

On the algebraic construction of sparse multilevel approximations of elliptic tensor product problems

H. Harbrecht and P. Zaspel

Departement Mathematik und Informatik
Fachbereich Mathematik
Universität Basel
CH-4051 Basel

Preprint No. 2018-01
January 2018

www.math.unibas.ch

On the algebraic construction of sparse multilevel approximations of elliptic tensor product problems

H. Harbrecht* P. Zaspel*

January 31, 2018

Abstract

We consider the solution of elliptic problems on the tensor product of two physical domains as e.g. present in the approximation of the solution covariance of elliptic partial differential equations with random input. Previous sparse approximation approaches used a geometrically constructed multilevel hierarchy. Instead, we construct this hierarchy for a given discretized problem by means of the algebraic multigrid method (AMG). Thereby, we are able to apply the sparse grid combination technique to problems given on complex geometries and for discretizations arising from unstructured grids, which was not feasible before. Numerical results show that our algebraic construction exhibits the same convergence behaviour as the geometric construction, while being applicable even in black-box type PDE solvers.

1 Introduction

The solution of elliptic problems on tensor products of a polygonally bounded domain $\Omega \subset \mathbb{R}^d$ with e.g. $d = 2, 3$ given by

$$\begin{aligned}(\Delta \otimes \Delta)u &= f \quad \text{on } \Omega \times \Omega, \\ u &= 0 \quad \text{on } \partial(\Omega \times \Omega),\end{aligned}$$

is an important high-dimensional problem. As an example, this problem shows up in the estimation of the output covariance of an elliptic partial differential equation with random input data that is given on a domain Ω , see [11, 13, 18, 19] for example. The problem becomes high-dimensional since the dimensionality of the elliptic problem on Ω is doubled. In case of real-world problems in $d = 3$, we end up solving a six-dimensional problem, which might become prohibitively expensive.

Recently, there have been developments to overcome this strong limitation. These developments are based on the introduction of a geometrically

*Departement für Mathematik und Informatik, Universität Basel, Schweiz

constructed multilevel frame, i.e. a hierarchy of discretizations of the elliptic problem on Ω . The multilevel frame gives rise to a sparse approximation with respect to the interaction of the involved domains in $\Omega \times \Omega$ [14, 18, 19]. It has been shown that the sparse approximation allows to solve the tensor product problem at a computational complexity that stays essentially (i.e. up to a poly-logarithmic factor) proportional to the number of degrees of freedom to discretize the domain Ω . In a more recent work by one of the authors [13], it has been shown that the sparse approximation can equivalently be replaced by the sparse grid combination technique [4, 8, 10, 15]. This further reduces the computational work and facilitates the implementation.

However, the currently available geometric construction of the multilevel hierarchy imposes limitations on the discretization in context of real-world problems. First, the coarsest mesh in the hierarchy of discretizations has to fully represent the boundary of the geometry Ω . This either limits the types of geometry to consider or the computational efficiency (in case even the coarsest mesh has to be fine at the boundary). Second, the use of a fully unstructured mesh becomes barely possible, since we are missing a coarsening strategy for such a mesh.

This work introduces *algebraically* constructed multilevel hierarchies [7, 9, 23] for the solution of elliptic problems on tensor product domains. While previous works [13, 14] first constructed the multilevel hierarchy of meshes and then discretized the problem by finite elements, the new approach first discretizes the problem on Ω on the finest (potentially unstructured) mesh and then constructs coarser versions of the linear system resulting from the fine discretization. The coarser problems are generated using algebraic coarsening known from the classical *Ruge-Stüben algebraic multigrid* (AMG) [17, 20]. The algebraic construction of multilevel hierarchies for frames has been previously discussed in context of optimal complexity solvers for elliptic problems in [23]. However, it has not been applied in the context of sparse approximation yet. Note that, by construction, our new approach allows us to overcome both the limitations in presence of complex geometries and the requirements on the structure of the mesh. Moreover, it perfectly fits into the context of black-box type PDE solvers.

