No. 668 **August 2023**

Numerical Analysis of a Time-Simultaneous Multigrid Solver for Stabilized Convection-Dominated Transport Problems in 1D

W. Drews, S. Turek, C. Lohmann

ISSN: 2190-1767

Numerical Analysis of a Time-Simultaneous Multigrid Solver for Stabilized Convection-Dominated Transport Problems in 1D

Wiebke Drews*, Stefan Turek, Christoph Lohmann

Institute of Applied Mathematics (LSIII), TU Dortmund University,
Vogelpothsweg 87, D-44227 Dortmund, Germany

* *Corresponding author:* wiebke.drews@math.tu-dortmund.de,
stefan.turek@math.tu-dortmund.de,
christoph.lohmann@math.tu-dortmund.de

Abstract

The work to be presented focuses on the convection-diffusion equation, especially in the regime of small diffusion coefficients, which is solved using a time-simultaneous multigrid algorithm closely related to multigrid waveform relaxation. For spatial discretization we use linear finite elements, while the time integrator is given by e.g. the Crank-Nicolson scheme. Blocking all time steps into a global linear system of equations and rearranging the degrees of freedom leads to a space-only problem with vector-valued unknowns for each spatial node. Then, common iterative solution techniques, such as the GMRES method with block Jacobi preconditioning, can be used for the numerical solution of the (spatial) problem and allow a higher degree of parallelization in space. We consider a time-simultaneous multigrid algorithm, which exploits space-only coarsening and the solution techniques mentioned above for smoothing purposes. By treating more time steps simultaneously, the dimension of the system of equations increases significantly and, hence, results in a larger number of degrees of freedom per spatial unknown. This can be used to employ parallel processes more efficiently. In numerical studies, the iterative multigrid solution of a problem with up to thousands of blocked time steps is analyzed in 1D. For the special case of the heat equation, it is well known that the number of iterations is bounded above independently of the number of blocked time steps, the time step size, and the spatial resolution. Unfortunately, convergence issues arise for the multigrid solver in convection-dominated regimes. In the context of the standard Galerkin method if the diffusion coefficient is small compared to the grid size and the magnitude of the velocity field, stabilization techniques are typically used to remove artificial oscillations in the solution. However, in our setting, special higher-order variational multiscale-type stabilization methods are discussed, which simultaneously improve the convergence behavior of the iterative solver as well as the smoothness of the numerical solution without significantly perturbing the accuracy.

Keywords. Convection-diffusion equations; Multigrid waveform relaxation; Variational multiscale methods

1 Introduction

The interest in massively parallel computing has grown rapidly due to the ever-increasing number of processors, or cores, in modern hardware architectures. To fully exploit the potential of these computers and actually reduce the run times of applications, it is essential to develop algorithms whose computational tasks can be performed in parallel on the different processors. More than 50 years ago, first investigations on parallel-in-time methods for the numerical solution of time-dependent partial differential equations were published. While initial value problems are typically solved numerically using methods that operate sequentially in time, the new algorithms solve the problem for all time steps simultaneously, providing increased parallelization capabilities that are otherwise limited by the spatial resolution. A general introduction and an overview of time parallel time integration methods can be found in [Gan15; OS20]. In this context, the Multigrid Waveform Relaxation (WRMG) method developed by Lubich and Ostermann [LO87] is a space-time multigrid method based on waveform relaxation and belongs to the iterative approaches among the time-parallel methods. This solution strategy is characterized by the fact that it is applied to the evolution equation before it is discretized in time. Recently, this approach was motivated in a different way in [Dün+21a] and referred to as a time-simultaneous multigrid method. Starting from a sequential problem already discretized in space and time, a global system of equations is set up in which all time steps are blocked so that it can be interpreted as a space-only problem for vector-valued unknowns. A geometric multigrid method with block Jacobi smoothing is applied to this linear system of equations and is designed to be highly parallelizable. While all time steps are treated simultaneously, the iterative multigrid solver allows parallelization in space. Some theoretical convergence results already exist for the WRMG method. For example, Janssen and Vandewalle [JV96a] proved that the asymptotic rate of convergence of the solver is the same as in the time-stepping approach if the problem is discretized in time using linear multistep methods. In another theoretical investigation, bounds on the spectral norm were derived using related circular matrices [Not22]. Furthermore, a Fourier analysis was exploited to analyze the time-simultaneous two-grid algorithm using a damped Jacobi (waveform relaxation) smoother: In case of the one-dimensional heat equation on a uniform grid, it was shown that the spectral norm of the iteration matrix is uniformly bounded if the one-step theta scheme is used for time integration [LDT22].

The application of many parallel-in-time methods to convection-dominated transport problems presents difficulties with respect to the parallel efficiency of the solver. Studies on higher-order hyperbolic problems and corresponding limitations were published in [FC03] for time-decomposed parallel time-integrators and for parareal, e.g., in [Bal05]. Recent studies on the multigrid reduction-in-time (MGRIT) algorithm applied to constant-wave-speed linear advection problems with an alternative coarse grid operator show fast solver convergence for various method-of-lines discretizations and a speed up compared to the sequential time-stepping method [De +21; De +23]. Moreover, optimized transmission conditions for the Schwarz waveform relaxation has been studied for the scope of convection-diffusion problems in, e.g., [GH07; DT22].

Considering these convection-dominated transport problems, it is additionally well known from discretization side that the Galerkin finite element solution is polluted by spurious artifacts. To reduce these oscillations, various stabilization techniques are proposed in the literature, e.g., see [Qua13, Sec. 12.8]. These include strongly consistent methods (such as GLS, SUPG), the introduction of artificial diffusion, or a decentralized discretization of the convection term based on a Petrov-Galerkin approach.

The stabilization to be considered in this work is a (fully implicit) variational multiscale (VMS) type method, which was originally proposed by [Hug95] and is adapted in [JKL06] and [Lay02]. Modification of the variational form of the underlying problem by adding a diffusive term and removing low frequency diffusion within the VMS context can improve the accuracy of the numerical solution compared to the one based on an artificial diffusion stabilization. In our context, however, the stabilization technique is used not only for accuracy reasons, but also for the improvement of the time-simultaneous multigrid solver under consideration. The main idea for its use is to perturb the system by higher order diffusion for preservation of the high accuracy, but better convergence behavior.

In this paper, the mentioned time-simultaneous multigrid algorithm is first investigated for the d-dimensional convection-diffusion equation in Sec. 2. In addition to the special case of the heat equation, our numerical studies focus on convection-dominated problems in 1D. As a second part of this work, we

introduce a higher-order VMS-type stabilization technique and numerically study this strategy in 1D in combination with the introduced time-simultaneous method in Sec. 3. Finally, Sec. 4 summarizes the results and offers considerations for future research.

2 Time-simultaneous multigrid method

The algorithm to be presented in this chapter is a geometric multigrid method in space and aims to be a highly parallelizable solution strategy to numerically solve the unsteady convection-diffusion equation. For this purpose, all time steps are considered simultaneously for each spatial grid point, allowing parallelization in space in a straightforward manner. We first introduce the spatial discretization of the d-dimensional problem under consideration and then formulate the linear system of equations to be solved by the time-simultaneous method. Afterwards, the time-simultaneous multigrid algorithm, which is highly related to the multigrid waveform relaxation [LO87; JV96a], is introduced. In numerical studies, the performance of the solver is analyzed in one spatial dimension to quickly obtain representative results.

2.1 Discretization

We consider the d-dimensional convection-diffusion problem: Find $u : \Omega \times (0,T) \to \mathbb{R}$ such that

$$\begin{aligned} \partial_t u(x,t) - \varepsilon \Delta u(x,t) + \boldsymbol{v}(x,t) \cdot \nabla u(x,t) &= f(x,t) && (x,t) \in \Omega \times (0,T) \\ u(x,t) &= g_D(x,t) && (x,t) \in \Gamma_D \times (0,T) \\ u(x,0) &= u^0(x) && x \in \Omega \end{aligned} \tag{1}$$

where $T > 0$ denotes the final time and $\Omega \subset \mathbb{R}^d, d \in \{1,2,3\}$, is the spatial domain on whose boundary $\partial\Omega$ homogeneous Dirichlet boundary values (i.e. $g_D = 0$) are imposed for simplicity. The velocity field and the right hand side are given by $\boldsymbol{v} : \Omega \times (0,T) \to \mathbb{R}^d$ and $f : \Omega \times (0,T) \to \mathbb{R}$, while $\varepsilon \geq 0$ is a constant diffusion coefficient.

Let $(\cdot,\cdot)$ denote the $L^2(\Omega)$-inner product. The solution $u : (0,T) \to V$ of the variational formulation for (1) satisfies the initial condition $u(0,\boldsymbol{x}) = u^0(\boldsymbol{x}) \in V = H^1(\Omega)$ and

$$(\partial_t u, \varphi) + \varepsilon(\nabla u, \nabla \varphi) + (\boldsymbol{v} \cdot \nabla u, \varphi) = (f, \varphi) \quad \forall \varphi \in V. \tag{2}$$

Problem (2) is discretized using the subspace $V_h \subset V$ of linear finite elements (FE) defined on the triangulation $\mathcal{T}_h$. Then $u_h : (0,T) \to V_h$ satisfies

$$(\partial_t u_h, \varphi_h) + \varepsilon(\nabla u_h, \nabla \varphi_h) + (\boldsymbol{v} \cdot \nabla u_h, \varphi_h) = (f, \varphi_h) \quad \forall \varphi_h \in V_h, \tag{3}$$

leading to the semi-discrete formulation in matrix form

$$\boldsymbol{M}_h \partial_t \boldsymbol{u}_h(t) + \varepsilon \boldsymbol{L}_h \boldsymbol{u}_h(t) + \boldsymbol{K}_h \boldsymbol{u}_h(t) = \boldsymbol{f}_h(t), \tag{4}$$

where $\boldsymbol{M}_h, \boldsymbol{L}_h, \boldsymbol{K}_h \in \mathbb{R}^{N \times N}$ are the mass, diffusion, and convection matrices, respectively. The discretized right hand side is given by $\boldsymbol{f}_h \in \mathbb{R}^N$ for $N \in \mathbb{N}$ spatial degrees of freedom. All occurring integrals are approximated using the Trapezoidal rule. This leads to a diagonal mass matrix $\boldsymbol{M}_h$ and problem (4) is equivalent to the well known second order finite difference (FD) discretization of (1) in case of an equidistant triangulation in one dimension and a constant velocity field.

