

A parallel space-time multigrid method for the eddy-current equation

Martin Neumüller

Martin Schwalsberger

DK-Report No. 2019-09

11 2019

A-4040 LINZ, ALTENBERGERSTRASSE 69, AUSTRIA

Supported by

Austrian Science Fund (FWF)



FWF

Der Wissenschaftsfonds.

Upper Austria



Editorial Board: Bruno Buchberger
Evelyn Buckwar
Bert Jüttler
Ulrich Langer
Manuel Kauers
Peter Paule
Veronika Pillwein
Silviu Radu
Ronny Ramlau
Josef Schicho

Managing Editor: Diego Dominici

Communicated by: Ulrich Langer
Ronny Ramlau

DK sponsors:

- **Johannes Kepler University Linz (JKU)**
- **Austrian Science Fund (FWF)**
- **Upper Austria**

A parallel space-time multigrid method for the eddy-current equation¹

M. Neumüller² M. Schwalsberger³

October 30, 2019

Abstract

We expand the applicabilities and capabilities of an already existing space-time parallel method based on a block Jacobi smoother. First we formulate a more detailed criterion for spatial coarsening, which enables the method to deal with unstructured meshes and varying material parameters. Further we investigate the application to the eddy-current equation, where the non-trivial kernel of the curl operator causes severe problems. This is remedied with a new nodal auxiliary space correction. We proceed to identify convergence rates by local Fourier analysis and numerical experiments. Finally, we present a numerical experiment which demonstrates its excellent scaling properties.

1 Introduction

Around 2004 processor speed ceased to increase due to physical limits. Therefore we have experienced exponential growth in the number of processing units instead. This manifests in modern processing clusters featuring 1000 to more than ten million CPUs as of 2019⁴. Meanwhile applications that are parallel in space frequently fail to efficiently utilize more than a few thousand CPUs (depending on problem size). To utilize this processing power massively parallel algorithms are required. For time dependent simulations the time dimension offers potential for further parallelization, however the classical time-stepping concept is fundamentally sequential. Primarily for parabolic evolution equations time-parallel methods have been developed to break this "curse of sequentiality". After the first concept for ordinary differential equations more than 55 years ago [24], a whole family of time-parallel methods of all kinds has blossomed, see [9] for a historical overview.

Despite still being sequential in time the method [12] by Hackbusch laid some groundwork for time-parallel multigrid schemes and defied the notion of exactly solving the

¹The research was funded by the Austrian Science Fund (FWF) under the grant W1214-N15, project DK8.

²Herzogsdorf, Austria, E-mail: martin@neumuller.at

³Doctoral Program "Computational Mathematics", Johannes Kepler University Linz, Altenberger Straße 69, A-4040 Linz, Austria, E-mail: martin.schwalsberger@dk-compmath.jku.at

⁴<https://www.top500.org/>, accessed 2019

problem forward in time. Lubich and Ostermann took a different approach by presenting a waveform relaxation type space-time multigrid method [19]. An in depth discussion of waveform relaxation methods can be found in the monograph [38]. A further variant described as space-time concurrent multigrid waveform relaxation with cyclic reduction can be found in [16].

A contemporary method similar to ours is multigrid reduction in time [7] which has been actively expanded by a team centered around the Lawrence Livermore National Laboratory [7, 6, 8, 18, 11, 27]. Further results can be found on the XBraid project website⁵. Another very active group is the Jülich Forschungszentrum with diverse publications about time-parallel methods [33, 31, 34, 32, 36, 40, 26, 35, 20, 17, 25, 30, 3, 29, 4].

A method, which can be described as the precursor to the next method, is space-time multigrid with point-wise relaxation [15]. The method used in this paper [10] replaces the point-wise relaxation with block relaxation. Thus it can be described as space-time multigrid with block relaxation. Recent publications demonstrated that this method can also be applied to stochastic differential equations [23], and in combination with isogeometric analysis methods [14]. This motivated us to investigate the application of our method to the eddy-current equation. Furthermore we generalize concepts to expand our scope towards more complex problems.

The eddy-current equations are used to study low frequency electromagnetic applications, like electric motors/generators, induction furnaces, non-destructive testing, and many more. In practice, time-harmonic approaches are often preferred over transient models due to lower computational costs. However time-harmonic approaches struggle with nonlinear materials (saturation effects) and moving geometry. As a stepping stone to these more complex problems, we only consider the linear case with constant geometry in this paper.

The remainder of this paper is organized as follows. In Section 2 the Space-time-parallel method is introduced and the concept of a coarsening rule is presented and expanded. Section 3 applies the method to the eddy-current equation and pinpoints the core problem of the naive approach. We continue to determine a remedy to this problem - namely a "nodal auxiliary space correction"-, which can be used to define a hybrid smoother. In Section 4, we perform a local Fourier analysis of the introduced hybrid smoother for the 2D curl-curl equation. In Section 5, we mimic the local Fourier analysis with numerical experiments for the 3D eddy-current equation. We further investigate the validity of a new generalized coarsening rule. Weak scaling results of the improved method are given in Section 6. Conclusions are drawn in Section 7 and we discuss possible implications of our findings for related models and methods.

⁵github.com/XBraid/xbraid/wiki/Project-Publications

2 The Method

2.1 The Space-Time Parallel Method

For simplicity we reiterate the space-time parallel multigrid method presented in [10] and [28] only for the case of the implicit Euler scheme with uniform step width $\tau = T/m$. It should be noted, that we are not concerned with the approximation quality, but we only want to solve a given discretized problem. After the usual discretization steps for a PDE in space and time (see Chapter 3 for an example) we need to solve an equation of the form

$$-M_h \underline{u}_h^k + (M_h + \tau K_h) \underline{u}_h^{k+1} = \tau \underline{f}_h^{k+1}, \quad k = 0, \dots, m-1, \quad M_h \underline{u}_h^0 = \underline{u}_{0,h}. \quad (1)$$

By defining $A_{\tau,h} := M_h + \tau K_h$ we can rewrite the time-stepping scheme as a huge linear system

$$\begin{pmatrix} A_{\tau,h} & & & & \\ -M_h & A_{\tau,h} & & & \\ & & \ddots & & \\ & & & \ddots & \\ & & & & -M_h & A_{\tau,h} \end{pmatrix} \begin{pmatrix} u_h^1 \\ u_h^2 \\ \vdots \\ u_h^m \end{pmatrix} = \begin{pmatrix} \tau f_h^1 + u_{0,h} \\ \tau f_h^2 \\ \vdots \\ \tau f_h^m \end{pmatrix}.$$

