -cycle variants. Dashed lines are results for using direct spatial coarsening and solid lines represent applying the delayed approach. The variants with FCF -relaxation show linear convergence behavior, and there is effectively no difference between both spatial coarsening strategies. In contrast, convergence of MGRIT with F -relaxation stagnates after five iterations. We assume this behavior comes from the scalar-valued unknowns of our system, which are injected between the spatial grids. The additional CF -relaxation in the FCF -relaxation scheme improves the spatial interpolation and fixes this problem.

Comparing the runtimes of the different MGRIT variants with and without spatial coarsening in Tables 1 and 2, the two-level algorithms show the best speedup, decreasing the total runtime by up to a factor of about 2.72, followed

Table 1 Results for the linear induction machine discretized on a space-time grid of size $69,384 \times 2^{15}$ in terms of iterations, setup, and solve time on 256 processors for the different MGRIT variants without spatial coarsening

MGRIT variant	Iterations	Setup time (s)	Solve time (s)	Total time (s)
Two-level cycles with F -relax.	12	9523	166,265	175,788
Two-level cycles with FCF -relax.	7	9515	125,149	134,664
Five-level V-cycles with FCF -relax.	8	3293	77,146	80,439
Five-level F-cycles with F -relax.	12	2396	85,117	87,513
Five-level F-cycles with FCF -relax.	7	3291	87,660	90,951

Table 2 Results similar to those of Table 1 but with spatial coarsening; $\times$ indicates no convergence to the desired tolerance in less than $N_t/m = 256$ iterations and speedup is measured relative to the same MGRIT variant without spatial coarsening

MGRIT variant	SC strategy	Iterations	Setup time	Solve time	Total time	Speedup w.r.t. no SC
Two-level cycles with F -relax.	Direct	$\times$	–	–	–	–
Two-level cycles with FCF -relax. (Ω_2)	Direct	7	2444 s	73,713 s	76,157 s	1.77
Two-level cycles with FCF -relax. (Ω_3)	Direct	7	685 s	48,866 s	49,551 s	2.72
Five-level V-cycles with FCF -relax.	Delayed	8	1664 s	73,022 s	74,686 s	1.08
	Direct	8	316 s	66,233 s	66,549 s	1.21
Five-level F-cycles with F -relax.	Delayed	$\times$	–	–	–	–
	Direct	$\times$	–	–	–	–
Five-level F-cycles with FCF -relax.	Delayed	7	1665 s	74,216 s	75,881 s	1.20
	Direct	7	317 s	60,372 s	60,689 s	1.50

Fig. 8 Convergence results for the different five-level F -cycle variants with direct (dashed lines) and delayed (solid lines) spatial coarsening

by the five-level F -cycle with FCF -relaxation and direct spatial coarsening, which is faster than the same variant without spatial coarsening by a factor of about 1.5. Note that, for this problem, the use of direct spatial coarsening does not degrade MGRIT convergence and, thus, choosing this more aggressive coarsening strategy leads to higher speedups than applying delayed spatial coarsening. Applying a more aggressive spatial coarsening strategy also in the two-level method with FCF -relaxation similarly allows for a larger speedup over the same algorithm without spatial coarsening. Using the spatial mesh Ω_3 instead of Ω_2 on the coarse time

grid, reduces the runtime from 76,157 seconds to 49,551 seconds and, thus, increases the speedup from a factor of about 1.77 to a factor of about 2.72. However, the scalability of the two-level method is limited by the size of the coarse-grid problem, and spatial coarsening can lead to divergence in the nonlinear setting, as will be seen in Sect. 5.2.

4.3 Discussion

Summarizing, the MGRIT algorithm is applied successfully to the linear material model of an induction machine with a pulsed voltage source. With our choice of spatial grid transfer operators, spatial coarsening can be added for all variants with FCF -relaxation, while the methods with F -relaxation methods do not converge to the time-stepping solution in less than N_t/m iterations. A better choice could resolve this behavior. For our problem, we see no difference in the direct and delayed strategy, and, therefore, the direct approach should be used for reducing the runtime. The speedup by adding spatial coarsening corresponds directly to the amount of work on the coarser grids, i.e., the two-level algorithm has the best speedup, followed by the F -cycle.