Discretization in time using the Crank-Nicolson (CN) scheme results in the discrete sequential form

$$\begin{gathered} \boldsymbol{A}_I \boldsymbol{u}_h^m + \boldsymbol{A}_E \boldsymbol{u}_h^{m-1} = \boldsymbol{f}^m, \quad m = 1, ..., K \\ \text{for } \boldsymbol{A}_I \coloneqq \boldsymbol{M}_h + \frac{1}{2}\delta t(\varepsilon \boldsymbol{L}_h + \boldsymbol{K}_h), \quad \boldsymbol{A}_E \coloneqq -\boldsymbol{M}_h + \frac{1}{2}\delta t(\varepsilon \boldsymbol{L}_h + \boldsymbol{K}_h), \quad \boldsymbol{f}^m \coloneqq \frac{1}{2}\delta t \left(\boldsymbol{f}_h^m + \boldsymbol{f}_h^{m-1}\right), \end{gathered} \tag{5}$$

where $\boldsymbol{u}_h^0 \in \mathbb{R}^N$ is a suitable approximation of u^0, δt is the time step size and $K \in \mathbb{N}$ denotes the number of time steps. The solution technique described below can be applied similarly to non-equidistant grids and to other time stepping methods like the implicit Euler method or Runge-Kutta methods.

So far, common discretization techniques have been presented. We now use algebraic transformations to obtain a global system matrix with a specific structure and construct the time-simultaneous multigrid method. Considering equation (5), using $\boldsymbol{u}^m$ as a shorthand notation for $\boldsymbol{u}_h^m$ and blocking all K time steps in a global linear system of equations results in

$$\underbrace{\begin{pmatrix} \boldsymbol{A}_I & & & \\ \boldsymbol{A}_E & \boldsymbol{A}_I & & \\ & \ddots & \ddots & \\ & & \boldsymbol{A}_E & \boldsymbol{A}_I \end{pmatrix}}_{\in \mathbb{R}^{NK \times NK}} \underbrace{\begin{pmatrix} \boldsymbol{u}^1 \\ \boldsymbol{u}^2 \\ \vdots \\ \boldsymbol{u}^K \end{pmatrix}}_{\in \mathbb{R}^{NK}} = \underbrace{\begin{pmatrix} \boldsymbol{f}^1 - \boldsymbol{A}_E \boldsymbol{u}^0 \\ \boldsymbol{f}^2 \\ \vdots \\ \boldsymbol{f}^K \end{pmatrix}}_{\in \mathbb{R}^{NK}},$$

Here, all degrees of freedom are sorted in a time-major ordering. The degrees of freedom are then rearranged so that all unknowns associated with one spatial node can be blocked into a macro degree of freedom as follows:

$$\begin{aligned} &(u_1^1, u_2^1, \dots u_N^1, u_1^2, u_2^2, \dots, u_N^2, \dots, u_1^K, u_2^K, \dots, u_N^K)^\top && \text{time-major ordering} \\ &\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad \downarrow \\ \boldsymbol{u} \coloneqq &(u_1^1, u_1^2 \dots, u_1^K, u_2^1, u_2^2, \dots, u_2^K, \dots, u_N^1, u_N^2, \dots, u_N^K)^\top && \text{space-major ordering} \end{aligned}$$

By doing the same for the system matrix as well as for the right-hand side, we obtain the following *space-only problem* with vector-valued unknowns for each spatial node:

$$\underbrace{\begin{pmatrix} \square & \square & & \\ \square & \square & \ddots & \\ & \ddots & \ddots & \square \\ & & \square & \square \end{pmatrix}}_{=: \boldsymbol{S} \in \mathbb{R}^{NK \times NK}} \boldsymbol{u} = \boldsymbol{f} \quad \text{with block matrix entries} \quad \begin{pmatrix} * & & & \\ * & * & & \\ & \ddots & \ddots & \\ & & * & * \end{pmatrix} \in \mathbb{R}^{K \times K} \tag{6}$$

The tridiagonal block structure of $\boldsymbol{S}$ stems from the sparsity pattern of matrices $\boldsymbol{A}_I$ and $\boldsymbol{A}_E$ of the one-dimensional problem considered here for simplicity, while each block entry is a lower bidiagonal matrix due to the use of the Crank-Nicolson time integrator. Obviously, this new system matrix has still the same dimension as the system matrix before. In what follows, we present a time-simultaneous solution algorithm for $\boldsymbol{S}\boldsymbol{u} = \boldsymbol{f}$, which highly exploits the special structure of the system matrix.

2.2 Solution strategy

In this section, we apply a geometric multigrid method in space to the constructed system (6), which is designed to be highly parallelizable on modern hardware architectures. The iterative solver allows parallelization in space while all time steps are treated simultaneously for each spatial grid point.

In general, a multigrid solver is based on a hierarchy of mesh levels, which are used to reduce different modes of the error. More precisely, the idea of a two-grid algorithm is to start with a fine grid, where a smoother performs a number of smoothing steps to dampen highly oscillating error components. The remaining smooth part of the error is then approximated on a coarser grid and used to update the solution on the fine mesh. Performing this procedure iteratively results in a very efficient solution strategy, if the coarse grid problem is recursively approximated using the same technique [Bra07, Sec. 5.1].

The time-simultaneous multigrid algorithm to be presented makes use of the same multigrid components, i.e., smoothing and coarse-grid correction, but is now applied to the space-only system (6) with vector-valued unknowns. For a more detailed description, the following sketch of the algorithm gives a brief overview of the i-th iteration of the two-grid case, which can be extended to the multigrid case.

Time-simultaneous two-grid algorithm (i-th iteration).

Let $\boldsymbol{x}^{(i)} \in \mathbb{R}^{NK}$ be a given initial guess.

1. **Pre-Smoothing:** $\boldsymbol{x}^{(i,1)} = \text{ smoother}(\boldsymbol{x}^{(i)}, \boldsymbol{D}, \nu_1)$
   - ν_1 iterations
   - preconditioner $\boldsymbol{D}$
2. **Coarse-Grid Correction:**
   - compute residual $\quad \boldsymbol{d} = \boldsymbol{f} - \boldsymbol{S}\boldsymbol{x}^{(i,1)}$
   - restrict residual $\quad \overline{\boldsymbol{f}} = \boldsymbol{R}\boldsymbol{d}$
   - solve coarse grid problem $\quad \overline{\boldsymbol{S}}\overline{\boldsymbol{x}}^{(i)} = \overline{\boldsymbol{f}}$
   - solution update $\quad \boldsymbol{x}^{(i,2)} = \boldsymbol{x}^{(i,1)} + \boldsymbol{P}\overline{\boldsymbol{x}}^{(i)}$
3. **Post-Smoothing:** $\boldsymbol{x}^{(i+1)} = \text{ smoother}(\boldsymbol{x}^{(i,2)}, \boldsymbol{D}, \nu_2)$
   - ν_2 iterations
   - preconditioner $\boldsymbol{D}$

On the fine grid with N spatial unknowns in each time step, a smoother is applied to a given initial guess to smooth high-frequency error components in the first place. Then the coarse grid correction is performed. It consists of computing the residual on the fine grid, which is transferred to the coarse grid using the restriction operator $\boldsymbol{R}$. With $\overline{N}$ spatial nodes on this coarse level, the system matrix $\overline{\boldsymbol{S}} \in \mathbb{R}^{\overline{N}K \times \overline{N}K}$ is discretized in the same way as $\boldsymbol{S}$ was constructed before. The coarse grid solution $\overline{\boldsymbol{x}}^{(i)} \in \mathbb{R}^{\overline{N}K}$ is then computed exactly. By prolongating this solution to the fine grid using the prolongation operator $\boldsymbol{P}$, the pre-smoothed solution can be corrected. Finally, some post-smoothing can be performed on the fine mesh to again dampen highly oscillating error components. The preconditioner $\boldsymbol{D}$ involved in the definition of the smoother should be designed to be efficiently applicable and will be discussed in what follows.

Smoothing We consider the iterative (damped) block Jacobi method

$$\boldsymbol{x} \mapsto \boldsymbol{x} + \omega \boldsymbol{D}^{-1}(\boldsymbol{f} - \boldsymbol{S}\boldsymbol{x})$$

applied to the linear system of equations (6) using a damping parameter $\omega \in (0, 1]$, where the block Jacobi preconditioner is the block diagonal of the system matrix $\boldsymbol{S}$

$$\boldsymbol{D} \coloneqq \begin{pmatrix} \square & & \\ & \ddots & \\ & & \square \end{pmatrix} \quad \text{with block matrix entries} \quad \underbrace{\begin{pmatrix} * & & & \\ * & * & & \\ & \ddots & \ddots & \\ & & * & * \end{pmatrix}}_{\in \mathbb{R}^{K \times K}}.$$

This lower triangular matrix is invertible because the diagonal entries of this preconditioner correspond to the ones of the matrix defined by the following Kronecker product

$$\boldsymbol{A}_I \otimes \boldsymbol{I}_K = (\boldsymbol{M}_h + \tfrac{1}{2}\delta t(\varepsilon \boldsymbol{L}_h + \boldsymbol{K}_h)) \otimes \boldsymbol{I}_K$$

and do not vanish, since the diagonal entries of $\boldsymbol{M}_h, \boldsymbol{L}_h$ and $\boldsymbol{K}_h$ are positive. The individual blocks of $\boldsymbol{D}$ provide a high degree of parallelization, because they can be considered independently of each other and only couple the (temporal) degrees of freedom associated with a single spatial node. Its lower bidiagonal structure makes it easy to solve the resulting linear systems of equations by other appropriate direct or iterative approaches.

For the construction of an efficient solution strategy, the block Jacobi method is not used as a standalone solver. Instead, it can be embedded into the multigrid algorithm as a smoother to attenuate high-frequency error components. In combination with the benefits of a Krylov subspace method, the smoother is finally given by the GMRES method as introduced in [SS86] exploiting the preconditioner $\boldsymbol{D}$. Using this scheme, ν_1 and $\nu_2 \geq 0$ pre- and post-smoothing steps are performed by the smoother.

Transfer between space-time grids Intergrid transfer operators are necessary to exploit the coarse grid correction in the multigrid method. The idea is to use common coarsening techniques in space while the temporal mesh stays fixed throughout the whole algorithm. The two-grid case is considered again, where h and $H = 2h$ are the mesh sizes of the fine level and the coarse level, respectively. To restrict the residual from the fine grid $\mathcal{T}_h$ to the coarse grid $\mathcal{T}_H$, the restriction operator

$$\boldsymbol{R} \coloneqq \boldsymbol{R}_h^H \otimes \boldsymbol{I}_K = \begin{pmatrix} \frac{1}{2} & 1 & \frac{1}{2} & & & & \\ & & \frac{1}{2} & 1 & \frac{1}{2} & & \\ & & & \ddots & \ddots & \ddots & \\ & & & & \frac{1}{2} & 1 & \frac{1}{2} \end{pmatrix} \otimes \boldsymbol{I}_K$$

has to be applied, which is defined as the Kronecker product of the canonical restriction operator $\boldsymbol{R}_h^H \in \mathbb{R}^{\overline{N} \times N}$ of the FE space and the identity matrix $\boldsymbol{I}_K \in \mathbb{R}^{K \times K}$. Similarly, the prolongation operator $\boldsymbol{P} = \boldsymbol{R}^\top$ projects an FE solution from the coarse grid to the fine grid, by means of an interpolation for piecewise linear functions [Bra07, Sec. 5.1].