This system can be compactly written as

$$L_{\tau,h} \underline{x} = \underline{f}.$$

The standard smoother in this method is a block-Jacobi preconditioned Richardson scheme with $\omega = 0.5$

$$\underline{x}^{k+1} = \underline{x}^k + \omega \underbrace{(I \otimes A_{\tau,h}^{-1})}_{=\text{diag}(A_{\tau,h}^{-1}, \dots, A_{\tau,h}^{-1})} (\underline{f} - L_{\tau,h} \underline{x}^k). \quad (2)$$

The overall method consists of pre-smoothing, a recursive coarse time-grid correction and post-smoothing. This V-cycle over the time meshes is guaranteed to geometrically converge, as long as the application of $A_{\tau,h}^{-1}$ is sufficiently well approximated and $M_h^{-1} K_h$ is positive semidefinite with real eigenvalues.

The presented smoother is simultaneously parallel in space and time. Because all operations are a Kronecker product of a time operation and a spatial operation, all spatial operations can be executed with the usual space-parallel approaches. On top of this also the time mesh can be subdivided among groups of processors. Only the application of $L_{\tau,h}$ requires forward communication for neighboring time nodes.

2.2 Spatial Coarsening

For performance reasons it is desirable to use not only a coarser time-mesh for the coarse grid correction, but also a coarser spatial mesh. For example without spatial

coarsening the coarse grid correction is responsible for approximately half the computational effort and is the less parallel part of the algorithm. In 3D the use of a coarser spatial mesh approximately reduces the effort for the coarse grid correction by a factor of 1/8 almost doubling the overall speed of the method.

However various inhibitors can prevent proper convergence rates for spatial coarsening. The "degree of anisotropy" in the discrete problem (denoted by λ , coined in [15]) is very significant for the behavior of the method. Assuming properly working spatial solvers the degree of anisotropy is indeed the only potential inhibitor for the heat equation. In fact if λ is above a certain threshold and no other inhibitor is active, then spatial coarsening has no negative effect on the convergence rate. Therefore we use spatial coarsening only if the *coarsening rule* is fulfilled

$$\lambda := \frac{\tau}{h^2} \geq \lambda_{crit},$$

with τ being the time step size and h the spatial mesh size. The threshold λ_{crit} depends on the used time-stepping scheme, the spatial discretization method, the smoothing steps, the dimension and the underlying differential equation. This rule needs to be generalized for practical applications to generic parabolic PDEs with a generic differential operator D_x

$$\alpha \frac{\partial u}{\partial t} + D_x^* \beta D_x u = f.$$

By a variable substitution argument we can account for the parameters α, β and by defining a local degree of anisotropy we can account for variation of the values over the spatial domain. The final coarsening rule is defined as

$$\lambda(\mathcal{T}) := \min_{x \in \mathcal{T}} \frac{\beta(x)\tau}{\alpha(x)h_{\mathcal{T}}^2} \geq \lambda_{crit} \quad \forall \mathcal{T} \in T_h(\Omega). \quad (3)$$

So for each element of the spatial mesh $T_h(\Omega)$ this condition needs to be satisfied for spatial coarsening. However in practice with piecewise constant parameters and other assumptions simplifications for checking this rule can be made as explained in [28]. To measure $h_{\mathcal{T}}$ we decide to use the length of the longest edge of the tetrahedron, whereas equivalent choices would be reflected in the value of λ_{crit} .

The determination of λ_{crit} in each case can be done either by a local Fourier analysis (LFA) or by numerical experiments. Both of these approaches are demonstrated in this paper.

3 Eddy-Currents

After the successful application of this scheme to the heat equation [10], we attempt to solve the eddy-current equation

$$\sigma(x) \frac{\partial \underline{u}}{\partial t} + \text{curl } \mu^{-1}(x) \text{curl } \underline{u} = \underline{f}, \quad \text{in } \Omega \times (0, T). \quad (4)$$

under the regularization assumption $\sigma(x) \geq 1[\frac{S}{m}]$ to avoid dealing with differential algebraic equations for now. We also always use homogeneous Dirichlet ($u \times n = 0$, on $\partial\Omega \times (0, T)$) or Neumann ($\mu^{-1}\text{curl}(u) \times n = 0$, on $\partial\Omega \times (0, T)$) boundary conditions and a homogeneous initial condition ($u = E_0$, for $t = 0$, $x \in \Omega$). Applying an implicit Euler scheme and lowest order Nedelec elements [21] to the semi-variational formulation

$$\frac{\partial}{\partial t} \int_{\Omega} \sigma \underline{u}(t) \cdot \underline{v} dx + \int_{\Omega} \mu^{-1} \text{curl } \underline{u}(t) \cdot \text{curl } \underline{v} dx = \int_{\Omega} \underline{f}(t) \cdot \underline{v} dx \quad \forall \underline{v} \in H(\Omega, \text{curl})$$

yields the discrete time-stepping scheme (1), where the occurring matrices are defined by the basis (ϕ_i) of the Nedelec finite elements

$$(M_h)_{i,j} := \int_{\Omega} \sigma \phi_i \cdot \phi_j dx, \quad (K_h)_{i,j} := \int_{\Omega} \mu^{-1} \text{curl } \phi_i \cdot \text{curl } \phi_j dx$$

3.1 Troubles with Spatial Coarsening

As $M_h^{-1}K_h$ is positive semidefinite and has real eigenvalues, it is not surprising, that the pure time multigrid method without spatial coarsening works reasonably well with convergence rates mostly below 0.5. For the idealized case of an exact solver for $A_{\tau,h}$ this rate is proven by inspecting the test equation (5) as in [22, ch.4.2]. However for the method as it is we come to the conclusion: *A coarse grid correction with spatial coarsening is completely ineffective regardless of λ .* This indicates another inhibitor for spatial coarsening, which will turn out to be the non-trivial kernel of the curl operator.

This puzzling finding made us question, if the implementation was really bug-free, while the explanation to this problem is in fact purely mathematical.