5 Full nonlinear model

In this section, we use the results from Sect. 4.2 of different MGRIT variants compared with one another for studying the application of MGRIT to the full nonlinear model of the induction machine “im_3_kw”, i.e., with a state dependent reluctivity ν which models magnetic saturation. In particular, we are interested in effects of spatial coarsening in this setting as well as in the parallel performance of MGRIT. Furthermore, we compare MGRIT with sequential time stepping.

5.1 Numerical parameters

To limit computational costs in the full nonlinear setting, we perform experiments on a smaller space-time domain. More precisely, we consider Ω_2 as the spatial grid and choose $\Delta t = 2^{-20}$ and $N_t = 10,753$ time steps, resulting in a final time $t_f \approx 0.01$ s. Convergence of MGRIT is measured based on the relative change in values of joule losses at C-points of two consecutive iterations. MGRIT iterations are stopped when the maximum norm of the relative difference of two successive iterates at C-points is below 10^{-2} , i.e., when the maximum relative change in joule losses is smaller than 1% for all C-points. Note that this stopping criterion is much looser compared to that in Sect. 4, but it corresponds to a choice made for prototyping of induction machines in a realistic setting. However, in contrast to using a convergence tolerance based on the space-time residual, it is not guaranteed that MGRIT converges to the discrete time-stepping solution. For the MGRIT variants considered in this section, we have checked that, if convergence is observed, iterates indeed converge to the discrete time-stepping solution.

In the full nonlinear setting, each application of the time propagator Φ involves an iterative nonlinear solve. GetDP, called by the time stepper routine of our MGRIT implementation, uses Newton’s method with damping for the solution of the nonlinear problems. As convergence of the Newton solver can depend on the initial guess, in Sect. 5.2, we consider not only the effect of spatial coarsening on MGRIT convergence, but also on the convergence of the Newton solver. We then use these results in Sect. 5.3 to choose a set of MGRIT variants for a strong scaling study and comparison to sequential time stepping.

5.2 Effects of spatial coarsening

The results in Sect. 4.2 show that adding spatial coarsening in MGRIT is beneficial in the case of FCF -relaxation, whereas MGRIT with F -relaxation does not converge when spatial coarsening is added. We now consider two- and five-level MGRIT variants with FCF -relaxation and compare the performance of these methods without and with direct spatial coarsening using the spatial mesh Ω_3 on coarse temporal

grids. Additionally, we look at the performance of two-level MGRIT with F -relaxation and without spatial coarsening. All tests are performed using 256 processors. Choosing again the non-uniform temporal coarsening strategy described in Sect. 4.1, a coarsening factor of 42 is used on the first level, and factor-four coarsening is applied on all coarse levels.

Table 3 shows iteration counts and runtimes of the different MGRIT variants, with timings split up again into setup and solve times. No results are given for the two-level method with FCF -relaxation and with spatial coarsening, since during the F -relaxation step of the first MGRIT iteration, the Newton solver does not converge for at least one time step and, thus, the algorithm cannot be applied in this setting. In contrast to the setting of linear material laws, here, an initial guess that is in the region of convergence of the Newton solver is crucial. Applying a nested iterations strategy in the two-level setting, the initial guess at the C-points of the fine grid are given by the interpolated solutions of the coarse grid. As no values are given at F-points, F -relaxation is then carried out using the solution from the previous time point as the initial guess, i.e., for the first F-point in each F-interval, the solution at the preceding C-point is used as the initial guess. After solving the corresponding nonlinear problem at the first F-point in each interval, the solution is taken as the initial guess for the following F-point, and so on. However, this strategy appears to result in a poor initial guess for the Newton solver at some time points. There are several difficulties for obtaining a good initial guess, including the use of spatial coarsening, the discontinuities of the right-hand side of the problem, as well as the different time-step sizes on the temporal grids in the multigrid hierarchy. Improving the initial guess requires further investigation of strategies for tackling these challenges. This is a topic of future work.

Looking at the total runtimes of the different variants without spatial coarsening, the five-level V-cycle algorithm is fastest, followed by the five-level F-cycle variant which is about a factor of 1.5 times slower than considering V-cycles. Furthermore, five-level F-cycles are already about twice as fast as the two-level method. Adding spatial coarsening in the multilevel schemes, we can benefit over the two-level algorithm even more. Considering five-level V- and F-cycles with spatial coarsening, the factor in comparison with the runtime of the two-level method can be increased from 1.9 or 2.8 when applying F- or V-cycles without spatial coarsening, respectively, to a factor of about 3.8 or 4.4 when spatial coarsening is added.