In general, the two-grid algorithm described above can be easily transferred to a multigrid method by recursively applying the two-grid idea to approximately solve the coarse grid problem. In this case, only the discrete problem on the coarsest mesh has to be solved exactly. In the numerical multigrid examples below, we mostly focus on the well known V- and F-cycle as described in [Bra07, Sec. 5.1, Sec. 5.4].

Overall, the presented time-simultaneous multigrid algorithm can be interpreted as a variant of multigrid waveform relaxation (WRMG) as introduced in [LO87] when using the Jacobi smoother and the same time-stepping method. This interpretation has already been discussed in [Dün+21b]. In this context, we therefore refer to existing literature on convergence analysis of the WRMG method in [VH95; JV96b; Not22; LDT22].

2.3 Numerical studies

We now numerically investigate the solution behavior of the time-simultaneous multigrid algorithm for the convection-diffusion equation (1) in 1D as a representative of similar results in higher dimensions. For this purpose, various combinations of the velocity field and the diffusion coefficient are considered, including the special case of the heat equation for $v = 0$ and convection-dominated transport problems for $v = 1$ and small ε. Unless otherwise specified, we consider the manufactured solution

$$u(x,t) = \exp\left(-\eta(\tfrac{1}{2} - x + \tfrac{1}{4}\sin(\tfrac{\pi}{2}t))^2\right)\sin(\pi x), \quad (x,t) \in (0,1) \times (0,T), \tag{7}$$

where the parameter $\eta = 100$ characterizes the steepness of u and is chosen to keep the spatial and temporal error in balance. The peak of the function oscillates periodically in time and follows the shape of a sine curve, as shown for certain time instances in Fig. 1. This exact solution satisfies homogenous Dirichlet boundary conditions while the initial data is given by

$$u(x,0) = \exp(-\eta(\tfrac{1}{2} - x)^2)\sin(\pi x)$$

coinciding with $u(x,t)$ for $t = 0, 2, 4, \ldots$. As mentioned in Section 2.1, we discretize the problem in space using an equidistant mesh and linear finite elements with quadrature-based mass lumping, while the Crank-Nicolson scheme with a fixed time step size δt is employed for time integration. This results in second order accuracy in space and time. In addition, a first-order upwind discretization is considered to

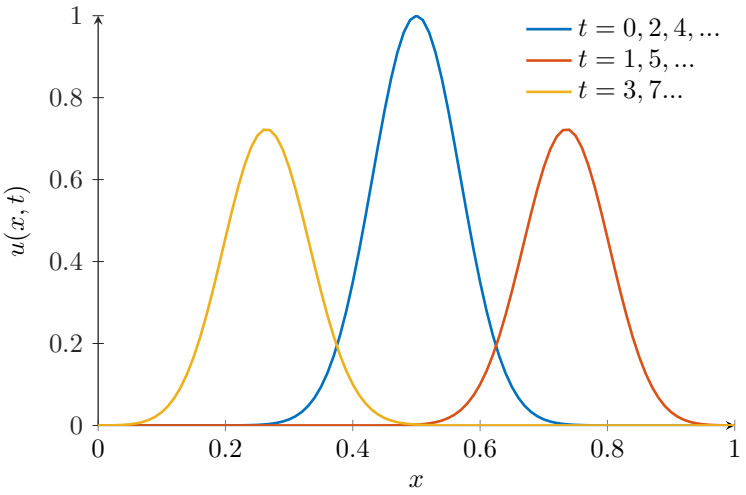

Figure 1: Smooth solution (7) on $\Omega = (0, 1)$ for certain time steps.

further illustrate the impact of the findings. This discretization technique is a special Petrov-Galerkin method and can be interpreted in the FE context as the Galerkin method supplemented by first-order artificial diffusion [QV94, Sec. 8.2.2].

In the following studies, parameters like the number of blocked time steps K, the time increment δt as well as the mesh size h are varied. In the context of the multigrid algorithm, the mesh size $h = 2^{-l}$ describes the resolution on the fine level l. In addition to the multigrid case (MG), where the coarse level is always chosen to be level 1, we also investigate the two-grid case (TG), using the coarse level $l - 1$. In this work, the smoother is always given by the GMRES method with block Jacobi preconditioning and performs $\nu_1 = \nu_2 = 4$ pre- and post-smoothing steps. We summarize the total number of iterations which are required to reduce the norm of the residual for the initial guess $\boldsymbol{x}^{(0)} = 0$ by a factor of 10^{-8}, while the maximum number of iterations is set to 100.

2.3.1 Heat equation

First, we take a brief look at some time-simultaneous multigrid results for the special case of the one-dimensional heat equation, which is equivalent to the partial differential equation (1) with the fixed velocity field $v = 0$. As known from theory for the time-simultaneous two-grid algorithm with a damped block Jacobi smoother, the spectral norm of the iteration matrix is uniformly bounded above by a value smaller than 1, which is independently of the mesh size, the time step size and the number of time steps [LDT22]. In Fig. 2, we investigate the V-cycle with the preconditioned GMRES smoother, where the number of iterations are plotted for different values of K. For a fixed ratio between the spatial and temporal resolutions, the number of iterations is indeed bounded above for the different numbers of time steps and does not depend on the fine mesh size. The upper bound is even independent of the time step size δt for sufficiently large K. In this special case, the number of iterations is bounded above by a value of 5.

Therefore, we can block many time steps simultaneously without increasing the number of iterations. The resulting linear system of equations can then be solved efficiently in parallel, due to the fact that the application of the preconditioner provides a decomposition into N independent local systems, which are each large enough to be solved on a single processor.

2.3.2 Convection-dominated transport problems

In what follows, we focus on the convection-diffusion equation by setting the velocity field to $v = 1$ and, thus, use the Galerkin discretization for the convective part as well. Due to the one-dimensional space and

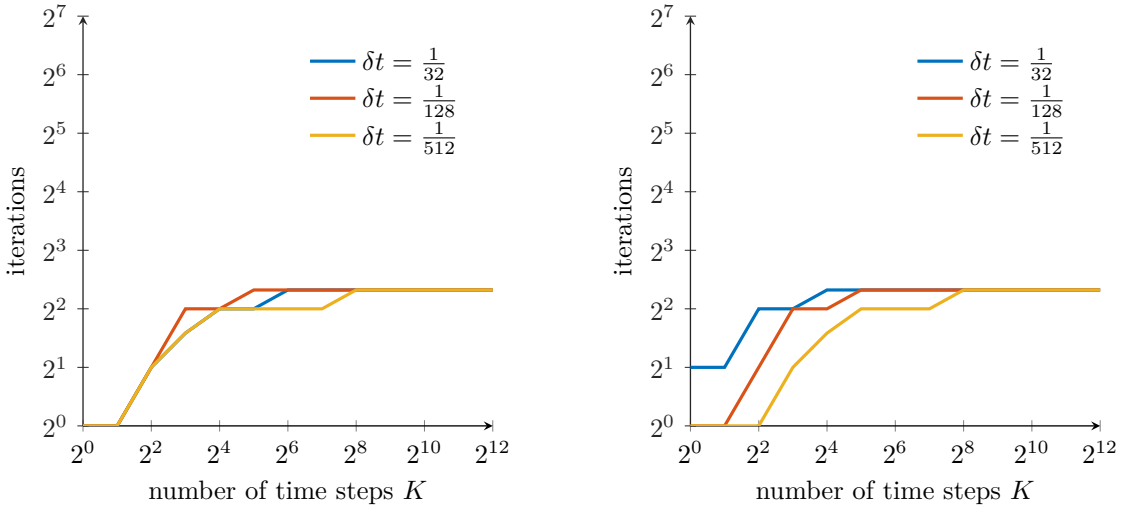

(a) $h = \delta t$. (b) $h = \frac{1}{128}$.

Figure 2: Number of iterations for multigrid method using V-cycle in case of the smooth solution, $v = 0$, and $\varepsilon = 10^{-2}$.

the use of mass lumping, the discretized convective term corresponds to the second order central difference operator for the first derivative in the context of finite differences. Furthermore, we also consider a lower order discretization at this point: The following tests are additionally investigated for the convection term discretized by the first-order upwinding technique as mentioned in [QV94, Sec. 8.2.2].

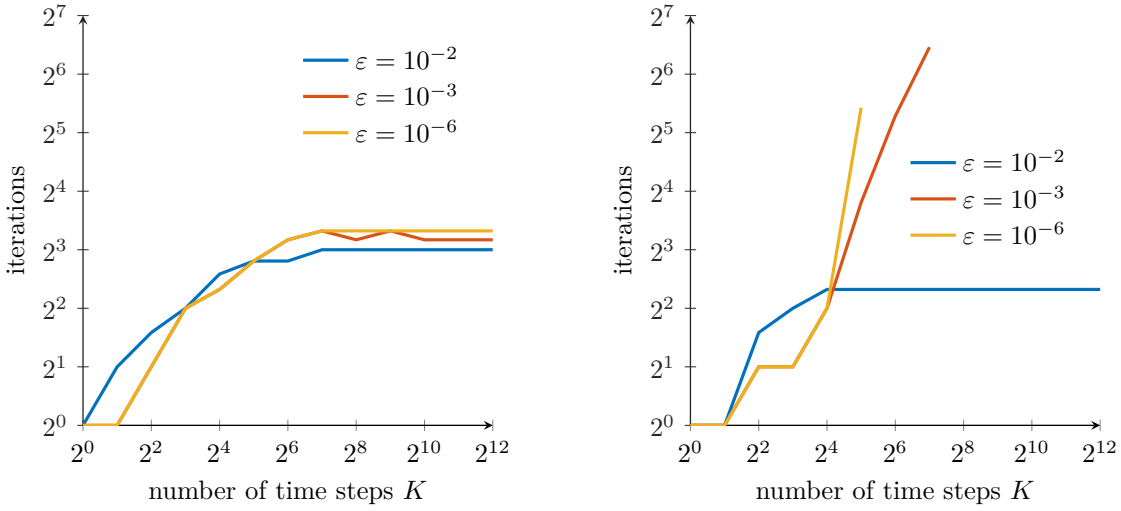

(a) Upwind discretization of convective term. (b) Galerkin discretization of convective term.

Figure 3: Number of iterations for multigrid method using V-cycle in case of the smooth solution, $v = 1$, and fixed $h = \delta t = \frac{1}{128}$.

Fig. 3 summarizes the multigrid results for a fixed fine level and time step size while different values of the diffusion coefficient are investigated. For the upwind approach and $\varepsilon = 10^{-2}$, it can be observed that the number of iterations stays bounded above even for a large time interval under consideration.