3.2 An Informal Explanation

This subsection explains the underlying mechanics in an intuitive, informal way. The basis for this explanation is the test equation

$$u' + \lambda u = f \quad \lambda \geq 0. \quad (5)$$

The case $\lambda = 0$ is the worst case for the convergence of our time parallel scheme assuming some fixed step size. For $\lambda \gg 1$ the smoother converges fast without a coarse grid correction. The semi-discretized eddy-current equation

$$\underline{u}' + M_h^{-1}K_h \underline{u} = \underline{f}$$

can be diagonalized by the Eigenvectors of $M_h^{-1}K_h$, which leads with to the test equation with the eigenvalues in place of λ . This implies that the smoother itself will dampen error components with high eigenvalues. As shown by the local Fourier

analysis for differential operators [37], high eigenvalues are associated with high spatial frequencies. Dampening those high spatial frequencies smooths the error, which enables spatial coarsening during the coarse grid correction.

For simply connected domains, the eddy-current equation can be viewed through the scope of the Helmholtz decomposition [39]

$$H(\Omega, \text{curl}) = \nabla H^1(\Omega) \oplus \text{curl } H(\Omega, \text{curl}).$$

By inserting the solenoidal part ($u_s \in \text{curl } H(\text{curl})$) into the eddy-current equation (4) without parameters and by using the identities $\text{curl curl} = \nabla \text{div} - \Delta$, $\text{div curl} = 0$ we get the resulting equation

$$\frac{\partial u_s}{\partial t} - \Delta u_s = f_s.$$

Analogously inserting the nodal part ($u_n \in \nabla H^1$) and using $\text{curl} \nabla = 0$ results in

$$\frac{\partial u_n}{\partial t} = f_n. \quad (6)$$

We can conclude from this, that the solenoidal part behaves like a heat equation and should be handled well by our method. In contrast the nodal part is equivalent to the worst case of the test equation 5. As a result high frequencies of the gradient part are practically not dampened by the smoother, thus there is no spatial smoothing at all, such that coarsening in space renders the coarse grid correction useless.

However equation (6) can be solved by just integrating the right hand side. Conceptually the error in the nodal component can be calculated by the operator

$$\int_0^s dt \nabla_x \Delta_x^{-1} \text{div}_x (f - \frac{\partial u}{\partial t}). \quad (7)$$

The interpretation of this procedure is as follows.

- $f - \frac{\partial u}{\partial t}$: Calculate the residual.
- div_x : Eliminate the solenoidal part.
- $\nabla_x \Delta_x^{-1}$: Revert div_x for the nodal part.
- $\int_0^s dt$: Integrate the equation.

This can be seen as an expansion of the method [13] with an additional time component. In this notation the nodal auxiliary space correction by R. Hiptmayr would be denoted by

$$\nabla(\text{div } \kappa \nabla)^{-1} \text{div}(f - (\kappa I + \text{curl } \mu^{-1} \text{curl})u).$$

3.3 Nodal Auxiliary Space Correction

With the previous insights in mind we can properly implement the nodal auxiliary correction method. If Nedelec elements of any order are used, an appropriate $H^1(\Omega)$ conforming finite element space can be chosen according to [42]. We define G as the discrete gradient matrix from the $H^1(\Omega)$ FE-space to the Nedelec FE-space and K_n as the nodal Laplace stiffness matrix for the operator $\text{div } \sigma \nabla$. If we also define Σ_t as the operator, which sums up all previous and the current time step

$$(\Sigma_t \underline{y})_i := \sum_{j=1}^i y_j,$$

then the correction can be expressed as

$$\underline{x} \leftarrow \underline{x} + G \Sigma_t K_n^{-1} G^T (\underline{f} - L_{\tau,h} \underline{x}). \quad (8)$$

In our notation the spatial operators above are applied to each time step separately, so we can omit the usual tensor products ($I \otimes \dots$) for more clarity. Furthermore through a time-multigrid approach the already trivial operator Σ_t could be made fully time parallel.

4 Hybrid Smoother

The aforementioned auxiliary nodal correction (8) denoted by A can be combined with the standard smoother (2) denoted by S to form a hybrid smoother. Like matrices the application is read from right to left, so SA is one application of A followed by an application of S . We restrict ourselves to have the same smoother during pre- and post smoothing with reversed order. The following local Fourier analysis indicates, that more than one application of A in the hybrid smoother has no apparent benefits, therefore we restrict ourselves to using A just once.

4.1 Local Fourier Analysis

The local Fourier analysis is based on the LFA for the 1D heat equation performed in [10] and extended with LFA symbols for the 2D curl-curl equation from [5]. The implementation is done in Mathematica [41]. The full methodology and more detailed results can be found in [28, chp. 5]. Additionally the nodal auxiliary correction is implemented. In the LFA we assume exact solutions to the spatial problems (implicit Euler step and Laplace problem), however the numerical experiments in chapter 5 show, that this is not required. Further we assume, that the coarse grid correction is solved exactly.

In Figure 1 we observe, how the addition of the auxiliary nodal correction introduces spatial smoothing, whereas the standard smoother features no spatial smoothing at

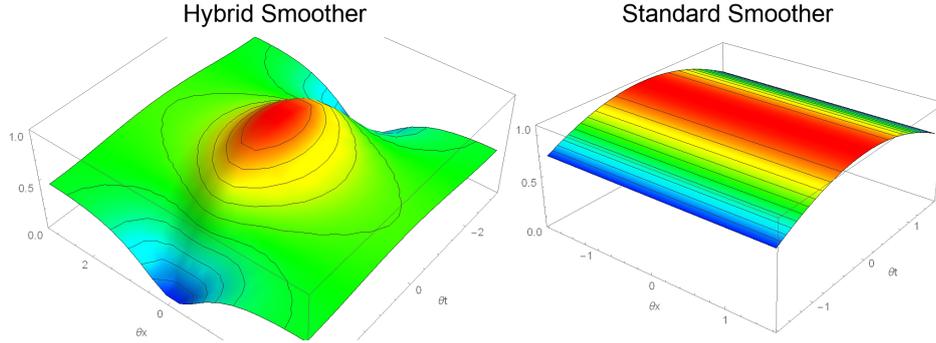


Figure 1: Comparison between the hybrid smoother SAS and the standard smoother S . Spectral radius for each frequency is plotted with $\theta_y = 0$, $\lambda = 1$. Only the hybrid smoother has spatial smoothing properties.

all regardless of λ (3). This reaffirms our observations in Section 3.2 .