5.3 Parallel results

Now we present strong scaling results of the five convergent MGRIT variants considered in Table 3. We perform parallel tests using between eight and 256 processors for parallelization in time. The temporal coarsening strategy in all runs is

Table 3 Number of iterations, setup, solve and total time on 256 processors of various MGRIT variants applied to the full nonlinear model of the induction machine “im_3_kw” discretized on a space-time grid of size $17,496 \times 10,753$

MGRIT variant	Spatial coarsening	Iterations	Setup time	Solve time	Total time
Two-level cycles with F -relax.	No	4	15,054 s	74,150 s	89,204 s
Two-level cycles with FCF -relax.	Yes	X	–	–	–
Five-level V-cycles with FCF -relax.	No	3	5870 s	25,509 s	31,379 s
	Yes	3	1297 s	18,763 s	20,060 s
Five-level F-cycles with FCF -relax.	No	3	5861 s	41,138 s	46,999 s
	Yes	3	1301 s	22,430 s	23,731 s

X indicates no convergence of at least one nonlinear spatial solve within GetDP

fixed at applying a coarsening factor of 42 on the finest grid and factor-four coarsening on all coarse grids. Additionally, we are interested in the benefits over sequential time stepping.

Figure 9 shows total runtimes of the different MGRIT variants as a function of the number of processors, as well as the time-to-solution of the sequential block forward solve for reference purposes. For the multilevel variants, runtimes are shown for using spatial coarsening (dashed lines) and without spatial coarsening (solid lines). The runtime of all multilevel variants using eight processors as well as the time-to-solution of sequential time stepping is about five to six days, while using eight-way parallelism in the two-level method results in a runtime of only about four days. However, the multilevel variants show better strong parallel scaling and, thus, lead to faster time-to-solution at large processor counts than the two-level method. Increasing the number of processors to 256, the total runtime can be reduced to about 8.72 hours or 5.57 hours, when considering MGRIT V-cycles without or with spatial coarsening, respectively. While F -cycles with spatial coarsening allow for a similar reduction, F -cycles without spatial coarsening using 256-way parallelism already take about half a the two-level variant reduces the runtime to about one day.

Table 4 details speedups and parallel efficiencies for the MGRIT variants. The speedup is computed relative to sequential time stepping, representing the algorithm with the minimum runtime on one processor, i. e., the speedup using p processors is given by $S_p = T_{TS}/T_{MGRIT}(p)$. The parallel efficiency is measured as S_p/p for p processors. Using multilevel methods, we can benefit more over sequential time stepping at larger processor counts than with the two-level algorithm. For example, while on 16 processors the speedup is similar for all methods, using 32 processors results in a speedup of up to a factor of 3.6 for a five-level variant, whereas the two-level method yields only a speedup of a factor of 2.8. Increasing the number of processors to 256, with the two-level method we achieve a speedup of a factor of about 5. In contrast, multilevel variants at least double the speedup factor. Considering spatial coarsening, leads to

the best speedup with the V-cycle algorithm being nearly 22 times faster than the sequential time stepping method. Moreover, the use of spatial coarsening in MGRIT F -cycle is crucial for improving the parallel scalability for more than 64 processors, since this allows reducing computations on coarse levels that dominate over communication costs. While speedups and efficiencies degrade for F -cycles without spatial coarsening, the degradation is only modest when adding spatial coarsening. Note that we bind the temporal coarsening strategy to the maximum number of available processors and, thus, faster runtimes on eight to 128 processors may be possible by adjusting the temporal coarsening strategy to the number of processors.

6 Conclusions

In this paper, the MGRIT algorithm is applied to a model of a realistic two-dimensional electrical machine with a pulsed excitation. Using the non-intrusive character of MGRIT for an existing model of an induction machine, performance of MGRIT is dominated by the cost of the time-stepping routine which carries out a nonlinear spatial solve for a given time step. To reduce the cost of spatial solves on coarse grids, the use of spatial coarsening is investigated. We demonstrate that spatial coarsening is a powerful option that can significantly reduce the runtime of the algorithm, but it also comes along with new challenges and problems, e.g., degradation in MGRIT convergence, as also observed in [27] for the p -Laplacian, and divergence of the Newton solver due to a poor initial guess. These are topics of future work.