Since we study the convection-dominated case in particular, smaller values of ε are also examined. In this case, we still find similar upper bounds for the number of iterations. Overall, these observations show a similar behavior to the studies of the heat equation, see Section 2.3.1. However, the number of iterations immediately reaches its maximum of 100 if we decrease the value of ε and treat more time steps simultaneously for the Galerkin discretization of the convective term as presented in Fig. 3b.

More precisely, the time-simultaneous multigrid algorithm converges fast if the diffusion coefficient is chosen sufficiently large. This also holds true for convection-dominated problems using $\varepsilon = 10^{-3}, 10^{-6}$ and a maximum number of blocked time steps of 16. However, in this convection-dominated regime, the number of iterations required to solve the global linear system of equations increases significantly when more time steps K are blocked. The same results can be observed for simultaneous two-grid approaches.

Next, we highlight the quality of the solution by focusing on the Heaviside step function as the exact solution with periodic boundary conditions, velocity field $v = 1$, and diffusion coefficient $\varepsilon = 0$. Fig. 4 compares the exact solution at final time $T = 1$ with the numerical solutions obtained using either the upwind scheme or the Galerkin discretization for the convective term. While strong oscillations can be observed in the solution of the second order approach, the upwind scheme provides a highly diffusive approximation.

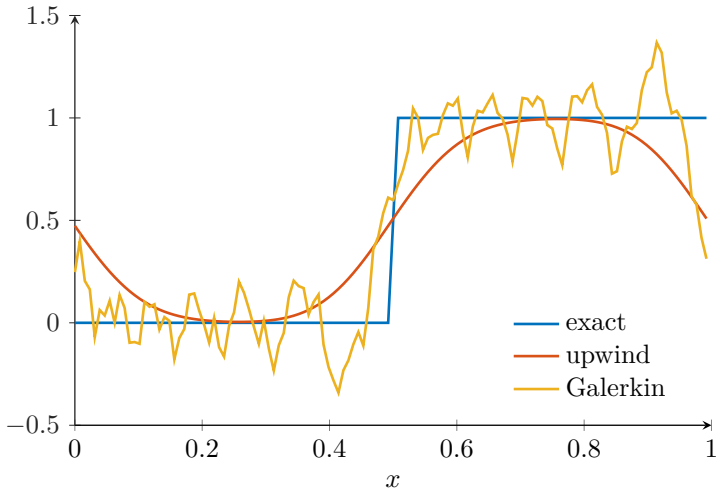

Figure 4: Heaviside step function at final time $T = 1$ in the case of $h = \delta t = \frac{1}{128}$ and corresponding numerical solutions.

The instabilities and oscillations, which occur in case of unstabilized Galerkin finite elements, are well known and can obviously be avoided using the upwind scheme. However, this approach is very diffusive and only leads to first-order accuracy $\mathcal{O}(h)$ [JS08; QV94, Sec. 8.3]. To preserve the higher order of the Galerkin discretization and especially improve the convergence of the solver even if a large number of time steps is blocked, we next introduce some stabilization based on the variational multiscale method. While accelerating the convergence of the solver is the main focus of this modification, we will observe that oscillations in the numerical solution are also damped as a positive side effect in this work.

3 Variational multiscale stabilization

When using the Galerkin approximation of the convection-diffusion equation (1), the numerical solution might be inaccurate and polluted by unphysical oscillations, if the diffusion coefficient ε is small compared to the mesh size h and the magnitude of the velocity field $\boldsymbol{v}$ [QV94, Sec. 8.3]. This problem was also observed in the numerical examples above and calls for some stabilization that attempts to reduce instabilities and artifacts in the numerical solution. In this paper we focus on the variational multiscale (VMS) method,

first introduced for this purpose in [Hug95], and in particular presented as a projection-based extension in [Lay02] and [JKL06]. The stabilization technique is presented below based on the d-dimensional problem under consideration. This is followed by an interpretation in the context of finite differences and numerical examples to discuss the influence of the method on the accuracy of the solution and the performance of the multigrid solver in 1D.

3.1 Definition

We again consider the d-dimensional convection-diffusion equation as presented in (1), where the velocity field $\boldsymbol{v}(x,t) = \boldsymbol{v} \in \mathbb{R}^d$ and the diffusion coefficient $\varepsilon > 0$ are given. The VMS method under investigation introduces an additional diffusive term and removes low frequency diffusion by means of the divergence of a recovered gradient approximation to the variational form introduced in (3). Let $(V_h)^d$ denote a d-dimensional vector-valued finite element subspace of $(L^2(\Omega))^d$ and $\alpha_{add} \geq 0$ be the constant stabilization parameter to be defined later. Then the solution $(u_h, \mathbf{g}_h)$ is sought so that

$$\begin{aligned}(\partial_t u_h, \varphi_h) + \varepsilon(\nabla u_h, \nabla\varphi_h) + (\boldsymbol{v}\cdot\nabla u_h, \varphi_h) + \alpha_{add}(\nabla u_h, \nabla\varphi_h) - \alpha_{add}(\mathbf{g}_h, \nabla\varphi_h) &= (f, \varphi_h) \quad \forall\varphi_h \in V_h,\\ (\mathbf{g}_h - \nabla u_h, \boldsymbol{\psi}_h) &= 0 \qquad\quad \forall\boldsymbol{\psi}_h \in (V_h)^d\end{aligned}$$

is satisfied, where $\mathbf{g}_h : (0,T) \to (V_h)^d$ corresponds to the projected gradient of u_h and ensures that the introduced stabilization term vanishes in the continuous problem. In contrast to the projection-based VMS stabilization technique as published in [Lay02] and [JKL06], the gradient is approximated using the same FE space defined on the triangulation $\mathcal{T}_h$. This procedure is also considered in [Loh+17, Sec. 5].

After discretization in space, the problem in matrix form reads

$$\begin{aligned}\boldsymbol{M}_h\partial_t\boldsymbol{u}_h(t) + \varepsilon\boldsymbol{L}_h\boldsymbol{u}_h(t) + \boldsymbol{K}_h\boldsymbol{u}_h(t) + \alpha_{add}\boldsymbol{L}_h\boldsymbol{u}_h(t) - \alpha_{add}\boldsymbol{B}_h^\top\mathbf{g}_h(t) &= \boldsymbol{f}_h(t),\\ \boldsymbol{N}_h\mathbf{g}_h(t) - \boldsymbol{B}_h\boldsymbol{u}_h(t) &= 0,\end{aligned}$$

where $\boldsymbol{B}_h \in \mathbb{R}^{N\times N}$ is the discrete counterpart of the gradient and $\boldsymbol{N}_h$ is the mass matrix corresponding to the vector-valued finite element space $(V_h)^d$. We eliminate the second equation by substituting $\mathbf{g}_h(t)$ into the first equation. This results in the system to be solved

$$\boldsymbol{M}_h\partial_t\boldsymbol{u}_h(t) + \varepsilon\boldsymbol{L}_h\boldsymbol{u}_h(t) + \boldsymbol{K}_h\boldsymbol{u}_h(t) + \alpha_{add}\boldsymbol{W}_h\boldsymbol{u}_h(t) = \boldsymbol{f}_h(t), \tag{8}$$

where the stabilization matrix $\boldsymbol{W}_h \coloneqq \boldsymbol{L}_h - \boldsymbol{B}_h^\top\boldsymbol{N}_h^{-1}\boldsymbol{B}_h$ can be explicitly determined for a diagonal mass matrix $\boldsymbol{N}_h$. As before, the Crank-Nicolson scheme is used as the time integrator. Then the discrete counterpart of (8) is given by

$$\boldsymbol{A}_I\boldsymbol{u}_h^m + \boldsymbol{A}_E\boldsymbol{u}_h^{m-1} + \tfrac{\delta t}{2}\alpha_{add}\boldsymbol{W}_h\boldsymbol{u}_h^m + \tfrac{\delta t}{2}\alpha_{add}\boldsymbol{W}_h\boldsymbol{u}_h^{m-1} = \boldsymbol{f}^m, \quad m = 1, ..., K \tag{9}$$

as a straightforward extension of (5). Generally, the sparsity pattern of the stabilization matrix $\boldsymbol{W}_h$ is more dense than the one of $\boldsymbol{A}_I$ and $\boldsymbol{A}_E$ due to the fact that the multiplication of FE matrices is involved in the definition of $\boldsymbol{W}_h$. Thus, stabilizing the system comes at the expense of a more complex iterative solution strategy. We will come back to this observation when discussing the matrix structures in more detail. Finally, in this work we treat the stabilization fully implicitly to reduce the computational effort. This results in the final formulation of the stabilized FE discretization to (1)

$$\boldsymbol{A}_I\boldsymbol{u}_h^m + \boldsymbol{A}_E\boldsymbol{u}_h^{m-1} + \delta t\alpha_{add}\boldsymbol{W}_h\boldsymbol{u}_h^m = \boldsymbol{f}^m, \quad m = 1, ..., K. \tag{10}$$

In the next section, we will argue that this fully implicit treatment is reasonable and actually does not reduce the accuracy of the numerical solution. We will then discuss the level-dependent choice of the stabilization parameter α_{add} in the context of the multigrid approach and study the stabilization numerically.

3.2 Interpretation using finite differences in 1D

To explore the fully implicit treatment of $\boldsymbol{W}_h$ in more detail, we first consider the time discretization of the stabilization.