Consequently we measure the maximal spectral radius over all frequencies to get an upper bound for convergence rates. The three investigated methods use either coarsening in space ("Semi-Space"), in time ("Semi-Time"), or both ("Full"). To relate to the coarsening rule, we vary λ over the horizontal axis. The results in Figure 2 confirm, that with the auxiliary smoother in place using a coarser spatial mesh is possible again. Furthermore we can observe, that the "Full" method behaves identical to the "Semi-Time" method for big enough λ and deteriorates for lower λ . Therefore it makes sense to use the coarsening rule (3) with $\lambda_{crit} \approx 0.9$. Not only does the addition of A allow spatial coarsening, but it also improves the convergence rates for "Semi-Time", which would otherwise converge with a worse rate independent of λ .

Different orders like SSA or ASS typically perform slightly worse than SAS , therefore we prefer putting A in between S . The convergence rates of SA or AS squared are approximately the same as for SAS , which means that two V-cycles with SA have the same impact as one V-cycle with SAS . This means, that in this setting SA is strictly inferior to SAS . On the other end $SSASS$ does not manage to square the convergence rate compared to SAS (0.2 vs. 0.3 for low λ) except for $\lambda > 1000$. In the case of low λ we use "Semi-Time", so the coarse grid correction is responsible for half the computational effort and $0.3^{1.337} = 0.2$, meaning one V-cycle of $SSASS$ is worth $4/3$ V-cycles of SAS . If $SSASS$ was used only on the finest grid, then this would mean approximately $(5/3 + 1)/2 = 4/3$ times the computational effort. This idealized rough estimate could set $SSASS$ on par with SAS in certain cases, maybe $SSAS$ could also be a good choice. In conclusion we will focus on SAS as a solid choice.

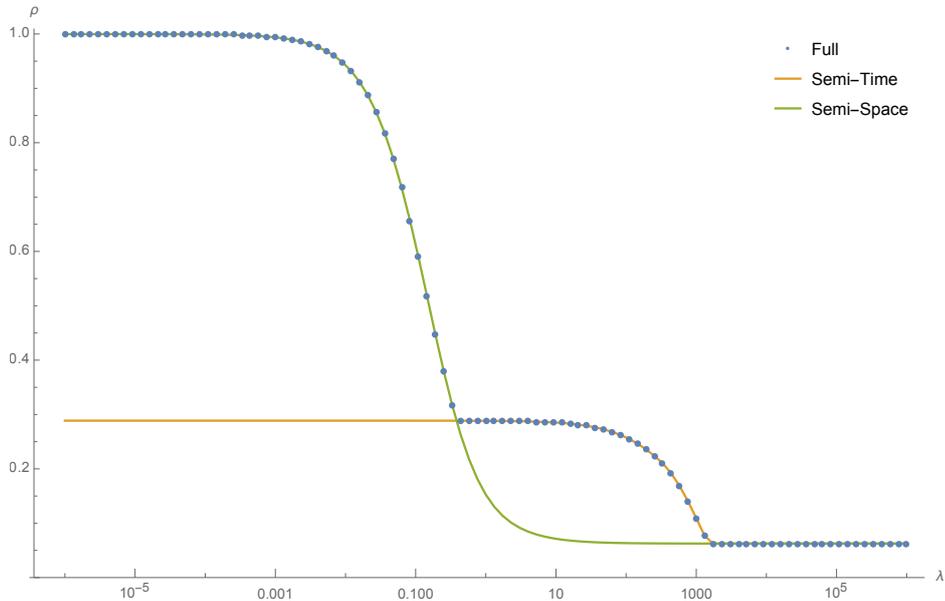


Figure 2: Convergence rates for the two-grid cycle with the *SAS* smoother as a function of λ .

5 Numerical Experiments

Our practical implementation is based on the implementation [10] and uses MFEM [1] and HYPRE [2] to deal with the arising spatial problems. The implementation is parallel in time, while everything except the auxiliary method is parallel in space, which was only prevented by time constraints. Therefore it is close to being parallel in space and time.

The modification of the existing multigrid framework to support a different spatial problem was seamless due to the non-intrusive design of the method. Specifically only the application of spatial prolongation/restriction, the discrete gradient matrix and functionality for A_τ , M_h , K_n needs to be provided through an interface. For example swapping out the eddy-current equation for a heat equation is a matter of minutes and matrix-free methods could be incorporated as well. The multigrid method itself fully relies on the interface for applying and solving spatial operators. The only hard-coded part is the used vector class, which could also be replaced by a template-based approach. There are other methods which aim for even less intrusiveness [REF reduced integration], but we are confident, that for many preexisting applications the required interface could be implemented with ease.

Space/Time level	9	8	7	6	5	4	3	2	1	0
2	x									
1		x	x	x	x	x	x	x	x	x
0										

Table 1: Coarsening strategy example for the "Full" method test case.

5.1 Determining Convergence Rates and the Parameter λ_{crit}

As we are not aware of any available local Fourier analysis for the 3D curl-curl equation, we recreate Figure 2 by using our implementation of the multigrid method with a *SAS* smoother. In addition we want to showcase how a fitting λ_{crit} can be determined with nothing but a working implementation. To determine convergence rates the power iteration method was applied to a test case. This means, that we apply up to 200 iterations of the method to a random initial guess (values between 0 and 1) with a 0 right hand side. During this process we measure the reduction of the ℓ^2 norm of the vector, which converges to the spectral radius of the method. The maximal occurring error reduction factor is considered as convergence rate, because these factors are almost monotonically increasing. In turn means, that our method has no worse convergence during a pre-asymptotic phase. Because we use no rescaling, the iteration is stopped, if the total error reduction gets below 10^{-100} to avoid a double underflow.

We compare the time multigrid method without coarsening in space ("Semi-time") with a variant, which uses just two levels of the spatial mesh and uses full coarsening on the highest level ("Full") as in Table 1. The reason is that the local Fourier analysis assumes the coarse grid correction to be solved exactly (two-grid cycle), so we mimic this by not using further spatial coarsening within the correction. For "Full" we want to use the spatial coarsening right on the highest level, because potentially bad convergence rates on lower levels are obscured on the highest level, where the convergence rates are measured.

The basic mesh consists of two tetrahedra, which are uniformly refined several times. The problem features homogeneous Neumann boundary conditions. We compare the results for 2, 3, 4, 5, 6 uniform refinements, so we can observe convergent behavior with an increasingly fine mesh. The spatial degrees of freedom range from 230 to 630240, with 512 time steps. The material parameters σ, μ are set to 1 and λ is varied by using different values for τ .