Strong scaling results show significant speedups compared to sequential time stepping. For a full nonlinear model of an induction machine and using moderate numbers of processors, we demonstrate that MGRIT allows reducing the simulation time from several days to only a few hours. While MGRIT is a non-intrusive approach, there is still a very large parameter space for the algorithm, such as the choice of the level structure, temporal and spatial coarsening strategies, and so on. Especially if computational resources are limited

Fig. 9 Total time-to-solution for the nonlinear induction machine “im_3_kw” using different MGRIT variants and sequential time stepping. Solid lines are runtimes without spatial coarsening, and dashed lines represent results with spatial coarsening. The dotted line shows the runtime of time stepping on one processor for reference purposes

Table 4 Speedup and efficiency of different MGRIT variants using various number of processors

p	8	16	32	64	126	256
<i>Two-level cycles with F-relax</i>						
Speedup	1.22	1.90	2.80	3.59	4.39	4.93
Efficiency	15.22%	11.85%	8.74%	5.60%	3.43%	1.93%
<i>Five-level V-cycle with FCF-relax.</i>						
Speedup	0.96	1.74	3.22	5.43	9.19	14.02
Efficiency	11.94%	10.86%	10.07%	8.48%	7.18%	5.48%
<i>Five-level F-cycle with FCF-relax.</i>						
Speedup	0.89	1.60	2.86	4.54	6.92	9.36
Efficiency	11.14%	10.02%	8.93%	7.10%	5.41%	3.66%
<i>Five-level V-cycle with FCF-relax. and SC</i>						
Speedup	1.04	1.90	3.62	6.40	12.01	21.93
Efficiency	12.96%	11.86%	11.31%	9.99%	9.44%	8.57%
<i>Five-level F-cycle with FCF-relax. and SC</i>						
Speedup	1.02	1.86	3.51	6.08	11.06	18.54
Efficiency	12.74%	11.63%	10.96%	9.50%	8.64%	7.24%

The speedup is given relative to the time-to-solution for the sequential time-stepping on one processor. Parallel efficiency is measured as S_p/p , where S_p is the speedup for p processors

to small or moderate numbers of processors, good parameter choices are crucial for obtaining the best possible speedup over sequential time stepping. Future work will focus on developing a model for optimizing MGRIT performance for given numbers of processors and time points.

Acknowledgements The work is supported by the Excellence Initiative of the German Federal and State Governments, the Graduate School of Computational Engineering at TU Darmstadt, and the BMBF (project PASIROM; grants 05M18RDA and 05M18PXB). The authors would like to thank Iryna Kulchytska-Ruchka for providing code for GetDP simulations of the “im_3_kw” model.

Funding Open Access funding enabled and organized by Projekt DEAL.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article’s Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article’s Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