Time discretization In Section 2.1, we introduced the Crank-Nicolson scheme for time integration of the original semi-discrete formulation. Since the stabilization term is treated fully implicitly, we now investigate its effect on the order of accuracy starting from the discrete form (10), where $\boldsymbol{W}_h$ is added to the unknown of the m-th time step with time step size δt. For this purpose, we algebraically transform the problem at hand into the following form:

$$\begin{aligned} & \boldsymbol{A}_I \boldsymbol{u}_h^m + \boldsymbol{A}_E \boldsymbol{u}_h^{m-1} + \delta t \alpha_{add} \boldsymbol{W}_h \boldsymbol{u}_h^m = \boldsymbol{f}^m \\ \Leftrightarrow \quad & \hat{\boldsymbol{A}}_I \boldsymbol{u}_h^m + \hat{\boldsymbol{A}}_E \boldsymbol{u}_h^{m-1} = \boldsymbol{f}^m \end{aligned} \tag{11}$$

where the matrices $\hat{\boldsymbol{A}}_I$ and $\hat{\boldsymbol{A}}_E$ are given by

$$\begin{aligned} \hat{\boldsymbol{A}}_I &\coloneqq \quad (\boldsymbol{M}_h + \alpha_{add} \tfrac{\delta t}{2} \boldsymbol{W}_h) + \tfrac{\delta t}{2} (\varepsilon \boldsymbol{L}_h + \boldsymbol{K}_h + \alpha_{add} \boldsymbol{W}_h), \\ \hat{\boldsymbol{A}}_E &\coloneqq -(\boldsymbol{M}_h + \alpha_{add} \tfrac{\delta t}{2} \boldsymbol{W}_h) + \tfrac{\delta t}{2} (\varepsilon \boldsymbol{L}_h + \boldsymbol{K}_h + \alpha_{add} \boldsymbol{W}_h). \end{aligned}$$

Obviously, this is nothing else than the Crank-Nicolson discretization of

$$\begin{aligned} & (\boldsymbol{M}_h + \alpha_{add} \tfrac{\delta t}{2} \boldsymbol{W}_h) \partial_t \boldsymbol{u}_h(t) + (\varepsilon \boldsymbol{L}_h + \boldsymbol{K}_h + \alpha_{add} \boldsymbol{W}_h) \boldsymbol{u}_h(t) = \boldsymbol{f}_h(t) \\ \Leftrightarrow \quad & \boldsymbol{M}_h \partial_t \boldsymbol{u}_h(t) + (\varepsilon \boldsymbol{L}_h + \boldsymbol{K}_h) \boldsymbol{u}_h(t) + \alpha_{add} \boldsymbol{W}_h (\tfrac{\delta t}{2} \partial_t \boldsymbol{u}_h(t) + \boldsymbol{u}_h(t)) = \boldsymbol{f}_h(t), \end{aligned} \tag{12}$$

and therefore guarantees that the order of convergence of the numerical approximation is not reduced if the semi-discrete solution to (12) converges to the exact solution of (1) with second order in space. This derivation motivates the more general variational formulation of the VMS stabilization in the d-dimensional case

$$\begin{aligned} & (\partial_t u_h, \varphi_h) + \varepsilon (\nabla u_h, \nabla \varphi_h) + (\boldsymbol{v} \cdot \nabla u_h, \varphi_h) \\ & \qquad + \alpha_{add} [(\nabla [u_h + \tfrac{\delta t}{2} \partial_t u_h], \nabla \varphi_h) - (\mathbf{g}_h, \nabla \varphi_h)] = (f, \varphi_h) \quad \forall \varphi_h \in V_h, \\ & \qquad (\mathbf{g}_h - \nabla [u_h + \tfrac{\delta t}{2} \partial_t u_h], \boldsymbol{\psi}_h) = 0 \quad \forall \boldsymbol{\psi}_h \in (V_h)^d. \end{aligned} \tag{13}$$

Discretization in time using the Crank-Nicolson scheme results in a fully implicit stabilization term and, hence, reduces the numerical complexity compared to problem (9). Next, we focus on the spatial accuracy of (12) and, for this, consider an FD interpretation of the stabilization matrix under investigation.

Space discretization Using one-dimensional linear finite elements, a uniform grid and quadrature-based mass lumping, both mass matrices $\boldsymbol{M}_h$ and $\boldsymbol{N}_h$ correspond to a scaled identity matrix. More precisely, the matrices under consideration are given by

$$\boldsymbol{N}_h = \boldsymbol{M}_h = \begin{pmatrix} \ddots & \ddots & \ddots & & \\ & 0 & h & 0 & \\ & & \ddots & \ddots & \ddots \end{pmatrix}, \qquad \boldsymbol{B}_h = \frac{1}{2} \begin{pmatrix} \ddots & \ddots & \ddots & & \\ & -1 & 0 & 1 & \\ & & \ddots & \ddots & \ddots \end{pmatrix}$$

neglecting the boundary conditions, so that the second part of the stabilization matrix is given by

$$\tilde{\boldsymbol{L}}_h \coloneqq \boldsymbol{B}_h^\top \boldsymbol{N}_h^{-1} \boldsymbol{B}_h = -\frac{1}{4h} \begin{pmatrix} \ddots & \ddots & \ddots & \ddots & \ddots & & \\ & 1 & 0 & -2 & 0 & 1 & \\ & & \ddots & \ddots & \ddots & \ddots & \ddots \end{pmatrix}.$$

Using the definition of the tridiagonal matrix $\boldsymbol{L}_h$ as the negative of the discrete Laplacian and already known from the diffusive part, we conclude that the stabilization matrix $\boldsymbol{W}_h$ reads

$$\boldsymbol{W}_h \coloneqq \boldsymbol{L}_h - \tilde{\boldsymbol{L}}_h = -\frac{1}{h} \begin{pmatrix} \ddots & \ddots & \ddots & & \\ & 1 & -2 & 1 & \\ & & \ddots & \ddots & \ddots \end{pmatrix} + \frac{1}{4h} \begin{pmatrix} \ddots & \ddots & \ddots & \ddots & \ddots & & \\ & 1 & 0 & -2 & 0 & 1 & \\ & & \ddots & \ddots & \ddots & \ddots & \ddots \end{pmatrix}.$$

To interpret this matrix in the context of finite differences, we multiply equation (12) by $\boldsymbol{M}_h^{-1} = h^{-1}\boldsymbol{I}$. This results in the FD matrices $\boldsymbol{L}_h^* \coloneqq \boldsymbol{M}_h^{-1}\boldsymbol{L}_h$ and $\tilde{\boldsymbol{L}}_h^* \coloneqq \boldsymbol{M}_h^{-1}\tilde{\boldsymbol{L}}_h$, which are equivalent to discrete diffusion operators applied to two different mesh sizes h and $2h$. In summary, the FD stabilization matrix is pentadiagonal and represented by

$$\boldsymbol{W}_h^* \coloneqq \boldsymbol{M}_h^{-1}\boldsymbol{W}_h = \boldsymbol{L}_h^* - \tilde{\boldsymbol{L}}_h^* = \frac{1}{(2h)^2}\begin{pmatrix} \ddots & \ddots & \ddots & \ddots & \ddots & & \\ & 1 & -4 & 6 & -4 & 1 & \\ & & \ddots & \ddots & \ddots & \ddots & \ddots \end{pmatrix} \in \mathbb{R}^{N\times N}.$$

Adding this stabilization to the original problem leads to a larger bandwidth of the system matrix and therefore increases the cost of the iterative solution strategy. The fully implicit treatment of the stabilization, as introduced at the beginning of this section, minimizes this additional effort.

A brief look at the Taylor expansion of the two central difference quotients

$$\begin{aligned} \boldsymbol{L}_h^*\boldsymbol{u}_h &\sim \frac{1}{h^2}(\ h^2 u_{xx} + \quad Ch^4 u_{xxxx} + O(h^6)) \\ \tilde{\boldsymbol{L}}_h^*\boldsymbol{u}_h &\sim \frac{1}{4h^2}(4h^2 u_{xx} + 2^4 Ch^4 u_{xxxx} + O(h^6)) \\ \Rightarrow \quad \boldsymbol{W}_h^* &\coloneqq (\boldsymbol{L}_h^* - \tilde{\boldsymbol{L}}_h^*)\boldsymbol{u}_h \sim \tilde{C}h^2 u_{xxxx} + O(h^4), \end{aligned} \tag{14}$$

where $\tilde{C} \coloneqq -3C$, illustrates that the stabilization matrix $\boldsymbol{W}_h^*$ corresponds to a scaled FD discretization of u_{xxxx}, while the factor h^2 guarantees second order of accuracy in space.

The d-dimensional stabilization problem was already given at the beginning of this section in (13) where the time discretization was considered. This problem corresponds in one dimension to the continuous case of the modified convection-diffusion equation

$$\partial_t u - \varepsilon u_{xx} + v u_x + \alpha_{add}\tilde{C}h^2[u + \tfrac{\delta t}{2}\partial_t u]_{xxxx} = f. \tag{15}$$

Therefore, the stabilization is a perturbation of the continuous problem of order h^2 and the solution of (11) converges to the exact solution of (1) with second order. Although (15) is continuous in space, the fixed factor h^2 corresponds to the grid used to discretize the problem at hand. This is the main difference to the d-dimensional form, which is not derived from an FD point of view, and is a crucial aspect in the next section, where the choice of the stabilization parameter is discussed.

3.3 Choice of stabilization parameter

At this point, we motivate our level-dependent choice of the stabilization parameter α_{add} for the time-simultaneous multigrid approach. For this purpose, we consider the situation of the multigrid algorithm, where the mesh sizes of a fine level and some coarser level are given by h and H, respectively. The time step size δt stays constant for all levels since coarsening is applied only in space. According to the FD interpretation mentioned above, the stabilized one-dimensional problem corresponding to the coarse level then reads

$$\partial_t u - \varepsilon u_{xx} + v u_x + \alpha_{add}\tilde{C}H^2[u + \tfrac{\delta t}{2}\partial_t u]_{xxxx} = f \tag{16}$$

and, hence, differs from the problem of the fine level given in (15), where the scaling of the stabilization term still depends on the mesh size h of the fine level. However, the aim is to solve the same continuous problem on each level within the multigrid algorithm, i.e., the fine grid problem (15). By the choice of $\alpha_{add} = \alpha\left(\frac{h}{H}\right)^2$ in (16), this requirement is satisfied due to the fact that

$$\begin{aligned} &\partial_t u - \varepsilon u_{xx} + v u_x + \alpha\left(\frac{h}{H}\right)^2 \tilde{C}H^2[u + \tfrac{\delta t}{2}\partial_t u]_{xxxx} = f \\ \Leftrightarrow \quad &\partial_t u - \varepsilon u_{xx} + v u_x + \alpha\tilde{C}h^2[u + \tfrac{\delta t}{2}\partial_t u]_{xxxx} = f. \end{aligned} \tag{17}$$

In the special case of the fine level with mesh size h, this parameter simplifies to $\alpha_{add} = \alpha$, so that the continuous problem resulting from the coarse level (16) and the fine level (15) coincide. Using some parameters $\alpha \geq 0, \gamma \geq 0$, we introduce a more general definition of the level-dependent stabilization parameter

$$\alpha_{add} \coloneqq \alpha \left(\frac{h}{H} \right)^{\gamma},$$

which is exploited to scale the stabilization term on the coarser level with mesh size H. While a reasonable value for α will be evaluated in the numerical studies, one possible choice of γ was already motivated above: According to (17) the same continuous problem is solved on each level for $\gamma = 2$. In this case, the value of α_{add} decreases for a larger mesh size H so that less stabilization is added on coarser levels. Corresponding effects can also be observed in the following numerical studies of the time-simultaneous multigrid method. Another obvious choice resulting from the d-dimensional stabilized problem (13) without the FD interpretation is given by $\gamma = 0$, which results in a level-independent value of α_{add}. The intermediate state $\gamma = 1$ of both derivations does not guarantee to solve the same problem in the continuous, but keeps the stabilization parameter larger on the coarser levels than with $\gamma = 2$. In the following section, we study the stabilization technique and the choice of different stabilization parameters numerically.