For solving the spatial problems $A_{\tau,h}, K_n$ we rely on 3 iterations of the preconditioned conjugate gradient method. The preconditioners are algebraic multigrid methods from HYPRE [2]. BoomerAMG was used for the Laplace problem and AMS for $A_{\tau,h}$. Quite remarkably, turning off the gradient-smoothing in AMS has no impact on convergence, because the gradient part is handled by the auxiliary correction. At the same time

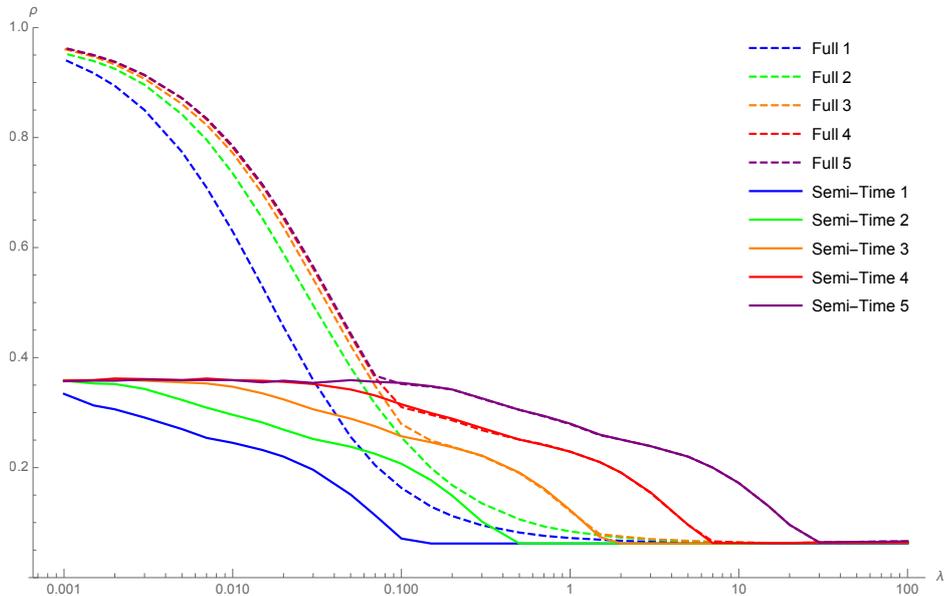


Figure 3: Convergence rates for the two-grid cycle with *SAS* smoother as a function of λ . Comparison between "Full" and "Semi-time" for different numbers of uniform refinements.

this cuts down total computation times by roughly 30%.

Our main goal is to confirm, that the method behaves similar to the local Fourier analysis for the 2D curl-curl equation. Furthermore we want to determine the value of λ_{crit} in this setting by observing, where the two methods "Full" and "Semi-time" start to differ.

In Figure 3 we can observe that the important parts of the curves converge with an increasing number of refinements. The only non-convergent feature is the transition from the high λ region with a convergence rate of 2^{-4} to the plateau of the "Semi-time" method with a convergence rate of 0.36. This means, that our method performs *better than expected* in this regard. On a side note the smoother on its own shifts to the right in the same fashion and is a suitable solver for sufficiently high λ and small problems. In the case of the heat equation the smoother never converges on its own and none of the shifting happens.

Figure 4 omits all cases except the most refined one for the sake of a less cluttered plot. The point where the "Full" method departs from the "Semi-time" method can be located at 0.1. When the diagonal line is extrapolated, it even appears to be around 0.075. This means that a choice of $\lambda_{crit} = 0.1$ will safely guarantee the optimal coarsening strategy to be selected. This holds for the case of the 3D eddy-current equation, lowest order Nedelec elements, implicit Euler time stepping, *SAS* smoother,

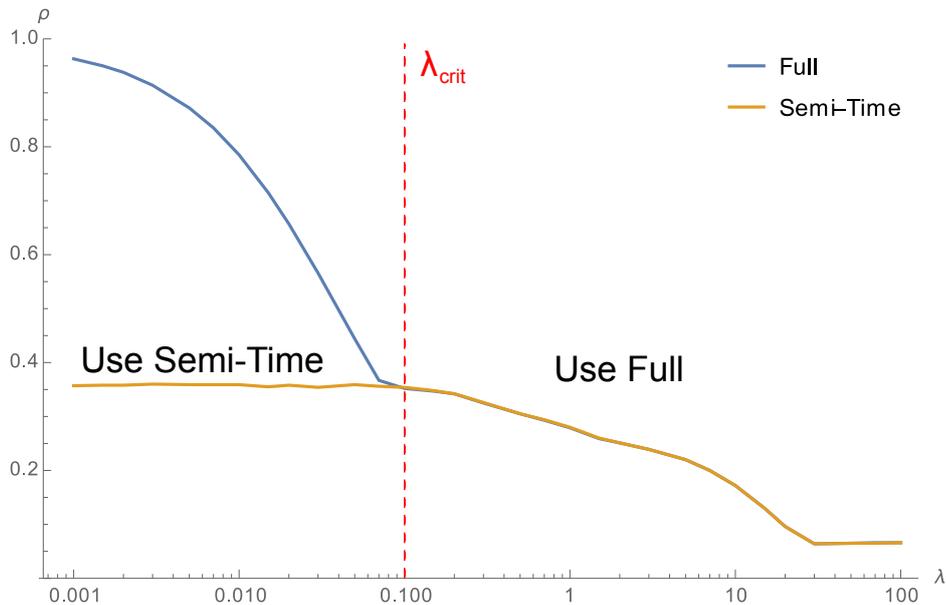


Figure 4: Convergence rates for the two-grid cycle with *SAS* smoother as a function of λ . Only the most refined cases with 630240 DoF in space. λ_{crit} determined as 0.1.

h = longest edge.

5.2 Coarsening Rule for Locally Different Parameters

In practical applications vastly different material parameters are very common. For example the ratio $\frac{\mu^{-1}}{\sigma}$ can differ by a factor of ≈ 30 between iron and copper, while between iron and weak conductors this gap can be way bigger. The limit case of perfect insulators ($\sigma = 0$) would even cause an infinitely large local degree of anisotropy λ . This motivates the question, if the local nature of the coarsening rule (3) can indeed handle different materials. On a side note we are not concerned, if the preconditioner can handle jumping coefficients, because our space-time parallel method can use any suitable method to solve the spatial problems.