1. Pyrhönen, J., Jokinen, T., Valéria, H.: Design of Rotating Electrical Machines. Wiley, New York (2008)
2. Jackson, J.D.: Classical Electrodynamics, 3rd edn. Wiley, New York (1999)
3. Takahashi, Y., Tokumasu, T., Fujiwara, K., Iwashita, T., Nakashima, H.: Parallel TP-EEC method based on phase conversion for time-periodic nonlinear magnetic field problems. *IEEE Trans. Magn.* **51**(3), 1 (2015)
4. Takahashi, Y., Fujiwara, K., Iwashita, T., Nakashima, H.: Parallel finite-element method based on space-time domain decomposition for magnetic field analysis of electric machines. *IEEE Trans. Magn.* **55**(6), 1–4 (2019)
5. Bermúdez, A., Gómez, D., Piñeiro, M., Salgado, P.: A novel numerical method for accelerating the computation of the steady-state in induction machines. *Comput. Math. Appl.* (2019). In Press
6. Lions, J.L., Maday, Y., Turinici, G.: Résolution d'EDP par un schéma en temps "pararéel". *C. R. l'Académie des Sci. Série I. Math.* **332**(7), 661–668 (2001)
7. Falgout, R.D., Friedhoff, S., Kolev, T.V., MacLachlan, S.P., Schroder, J.B.: Parallel time integration with multigrid. *SIAM J. Sci. Comput.* **36**(6), C635–C661 (2014)
8. Schöps, S., Niyonzima, I., Clemens, M.: Parallel-in-time simulation of eddy current problems using parareal. *IEEE Trans. Magn.* **54**(3), 1–4 (2018)
9. Gander, M.J., Kulchytska-Ruchka, I., Niyonzima, I., Schöps, S.: A new parareal algorithm for problems with discontinuous sources. *SIAM J. Sci. Comput.* **41**(2), B375–B395 (2019)
10. Bast, D., Kulchytska-Ruchka, I., Schöps, S., Rain, O.: Accelerated steady-state torque computation for induction machines using parallel-in-time algorithms. *IEEE Trans. Magn.* **56**(2), 1–9 (2019)
11. Friedhoff, S., Hahne, J., Schöps, S.: Multigrid-reduction-in-time for eddy current problems. *PAMM* **19**(1), e201900262 (2019)
12. Friedhoff, S., Hahne, J., Kulchytska-Ruchka, I., Schöps, S.: Exploring parallel-in-time approaches for eddy current problems. In: Faragó, I., Izsák, F., Simon, P.L. (eds.) *Progress in Industrial Mathematics at ECMI 2018*, The European Consortium for Mathematics in Industry, vol. 30. Springer, Berlin (2020)
13. Falgout, R.D., Friedhoff, S., Kolev, T.V., MacLachlan, S.P., Schroder, J.B., Vandewalle, S.: Multigrid methods with space-time concurrency. *Comput. Vis. Sci.* **18**(4), 123–143 (2017)
14. Henthenthaler, A., Nordsletten, D., Röhrle, O., Schroder, J.B., Falgout, R.D.: Convergence of the multigrid reduction in time algorithm for the linear elasticity equations. *Numer. Linear Algebra Appl.* **25**(3), e2155 (2018). 18
15. Lecouvez, M., Falgout, R.D., Woodward, C.S., Top, P.: A parallel multigrid reduction in time method for power systems. In: 2016 IEEE Power and Energy Society General Meeting (PESGM), pp. 1–5 (2016)
16. Schroder, J.B., Falgout, R.D., Woodward, C.S., Top, P., Lecouvez, M.: Parallel-in-time solution of power systems with scheduled events. In: 2018 IEEE Power Energy Society General Meeting (PESGM), pp. 1–5 (2018)
17. Ruprecht, D.: Convergence of parareal with spatial coarsening. *PAMM* **14**(1), 1031–1034 (2014)
18. Lunet, T., Bodart, J., Gratton, S., Vasseur, X.: Time-parallel simulation of the decay of homogeneous turbulence using parareal with spatial coarsening. *Comput. Vis. Sci.* **19**(1), 31–44 (2018)
19. Hackbusch, W.: Parabolic multi-grid methods. *Comput. Methods Appl. Sci. Eng.* **VI**, 189–197 (1984)
20. Vandewalle, S., Horton, G.: Multicomputer-Multigrid Solution of Parabolic Partial Differential Equations. In: Hemker, P., Wesseling, P. (eds.) *Multigrid Methods IV*, ISNM International Series of Numerical Mathematics, vol. 116, pp. 97–109. Birkhäuser, Basel (1994)
21. Gander, M.J., Neumüller, M.: Analysis of a new space-time parallel multigrid algorithm for parabolic problems. *SIAM J. Sci. Comput.* **38**(4), A2173–A2208 (2016)
22. Franco, S.A.R., Gaspar, F.J., Pinto, M.A.V., Rodrigo, C.: Multigrid method based on a space-time approach with standard coarsening for parabolic problems. *Appl. Math. Comput.* **317**(Supplement C), 25–34 (2018)
23. Emmett, M., Minion, M.L.: Toward an efficient parallel in time method for partial differential equations. *Commun. Appl. Math. Comput. Sci.* **7**, 105–132 (2012)
24. Minion, M.L., Speck, R., Bolten, M., Emmett, M., Ruprecht, D.: Interweaving PFASST and parallel multigrid. *SIAM J. Sci. Comput.* **37**, S244–S263 (2015)
25. Speck, R., Ruprecht, D., Emmett, M., Bolten, M., Krause, R.: A space-time parallel solver for the three-dimensional heat equation. In: *Parallel Computing: Accelerating Computational Science and Engineering (CSE)*, *Advances in Parallel Computing*, vol. 25, pp. 263–272. IOS Press (2014)
26. Ruprecht, D., Speck, R., Emmett, M., Bolten, M., Krause, R.: Poster: Extreme-scale space-time parallelism. In: *Proceedings of the 2013 Conference on High Performance Computing Networking, Storage and Analysis Companion, SC '13 Companion* (2013)
27. Falgout, R.D., Manteuffel, T.A., O'Neill, B., Schroder, J.B.: Multigrid reduction in time for nonlinear parabolic problems: a case study. *SIAM J. Sci. Comput.* **39**(5), S298–S322 (2017)
28. Howse, A., Sterck, H., Falgout, R., MacLachlan, S., Schroder, J.: Parallel-in-time multigrid with adaptive spatial coarsening for the linear advection and inviscid burgers equations. *SIAM J. Sci. Comput.* **41**, A538–A565 (2019)
29. Ries, M., Trottenberg, U., Winter, G.: A note on MGR methods. *Linear Algebra Appl.* **49**, 1–26 (1983)
30. Nielsen, A.S., Brunner, G., Hesthaven, J.S.: Communication-aware adaptive Parareal with application to a nonlinear hyperbolic system of partial differential equations. *J. Comput. Phys.* **371**, 483–505 (2018)
31. Falgout, R.D., Lecouvez, M., Woodward, C.S.: A parallel-in-time algorithm for variable step multistep methods. *J. Comput. Sci.* **37**, 101029 (2019). 12
32. Stüben, K.: An introduction to algebraic multigrid. *Multigrid*, pp. 413–523. Academic press, London (2001)
33. Brandt, A.: Multi-level adaptive solutions to boundary-value problems. *Math. Comput.* **31**(138), 333–390 (1977)
34. Falgout, R.D., Katz, A., Kolev, T.V., Schroder, J.B., Wissink, A., Yang, U.M.: Parallel time integration with multigrid reduction for a compressible fluid dynamics application. Technical report, llnl-jrnl-663416, Lawrence Livermore National Laboratory (2015)
35. Kronsjö, L., Dahlquist, G.: On the design of nested iterations for elliptic difference equations. *Nord. Tidskr. Informationsbehandling* **12**, 63–71 (1972)
36. Kronsjö, L.: A note on the "nested iterations" methods. *Nord. Tidskr. Informationsbehandling* **15**(1), 107–110 (1975)
37. Schöps, S., De Gersem, H., Weiland, T.: Winding functions in transient magnetoquasistatic field-circuit coupled simulations. *COMPEL Int. J. Comput. Math. Electr. Electron. Eng.* **32**, 2063–2083 (2013). <https://doi.org/10.1108/COMPEL-01-2013-0004>
38. Ferreira da Luz, M.V., Dular, P., Sadowski, N., Geuzaine, C., Bastos, J.P.A.: Analysis of a permanent magnet generator with dual formulations using periodicity conditions and moving band. *IEEE Trans. Magn.* **38**(2), 961–964 (2002)
39. Bartel, A., Baumanns, S., Schöps, S.: Structural analysis of electrical circuits including magnetoquasistatic devices. *Appl. Numer. Math.* **61**, 1257–1270 (2011)