3.4 Numerical studies

This section focuses on the qualitative and quantitative effects on the solution behavior of the stabilization technique introduced above. For this purpose, the influence on the accuracy of the solution as well as the performance of the time-simultaneous multigrid solver are investigated. As before, we again use linear finite elements for discretization in space in 1D, while the time integrator is given by the Crank-Nicolson scheme.

3.4.1 Accuracy of the solution

Our study on the quality of the solution computed by the stabilized method consists of two parts, which differ mainly by the smoothness of the considered exact solution. The choice of the parameter γ can be neglected at this point because we are only interested in the fine grid solution, which is not effected by γ due to the definition of α_{add}.

Heaviside step function By considering the Heaviside step function and the coefficients $v = 1$ as well as $\varepsilon = 0$, we focus on the quality of the numerical solution as already discussed in Section 2.3.2. We observed oscillations for the Galerkin and hence unstabilized discretization of the convective term, while the first-order upwind scheme provided a highly diffusive result. Since the higher-order stabilization is studied in combination with the Galerkin discretization, the corresponding solution of the unstabilized case, polluted by artificial oscillations, is shown again in Fig. 5. Furthermore, this figure shows the numerical approximation of the stabilized configuration at the final time $T = 1$. In this case, the artifacts of the unstabilized solution are smoothed by the stabilization using $\alpha = 0.1$, which seems to be an appropriate choice, as we will see below.

Order of convergence We now come back to the smooth solution (7) introduced in the numerical studies of the time-simultaneous algorithm in Section 2.3 and focus on the order of convergence, which was an important aspect for the choice of the stabilization technique. In the following investigations the convection-dominated region with $\varepsilon = 10^{-3}$ and velocity field $v = 1$ is considered. In Table 1, we summarize the error for the final time $T = 2$ in the discrete L_2-norm for different values of the stabilization parameter α. The error is reduced by a factor of 4 when the mesh size and the time step size are both halved, no matter how the stabilization parameter is chosen. Since we consider linear finite elements in space and the Crank-Nicolson scheme in time, this confirms the theoretical expectations with and without stabilizing the problem.

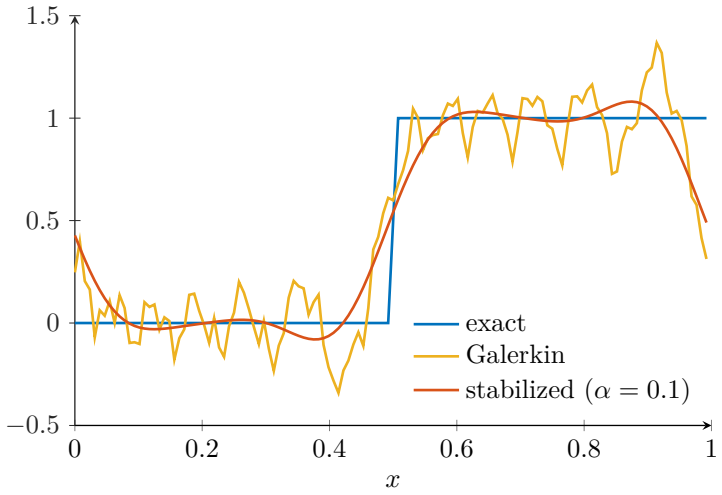

Figure 5: Heaviside step function at final time $T = 1$ in the case of $h = \delta t = \frac{1}{128}$ and corresponding numerical solution compared to stabilized discretization.

Table 1: Discrete L_2-error at final time $T = 2$ in the case of the smooth solution, $v = 1$, and $\varepsilon = 10^{-3}$.

$h = \delta t$	$\alpha = 0$	10^{-2}	10^{-1}	10^{0}	10^{1}
1/64	$1.8 \cdot 10^{-3}$	$1.9 \cdot 10^{-3}$	$6.4 \cdot 10^{-3}$	$4.9 \cdot 10^{-2}$	$1.6 \cdot 10^{-1}$
1/128	$4.6 \cdot 10^{-4}$	$4.8 \cdot 10^{-4}$	$1.7 \cdot 10^{-3}$	$1.5 \cdot 10^{-2}$	$8.8 \cdot 10^{-2}$
1/256	$1.1 \cdot 10^{-4}$	$1.2 \cdot 10^{-4}$	$4.2 \cdot 10^{-4}$	$4.0 \cdot 10^{-3}$	$3.5 \cdot 10^{-2}$
1/512	$2.9 \cdot 10^{-5}$	$3.0 \cdot 10^{-5}$	$1.0 \cdot 10^{-4}$	$1.0 \cdot 10^{-3}$	$9.9 \cdot 10^{-3}$
1/1024	$7.2 \cdot 10^{-6}$	$7.6 \cdot 10^{-6}$	$2.6 \cdot 10^{-5}$	$2.5 \cdot 10^{-4}$	$2.5 \cdot 10^{-3}$

So far, we did not argue how to choose the stabilization parameter. This can be now done by considering Table 1 more precisely: First of all, we notice that the accuracy of the solution deteriorates as α increases. For example, for $\alpha = 10^{-1}$, the numerical solution is approximately as accurate as in the case of $\alpha = 0$ for two times larger time increments and mesh sizes, i.e., we lose one level of mesh refinement, which we assume to be acceptable. Therefore, as a first observation, choosing α not too large is reasonable from the point of view of accuracy. We will come back to this when discussing the convergence behavior of the multigrid solver below, due to the fact that an accurate solution is an important aspect of the stabilization method.

3.4.2 Performance of the solver

In the previous studies on the time-simultaneous method, a slow iterative convergence behavior was observed for the convection-dominated case when the diffusion coefficient was at most $\varepsilon = 10^{-3}$, the velocity field was set to 1, and 64 or more time steps were treated simultaneously. Therefore, we now focus on the effect of the stabilization method for this parameter setting and, especially, discuss the influence of the stabilization parameter α. In addition to the convergence behavior of the iterative solver, which is illustrated by the number of iterations necessary for the time-simultaneous method to solve the problem at hand, the behavior of the error is also considered for accuracy reasons. Furthermore, we are interested in finding criterions how to choose the stabilization parameter α: While the number of required iterations should not grow arbitrarily for different numbers of blocked time steps K, the error is intended to be as close as possible to the one of the Galerkin approximation for smooth solutions.

We first focus on the two-grid solution algorithm and then discuss the influence of its multigrid extension on the choice of the stabilization parameter.

Two-grid algorithm To illustrate how to read the subsequent figures, we explain the context of Fig. 6a in detail: The number of iterations for the two-grid algorithm for three different numbers of blocked time steps K are shown in blue lines while the stabilization parameter α varies between 10^{-3} and 10^{1}. The results for the smallest value of $\alpha = 10^{-3}$ are in good agreement with the unstabilized ones, since the convergence behavior is very similar in both cases. This means in this case that the solver does not converge within the maximum number of iterations. However, the solver converges significantly faster when more stabilization is incorporated into the system. For stabilization with $\alpha = 10^{-1}$, the lines for the different values of K meet, i.e., the algorithm converges independently of the number of time steps treated simultaneously and requires only a few iterations ($< 2^4$) to solve the problem at hand. At the same time, and for this $\alpha = 10^{-1}$, the errors are about ("only") four times larger than those without stabilization, corresponding to a loss of accuracy of one mesh level. This can be observed in the same figure, where the discrete L_2-error normalized with respect to the error for $\alpha = 10^{-3}$ is plotted in red. Again, the errors for $\alpha = 10^{-3}$ are very close to the ones corresponding to $\alpha = 0$, which do not exploit any stabilization at all. Overall, the results presented in Fig. 6a illustrate the desired effect of stabilization with $\alpha = 10^{-1}$ on the two-grid solver. The associated loss of accuracy seems to be acceptable if the algorithm recovers the original quality of the solution on a finer resolution while requiring significantly less iterations.

Next, the results for a finer time step size are summarized in Fig. 6b. In this case, the convergence behavior improves slightly for smaller values of K and without stabilization (on the left of the x-axis). However, the remaining convergence issues can be further reduced by the stabilization, so that the choice of $\alpha = 10^{-1}$ might still be beneficial.

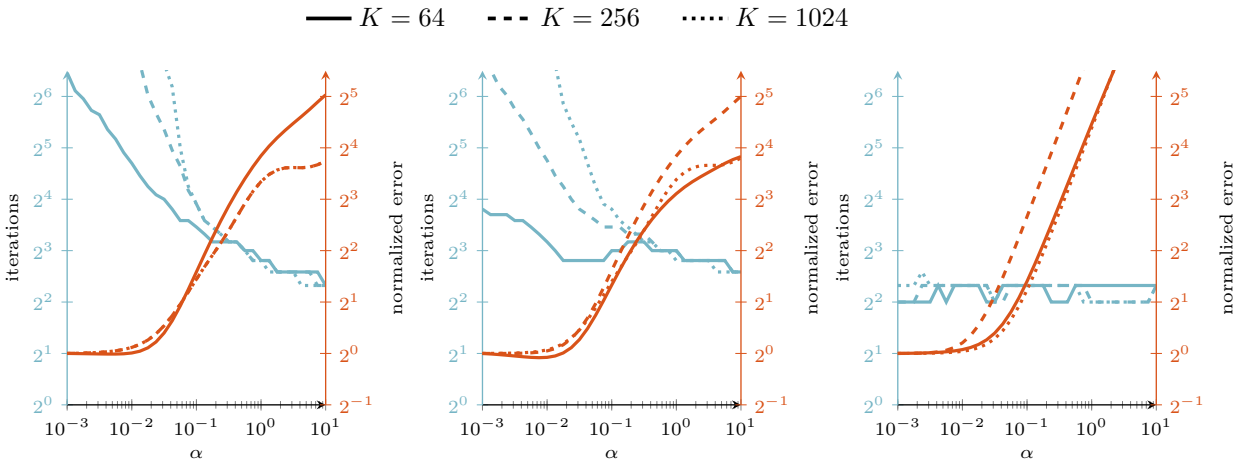

(a) $h = \delta t = \frac{1}{32}$. (b) $h = \frac{1}{32}, \delta t = \frac{1}{128}$. (c) $h = \frac{1}{512}, \delta t = \frac{1}{128}$.

Figure 6: Number of iterations and normalized error for two-grid method in case of the smooth solution, $v = 1, \varepsilon = 10^{-3}$, and stabilization parameter $\gamma = 2$.

Finally, the case of a finer spatial mesh is presented in Fig. 6c. We observe that a finer level improves the convergence behavior even without stabilization. Although stabilization would not be necessary for this configuration, the figure illustrates that stabilization does not worsen the convergence behavior.

In summary, the stabilization parameter α should not be chosen too large to achieve accurate solutions, while small values of α might not sufficiently improve the performance of the iterative solver. In this trade-off situation, the choice of $\alpha_1 := 10^{-1}$ as an upper bound for α seems to be a good compromise, as illustrated in Fig. 6.

