Robust Multigrid for Isogeometric Analysis Based on Stable Splittings of Spline Spaces

Clemens Hofreither, Stefan Takacs

ROBUST MULTIGRID FOR ISOGEOMETRIC ANALYSIS BASED ON STABLE SPLITTINGS OF SPLINE SPACES

CLEMENS HOFREITHER[*] AND STEFAN TAKACS[†]

Abstract. We present a robust and efficient multigrid method for single-patch isogeometric discretizations using tensor product B-splines of maximum smoothness. Our method is based on a stable splitting of the spline space into a large subspace of "interior" splines which satisfy a robust inverse inequality, as well as one or several smaller subspaces which capture the boundary effects responsible for the spectral outliers which occur in Isogeometric Analysis. We then construct a multigrid smoother based on an additive subspace correction approach, applying a different smoother to each of the subspaces. For the interior splines, we use a mass smoother, whereas the remaining components are treated with suitably chosen Kronecker product smoothers or direct solvers.

We prove that the resulting multigrid method exhibits iteration numbers which are robust with respect to the spline degree and the mesh size. Furthermore, it can be efficiently realized for discretizations of problems in arbitrarily high geometric dimension. Some numerical examples illustrate the theoretical results and show that the iteration numbers also scale relatively mildly with the problem dimension.

Key words. Isogeometric Analysis, Multigrid methods, B-splines, Stable splittings, Subspace correction methods

AMS subject classifications. 65N55, 65N30, 65F08, 65D07

1. Introduction. Isogeometric Analysis (IgA) is a method for the numerical solution of partial differential equations (PDEs) introduced in the seminal paper [18] which has since attracted a sizable research community. Spline spaces, such as spaces spanned by tensor product B-splines or NURBS, are commonly used for geometry representation in industrial CAD systems. The foundational idea in IgA is to use such spline spaces both for the representation of the computational domain and for the discretization of the quantities of interest when solving a PDE. The overall goal is to create a tighter integration between geometric design and analysis.

There is a need for efficient solvers for the large, sparse linear systems which arise when applying isogeometric discretizations to boundary value problems. By now, most established solution strategies known from the finite element literature have been applied in one way or another to IgA: among these, direct solvers [2], non-overlapping and overlapping domain decomposition methods [19, 4, 5, 6], and multilevel and multigrid methods [1, 11, 17, 10, 15]. A recent contribution [20] constructs preconditioners based on fast solvers for Sylvester equations. The above list is certainly not comprehensive.

In IgA, we typically encounter as discretization parameters the mesh size and the spline degree. In the early IgA solver literature, the focus was on translating solvers from the finite element world to IgA with minimal adaptations. As a rule, it was found that such an approach results in methods that work well for low spline degrees, but deteriorate in performance as the degree is increased; often dramatically so. This motivated the search for IgA solvers that are robust not only with respect to the mesh size (which is often easy to achieve), but also with respect to the spline degree.

---

[*]Department of Computational Mathematics, Johannes Kepler University Linz, Austria (chofreither@numa.uni-linz.ac.at).

[†]Johann Radon Institute for Computational and Applied Mathematics (RICAM), Austrian Academy of Sciences (stefan.takacs@ricam.oeaw.ac.at).

Within the class of multigrid methods for IgA, advances towards a robust method were made using two approaches. In [9], a careful analysis of the symbol of isogeometric stiffness matrices served as the basis for the construction of multigrid methods. This theoretical approach is somewhat related to the technique known as Local Fourier Analysis (LFA) in the multigrid literature (see, e.g., [22]). It appears that the method presented in [9] is roughly comparable to the one studied in [16], which uses mass matrices as multigrid smoothers, an approach itself motivated by LFA. For both methods, an increase in the number of smoothing steps, roughly linearly with the spline degree, is required in order to maintain robust convergence. They can thus not be considered totally robust and efficient in the strict meaning that we will use in the present work.

A second approach towards robust and efficient multigrid was presented in [15]. Based on a robust inverse inequality and approximation error estimate in a large subspace of maximally smooth spline spaces derived in [21], it was shown that mass matrices can be used as robust smoothers in this large subspace. For the remaining, relatively few degrees of freedom, a low-rank correction was constructed. (These degrees of freedom are associated with the boundary of the domain and cannot be captured by LFA, which assumes periodic boundary conditions.) This approach resulted in a provably robust and efficient multigrid method for two-dimensional problems with splines of maximum smoothness. It was however not clear how to extend this approach efficiently to three and higher dimensions.

The present work can be viewed as a continuation of [15]. Based on the theoretical results from [21], we construct a splitting of the tensor product spline space into a large, regular interior part and several smaller spaces which capture boundary effects. The splitting is L_2-orthogonal and H^1-stable with respect to both the mesh size and the spline degree. This stability enables us to construct a multigrid smoother based on an additive subspace correction approach, applying a different smoother in each of the subspaces. In the regular interior subspace, we use a mass smoother. In the other subspaces, we construct smoothers which exploit the particular structure of the subspaces while still permitting an efficient application through a Kronecker product representation. In one small subspace associated with the corners of the domain, we apply a direct solver.

Unlike the low-rank correction approach from [15], the subspace correction approach generalizes easily to three-dimensional problems, and indeed to problems of arbitrary space dimension. We show that the method converges robustly with respect to mesh size and spline degree, and that one iteration is asymptotically not more expensive than an application of the stiffness matrix. The result is a quasi-optimal solution method for problems of arbitrary space dimensions.

It appears that the stable splitting of the tensor product spline space presented in Section 3 is an interesting theoretical result in its own right. It may have future applications to other aspects of IgA beyond the one presented here.

The remainder of the paper is organized as follows. In Section 2, we introduce the needed spline spaces and present an isogeometric model problem. We also present an algorithmic multigrid framework and an abstract convergence result which forms the basis of our later analysis. In Section 3, we derive the main new theoretical result used in our construction: the L_2-orthogonal and H^1-stable splitting of the spline space into a large, regular interior part and smaller spaces which capture boundary effects. In Section 4, we use this space splitting to construct a multigrid smoother based on the idea of additive subspace correction and show that it results in a robust solver. In Section 5, we present details on the computational realization of the proposed

smoother and show that it permits an efficient implementation in arbitrary space dimensions. In Section 6, we present numerical experiments which demonstrate the performance of the proposed method in practice.

2. Preliminaries.

2.1. Spline spaces and B-splines. Consider a subdivision of the interval $(0,1)$ into $m \in \mathbb{N}$ intervals of length $h = 1/m$. We introduce the spline space of degree $p \in \mathbb{N}$ with maximum smoothness,

$$S := \{u \in C^{p-1}(0,1) : u|_{((j-1)h, jh)} \in \mathcal{P}^p \quad \forall j = 1, \ldots, m\},$$

where $C^{p-1}(0,1)$ is the space of all $p-1$ times continuously differentiable functions on $(0,1)$ and $\mathcal{P}^p$ is the space of all polynomials of degree at most p. We have $n := \dim S = m + p$. As a basis for S, we use the normalized (i.e., satisfying a partition of unity; cf. [8]) B-splines with an open knot vector. In higher dimensions $d > 1$, we introduce the space of tensor product splines (cf. [8])

$$S^d := S \otimes \ldots \otimes S$$

defined over $(0,1)^d$ with $\dim S^d = n^d$ and the corresponding tensor product B-spline basis. For notational convenience, we assume that the same spline space S is used in each of the d coordinate directions. Both our construction and our analysis are however straightforward to generalize to the case where different spline spaces are used in different coordinate directions.

2.2. Isogeometric model problem. Let $\Omega = (0,1)^d$ with $d \in \mathbb{N}$. As a model problem, we consider a pure Neumann boundary value problem for the PDE $-\Delta u + u = f$. The variational formulation reads: find $u \in H^1(\Omega)$ such that

$$a(u,v) = \langle f, v \rangle \qquad \forall v \in H^1(\Omega), \tag{1}$$

where

$$a(u,v) = \int_\Omega (\nabla u \cdot \nabla v + uv) \, dx \qquad \forall u, v \in H^1(\Omega) \tag{2}$$

and f is a linear functional on $H^1(\Omega)$. We will sometimes refer to the operator $A : H^1(\Omega) \to H^1(\Omega)'$ given by $Av = a(v, \cdot)$, where $H^1(\Omega)'$ denotes the continuous dual. Note that $\|v\|_A^2 = a(v,v) = \|v\|_{H^1(\Omega)}^2$.