As it is well-known, a full theory for algebraic multigrid methods, especially in the multilevel context and on unstructured grids, is still to be developed. Nevertheless, this technique is extremely popular as solver in real-world applications and, usually, empirically shows the same performance as geometric multigrid. This work follows the same spirit and focuses on the formal construction and the empirical analysis of the resulting numerical method. Thereby, we are able to match the convergence results available for geometrically constructed sparse approximations, while being able to apply this approach to complex geometries and unstructured grids in a black-box fashion.

In Section 2, the algebraic multilevel construction is outlined. This con-

struction is introduced to the tensor product problem with sparse approximation and the sparse grid combination technique in Section 3. Section 4 briefly discusses the implementation. In Section 5, we give a series of numerical examples with empirical error analysis. Finally, Section 6 summarizes this work.

2 Algebraic multilevel constructions

In our algebraic construction, we aim at replacing classical multilevel discretizations for elliptic partial differential equations by a purely matrix-based construction. That is, we consider an elliptic partial differential equation

$$\begin{aligned} -\Delta u &= f & \text{on } \Omega \\ u &= 0 & \text{on } \partial\Omega \end{aligned} \tag{1}$$

on a polygonally bounded domain $\Omega \subset \mathbb{R}^d$. This problem has been discretized by some method on a discretization level J , leading to a system of linear equations

$$\mathbf{A}_J \mathbf{u}_J = \mathbf{f}_J, \tag{2}$$

where $\mathbf{A}_J \in \mathbb{R}^{N_J \times N_J}$ is an M -matrix and $\mathbf{u}_J, \mathbf{f}_J \in \mathbb{R}^{N_J}$. In case of the discretization by finite elements, \mathbf{A}_J corresponds to the stiffness matrix and \mathbf{f}_J is the load vector, obtained by, for example, using the mass matrix \mathbf{M}_J and interpolation. Moreover, we identify each variable $u_{J,i}$ in $\mathbf{u}_J = (u_{J,1} \dots u_{J,N_J})^\top$ by its index i and introduce the corresponding index set $\mathcal{D}_J := \{1, \dots, N_J\}$ for discretization level J .

2.1 Multilevel hierarchy of discretized problems

The objective is to construct from (2) a hierarchy of systems of linear equations

$$\mathbf{A}_j \mathbf{u}_j = \mathbf{f}_j, \quad j = 0, \dots, J, \tag{3}$$

which are similar to discretizations on different geometric refinement levels. Especially, we intend to do this in a purely matrix-based, i.e. algebraic, way by using coarsening and transfer operators from algebraic multigrid (AMG) [20]. To this end, we first introduce a construction method for a hierarchy of variable sets

$$\mathcal{D}_0 \subset \mathcal{D}_1 \subset \dots \subset \mathcal{D}_J \tag{4}$$

of sizes

$$N_0 \leq N_1 \leq \dots \leq N_J.$$

Algorithm 1 Standard coarsening algorithm [21]

Require: level j

```

1: function AMGSTANDARDCOARSENING
2:    $F_j := \emptyset, \mathcal{D}_{j-1} := \emptyset, U_j := \mathcal{D}_j$ 
3:   for  $i \in U_j$  do
4:      $\lambda_j(i) := \left| \mathcal{S}_j(i)^\top \cap U_j \right| + 2 \left| \mathcal{S}_j(i)^\top \cap F_j \right|$ 
5:     while  $\exists i$  s.th.  $\lambda_j(i) \neq 0$  do
6:       find  $i_{\max} := \arg \max_i \lambda_j(i)$ 
7:        $\mathcal{D}_{j-1} := \mathcal{D}_{j-1} \cup \{i_{\max}\}$ 
8:        $U_j := U_j \setminus \{i_{\max}\}$ 
9:       for  $j \in (\mathcal{S}_j(i)^\top \cap U_j)$  do
10:         $F_j := F_j \cup \{j\}$ 
11:         $U_j := U_j \setminus \{j\}$ 
12:       for  $i \in U_j$  do
13:         $\lambda_j(i) := \left| \mathcal{S}_j(i)^\top \cap U_j \right| + 2 \left| \mathcal{S}_j(i)^\top \cap F_j \right|$ 
14:   return  $\mathcal{D}_{j-1}, F_j$ 