The previous results were computed using quadrature-based mass lumping, which allowed us to exploit an FD interpretation for the analysis of the stabilization. Furthermore, the employed preconditioner $\boldsymbol{D}$ becomes exact as $\delta t \to 0$, which is not satisfied anymore when all integrals are computed exactly. In Fig. 7, the behavior of the stabilization is shown for the discretization with the consistent mass matrix $\boldsymbol{M}_h^c$ in front of the time derivate, i.e., $\boldsymbol{M}_h^c \partial_t \boldsymbol{u}_h(t)$, while mass lumping is still exploited for the computation of $\boldsymbol{N}_h$ in (8). The number of iterations is compared for the lumped and consistent mass matrices $\boldsymbol{M}_h$ and $\boldsymbol{M}_h^c$ in Fig. 7a and Fig. 7b illustrates the corresponding discrete L_2-errors. A vertical line marks the value of $\alpha_1 = 10^{-1}$, which was derived in previous investigations as an appropriate balance between accuracy and convergence behavior in case of the lumped mass matrix. Especially, the number of iterations stayed bounded above independently of the number of blocked time steps K. However, the replacement of the lumped mass matrix by the consistent one further reduces the numerical effort to solve the system at hand by a factor of 2. In this case, the number of iterations already stays bounded independently of K for a smaller value of α. The new upper bound α_2 is marked by another vertical line in the figure and indicates that the number of iterations is still lower than the ones using the lumped mass matrix for the stabilization parameter α_1. On the other hand, the errors are approximately the same for α_1 and both choices of the mass matrix. In case of $\boldsymbol{M}_h^c$, even smaller values of α are acceptable resulting in an improved accuracy of the numerical approximation. As a conclusion, by using the consistent mass matrix we can achieve a similar error for α_1, but with less numerical effort. On the other hand, there is the possibility to choose α even smaller for more accurate solutions while still achieving iteration numbers that are independent of the number of blocked time steps. Due to this observation, the following investigations are performed using the consistent mass matrix and especially focusing on $\alpha = 10^{-1}$.

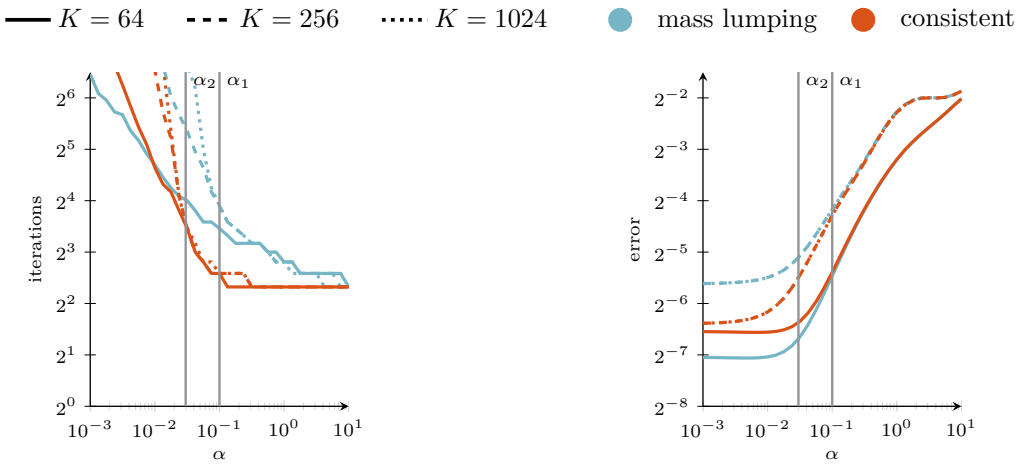

(a) Number of iterations.

(b) Absolute error.

Figure 7: Number of iterations and discrete L_2-error for two-grid method in case of the smooth solution, $v = 1, \varepsilon = 10^{-3}$, stabilization parameter $\gamma = 2$, and $h = \delta t = \frac{1}{32}$. Vertical lines mark $\alpha_1 = 0.1$ and $\alpha_2 = 0.03$.

Multigrid algorithm Time-simultaneous multigrid results are shown in Fig. 8 and illustrate the number of iterations and the normalized error behavior in the same way as in the two-grid analysis above. However, the behavior of the discrete L_2-error is slightly different, due to the use of the consistent mass matrix instead of its lumped counterpart. The discrepancies have already been discussed in the previous section, and remain valid here, since we are still solving the same problem (on the finest level) in both the TG and MG approaches. Therefore, we focus mainly on the convergence behavior of the iterative solver in the following investigations. In Figs 8a and 8c, the fine mesh size and the time step size δt coincide again. We first focus on the MG results for $\gamma = 2$ presented in Fig. 8c, where finer resolutions are considered. From

the TG findings, we concluded that the choice of $\alpha = 10^{-1}$ results in a good balance between accuracy of the solution and performance of the solver. Looking at the MG results stabilized with this value of α, we notice that the number of iterations for different values of K can still be improved compared to the unstabilized case, but increases when more time steps are treated simultaneously. The reason for this might be that the stabilization on a coarse grid seems not to be sufficient for uniform convergence behavior. Thus, the stabilization parameter α has to be chosen much larger in these cases leading to a loss of accuracy again. In Section 3.3, we already mentioned that the parameter $\gamma = 2$ implies less stabilization on coarser levels, which occur especially in the MG case where level 1 corresponds to the coarsest mesh. To keep the stabilization parameter $\alpha_{add} \coloneqq \alpha \left(\frac{h}{H}\right)^{\gamma}$ larger on the coarser levels while still using the same value on the fine level, $\gamma = 1$ is additionally considered and shown in dark blue in the same figure. For this setup, there is indeed an improvement, due to the fact that the number of iterations is again bounded for different values of K and the stabilization parameter $\alpha = 10^{-1}$. This behavior can also be observed in Figs 8a and 8b, where a coarser mesh size and time step sizes are investigated. In summary, the stabilized MG results are comparable to those of the two-grid algorithm, but the numbers of iterations are slightly higher.

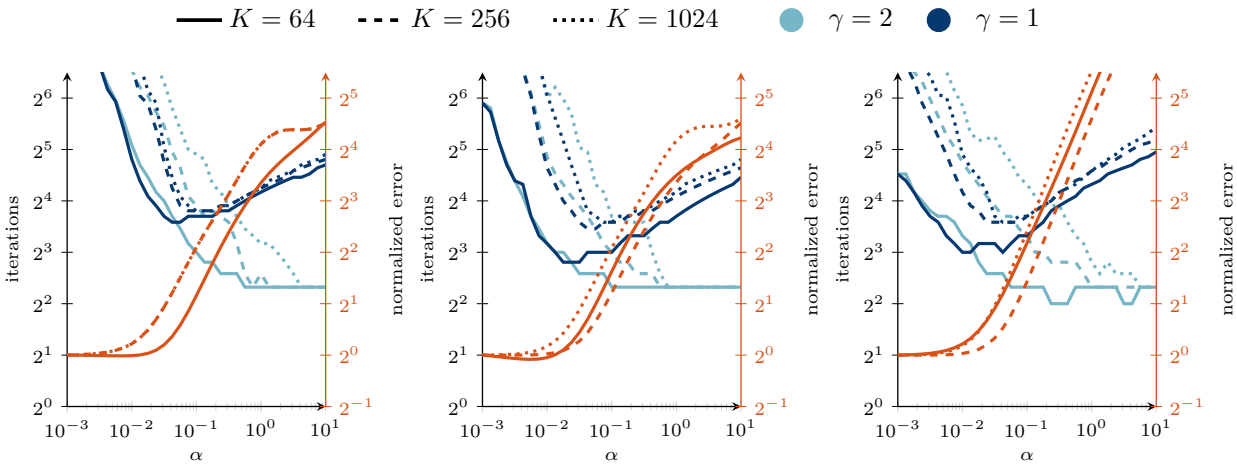

(a) $h = \delta t = \frac{1}{32}$. (b) $h = \frac{1}{32}, \delta t = \frac{1}{128}$. (c) $h = \delta t = \frac{1}{128}$.

Figure 8: Number of iterations and normalized error for multigrid method using F-cycle in case of the smooth solution, $v = 1, \varepsilon = 10^{-3}$, and the consistent mass matrix.

These observations can also be confirmed with the help of Table 2b, which additionally contains results for the stabilization parameter $\alpha = 10^{-1}$ and, in particular, $\gamma = 0$. In certain multigrid configurations, especially for significantly larger values of K, the choice of $\gamma = 0$ provides convergence rates that are uniformly bounded above, while the solver does not converge within the maximum number of iterations for $\gamma = 1$ and $\gamma = 2$. In contrast to that, the two-grid solver provides the lowest numbers of iterations bounded above for $\gamma = 2$. Corresponding results can be found in Table 2a.

Finally, we highlight two further scenarios. The first setup in Fig. 9a deals with difficulties that occur when the fine mesh size is much smaller than the time step size δt. In this case, neither the choice of $\gamma = 2$ nor $\gamma = 1$ will give the desired K-independent convergence behavior for $\alpha = 10^{-1}$. Choosing a slightly larger α would satisfy this criterion for $\gamma = 1$. However, the number of iterations are still quite large, while the error is relatively small. Since the choice of a time step size larger than the mesh resolution is not necessarily physically reasonable for convection-dominated problems, the practical relevance of this consideration remains questionable. Even in the multigrid approach, the case where the spatial mesh size is much smaller than the time step size does not arise, since coarsening is only performed in space, but could occur for other coarsening strategies (in time). The second setup considered in Fig. 9b fixes the final time $T = 2$, but increases the number of time steps K for smaller and smaller time step sizes δt. In

Table 2: Number of iterations in case of the smooth solution, $v = 1, \varepsilon = 10^{-6}$, the consistent mass matrix, and stabilization parameter $\alpha = 10^{-1}$. A dash "-" indicates that the solver did not converge within the maximum number of 100 iterations.

(a) Two-grid

	$\gamma = 0$			$\gamma = 1$			$\gamma = 2$		
$K \diagdown h = \delta t$	1/32	1/128	1/512	1/32	1/128	1/512	1/32	1/128	1/512
1	4	2	2	4	2	2	3	1	2
4	6	5	4	5	4	4	4	3	3
16	10	10	6	8	7	4	5	3	3
64	20	21	8	13	11	6	6	5	3
256	28	29	21	14	13	11	6	5	4
1024	28	28	24	14	12	12	6	5	5
4096	28	26	23	14	12	12	6	5	5
16384	28	25	21	14	11	11	6	5	4

(b) Multigrid (F-cycle)

	$\gamma = 0$			$\gamma = 1$			$\gamma = 2$		
$K \diagdown h = \delta t$	1/32	1/128	1/512	1/32	1/128	1/512	1/32	1/128	1/512
1	1	2	2	3	1	1	2	1	2
4	6	5	6	5	4	4	4	3	3
16	10	10	14	8	7	9	5	3	3
64	20	20	19	13	11	8	9	5	4
256	28	34	33	15	15	12	19	12	8
1024	28	36	45	16	15	18	32	26	23
4096	29	35	48	29	21	19	-	51	42
16384	29	34	46	-	-	28	-	-	55

this case, the convergence behavior for both values of γ can be improved in the same way for the different numbers of blocked time steps. For $\alpha = 10^{-1}$, we observe the desired effect of the stabilization even in the case of the MG algorithm, while the solver does not converge within the maximum number of iterations in most unstabilized cases.