Discretizing (1) using tensor product splines, we seek $u_h \in S^d$ such that

$$a(u_h, v_h) = \langle f, v_h \rangle \qquad \forall v_h \in S^d. \tag{3}$$

We are interested in robust and efficient iterative solvers for the discrete problem (3). Here, by "robust" we mean that the number of iterations to solve the problem should stay uniformly bounded with respect to both the mesh size h and the spline degree p, and by "efficient" we mean that one iteration of the method should not be asymptotically more expensive than computing the product of the stiffness matrix with a vector. Combined, these properties allow us to solve (3) in quasi-optimal time.

In IgA, one introduces a bijective geometry map from Ω to the actual domain of interest in order to be able to treat more complicated computational domains. Basis functions on the transformed domain are defined by composing the basis functions on the reference domain with the inverse of the geometry map. Furthermore, one is often

interested in more general PDEs with varying and possibly matrix-valued coefficients. Discretizations for such more general problems can be preconditioned with a solver for the model problem (3), and the resulting condition number depends only on the geometry map and the coefficient functions, but not on discretization parameters like the mesh size h or the spline degree p. This principle has been widely used in the literature on IgA solvers (see, e.g., [9, 15]) and formalized in [20]. Therefore, a robust and efficient solver for the model problem (3) immediately yields robust and efficient solvers for a more general class of problems with "benign" geometry maps and mildly varying coefficients. This justifies the study of solvers for the model problem.

Three different refinement strategies for IgA discretizations were proposed in [18]: h-refinement (reducing the mesh size), p-refinement (increasing the spline degree), and the so-called k-refinement. The latter is unique to IgA and maintains the maximum possible smoothness C^{p-1} for the spline space of degree p. Already in [18], the favorable performance of k-refinement was observed, and it appears to be the most popular refinement strategy in the wider IgA literature. This motivates the study of solvers for spline spaces with maximum smoothness.

2.3. A multigrid method framework. Given a discretization space V and a coarse space $V_c \subset V$, we denote by $P : V_c \to V$ the canonical embedding. Let $A : V \to V'$ denote the operator in a (discretized) equation

$$Au = f$$

to be solved for $u \in V$. The corresponding coarse-space operator is given by $A_c := P'AP$. Furthermore, we assume that we are given a self-adjoint and positive definite smoothing operator $L : V \to V'$.

Given a previous iterate $u^{(k)}$, we let $u^{(k,0)} := u^{(k)}$ and perform $\nu \in \mathbb{N}$ *smoothing steps* given by

$$u^{(k,j)} := u^{(k,j-1)} + \tau L^{-1}(f - Au^{(k,j-1)}), \qquad j = 1, \ldots, \nu,$$

where $\tau > 0$ is a damping parameter. Then, we perform one *coarse-grid correction step* given by

$$u^{(k+1)} := u^{(k,\nu)} + PA_c^{-1}P'(f - Au^{(k,\nu)}).$$

Together, these updates describe one iteration $u^{(k)} \mapsto u^{(k+1)}$ of a *two-grid method*. Given an entire sequence of nested spaces $V_0 \subset \ldots \subset V_L = V$, we can replace the exact inversion of A_c in the coarse-grid correction step by one or two recursive applications of the two-grid method on the next coarser level V_{L-1}, and so on until we reach the coarsest level V_0, where an exact solver is used. Using one or two recursive iteration steps results in the *V-cycle* or the *W-cycle multigrid method*, respectively.

The following theorem is an abstract convergence result for the two-grid method with the abovementioned smoother. Its proof is given in [15, Theorem 3] and is based on a variant of the standard multigrid theory as developed by Hackbusch [14]. In [15, Theorem 4], it was shown that under the same assumptions also a W-cycle multigrid method converges.

THEOREM 1 ([15]). *Assume that there are constants C_A and C_I such that the inverse inequality*

$$(4) \qquad \qquad \|u\|_A^2 \le C_I \|u\|_L^2 \qquad \forall u \in V$$

and the approximation property for the A-orthogonal projector $T_c : V \to V_c$

$$(5) \qquad \qquad \|(I - T_c)u\|_L^2 \le C_A \|u\|_A^2 \qquad \forall u \in V$$

hold. Then the two-grid method converges for any choice of the damping parameter $\tau \in (0, C_I^{-1}]$ *and any number of smoothing steps* $\nu > \nu_0 := \tau^{-1} C_A$ *with rate* $q = \nu_0/\nu < 1$.

In particular, if C_A and C_I do not depend on the mesh size h and the spline degree p, then the two-grid method converges with a rate $q < 1$ which does not depend on h and p. In other words, the two-grid method is then robust.

In addition to properties (4) and (5), care must be taken that the smoother can be realized efficiently. In other words, it should be possible to apply the inverse L^{-1} with a computational cost which is roughly comparable to that for applying A.

3. Stable splittings of spline spaces. Consider first the univariate case, $d = 1$, with $\Omega = (0, 1)$. In [21], the subspace

$$S_0 := \left\{ u \in S : u^{(2l+1)}(0) = u^{(2l+1)}(1) = 0 \quad \forall l \in \mathbb{N}_0 \text{ with } 2l + 1 < p \right\}$$

of splines with vanishing odd derivatives of order less than p at the boundaries was introduced (denoted in [21] by $\widetilde{S}_{p,h}(\Omega)$). It is a large subspace of S in the sense that $\dim S_0 \geq \dim S - p$.

The subspace S_0 has the very desirable property of satisfying both a (first-order) approximation property and an inverse inequality, both with constants which are independent of the spline degree p. To formulate these results, let $Q_0 : L_2(\Omega) \to S_0$ denote the L_2-orthogonal projector into S_0, and let $\Pi_0 : H^1(\Omega) \to S_0$ denote the projector into S_0 which is orthogonal with respect to the scalar product

$$(u, v)_{H^1_\circ(\Omega)} := (\nabla u, \nabla v)_{L_2(\Omega)} + \frac{1}{|\Omega|} \left(\int_\Omega u(x) dx \right) \left(\int_\Omega v(x) dx \right).$$

We abbreviate the $L_2(\Omega)$-norm by $\| \cdot \|_0$, and the full $H^1(\Omega)$-norm and the seminorm by $\| \cdot \|_1$ and $| \cdot |_1$, respectively. Furthermore, we write c for a generic positive constant which does not depend on the mesh size h or the spline degree p.

THEOREM 2 ([21, Theorem 6.1]). *For any spline degree* $p \in \mathbb{N}$, *we have the inverse inequality*

$$|u|_1 \leq 2\sqrt{3} h^{-1} \|u\|_0 \qquad \forall u \in S_0.$$

THEOREM 3 ([21, Corollary 5.1], [15, Theorem 14]). *For any spline degree* $p \in \mathbb{N}$ *and any* $u \in H^1(\Omega)$, *we have the approximation error estimates*

$$\|(I - Q_0)u\|_0 \leq \sqrt{2} h |u|_1 \qquad and \qquad \|(I - \Pi_0)u\|_0 \leq \sqrt{2} h |u|_1.$$

Contrast these properties with the entire spline space S, which does satisfy a robust approximation property, but whose inverse inequality deteriorates with increasing degree p ([21]). On the other hand, a smaller space of only "interior" splines, built by discarding the p leftmost and p rightmost B-splines, does satisfy a robust inverse inequality but loses the approximation property.

We remark that the non-robustness of the inverse inequality in S is the root cause of the spectral "outliers" commonly observed when solving eigenvalue problems using IgA (cf. [3]). No such outliers appear in the space S_0.

3.1. A stable splitting in one dimension. Let $S_1 := S_0^{\perp_{L_2}}$ denote the L_2-orthogonal complement of S_0 in S. Consider the splitting of S into the direct sum

$$S = S_0 \oplus S_1 \quad \longleftrightarrow \quad u = Q_0 u + (I - Q_0)u$$

of S_0 and its complement, illustrated in Fig. 1. Due to orthogonality, we have

$$\|u\|_0^2 = \|Q_0 u\|_0^2 + \|(I - Q_0)u\|_0^2. \tag{6}$$

FIG. 1. *Bases for the space S_0 (left) and its orthogonal complement S_1 (right) for $p = 4$, $h = 1/20$. Here,* $\dim S_0 = 20$ *and* $\dim S_1 = 4$.

Crucially, we can prove that this splitting is stable also in the H^1-norm. This is a direct result of the space S_0 satisfying both an approximation property and an inverse inequality.

THEOREM 4. *For any spline $u \in S$, we have*

$$c^{-1}|u|_1^2 \leq |Q_0 u|_1^2 + |(I - Q_0)u|_1^2 \leq c|u|_1^2$$

and the corresponding result for the full H^1-norm.