40. Cortes Garcia, I., De Gersem, H., Schöps, S.: A structural analysis of field/circuit coupled problems based on a generalised circuit element. *Numer. Algorithms* **83**, 1–22 (2019)
41. Black, H.: *Modulation Theory*. Bell Telephone Laboratories Series. Van Nostrand, New York (1953)
42. Sun, J.: *Pulse-Width Modulation*, pp. 25–61. Springer, London (2012)
43. Gyselinck, J., Vandeveld, L., Melkebeek, J.: Multi-slice fe modeling of electrical machines with skewed slots-the skew discretization error. *IEEE Trans. Magn.* **37**, 3233–3237 (2001)
44. Geuzaine, C., Remacle, J.F.: Gmsh: A 3-d finite element mesh generator with built-in pre- and post-processing facilities. *Int. J. Numer. Methods Eng.* **79**(11), 1309–1331 (2009)
45. Gmsh: a three-dimensional finite element mesh generator with built-in pre- and post-processing facilities. <http://www.gmsh.info>
46. Hahne, J., Friedhoff, S.: PyMGRIT: Multigrid-reduction-in-time in Python v1.0 (2020). <https://github.com/pymgrit/pymgrit>. Release 1.0
47. Dular, P., Geuzaine, C., Henrotte, F., Legros, W.: A general environment for the treatment of discrete problems and its application to the finite element method. *IEEE Trans. Magn.* **34**(5), 3395–3398 (1998)
48. Geuzaine, C.: GetDP: a general finite-element solver for the de Rham complex. *PAMM* **7**(1), 1010603–1010604 (2007)
49. Getdp: A general environment for the treatment of discrete problems. <http://www.getdp.info>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.