4 Conclusion and Outlook

The time-simultaneous multigrid algorithm and its application to the one-dimensional heat equation was presented in the first part of this work. This investigation, as well as the performance of the solver for the convection-diffusion equation when the diffusion parameter is sufficiently large, illustrated that the rate of convergence is uniformly bounded above independently of the number of blocked time steps. Since convergence issues arise for the iterative solution strategy and convection-dominated problems, we introduced a variational multiscale-type stabilization technique of higher order, which intends to improve the convergence behavior of the multigrid solver. The convergence behavior of the time-simultaneous two-grid algorithm could be extremely improved when using the stabilized system. In most cases, even the number of iterations of the multigrid algorithm are bounded above, while still leading to second order of accuracy in space and time. To obtain those findings, the choice of the stabilization parameter is a crucial question. An FD interpretation of the stabilization was used to derive a level-dependent parameter for

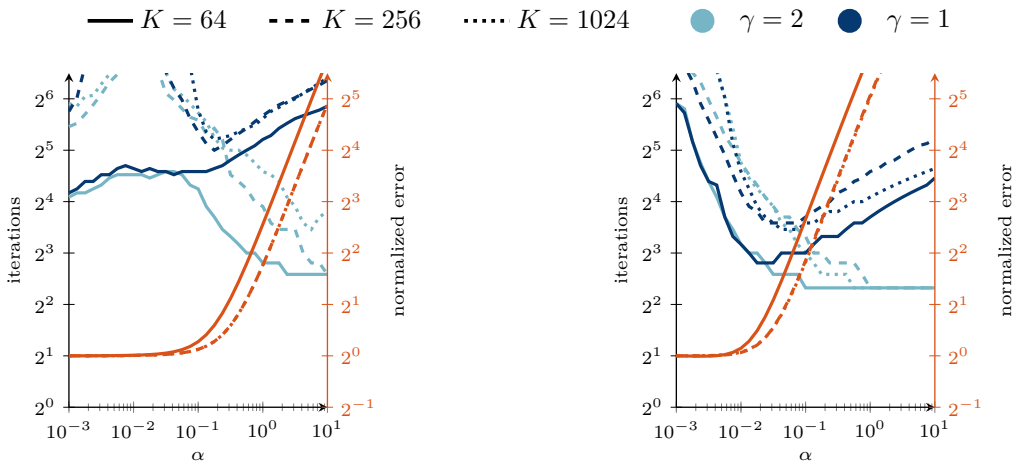

(a) Fixed grid sizes $\delta t > h$.

(b) Fixed final time $T = 2$.

Figure 10: Number of iterations for multigrid method using F-cycle in case of further setups with the smooth solution, $v = 1, \varepsilon = 10^{-3}$, and the consistent mass matrix.

which the number of iterations in the two-grid studies was small and bounded above. A similar behavior has been found in the numerical studies of the multigrid case when more stabilization is added on the coarser grids, which is accompanied by a neglect of the consistency of the coarse grid problem. Finally, as a side effect, it was observed that artifacts occurring in the solutions for the standard Galerkin discretization of convection-dominated problems can be smoothed by the stabilization.

Future investigations on the described stabilization technique and its application to the presented multigrid algorithm include extensions to 2D and 3D problems as well as studies on the computational efficiency of the time-simultaneous approach. To exploit even more parallelism, this time-simultaneous multigrid algorithm could be combined with other approaches, such as parareal, MGRIT, or another multigrid version of parareal [WT22], which would extend the method by working parallel in time. As a forthcoming stabilization technique, it is convenient to apply the stabilization in time rather than in space, since all time steps are treated simultaneously in a global system in this method. Overall, it is reasonable to explore especially the multigrid case for convection-dominated problems in more detail, since less stabilization on coarser levels does not always show the desired effect; for example, adaptive control of the stabilization parameter or level-dependent numbers of smoothing steps might offer potential for improvement.

Acknowledgements

This work was supported by the Federal Ministry of Education and Research (BMBF) through the project “StroemungsRaum” 16ME0706K, which is part of the initiative “Neue Methoden und Technologien für das Exascale-Höchstleistungsrechnen” (SCALEXA) that aims at tackling the various challenges of large-scale scientific computing.

References

[Bal05] G. Bal. “On the Convergence and the Stability of the Parareal Algorithm to Solve Partial Differential Equations”. In: vol. 40. Jan. 2005, pp. 425–432.

[Bra07] D. Braess. *Finite Elements: Theory, Fast Solvers, and Applications in Solid Mechanics*. 3rd ed. Cambridge University Press, 2007.

[De +21] H. De Sterck, R. D. Falgout, S. Friedhoff, O. A. Krzysik, and S. P. MacLachlan. “Optimizing multigrid reduction-in-time and Parareal coarse-grid operators for linear advection”. In: *Numerical Linear Algebra with Applications* 28.4 (Mar. 2021).

[De +23] H. De Sterck, R. D. Falgout, O. A. Krzysik, and J. B. Schroder. “Efficient Multigrid Reduction-in-Time for Method-of-Lines Discretizations of Linear Advection”. In: *Journal of Scientific Computing* 96.1 (2023).

[DT22] W. Dong and H. Tang. *Convergence Analysis of Waveform Relaxation Method to Compute Coupled Advection-Diffusion-Reaction Equations*. May 2022. arXiv: 2205.01708 [math.NA].

[Dün+21a] J. Dünnebacke, S. Turek, C. Lohmann, A. Sokolov, and P. Zajac. “Increased space-parallelism via time-simultaneous Newton-multigrid methods for nonstationary nonlinear PDE problems”. In: *The International Journal of High Performance Computing Applications* 35.3 (Apr. 2021), pp. 211–225.

[Dün+21b] J. Dünnebacke, S. Turek, P. Zajac, and A. Sokolov. “A Time-Simultaneous Multigrid Method for Parabolic Evolution Equations”. In: *Numerical Mathematics and Advanced Applications ENUMATH 2019*. Ed. by Fred J. Vermolen and Cornelis Vuik. Cham: Springer International Publishing, 2021, pp. 333–342.

[FC03] C. Farhat and M. Chandesris. “Time-decomposed parallel time-integrators: Theory and feasibility studies for fluid, structure, and fluid-structure applications”. In: *International Journal for Numerical Methods in Engineering* 58 (Nov. 2003), pp. 1397–1434.

[Gan15] M. J. Gander. “50 Years of Time Parallel Time Integration”. In: *Multiple Shooting and Time Domain Decomposition Methods*. Ed. by Thomas Carraro, Michael Geiger, Stefan Körkel, and Rolf Rannacher. Cham: Springer International Publishing, 2015, pp. 69–113.

[GH07] M. J. Gander and L. Halpern. “Optimized Schwarz Waveform Relaxation Methods for Advection Reaction Diffusion Problems”. In: *SIAM Journal on Numerical Analysis* 45.2 (2007), pp. 666–697.

[Hug95] T. J.R. Hughes. “Multiscale phenomena: Green’s functions, the Dirichlet-to-Neumann formulation, subgrid scale models, bubbles and the origins of stabilized methods”. In: *Computer Methods in Applied Mechanics and Engineering* 127.1 (1995), pp. 387–401.

[JKL06] V. John, S. Kaya, and W. Layton. “A two-level variational multiscale method for convection-dominated convection-diffusion equations”. In: *Computer Methods in Applied Mechanics and Engineering* 195.33 (2006), pp. 4594–4603.

[JS08] V. John and E. Schmeyer. “A two-level variational multiscale method for convection-dominated convection-diffusion equations”. In: *Computer Methods in Applied Mechanics and Engineering* 198 (2008), pp. 475–494.

[JV96a] J. Janssen and S. Vandewalle. “Multigrid Waveform Relaxation on Spatial Finite Element Meshes: The Continuous-Time Case”. In: *SIAM Journal on Numerical Analysis* 33.2 (1996), pp. 456–474.

[JV96b] J. Janssen and S. Vandewalle. “Multigrid Waveform Relaxation on Spatial Finite Element Meshes: The Discrete-Time Case”. In: *SIAM Journal on Scientific Computing* 17.1 (1996), pp. 133–155.

[Lay02] W. Layton. “A connection between subgrid scale eddy viscosity and mixed methods”. In: *Applied Mathematics and Computation* 133.1 (2002), pp. 147–157.

[LDT22] C. Lohmann, J. Dünnebacke, and S. Turek. “Fourier analysis of a time–simultaneous two–grid algorithm using a damped Jacobi waveform relaxation smoother for the one–dimensional heat equation”. In: *Journal of Numerical Mathematics* 30.3 (Sept. 2022), pp. 173–207.

[LO87] C. Lubich and A. Ostermann. “Multi–grid dynamic iteration for parabolic equations”. In: *BIT Numerical Mathematics* 27 (June 1987), pp. 216–234.

[Loh+17] C. Lohmann, D. Kuzmin, J. N. Shadid, and S. Mabuza. “Flux-corrected transport algorithms for continuous Galerkin methods based on high order Bernstein finite elements”. In: *Journal of Computational Physics* 344 (2017), pp. 151–186.

[Not22] Y. Notay. “Rigorous convergence proof of space-time multigrid with coarsening in space”. In: *Numerical Algorithms* 89.2 (2022), pp. 675–699.

[OS20] B. W. Ong and J. B. Schroder. “Applications of time parallelization”. In: *Computing and Visualization in Science* 23.11 (2020).

[Qua13] A. Quarteroni. *Numerical Models for Differential Problems*. 2nd. Springer Publishing Company, Incorporated, 2013.

[QV94] A. Quarteroni and A. Valli. *Numerical Approximation of Partial Differential Equations*. Springer Berlin, Heidelberg, 1994.

[SS86] Y. Saad and M. H. Schultz. “GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems”. In: *Siam Journal on Scientific and Statistical Computing* 7 (1986), pp. 856–869.

[VH95] S. Vandewalle and G. Horton. “Fourier mode analysis of the multigrid waveform relaxation and time-parallel multigrid methods”. In: *Computing* 54.4 (1995), pp. 317–330.

[WT22] L. Wambach and S. Turek. *Numerical studies of a multigrid version of the parareal algorithm*. Tech. rep. Ergebnisberichte des Instituts für Angewandte Mathematik, Nummer 650. Fakultät für Mathematik, TU Dortmund, Mar. 2022.