Proof. The left inequality follows from the Cauchy-Schwarz inequality with $c = 2$. For the right inequality, we observe that

$$|Q_0 u|_1 \leq |\Pi_0 u|_1 + |(\Pi_0 - Q_0)u|_1 \leq |u|_1 + ch^{-1} \left(\|(I - \Pi_0)u\|_0 + \|(I - Q_0)u\|_0 \right)$$

because of the triangle inequality, the stability of the $H_\circ^1$-projector Π_0 in the H^1-seminorm and the robust inverse inequality in S_0 (Theorem 2). With the approximation error estimate (Theorem 3) we obtain H^1-stability of the L_2-projector,

$$|Q_0 u|_1 \leq c|u|_1. \tag{7}$$

The desired result follows from (7) and

$$|(I - Q_0)u|_1 \leq |u|_1 + |Q_0 u|_1 \leq (1 + c)|u|_1.$$

The result for the full H^1-norm follows by adding the identity (6). □

3.2. A stable splitting in two dimensions. The two-dimensional tensor product spline space is given by $S^2 = S \otimes S$. Since the tensor product distributes over direct sums, we obtain the splitting

$$S^2 = (S_0 \otimes S_0) \oplus (S_0 \otimes S_1) \oplus (S_1 \otimes S_0) \oplus (S_1 \otimes S_1) = S_{00} \oplus S_{01} \oplus S_{10} \oplus S_{11}$$

with the abbreviations $S_{\alpha_1, \alpha_2} := S_{\alpha_1} \otimes S_{\alpha_2}$ for $\alpha_j \in \{0, 1\}$. A visualization of this splitting is shown in Fig. 2. Note that the shaded regions do not correspond to the supports of the function spaces; in fact, each of the subspaces has global support. However, the shaded regions roughly correspond to regions where the corresponding functions are "largest", and their areas roughly correspond to the space dimensions. In view of this, it makes sense to think of S_{00} as an "interior" space, of S_{01} and S_{10} as "edge" spaces, and of S_{11} as a "corner" space.

FIG. 2. *Visualization of the splitting in 2D.*

Again, we can prove that the splitting is H^1-stable. In the following, we let $M : S \to S'$, $K : S \to S'$ denote the operators in the univariate spline space associated with the bilinear forms

$$\langle Mu, v \rangle := \int_0^1 u(x)v(x)\,dx, \quad \langle Ku, v \rangle := \int_0^1 u'(x)v'(x)\,dx \quad \forall u, v \in S,$$

that is, the one-dimensional mass and stiffness operators, respectively. For any $(\alpha_1, \alpha_2) \in \{0, 1\}^2$, we furthermore introduce the abbreviations

$$Q_1 := I - Q_0 : S \to S_1, \qquad\qquad Q_{\alpha_1, \alpha_2} := Q_{\alpha_1} \otimes Q_{\alpha_2} : S^2 \to S_{\alpha_1, \alpha_2}$$
$$K_{\alpha_j} := Q'_{\alpha_j} K Q_{\alpha_j} : S_{\alpha_j} \to S'_{\alpha_j}, \qquad M_{\alpha_j} := Q'_{\alpha_j} M Q_{\alpha_j} : S_{\alpha_j} \to S'_{\alpha_j}.$$

As tensor products of $L_2(0, 1)$-orthogonal projectors, the projectors Q_{α_1, α_2} are $L_2(\Omega)$-orthogonal, as one easily verifies. Thus the splitting of S^2 given above is a direct sum of L_2-orthogonal subspaces, and we have

$$\|u\|_0^2 = \sum_{(\alpha_1, \alpha_2)} \|Q_{\alpha_1, \alpha_2} u\|_0^2, \tag{8}$$

where here and below sums over (α_1, α_2) are taken to run over the set $\{0, 1\}^2$.

THEOREM 5. *For any tensor product spline $u \in S^2$, we have*

$$c^{-1}|u|_1^2 \le \sum_{(\alpha_1, \alpha_2)} |Q_{\alpha_1, \alpha_2} u|_1^2 \le c|u|_1^2,$$

and the corresponding result for the full H^1-norm.

Proof. The left inequality follows by the Cauchy-Schwarz inequality. For the right one, fix $(\alpha_1, \alpha_2) \in \{0, 1\}^2$. The H^1-seminorm can be written using tensor products of one-dimensional operators as

$$|Q_{\alpha_1, \alpha_2} u|_1^2 = |Q_{\alpha_1, \alpha_2} u|_{K \otimes M}^2 + |Q_{\alpha_1, \alpha_2} u|_{M \otimes K}^2. \tag{9}$$

The first term can be rewritten, using the definitions and basic identities for tensor products of operators, as

$$|Q_{\alpha_1, \alpha_2} u|_{K \otimes M}^2 = \langle Q'_{\alpha_1, \alpha_2}(K \otimes M)Q_{\alpha_1, \alpha_2} u, u \rangle = \langle (K_{\alpha_1} \otimes M_{\alpha_2})u, u \rangle.$$

Due to orthogonality and Theorem 4, we have $M_0 + M_1 = M$ and $K_0 + K_1 \le cK$, where all summands are positive semidefinite operators. This implies that we can estimate, in the spectral sense, $K_{\alpha_1} \le cK$ and $M_{\alpha_2} \le M$, and we obtain

$$|Q_{\alpha_1, \alpha_2} u|_{K \otimes M}^2 \le c|u|_{K \otimes M}^2.$$

Treating the second term in (9) analogously, we obtain

$$|Q_{\alpha_1,\alpha_2}u|_1^2 \le c(|u|_{K\otimes M}^2 + |u|_{M\otimes K}^2) = c|u|_1^2.$$

The right inequality now follows by summing up over all (α_1,α_2). The result for the full H^1-norm follows by adding the identity (8).   □

3.3. Stable splitting in arbitrary dimensions. For any $d \in \mathbb{N}$, we define multiindices $\alpha \in \{0,1\}^d$ and generalize the notations from Section 3.2 in the straightforward way to higher dimensions. We obtain the splitting into the direct sum of 2^d subspaces

$$S^d = \bigoplus_\alpha S_\alpha, \qquad \text{where} \qquad S_\alpha = S_{\alpha_1} \otimes \ldots \otimes S_{\alpha_d}.$$

The L_2-orthogonal projectors into the subspaces are given by

$$Q_\alpha = Q_{\alpha_1} \otimes \ldots \otimes Q_{\alpha_d} : S^d \to S_\alpha.$$

As in the two-dimensional case, we can prove that this splitting is H^1-stable.

THEOREM 6. *For any d-dimensional tensor product spline $u \in S^d$, we have*

$$c^{-1}|u|_1^2 \le \sum_{\alpha=(0,\ldots,0)}^{(1,\ldots,1)} |Q_\alpha u|_1^2 \le c|u|_1^2$$

and the corresponding result for the full H^1-norm.

Proof. Completely analogous to Theorem 5.   □

4. Construction of a robust multigrid method. Recall that S was a univariate spline space of degree p and mesh size h. Let $S_c \subset S$ be the analogous coarse spline space with uniform mesh size $2h$. For the construction of our two-grid method in d dimensions in accordance with the framework introduced in Section 2.3, we let

$$V := S^d, \qquad V_c := (S_c)^d \subset V.$$

The prolongation $P : V_c \to V$ is the canonical embedding of the coarse tensor product spline space in the fine one. It can be represented as the d-fold tensor product of prolongations for the univariate spline spaces, $I : S_c \to S$.

The following result states that a robust approximation error estimate holds for the Galerkin projector to the coarse spline space. It was proved for $d = 1$ and $d = 2$ in [15]. We extend the proof to arbitrary dimensions in the Appendix.

LEMMA 7. *The A-orthogonal projector $T_c : S^d \to (S_c)^d$ satisfies the approximation error estimate*

$$\|(I - T_c)u\|_{L_2(\Omega)} \le ch\|u\|_A \qquad \forall u \in S^d$$

with a constant c which is independent of h and p (but may depend on d).

In the following subsections, we construct a smoother for the two-grid method on these nested spline spaces which leads to a robust and efficient iterative method.

4.1. A multigrid smoother based on subspace correction. In each of the 2^d subspaces $S_\alpha \subset S^d$, $\alpha \in \{0,1\}^d$, defined in Section 3.3, we prescribe a local, symmetric and positive definite smoothing operator $L_\alpha : S_\alpha \to S'_\alpha$. The overall smoothing operator is then given by the additive subspace operator

$$(10) \qquad L := \sum_\alpha Q'_\alpha L_\alpha Q_\alpha : S^d \to S^{d'},$$

from S^d to its dual $S^{d'}$, and its inverse has the form

$$L^{-1} = \sum_\alpha L_\alpha^{-1} Q'_\alpha : S^{d'} \to S^d.$$

The assumptions of Theorem 1 for L, and thus the convergence of the two-grid method with such a smoother, can be guaranteed under simple assumptions on the subspace operators L_α, as the following two lemmas show. The stability of the space splitting is crucial to both proofs. Although we do not explicitly use any results from the literature on subspace correction methods, we rely heavily on the ideas developed therein; cf., e.g., [23, 13].