We assign two different material parameters to the two original tetrahedra of the test mesh. After uniform refinements two regions with two distinct degrees of anisotropy $\lambda_1 < \lambda_2$ are given. As our method mostly benefits from increasing λ , it is sufficient to only deal with the smallest λ as reflected in the coarsening rule (3).

In contrast the method in [15] chooses to use semi coarsening in time if $\lambda > \lambda_{crit}$ and semi coarsening in space if $\lambda < \lambda_{crit}$. If this rule is violated both semi coarsening in time (when $\lambda < 2^{-4}\lambda_{crit}$) and semi coarsening in space (when $\lambda > 2^4\lambda_{crit}$) yield bad convergence rates (> 0.8). Therefore we are concerned, that in the case of two materi-

$\mu_2^{-1} \setminus \lambda$	2^0	2^1	2^2	2^3	2^4	2^5	2^6	0.01
1	0.121	0.19	0.23	0.252	0.372	0.533	0.688	0.77
10^1	0.083	0.159	0.213	0.242	0.34	0.498	0.662	0.753
10^2	0.083	0.158	0.212	0.242	0.34	0.5	0.664	0.755
10^3	0.083	0.158	0.212	0.242	0.34	0.5	0.665	0.756
10^4	0.083	0.158	0.212	0.243	0.34	0.501	0.665	0.756
10^5	0.083	0.157	0.212	0.243	0.34	0.501	0.665	0.756
10^6	0.083	0.157	0.213	0.242	0.34	0.501	0.665	0.756

Table 2: Maximum error reduction factor for the "Two λ " test case using the "Full" method. 10744 DoF in space. After an initial boost to convergence, the rates remain stable for increasing μ_2^{-1} . Again λ_{crit} appears to be located between 0.06125 and 0.125.

als with $\lambda_1 \leq 2^{-5}$, $\lambda_2 \geq 2^3$ neither coarsening approach would yield useful convergence rates. Vastly different λ_i can also arise during adaptive refinements, where the size of the elements can differ significantly.

For this "Two λ " test case we use mostly the same setup as for the previous test case, but we only investigate the "Full" method, only use 4 uniform refinements, have a maximum of 100 iterations and vary $\mu^{-1} \geq 1$ on one of the two subdomains. This has no effect on the minimal λ , as it only increases λ_2 . The minimal λ is again manipulated through changing τ .

We can observe in Table 2, that only the minimal λ has a notable impact on convergence, while increasing the bigger λ_2 can only improve the convergence rate. With this and previous results [28] our method is one step closer to real applications with varying materials.

6 Parallel Performance

To investigate the parallel performance of the method, we perform a weak scaling test with respect to time. In contrast to other weak scaling tests we expand the time domain instead of refining it. This is done to keep the degree of anisotropy constant at $\lambda = 0.08$, which allows for instant spatial coarsening. For comparison we also run a test series without spatial coarsening.

- Spatial domain $(0, 1)^3$, 7768 DoF
- 2^k time-parallel processors
- Time domain $32 * 2^k$ time steps with size 0.016
- \Rightarrow Constant $\lambda = 0.512$

- $\lambda_{crit} = 0.15$
- Material parameters = 1
- Residual reduction tolerance 10^{-10}
- SAS smoother

The "Full" method always finished in 12 iterations, while the "Semi-Time" method needed 14 iterations. Without some of the idealizations assumed in the analysis it is very well possible, that the "Full" method has a marginally better convergence rate than the "Semi-Time" method. The corresponding run-times can be observed in Table 3. The Adjusted Ratio is the ratio of the runtimes multiplied by 14/12 to adjust for the iterations.

#Processors	1	2	4	8	16	32	64	128	256	512
Full	35.3	35.3	35.4	39.3	49.8	49.7	50.2	50.2	50.4	52.1
Semi-Time	74.6	75.8	77.4	87.6	112.3	114.1	115.6	117.6	118.9	120.8
Adjusted Ratio	0.552	0.543	0.534	0.523	0.517	0.508	0.507	0.498	0.495	0.503

Table 3: Weak scaling results in seconds for the "Full" and the "Semi-Time" method. Adjusted ratio is Full/Semi-Time*14/12

We can clearly observe how the "Full" method scales better with barely any increase in computation times in the upper range. The sudden increase around the 16 cores mark can be most likely attributed to the architecture of the cluster with a node size of 16 cores.

It would have been interesting to run the scaling tests on even bigger clusters, where possibly a more severe discrepancy between the two methods would surface.

7 Conclusion

This paper demonstrates the difficulties in solving the eddy-current equation with a space-time multigrid method using also spatial coarsening. The non-trivial kernel of the curl-operator can be identified as the cause of these problems, and a nodal auxiliary space correction with time integration is devised to restore good convergence rates. This procedure is analogous to the method in [13], but with an additional time component.

The local Fourier analysis of the original and the modified method confirms the effectiveness of the auxiliary correction, and in fact lead the authors of this paper to devise the auxiliary correction. From the gained insights, we expect similar problems to occur whenever the spatial operator features a non-trivial kernel. Furthermore,

we expect other space-time multigrid methods to experience similar problems and to profit from an implementation of the described universal correction (7). Even multigrid methods without spatial coarsening could benefit from this modification, because it also improves the convergence rates of our pure time-multigrid method.

We proceed to determine convergence rates and $\lambda_{crit} = 0.1$ with numerical tests. This result holds for the 3D eddy-current equation and one specific configuration of the method. The coarsening rule passes one basic test for robustness, which involves solving test problems with vastly different local degrees of anisotropy λ_i .

We finally demonstrate with the help of weak scaling tests, how spatial coarsening can improve parallel performance.

In conclusion the explored "nodal auxiliary space correction" is a highly useful tool to enhance space-time multigrid methods for the eddy-current equation.

Acknowledgment

The research was funded by the Austrian Science Fund (FWF) under the grant W1214-N15, project DK8. We further want to thank the RICAM institute for enabling computations on the RADON1 cluster and the people associated with the NUMA institute for helpful comments and insights. We want to especially thank Ulrich Langer for fruitful discussions. Thank also goes to Bianca Klammer for additional proofreading.