LEMMA 8. *Assume that for every* $\alpha \in \{0,1\}^d$, *we have*

$$(11) \qquad \langle Av_\alpha, v_\alpha \rangle \leq c \langle L_\alpha v_\alpha, v_\alpha \rangle \qquad \forall v_\alpha \in S_\alpha.$$

Then the subspace correction smoother satisfies

$$\langle Av, v \rangle \leq c \langle Lv, v \rangle \qquad \forall v \in S^d.$$

Proof. Due to Theorem 6 and (11), we have

$$\langle Av, v \rangle \leq c \sum_\alpha \langle AQ_\alpha v, Q_\alpha v \rangle \leq c \sum_\alpha \langle L_\alpha Q_\alpha v, Q_\alpha v \rangle = c \langle Lv, v \rangle. \qquad \square$$

LEMMA 9. *Assume that for every* $\alpha \in \{0,1\}^d$, *we have*

$$(12) \qquad \langle L_\alpha v_\alpha, v_\alpha \rangle \leq c \langle (A + h^{-2} M^d) v_\alpha, v_\alpha \rangle \qquad \forall v_\alpha \in S_\alpha,$$

where $M^d : S^d \to S^{d'}$ *is the mass operator in the tensor product spline space. Then the subspace correction smoother satisfies*

$$\|(I - T_c)v\|_L \leq c\|v\|_A \qquad \forall v \in S^d.$$

Proof. From (12), Theorem 6 and L_2-orthogonality, we obtain

$$\langle Lv, v \rangle \leq c \sum_\alpha \langle (A + h^{-2} M^d) Q_\alpha v, Q_\alpha v \rangle \leq c \langle (A + h^{-2} M^d) v, v \rangle.$$

Thus, it follows

$$\|(I - T_c)v\|_L^2 \leq c\|(I - T_c)v\|_A^2 + c h^{-2} \|(I - T_c)v\|_{M^d}^2 \leq c\|v\|_A^2,$$

where we used the stability of the coarse-grid projector and the coarse-grid approximation property Lemma 7. $\qquad \square$

4.2. Choice of the local smoothing operators. We now construct suitable local operators L_α which satisfy the assumptions of Lemma 8 and Lemma 9. In the two-dimensional case, the operator associated with the bilinear form (2) admits the representation

$$A = K \otimes M + M \otimes K + M \otimes M$$

in terms of the stiffness and mass operators for the univariate case. Restricting A to a subspace $S_\alpha = S_{\alpha_1, \alpha_2}$, we obtain

$$A_\alpha := Q'_\alpha A Q_\alpha = K_{\alpha_1} \otimes M_{\alpha_2} + M_{\alpha_1} \otimes K_{\alpha_2} + M_{\alpha_1} \otimes M_{\alpha_2}.$$

The inverse inequality in S_0 (Theorem 2) allows us to estimate

$$K_0 \leq \sigma M_0,$$

where $\sigma = 12h^{-2}$. We obtain subspace smoothers L_α by replacing K_0 by σM_0,

$$
\begin{aligned}
A_{00} &\leq (1 + 2\sigma) M_0 \otimes M_0 & &=: L_{00}, \\
A_{01} &\leq M_0 \otimes ((1 + \sigma) M_1 + K_1) & &=: L_{01}, \\
A_{10} &\leq ((1 + \sigma) M_1 + K_1) \otimes M_0 & &=: L_{10}, \\
A_{11} &= M_1 \otimes M_1 + K_1 \otimes M_1 + M_1 \otimes K_1 =: L_{11},
\end{aligned}
$$

where (11), the assumption of Lemma 8, holds by construction. It is easy to see that each L_α can be spectrally bounded from above by a constant times the matrix $Q'_\alpha (A + h^{-2} M \otimes M) Q_\alpha$, which proves the assumption (12) of Lemma 9. Using the statements of these two lemmas, Theorem 1 implies the two-grid convergence.

The same approach generalizes to higher dimensions, and we illustrate this in the three-dimensional setting. Here, we have

$$A = K \otimes M \otimes M + M \otimes K \otimes M + M \otimes M \otimes K + M \otimes M \otimes M.$$

Again, we define A_α as above and obtain the operators L_α by replacing K_0 by σM_0,

$$
\begin{aligned}
A_{000} &\leq (1 + 3\sigma) M_0 \otimes M_0 \otimes M_0 & &=: L_{000}, \\
A_{001} &\leq M_0 \otimes M_0 \otimes ((1 + 2\sigma) M_1 + K_1) & &=: L_{001}, \\
A_{010} &\leq M_0 \otimes ((1 + 2\sigma) M_1 + K_1) \otimes M_0 & &=: L_{010}, \\
A_{100} &\leq ((1 + 2\sigma) M_1 + K_1) \otimes M_0 \otimes M_0 & &=: L_{100}, \\
A_{011} &\leq M_0 \otimes ((1 + \sigma) M_1 \otimes M_1 + K_1 \otimes M_1 + M_1 \otimes K_1) & &=: L_{011}, \\
A_{110} &\leq ((1 + \sigma) M_1 \otimes M_1 + K_1 \otimes M_1 + M_1 \otimes K_1) \otimes M_0 & &=: L_{110}, \\
A_{101} &\leq K_1 \otimes M_0 \otimes M_1 + (1 + \sigma) M_1 \otimes M_0 \otimes M_1 + M_1 \otimes M_0 \otimes K_1 & &=: L_{101}, \\
A_{111} &= M_1 \otimes M_1 \otimes M_1 + K_1 \otimes M_1 \otimes M_1 + M_1 \otimes K_1 \otimes M_1 + M_1 \otimes M_1 \otimes K_1 =: L_{111}.
\end{aligned}
$$

We point out that, whereas L_{011} and L_{110} permit a tensor product factorization, the operator L_{101} cannot directly be factorized due to the ordering of the involved spaces. However, the tensor product space S_{101} is isomorphic to S_{011} by a simple swapping of the order of the involved tensor products. We exploit this in Section 5.3 below by a simple renumbering of the degrees of freedom in order to obtain an efficient method for inverting L_{101}.

It is clear that the rule of replacing K_0 by σM_0 in each operator A_α to obtain L_α extends directly to arbitrary dimension d. By the same arguments as above, we see that the resulting subspace correction smoother satisfies the assumptions of Lemma 8 and Lemma 9. Thus Theorem 1 shows that the resulting two-grid method converges robustly with respect to h and p. We summarize this in the following theorem.

THEOREM 10. *For any $d \in \mathbb{N}$, there exist choices for τ and ν, independent of h and p, such that the two-grid method in S^d with the smoother induced by the subspace operators L_α as constructed above converges with a rate $q < 1$ which does not depend on the grid size h or the spline degree p.*

The robust convergence of the W-cycle multigrid method follows using standard arguments, cf. [14, 15].

5. Computational realization. In Section 4, we have proposed a smoother and shown that it leads to a robust two-grid method. In this section, we provide details on the realization of the method and show that it permits an efficient implementation.

5.1. Computation of a basis for S_0 and S_1. In order to be able to work with the space S_0 and its orthogonal complement, we require bases for them. The aim of this subsection is to provide an algorithm for computing such bases as linear combinations of B-splines.

Recall that the univariate spline space S with m knot spans of width $h = 1/m$, degree p and maximum smoothness C^{p-1} has dimension $n = m + p$. Let

$$\mathcal{B} := \{\varphi_1, \ldots, \varphi_n\}$$

denote the normalized (i.e., satisfying a partition of unity, cf. [8]) B-spline basis of S. We have $\operatorname{supp} \varphi_j = [(j - p - 1)h, jh] \cap [0, 1]$. All interior B-splines

$$\mathcal{B}^I := \{\varphi_{p+1}, \ldots, \varphi_{n-p}\}$$

vanish with all their derivatives up to the $p - 1$st at the boundaries of the interval $[0, 1]$ and therefore lie in S_0. (Here and in the following we assume that $p + 1 \leq m$ such that $\mathcal{B}^I$ is nonempty.)

It remains to find linear combinations of the first and last p B-splines which complete $\mathcal{B}^I$ to a basis of S_0. Recall that $u \in S$ lies in S_0 iff

$$u^{(2l+1)}(0) = u^{(2l+1)}(1) = 0 \quad \forall l \in \mathbb{N}_0 \text{ with } 2l + 1 < p.$$

Consider first the left boundary. We need to satisfy $k := \lfloor p/2 \rfloor$ conditions on the derivatives of the splines. Let

$$\tilde{D} = \left(h^{2i-1} \varphi_j^{(2i-1)}(0) \right)_{i=1,\ldots,k,\ j=1,\ldots,p} \in \mathbb{R}^{k \times p}$$

denote the matrix of the relevant B-spline derivatives at 0, scaled with a suitable power of h in order to avoid numerical instabilities. We pad $\tilde{D}$ with $p - k$ zero rows to obtain a square matrix $D \in \mathbb{R}^{p \times p}$. Computing the singular value decomposition (SVD), we obtain

$$D = U \Sigma V^\top$$

with $U, V \in \mathbb{R}^{p \times p}$ orthogonal and $\Sigma \in \mathbb{R}^{p \times p}$ being the diagonal matrix of singular values in descending order. By construction, Σ contains k nonzero and $p - k$ zero singular values. Therefore, the rightmost $p - k$ columns of V span the kernel of D, and the linear combinations

$$\mathcal{B}_0^L := \left\{ \sum_{i=1}^{p} V_{i,j} \varphi_i : j = p - k + 1, \ldots, p \right\}$$

lie in S_0. By the analogous procedure at the right boundary, we compute a set $\mathcal{B}_0^R$ of $p - k$ linear combinations of the last p B-splines. Then, the functions in the set

$$\mathcal{B}_0 := \mathcal{B}_0^L \cup B^I \cup \mathcal{B}_0^R$$

are by construction linearly independent and lie in S_0. Since $n_0 := |\mathcal{B}_0| = n - 2k = \dim S_0$, we have

$$\operatorname{span} \mathcal{B}_0 = S_0.$$

In practice, we collect the coefficients in a sparse block diagonal matrix

$$P_0 = \begin{bmatrix} V^L[:, p - k + 1 : p] & & \\ & I_{n-2p} & \\ & & V^R[:, p - k + 1 : p] \end{bmatrix} \in \mathbb{R}^{n \times n_0},$$

where $V^L[:, p - k + 1 : p] \in \mathbb{R}^{p \times (p-k)}$ denotes the last $p - k$ columns of the matrix V computed for the left boundary, analogously V^R that for the right boundary, and I_d is the $d \times d$ identity matrix. Then clearly, splines in S_0 can be uniquely represented in terms of the B-spline basis as

$$u \in S_0 \quad \Longleftrightarrow \quad \exists \underline{u} \in \mathbb{R}^{n_0} : u = \sum_{j=1}^{n} (P_0 \underline{u})_j \varphi_j.$$

Since the SVD produces an orthonormal basis, collecting the remaining columns of V^L and V^R in a second sparse block matrix

$$P_\perp = \begin{bmatrix} V^L[:, 1 : k] & 0 \\ 0 & 0 \\ 0 & V^R[:, 1 : k] \end{bmatrix} \in \mathbb{R}^{n \times 2k}$$

satisfies $P_0^\top P_\perp = 0$. In fact, the columns of the concatenation $\begin{bmatrix} P_0 & P_\perp \end{bmatrix}$ form an orthonormal basis of $\mathbb{R}^n$. Let

$$P_1 := \underline{M}^{-1} P_\perp \in \mathbb{R}^{n \times 2k},$$

where $\underline{M}$ denotes the $\mathcal{B}$-mass matrix. Note that P_1 is no longer sparse. Furthermore, let $\underline{u} \in \mathbb{R}^{n_0}$ and $\underline{v} \in \mathbb{R}^{2k}$ with associated splines

$$u = \sum_{j=1}^{n} (P_0 \underline{u})_j \varphi_j, \qquad v = \sum_{j=1}^{n} (P_1 \underline{v})_j \varphi_j.$$

By construction, $u \in S_0$. We have

$$\langle u, v \rangle_{L_2(\Omega)} = \langle \underline{M} P_0 \underline{u}, \underline{M}^{-1} P_\perp \underline{v} \rangle = \underline{u}^\top P_0^\top P_\perp \underline{v} = 0.$$

Since this holds for all $u \in S_0$, v lies in the L_2-orthogonal complement of S_0. All in all, we have constructed basis representations or "prolongation matrices"

$$P_0 \in \mathbb{R}^{n \times (n-2k)}, \qquad P_1 = \underline{M}^{-1} P_\perp \in \mathbb{R}^{n \times 2k}$$

for S_0 and its L_2-orthogonal complement S_1, respectively.

For $d > 1$, we let $\alpha \in \{0, 1\}^d$ and introduce the Kronecker products

$$P_\alpha := P_{\alpha_1} \otimes \ldots \otimes P_{\alpha_d} \in \mathbb{R}^{n^d \times n_\alpha},$$

where $n_\alpha = \dim S_\alpha$, which represent bases for the spaces S_α in terms of the coefficients of linear combinations of the tensor product B-spline basis $\mathcal{B}^{\otimes d}$.

5.2. Implementation of the subspace correction smoother. For any $\alpha \in \{0,1\}^d$, the matrices P_α as defined in Section 5.1 describe a basis for S_α. Let $\underline{L}_\alpha \in \mathbb{R}^{n_\alpha \times n_\alpha}$ be the (symmetric and positive definite) matrix representation of $L_\alpha : S_\alpha \to S'_\alpha$ (as defined in Section 4.2) with respect to that basis. Then the matrix representation of

$$L^{-1} = \sum_\alpha L_\alpha^{-1} Q'_\alpha = \sum_\alpha I_{S_\alpha \to S^d} L_\alpha^{-1} I_{S^{d'} \to S'_\alpha} Q'_\alpha I_{S^{d'} \to S'_\alpha}$$

is given by

$$(13) \qquad \underline{L}^{-1} = \sum_\alpha P_\alpha \underline{L}_\alpha^{-1} P_\alpha^\top \underline{M} P_\alpha \underline{M}_\alpha^{-1} P_\alpha^\top = \sum_\alpha P_\alpha \underline{L}_\alpha^{-1} P_\alpha^\top,$$

where we used that the matrix representation of the embedding $I_{S_\alpha \to S^d}$ is P_α and the matrix representation of the L_2-projector Q_α is

$$\underline{M}_\alpha^{-1} P_\alpha^\top \underline{M}, \quad \text{where} \quad \underline{M}_\alpha = P_\alpha^\top \underline{M} P_\alpha.$$

Hence (13) can be used to implement the subspace correction smoother using only the prolongation matrices P_α and a fast method for applying $\underline{L}_\alpha^{-1}$. It is never necessary to explicitly apply the L_2-projectors Q_α. Furthermore, due to the use of additive subspace correction, the residual needs to be computed only once, and the individual subspace corrections may be done in parallel.

5.3. Inversion of the subspace operators. The final required algorithmic component is a fast method for applying the inverse of the local smoothing matrices $\underline{L}_\alpha \in \mathbb{R}^{n_\alpha \times n_\alpha}$. We illustrate this in the three-dimensional setting as described in Section 4.2, but the principles are the same regardless of dimension. A detailed discussion of the computational costs for arbitrary dimension is given in Section 5.4.

Interior space and face spaces. The interior space S_{000} and the face spaces $S_{001}, S_{010}, S_{100}$ contain the complement space S_1 as a factor space at most once, and thus the matrices associated with their smoothing operators can be represented as Kronecker products of three one-dimensional discretization matrices, e.g.,

$$\underline{L}_{000} = (1 + 3\sigma)\underline{M}_0 \otimes \underline{M}_0 \otimes \underline{M}_0, \quad \underline{L}_{001} = \underline{M}_0 \otimes \underline{M}_0 \otimes ((1 + 2\sigma)\underline{M}_1 + \underline{K}_1).$$

Here the symmetric matrices $\underline{M}_\beta, \underline{K}_\beta \in \mathbb{R}^{\dim S_\beta \times \dim S_\beta}$, $\beta \in \{0,1\}$, are the matrix representations of M_β and K_β, respectively, with respect to the bases described by P_β as computed in Section 5.1 above. For $\beta = 0$, $\underline{M}_\beta$ and $\underline{K}_\beta$ have dimension $\mathcal{O}(n)$ and bandwidth $\mathcal{O}(p)$, whereas for $\beta = 1$ they have dimension $\mathcal{O}(p)$ and are dense.

Since the Kronecker product can be inverted componentwise, we obtain, e.g.,

$$\underline{L}_{001}^{-1} = \underline{M}_0^{-1} \otimes \underline{M}_0^{-1} \otimes ((1 + 2\sigma)\underline{M}_1 + \underline{K}_1)^{-1}.$$

Instead of computing this (dense) inverse explicitly, we employ the algorithm described by de Boor [7] for computing the application of a Kronecker product of matrices to a vector, given only routines for applying the individual Kronecker factors. For the latter, we use Cholesky factorization.