References

- [1] MFEM: Modular finite element methods library. mfem.org.
- [2] HYPRE: Scalable linear solvers and multigrid methods. <https://computing.llnl.gov/projects/hypre-scalable-linear-solvers-multigrid-methods>, 2018.
- [3] BOLTEN, M., MOSER, D., AND SPECK, R. A multigrid perspective on the parallel full approximation scheme in space and time. *Numer. Linear Algebra Appl.* *24*, 6 (2017), e2110, 24.
- [4] BOLTEN, M., MOSER, D., AND SPECK, R. Asymptotic convergence of the parallel full approximation scheme in space and time for linear problems. *Numer. Linear Algebra Appl.* *25*, 6 (2018), e2208, 22.
- [5] BOONEN, T., VAN LENT, J., AND VANDEWALLE, S. Local Fourier analysis of multigrid for the curl-curl equation. *SIAM J. Sci. Comput.* *30*, 4 (2008), 1730–1755.
- [6] DOBREV, V. A., KOLEV, T., PETERSSON, N. A., AND SCHRODER, J. B. Two-level convergence theory for multigrid reduction in time (MGRIT). *SIAM J. Sci. Comput.* *39*, 5 (2017), S501–S527.
- [7] FALGOUT, R. D., FRIEDHOFF, S., KOLEV, T. V., MACLACHLAN, S. P., AND SCHRODER, J. B. Parallel time integration with multigrid. *SIAM J. Sci. Comput.* *36*, 6 (2014), C635–C661.
- [8] FALGOUT, R. D., MANTEUFFEL, T. A., O’NEILL, B., AND SCHRODER, J. B. Multigrid reduction in time for nonlinear parabolic problems: a case study. *SIAM J. Sci. Comput.* *39*, 5 (2017), S298–S322.
- [9] GANDER, M. J. 50 years of time parallel time integration. In *Multiple shooting and time domain decomposition methods*, vol. 9 of *Contrib. Math. Comput. Sci.* Springer, Cham, 2015, pp. 69–113.
- [10] GANDER, M. J., AND NEUMÜLLER, M. Analysis of a new space-time parallel multigrid algorithm for parabolic problems. *SIAM J. Sci. Comput.* *38*, 4 (2016), A2173–A2208.
- [11] GÜNTHER, S., GAUGER, N. R., AND SCHRODER, J. B. A non-intrusive parallel-in-time adjoint solver with the XBraid library. *Comput. Vis. Sci.* *19*, 3-4 (2018), 85–95.
- [12] HACKBUSCH, W. Parabolic multigrid methods. In *Computing methods in applied sciences and engineering, VI (Versailles, 1983)*. North-Holland, Amsterdam, 1984, pp. 189–197.

- [13] HIPTMAIR, R. Multigrid method for Maxwell’s equations. *SIAM J. Numer. Anal.* 36, 1 (1999), 204–225.
- [14] HOFER, C., LANGER, U., NEUMÜLLER, M., AND TOULOPOULOS, I. Time-multipatch discontinuous Galerkin space-time isogeometric analysis of parabolic evolution problems. *Electron. Trans. Numer. Anal.* 49 (2018), 126–150.
- [15] HORTON, G., AND VANDEWALLE, S. A space-time multigrid method for parabolic partial differential equations. *SIAM J. Sci. Comput.* 16, 4 (1995), 848–864.
- [16] HORTON, G., VANDEWALLE, S., AND WORLEY, P. An algorithm with polylog parallel complexity for solving parabolic partial differential equations. *SIAM J. Sci. Comput.* 16, 3 (1995), 531–541.
- [17] KREIENBUEHL, A., NAEGEL, A., RUPRECHT, D., SPECK, R., WITTUM, G., AND KRAUSE, R. Numerical simulation of skin transport using Parareal. *Comput. Vis. Sci.* 17, 2 (2015), 99–108.
- [18] LECOUCVEZ, M., FALGOUT, R. D., WOODWARD, C. S., AND TOP, P. A parallel multigrid reduction in time method for power systems. In *2016 IEEE Power and Energy Society General Meeting (PESGM)* (July 2016), pp. 1–5.
- [19] LUBICH, C., AND OSTERMANN, A. Multigrid dynamic iteration for parabolic equations. *BIT* 27, 2 (1987), 216–234.
- [20] MINION, M. L., SPECK, R., BOLTEN, M., EMMETT, M., AND RUPRECHT, D. Interweaving PFASST and parallel multigrid. *SIAM J. Sci. Comput.* 37, 5 (2015), S244–S263.
- [21] MONK, P. *Finite element methods for Maxwell’s equations*. Numerical Mathematics and Scientific Computation. Oxford University Press, New York, 2003.
- [22] NEUMÜLLER, M. *Space-Time Methods: Fast Solvers and Applications*. PhD thesis, Graz University of Technology, 2013.
- [23] NEUMÜLLER, M., AND THALHAMMER, A. A fully parallelizable space-time multilevel Monte Carlo method for stochastic differential equations with additive noise. *SIAM J. Sci. Comput.* 40, 3 (2018), C388–C400.
- [24] NIEVERGELT, J. Parallel methods for integrating ordinary differential equations. *Comm. ACM* 7 (1964), 731–733.
- [25] RUPRECHT, D., AND SPECK, R. Spectral deferred corrections with fast-wave slow-wave splitting. *SIAM J. Sci. Comput.* 38, 4 (2016), A2535–A2557.