Edge spaces. The spaces $S_{011}, S_{110}, S_{101}$ contain the complement space S_1 as a factor twice. In S_{011}, the matrix to be inverted has the form

$$\underline{L}_{011} = \underline{M}_0 \otimes ((1 + \sigma)\underline{M}_1 \otimes \underline{M}_1 + \underline{K}_1 \otimes \underline{M}_1 + \underline{M}_1 \otimes \underline{K}_1).$$

It again has Kronecker product structure and can be inverted using the algorithm described in the previous case. The same holds for S_{110}.

In the case of the space S_{101}, the associated matrix

$$\underline{L}_{101} = \underline{K}_1 \otimes \underline{M}_0 \otimes \underline{M}_1 + (1+\sigma)\underline{M}_1 \otimes \underline{M}_0 \otimes \underline{M}_1 + \underline{M}_1 \otimes \underline{M}_0 \otimes \underline{K}_1$$

does not permit a Kronecker product factorization due to the order of the involved spaces. However, by a simple renumbering of the degrees of freedom, S_{101} can be identified with S_{011}, and then $\underline{L}_{011}^{-1}$ can be applied as above.

Alternatively, the matrix $\underline{L}_{101}$ could be directly computed and inverted in its entirety using Cholesky factorization. This would exceed asymptotically (for $p \to \infty$) the computational costs derived in the following subsection, however this slowdown appears to be negligible in practice. For $d > 3$, this shortcut seems no longer viable.

Corner space. The space S_{111} is the tensor product of the three complement spaces and has dimension $\dim(S_1)^3 \leq p^3$. The associated matrix

$$\underline{L}_{111} = M_1 \otimes M_1 \otimes M_1 + K_1 \otimes M_1 \otimes M_1 + M_1 \otimes K_1 \otimes M_1 + M_1 \otimes M_1 \otimes K_1$$

is dense and is inverted by means of its Cholesky factorization.

5.4. Computational costs. We now study the computational complexity for applying the subspace correction smoother in the general d-dimensional setting. In our analysis, we ignore multiplicative constants which depend only on d. Repeatedly, we make use of the fact that the Cholesky factorization of a symmetric matrix of dimension N and bandwidth q can be computed in $\mathcal{O}(Nq^2)$ operations, and its inverse can then be applied in $\mathcal{O}(Nq)$ operations. If the matrix is not banded but dense, the factorization and inversion require $\mathcal{O}(N^3)$ and $\mathcal{O}(N^2)$ operations, respectively (cf. [12]).

By the renumbering of degrees of freedom described in Section 5.3, we can always rearrange the factor spaces such that we only need to consider spaces of the form

$$\underbrace{S_0 \otimes \ldots \otimes S_0}_{k \text{ times}} \otimes \underbrace{S_1 \otimes \ldots \otimes S_1}_{d-k \text{ times}}.$$

The smoothing matrices to be inverted, constructed as in Section 5.3, have the form

$$\underline{L}_{\{k,d-k\}} := \underbrace{\underline{M}_0 \otimes \ldots \otimes \underline{M}_0}_{k \text{ times}} \otimes \underline{X}_{d-k},$$

where $\underline{X}_j \in \mathbb{R}^{(\dim S_1)^j \times (\dim S_1)^j}$ is a dense, symmetric matrix. Recall that $\dim S_1 \leq p$.

Setup costs. The computation of the basis for S_0 and its L_2-orthogonal complement as described in Section 5.1 requires computing the SVD of two matrices of dimension $\mathcal{O}(p)$ as well as $\mathcal{O}(p)$ applications of the inverse of $\underline{M}$, which has dimension $n = m + p$ and bandwidth $\mathcal{O}(p)$, where m is the number of subintervals. The costs for this step are thus $\mathcal{O}(p^3 + np^2) = \mathcal{O}(p^3 + mp^2)$.

The one-dimensional mass matrix in S_0, $\underline{M}_0$, has dimension $\mathcal{O}(m)$ and bandwidth $\mathcal{O}(p)$ and thus requires $\mathcal{O}(mp^2)$ operations to factorize.

The matrices $\underline{X}_j$, $j = 1, \ldots, d$, are dense and therefore require $\mathcal{O}(p^{3j})$ operations to factorize.

The overall setup costs are therefore $\mathcal{O}(mp^2 + p^{3d})$.

Application costs. After factorization, the cost for applying the inverse $\underline{M}_0^{-1}$ is $\mathcal{O}(mp)$, and for $\underline{X}_j^{-1}$, it is $\mathcal{O}(p^{2j})$. To apply $\underline{L}_{\{k,d-k\}}^{-1}$ using the Kronecker product

algorithm from [7], we need to perform m^{d-1} applications of each of the k factors $\underline{M}_0^{-1}$ and m^k applications of $\underline{X}_j^{-1}$. Thus, the cost is $\mathcal{O}(km^d p + m^k p^{2(d-k)})$.

The inverse of $\underline{L}_{\{k,d-k\}}$ needs to be applied $\binom{d}{k}$ times since that is the number of multiindices $\alpha \in \{0,1\}^d$ which permute to $(0,\ldots,0,1,\ldots,1)$ with exactly k leading zeros. The binomial coefficient satisfies $\binom{d}{k} = \mathcal{O}(2^d/\sqrt{d})$ and in particular can be bounded from above by a constant which depends only on d. The overall cost for one application of the subspace correction smoother is then

$$\sum_{k=0}^{d} \binom{d}{k} \mathcal{O}(km^d p + m^k p^{2(d-k)}) = \mathcal{O}\left(m^d p + \max_{k=0,\ldots,d} m^k p^{2(d-k)}\right) = \mathcal{O}(m^d p + p^{2d}).$$

Overall costs. For $d \geq 2$, we have $mp^2 \leq m^2 + p^4 \leq m^d + p^{2d}$. Therefore, the overall costs for setting up and applying the smoother are bounded by

$$\mathcal{O}(m^d p + p^{3d}).$$

Assuming $p^2 \lesssim m$, the overall costs are asymptotically not more expensive than one application of the stiffness matrix, which has complexity $\mathcal{O}(n^d p^d) = \mathcal{O}(m^d p^d + p^{2d})$.

In a multigrid setting, assuming $\mathcal{O}(\log \mathfrak{m})$ levels with $m = \mathfrak{m}, \frac{\mathfrak{m}}{2}, \frac{\mathfrak{m}}{4}, \frac{\mathfrak{m}}{8} \ldots$ intervals per dimension, one obtains for $d \geq 2$ by summing up the overall costs of

$$\mathcal{O}(\mathfrak{m}^d p + (\log \mathfrak{m}) p^{3d}) \qquad \text{and} \qquad \mathcal{O}(\mathfrak{m}^d p + \mathfrak{m} p^{2d} + (\log \mathfrak{m}) p^{3d})$$

for smoothing in the V-cycle and the W-cycle, respectively. The full complexity including the costs for the exact coarse-grid solver and the intergrid transfers is asymptotically the same. Under mild assumptions on the relation between p and $\mathfrak{m}$, again the overall effort is asymptotically not higher than that for one application of the stiffness matrix.

6. Numerical experiments.

6.1. Experiments for the model problem.
We solve the problem (1), i.e.,

$$-\Delta u + u = f \quad \text{in } \Omega = (0,1)^d, \qquad \partial_n u = 0 \quad \text{on } \partial\Omega$$

for $d = 1, 2, 3$ with the right-hand side

$$(14) \qquad\qquad f(x) = d\pi^2 \prod_{j=1}^{d} \sin(\pi(x_j + \tfrac{1}{2})).$$

We perform a (tensor product) B-spline discretization using equidistant knot spans and maximum-continuity splines for varying spline degrees p. We refer to the coarse discretization with only one single interval as level $\ell = 0$ and perform uniform, dyadic refinement to obtain the finer discretization levels ℓ with $2^{\ell d}$ elements and $h_\ell = 2^{-\ell}$.

We set up a V-cycle multigrid method as described in Section 2.3 and using on each level the proposed smoother (10) as constructed in Section 4. We always use one pre- and one post-smoothing step with $\tau = 1$. The parameter σ was chosen as $\frac{1}{0.09} h^{-2}$ in 1D, $\frac{1}{0.18} h^{-2}$ in 2D, and $\frac{1}{0.19} h^{-2}$ in 3D. In each test, the coarsest grid was chosen in such a way that the spaces S_0 on each higher level are non-empty, i.e., such that the smoother is well-defined. We perform tests both using the V-cycle multigrid method and a conjugate gradient solver preconditioned with one V-cycle.