- [26] RUPRECHT, D., SPECK, R., AND KRAUSE, R. Parareal for diffusion problems with space- and time-dependent coefficients. In *Domain Decomposition Methods in Science and Engineering XXII* (Sep 2015), vol. 104 of *Lecture Notes in Computational Science and Engineering*, 22nd International Conference on Domain Decomposition Methods, Lugano (Switzerland), 16 Sep 2013 - 20 Sep 2013, Springer International Publishing Switzerland, pp. 3–10.
- [27] SCHRODER, J. B., FALGOUT, R. D., WOODWARD, C. S., TOP, P., AND LECOUCVEZ, M. Parallel-in-time solution of power systems with scheduled events. In *2018 IEEE Power Energy Society General Meeting (PESGM)* (Aug 2018), pp. 1–5.
- [28] SCHWALSBERGER, M. A space-time parallel multigrid method for the transient eddy-current equation. Master’s thesis, JKU Linz, 2018. <https://www.numa.uni-linz.ac.at/Teaching/Diplom/Finished/schwalsberger>.
- [29] SPECK, R. Parallelizing spectral deferred corrections across the method. *Comput. Vis. Sci.* 19, 3-4 (2018), 75–83.
- [30] SPECK, R., AND RUPRECHT, D. Toward fault-tolerant parallel-in-time integration with PFASST. *Parallel Comput.* 62 (2017), 20–37.
- [31] SPECK, R., RUPRECHT, D., EMMETT, M., MINION, M., BOLTEN, M., AND KRAUSE, R. A space-time parallel solver for the three-dimensional heat equation. In *Parallel Computing: Accelerating Computational Science and Engineering (CSE)* (Sep 2014), vol. 25 of *Advances in Parallel Computing*, International Conference on Parallel Computing, Munich (Germany), 10 Sep 2013 - 13 Sep 2013, IOS Press, pp. 263 – 272.
- [32] SPECK, R., RUPRECHT, D., EMMETT, M., MINION, M., BOLTEN, M., AND KRAUSE, R. A multi-level spectral deferred correction method. *BIT* 55, 3 (2015), 843–867.
- [33] SPECK, R., RUPRECHT, D., KRAUSE, R., EMMETT, M., MINION, M., WINKEL, M., AND GIBBON, P. A massively space-time parallel n-body solver. In *SC ’12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis* (Nov 2012), pp. 1–11.
- [34] SPECK, R., RUPRECHT, D., KRAUSE, R., EMMETT, M., MINION, M., WINKEL, M., AND GIBBON, P. Integrating an N-Body Problem with SDC and PFASST. vol. 98 of *Lecture Notes in Computational Science and Engineering*, Domain Decomposition Methods in Science and Engineering XXI, Rennes (France), 25 Jun 2012 - 29 Jun 2012, Springer International Publishing, pp. 637 – 645.

- [35] SPECK, R., RUPRECHT, D., MINION, M., EMMETT, M., AND KRAUSE, R. Inexact spectral deferred corrections. In *Domain Decomposition Methods in Science and Engineering XXII* (Sep 2015), vol. 104 of *Lecture Notes in Computational Science and Engineering*, 22nd International Conference on Domain Decomposition Methods, Lugano (Switzerland), 16 Sep 2013 - 20 Sep 2013, Springer International Publishing Switzerland, pp. 127–133.
- [36] STEINER, J., RUPRECHT, D., SPECK, R., AND KRAUSE, R. Convergence of Parareal for the Navier-Stokes Equations Depending on the Reynolds Number. In *Numerical Mathematics and Advanced Applications - ENUMATH 2013* (Cham, Aug 2015), vol. 103 of *Lecture Notes in Computational Science and Engineering*, European Numerical Mathematics and Advanced Applications, Lausanne (Switzerland), 26 Aug 2013 - 30 Aug 2013, Springer International Publishing, pp. 195 – 202.
- [37] TROTTEBERG, U., OOSTERLEE, C. W., AND SCHÜLLER, A. *Multigrid*. Academic Press, Inc., San Diego, CA, 2001. With contributions by A. Brandt, P. Oswald and K. Stüben.
- [38] VANDEWALLE, S. *Parallel multigrid waveform relaxation for parabolic problems*. Teubner Skripten zur Numerik. [Teubner Scripts on Numerical Mathematics]. B. G. Teubner, Stuttgart, 1993.
- [39] VON HELMHOLTZ, H. Über Integrale der hydrodynamischen Gleichungen, welche den Wirbelbewegungen entsprechen. *J. Reine Angew. Math.* 55, 3 (1858), 25–55.
- [40] WINKEL, M., SPECK, R., AND RUPRECHT, D. A high-order Boris integrator. *J. Comput. Phys.* 295 (2015), 456–474.
- [41] WOLFRAM RESEARCH, INC. Mathematica, Version 12.0. Champaign, IL, 2019.
- [42] ZAGLMAYR, S. *High Order Finite Element Methods for Electromagnetic Field Computation*. PhD thesis, JKU Linz, 2006. <https://www.numa.uni-linz.ac.at/Teaching/PhD/Finished/zaglmayr.de.shtml>.

Technical Reports of the Doctoral Program

“Computational Mathematics”

2019

- 2019-01** A. Seiler, B. Jüttler: *Approximately C^1 -smooth Isogeometric Functions on Two-Patch Domains* Jan 2019. Eds.: J. Schicho, U. Langer
- 2019-02** A. Jiménez-Pastor, V. Pillwein, M.F. Singer: *Some structural results on D^n -finite functions* Feb 2019. Eds.: M. Kauers, P. Paule
- 2019-03** U. Langer, A. Schafelner: *Space-Time Finite Element Methods for Parabolic Evolution Problems with Non-smooth Solutions* March 2019. Eds.: B. Jüttler, V. Pillwein
- 2019-04** D. Dominici, F. Marcellán: *Discrete semiclassical orthogonal polynomials of class 2* April 2019. Eds.: P. Paule, V. Pillwein
- 2019-05** D. Dominici, V. Pillwein: *A sequence of polynomials generated by a Kapteyn series of the second kind* May 2019. Eds.: P. Paule, J. Schicho
- 2019-06** D. Dominici: *Mehler-Heine type formulas for the Krawtchouk polynomials* June 2019. Eds.: P. Paule, M. Kauers
- 2019-07** M. Barkatou, A. Jiménez-Paster: *Linearizing Differential Equations Riccati Solutions as D^n -Finite Functions* June 2019. Eds.: P. Paule, M. Kauers
- 2019-08** D. Dominici: *Recurrence coefficients of Toda-type orthogonal polynomials I. Asymptotic analysis.* July 2019. Eds.: P. Paule, M. Kauers
- 2019-09** M. Neumüller, M. Schwalsberger: *A parallel space-time multigrid method for the eddy-current equation* Nov 2019. Eds.: U. Langer, R. Ramlau

Doctoral Program

“Computational Mathematics”

Director:

Dr. Veronika Pillwein
Research Institute for Symbolic Computation

Deputy Director:

Prof. Dr. Bert Jüttler
Institute of Applied Geometry

Address:

Johannes Kepler University Linz
Doctoral Program “Computational Mathematics”
Altenbergerstr. 69
A-4040 Linz
Austria
Tel.: ++43 732-2468-6840

E-Mail:

office@dk-compmath.jku.at

Homepage:

<http://www.dk-compmath.jku.at>