TABLE 1
Iteration numbers: unit interval (1D)

	$\ell \setminus p$	2	3	4	5	6	7	8	9	10	11	12	13	14
	9	33	34	34	33	33	33	32	31	31	31	28	28	29
V-cycle	8	33	34	34	32	33	33	31	30	30	31	28	28	27
	7	33	34	34	32	33	33	31	28	30	29	28	25	26
	9	13	13	13	13	13	13	13	13	12	12	12	12	12
PCG	8	13	13	13	13	13	13	12	12	12	12	12	12	11
	7	13	13	13	13	13	12	12	12	12	11	11	11	11

TABLE 2
Iteration numbers: unit square (2D).

	$\ell \setminus p$	2	3	4	5	6	7	8	9	10
	8	38	39	39	39	38	38	37	37	36
V-cycle	7	38	39	39	38	38	37	36	36	34
	6	38	38	38	37	37	35	34	34	32
	5	36	37	34	34	32	30	28	26	24
	8	14	14	14	14	14	14	14	14	13
PCG	7	14	14	14	14	14	14	14	13	13
	6	14	14	14	14	14	13	13	13	12
	5	14	14	13	13	13	12	11	11	10

TABLE 3
Iteration numbers: unit cube (3D).

	$\ell \setminus p$	2	3	4	5	6	7
	6	46	44	43	43	42	41
V-cycle	5	44	43	42	39	38	35
	4	39	36	32	29	25	23
	3	30	42	18	22	12	17
	6	17	16	15	15	15	15
PCG	5	17	16	15	15	14	13
	4	14	16	13	14	11	12
	3	12	13	9	10	7	8

The iteration numbers required to reduce the ℓ^2-norm of the initial residual by a factor of 10^{-8} for the 1D, 2D and 3D problem are given in Tables 1–3, respectively.

The method was implemented in C++ based on the G+SMO library[1] which is developed in the framework of the National Research Network "Geometry + Simulation" at Johannes Kepler University, Linz.

We observe that the iteration numbers are robust with respect to both the discretization level ℓ (and thus h) and the spline degree p. They do increase with the space dimension d, but this dependence, which we have not fully analyzed, appears

---

[1] http://www.gs.jku.at/gismo

to be relatively mild. In particular, the 2D iteration numbers are significantly lower than those obtained using the boundary-corrected mass smoother in [15].

6.2. Experiments for non-trivial computational domains. We perform experiments with varying, matrix-valued diffusion coefficients on the non-trivial geometries shown in Fig. 3. The geometry map for the quarter annulus in the two-dimensional example is described exactly with NURBS, that for the three-dimensional object with B-splines. On these objects, we solve

$$- \operatorname{div}(A(x)\nabla u(x)) = f(x) \qquad \text{in } \Omega$$

with Dirichlet boundary conditions $g(x)$ on Γ_D as indicated in Fig. 3 and homogeneous Neumann boundary conditions on the remaining part of the boundary. Furthermore, f is given by (14) and the diffusion coefficient is given by

$$A^{(2D)}(x) = \begin{pmatrix} 1 + x_1^2 & -x_1 x_2 \\ -x_1 x_2 & 1 + x_2^2 \end{pmatrix}, \quad A^{(3D)}(x) = \begin{pmatrix} 1 + x_1^2 & -\frac{1}{3}x_1 x_2 & -\frac{1}{3}x_1 x_3 \\ -\frac{1}{3}x_1 x_2 & 1 + x_2^2 & -\frac{1}{3}x_2 x_3 \\ -\frac{1}{3}x_1 x_3 & -\frac{1}{3}x_2 x_3 & 1 + x_3^2 \end{pmatrix}.$$
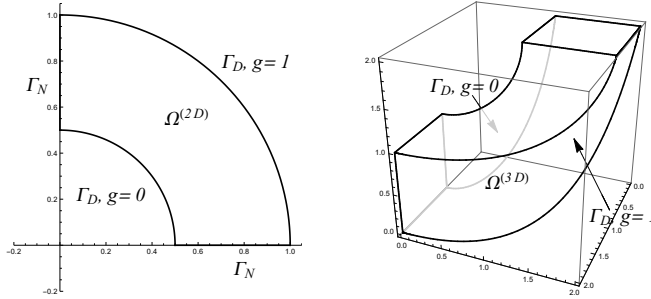

FIG. 3. *Computational domains for 2D and 3D example.*

Table 4 gives the iteration numbers for a conjugate gradient method, preconditioned with one V-cycle of the proposed multigrid solver, where the multigrid solver was set up as solver for the model problem $-\Delta u + u = f$ on the parameter domain.

TABLE 4
Iteration numbers for the nontrivial 2D (top) and 3D (bottom) domains as shown in Fig. 3.

$\ell \setminus p$	2	3	4	5	6	7	8	9	10
8	53	55	56	56	55	55	55	54	54
7	52	53	54	53	53	52	51	50	51
6	47	50	50	48	48	48	46	46	45
5	43	45	45	44	44	41	41	40	41

$\ell \setminus p$	2	3	4	5	6	7
5	87	90	91	90	89	90
4	73	76	76	79	81	83
3	55	61	66	67	72	75

Obviously, the condition number of the preconditioned system depends only on the geometry transformation, the diffusion coefficient and on the contraction number

of the multigrid method (as a solver for the model problem on the parameter domain). All of these quantities are independent of the grid size and the polynomial degree p. This is reflected in the numerical results, which are robust in those two parameters.

Appendix. The aim of this section is to prove Lemma 7, an approximation result for the coarse spline space Galerkin projector in d dimensions. It was shown in [15] for $d = 1$ and $d = 2$, and here we extend it to arbitrary dimensions by induction.

Before we give the proof, we need an auxiliary lemma which is a variant of the Aubin-Nitsche duality argument in a finite-dimensional Hilbert space V. By the choice of a suitable basis, we can identify V with $\mathbb{R}^n$, and operators A on V with matrices. We use this matrix representation implicitly in the following, and operations like $A^{1/2}$ and $A^\top$ are to be understood in the matrix sense.

LEMMA 11. *Let A and M be self-adjoint and positive definite linear operators on V, $T : V \to W \subset V$ an A-orthogonal projector, and $\theta > 0$. Then, the statements*

$$(15) \qquad \|Tu\|_M \leq \theta \|u\|_A \quad \forall u \in V \qquad and \qquad \|Tu\|_A \leq \theta \|u\|_{AM^{-1}A} \quad \forall u \in V$$

are equivalent.

Proof. We first observe that the statements in (15) are equivalent to

$$(16) \qquad \|M^{1/2}TA^{-1/2}\| \leq \theta \quad and \quad \|A^{1/2}TA^{-1}M^{1/2}\| \leq \theta,$$

respectively. Since T is self-adjoint in the scalar product $(\cdot, \cdot)_A$, $AT = T^\top A$ and further

$$(17) \qquad TA^{-1} = A^{-1}T^\top$$

hold. Using (17) as well as the self-adjointness of M and A, we obtain

$$\|M^{1/2}TA^{-1/2}\| = \|M^{1/2}A^{-1}T^\top A^{1/2}\| = \|(A^{1/2}TA^{-1}M^{1/2})^\top\| = \|A^{1/2}TA^{-1}M^{1/2}\|.$$

This proves that the two statements in (16) and, consequently, those in (15) are equivalent. $\square$

Proof of Lemma 7. Within this proof, we denote the dimensions explicitly and use a recursive representation,

$$M_1 := M, \qquad\qquad A_1 := K + M,$$
$$M_d := M_{d-1} \otimes M, \qquad A_d := A_{d-1} \otimes M + M_{d-1} \otimes K.$$

Furthermore we let T_d denote the A_d-orthogonal projector into $(S_c)^d$.

In [15], the desired result was proved for $d = 1$, namely,

$$(18) \qquad \|(I - T_1)u\|_{M_1} \leq ch\|u\|_{A_1} \qquad \forall u \in S.$$

By Lemma 11, this is equivalent to

$$(19) \qquad \|(I - T_1)u\|_{A_1} \leq ch\|u\|_{A_1 M_1^{-1} A_1} \qquad \forall u \in S.$$

Stability of the A_1-orthogonal projector means that

$$(20) \qquad \|(I - T_1)u\|_{A_1} \leq \|u\|_{A_1} \qquad \forall u \in S.$$

We now show the desired result using induction. Assume that we have already shown

$$\|(I - T_{d-1})u\|_{M_{d-1}} \le ch\|u\|_{A_{d-1}} \qquad \forall u \in S^{d-1} \tag{21}$$

for some $d > 1$. Using Lemma 11, this implies

$$\|(I - T_{d-1})u\|_{A_{d-1}} \le ch\|u\|_{A_{d-1}M_{d-1}^{-1}A_{d-1}} \qquad \forall u \in S^{d-1}. \tag{22}$$

Stability of the A_{d-1}-orthogonal projector means that

$$\|(I - T_{d-1})u\|_{A_{d-1}} \le \|u\|_{A_{d-1}} \qquad \forall u \in S^{d-1}. \tag{23}$$

Using equations (18)–(23) and the fact that the operator norm of a tensor product is the product of the individual operator norms, we obtain for all $u \in S^d$

$$\|(I - T_{d-1}) \otimes (I - T_1)u\|_{A_{d-1} \otimes M_1 + M_{d-1} \otimes A_1} \le ch\|u\|_{A_{d-1} \otimes A_1},$$
$$\|(I - T_{d-1}) \otimes Iu\|_{A_{d-1} \otimes M_1 + M_{d-1} \otimes A_1} \le ch\|u\|_{A_{d-1}M_{d-1}^{-1}A_{d-1} \otimes M_1 + A_{d-1} \otimes A_1},$$
$$\|I \otimes (I - T_1)u\|_{A_{d-1} \otimes M_1 + M_{d-1} \otimes A_1} \le ch\|u\|_{A_{d-1} \otimes A_1 + M_{d-1} \otimes A_1 M_1^{-1} A_1}.$$

Since $I - T_{d-1} \otimes T_1 = (I - T_{d-1}) \otimes I + I \otimes (I - T_1) - (I - T_{d-1}) \otimes (I - T_1)$, this implies using the triangle inequality

$$\|(I - T_{d-1} \otimes T_1)u\|_{A_{d-1} \otimes M_1 + M_{d-1} \otimes A_1}$$
$$\le ch\|u\|_{A_{d-1}M_{d-1}^{-1}A_{d-1} \otimes M_1 + A_{d-1} \otimes A_1 + M_{d-1} \otimes A_1 M_1^{-1} A_1}.$$

As the norm on the left-hand side is bounded from below by $\|\cdot\|_{A_d}$ and the norm on the right-hand side is bounded from above by $c\|\cdot\|_{A_d M_d^{-1} A_d}$, we further obtain

$$\|(I - T_{d-1} \otimes T_1)u\|_{A_d} \le ch\|u\|_{A_d M_d^{-1} A_d} \qquad \forall u \in S^d.$$

Both $T_{d-1} \otimes T_1$ and T_d are projectors into $(S_c)^d$. Since the latter projector produces the best approximation in the A_d-norm, we have

$$\|(I - T_d)u\|_{A_d} \le ch\|u\|_{A_d M_d^{-1} A_d} \qquad \forall u \in S^d,$$

which, by Lemma 11, is equivalent to the desired result

$$\|(I - T_d)u\|_{M_d} \le ch\|u\|_{A_d} \qquad \forall u \in S^d. \qquad \square$$

REFERENCES

[1] A. Buffa, H. Harbrecht, A. Kunoth, and G. Sangalli, *BPX-preconditioning for isogeometric analysis*, Computer Methods in Applied Mechanics and Engineering, 265 (2013), pp. 63–70, doi:10.1016/j.cma.2013.05.014.
[2] N. Collier, D. Pardo, L. Dalcin, M. Paszynski, and V. M. Calo, *The cost of continuity: A study of the performance of isogeometric finite elements using direct solvers*, Computer Methods in Applied Mechanics and Engineering, 213–216 (2012), pp. 353–361, doi:10.1016/j.cma.2011.11.002.

[3] J. Cottrell, A. Reali, Y. Bazilevs, and T. Hughes, *Isogeometric analysis of structural vibrations*, Computer Methods in Applied Mechanics and Engineering, 195 (2006), pp. 5257–5296, doi:10.1016/j.cma.2005.09.027. John H. Argyris Memorial Issue. Part II.

[4] L. B. da Veiga, D. Cho, L. Pavarino, and S. Scacchi, *Overlapping Schwarz methods for isogeometric analysis*, SIAM Journal on Numerical Analysis, 50 (2012), pp. 1394–1416, doi:10.1137/110833476.

[5] L. B. da Veiga, D. Cho, L. Pavarino, and S. Scacchi, *BDDC preconditioners for isogeometric analysis*, Mathematical Models and Methods in Applied Sciences, 23 (2013), pp. 1099–1142, doi:10.1142/S0218202513500048.

[6] L. B. da Veiga, L. F. Pavarino, S. Scacchi, O. B. Widlund, and S. Zampini, *Isogeometric BDDC preconditioners with deluxe scaling*, SIAM Journal on Scientific Computing, 36 (2014), pp. A1118–A1139, doi:10.1137/130917399.

[7] C. de Boor, *Efficient computer manipulation of tensor products*, ACM Transactions on Mathematical Software (TOMS), 5 (1979), pp. 173–182.

[8] C. de Boor, *A Practical Guide to Splines (revised edition)*, vol. 27 of Applied Mathematical Sciences, Springer, 2001.

[9] M. Donatelli, C. Garoni, C. Manni, S. Serra-Capizzano, and H. Speleers, *Robust and optimal multi-iterative techniques for IgA Galerkin linear systems*, Computer Methods in Applied Mechanics and Engineering, 284 (2014), pp. 230–264, doi:10.1016/j.cma.2014.06.001.

[10] M. Donatelli, C. Garoni, C. Manni, S. Serra-Capizzano, and H. Speleers, *Symbol-based multigrid methods for Galerkin B-spline isogeometric analysis*, Tech. Report TW650, Department of Computer Science, KU Leuven, July 2014, http://www.cs.kuleuven.be/publicaties/rapporten/tw/TW650.abs.html.

[11] K. P. S. Gahalaut, J. K. Kraus, and S. K. Tomar, *Multigrid methods for isogeometric discretization*, Computer Methods in Applied Mechanics and Engineering, 253 (2013), pp. 413–425, doi:10.1016/j.cma.2012.08.015.

[12] G. Golub and C. Van Loan, *Matrix Computations*, Johns Hopkins University Press, fourth ed., 2012.

[13] M. Griebel and P. Oswald, *On the abstract theory of additive and multiplicative Schwarz algorithms*, Numerische Mathematik, 70 (1995), pp. 163–180, doi:10.1007/s002110050115.

[14] W. Hackbusch, *Multi-Grid Methods and Applications*, Springer, Berlin, 1985.

[15] C. Hofreither, S. Takacs, and W. Zulehner, *A robust multigrid method for isogeometric analysis in two dimensions using boundary correction*, Computer Methods in Applied Mechanics and Engineering, (2016), doi:10.1016/j.cma.2016.04.003. Available online.

[16] C. Hofreither and W. Zulehner, *Mass smoothers in geometric multigrid for isogeometric analysis*, in Curves and Surfaces, J.-D. Boissonnat, A. Cohen, O. Gibaru, C. Gout, T. Lyche, M.-L. Mazure, and L. L. Schumaker, eds., vol. 9213 of Lecture Notes in Computer Science, Springer International Publishing, 2015, pp. 272–279, doi:10.1007/978-3-319-22804-4_20.

[17] C. Hofreither and W. Zulehner, *Spectral analysis of geometric multigrid methods for isogeometric analysis*, in Numerical Methods and Applications, I. Dimov, S. Fidanova, and I. Lirkov, eds., vol. 8962 of Lecture Notes in Computer Science, Springer International Publishing, 2015, pp. 123–129, doi:10.1007/978-3-319-15585-2_14.

[18] T. J. R. Hughes, J. A. Cottrell, and Y. Bazilevs, *Isogeometric analysis: CAD, finite elements, NURBS, exact geometry and mesh refinement*, Computer Methods in Applied Mechanics and Engineering, 194 (2005), pp. 4135–4195, doi:10.1016/j.cma.2004.10.008.

[19] S. K. Kleiss, C. Pechstein, B. Jüttler, and S. Tomar, *IETI – Isogeometric tearing and interconnecting*, Computer Methods in Applied Mechanics and Engineering, 247–248 (2012), pp. 201–215, doi:10.1016/j.cma.2012.08.007.

[20] G. Sangalli and M. Tani, *Isogeometric preconditioners based on fast solvers for the Sylvester equation*. ArXiv e-print 1602.01636. http://arxiv.org/abs/1602.01636, Feb. 2016.

[21] S. Takacs and T. Takacs, *Approximation error estimates and inverse inequalities for B-splines of maximum smoothness*, Mathematical Models and Methods in Applied Sciences, 26 (2016), pp. 1411–1445, doi:10.1142/S0218202516500342.

[22] U. Trottenberg, C. Oosterlee, and A. Schüller, *Multigrid*, Academic Press, London, 2001.

[23] J. Xu, *Iterative methods by space decomposition and subspace correction*, SIAM Review, 34 (1992), pp. 581–613, doi:10.1137/1034116.