```

In classical Ruge-Stüben AMG [17, 21], this is achieved by recursively splitting the set of variables \mathcal{D}_j on level j into a set of coarse and fine grid variables

$$\mathcal{D}_j = \mathcal{D}_{j-1} \cup \mathcal{F}_j.$$

Each fine grid variable is supposed to be in the neighborhood of an appropriate amount of strongly coupled coarse grid variables, where we define the *neighborhood* of a variable $i \in \mathcal{D}_j$ by

$$\mathcal{N}_j(i) := \{i' \in \mathcal{D}_j : i' \neq i, a_{j,ii'} \neq 0\},$$

where $\mathbf{A}_j = (a_{j,ii'})_{i,i'=1}^{N_j}$. That is, we consider neighborhoods between variables by reinterpreting the system matrix \mathbf{A}_j as the adjacency matrix of a graph with edges between nodes for each non-zero matrix entry. Moreover, the set of neighboring strongly negatively coupled variables of a variable i is

$$\mathcal{S}_j(i) := \left\{ i' \in \mathcal{N}_j(i) \mid -a_{j,ii'} \geq \epsilon_{str} \max_k |a_{j,ik}| \right\}$$

with a strength measure $0 < \epsilon_{str} < 1$. The *standard coarsening* procedure, cf. Algorithm 1 [21], builds an appropriate splitting $\mathcal{D}_j = \mathcal{D}_{j-1} \cup \mathcal{F}_j$ based on these considerations. It also involves the sets $\mathcal{S}_j(i)^\top$, which are given by

$$\mathcal{S}_j(i)^\top := \{i' \in \mathcal{D}_j : i \in \mathcal{S}_j(i')\}.$$

In order to define the hierarchy of linear systems (3), we further need a means to transfer information between two consecutive levels j and $j + 1$.

This is done by prolongation operators $\mathbf{P}_j^{j+1} \in \mathbb{R}^{N_{j+1} \times N_j}$ and restriction operators $\mathbf{P}_{j+1}^j \in \mathbb{R}^{N_j \times N_{j+1}}$. Prolongation and restriction are done in a purely algebraic way based on AMG. In *standard interpolation* [21], which is one possible type of algebraic prolongation, data given on a fine grid node $i \in \mathcal{F}_j$ are interpolated from the set of interpolatory variables

$$\mathcal{I}_j(i) := (\mathcal{D}_{j-1} \cap \mathcal{S}_j(i)) \cap \left(\bigcup_{i' \in \mathcal{F}_j \cap \mathcal{S}_j(i)} (\mathcal{D}_{j-1} \cap \mathcal{S}_j(i')) \right).$$

Thus, it is interpolated from strongly negatively coupled coarse grid points and all coarse grid points that are strongly negatively coupled to strongly negatively coupled fine grid points. The exact choice of prolongation / interpolation weights is known from literature [21]. Restriction is given as the transpose of the prolongation, i.e. $\mathbf{P}_{j+1}^j = \mathbf{P}_j^{j+1 \top}$.

Finally, we recursively define for $j = J - 1, \dots, 0$ the matrices and the right-hand sides involved in the hierarchy of linear systems (3) as

$$\mathbf{A}_j := \mathbf{P}_{j+1}^j \mathbf{A}_{j+1} \mathbf{P}_j^{j+1}, \quad \mathbf{f}_j := \mathbf{P}_{j+1}^j \mathbf{f}_{j+1}.$$

In order to achieve optimal complexity in AMG, coarser levels are constructed such that the *operator complexity*

$$C_A := \sum_j \frac{\eta(\mathbf{A}_j)}{\eta(\mathbf{A}_J)},$$

where $\eta(\mathbf{A}_J)$ is the number of non-zeros in \mathbf{A}_J , stays bounded by some constant independent of J . If standard interpolation and standard coarsening fail in achieving this, stronger or more aggressive versions such as *extended / multi-pass interpolation* and *aggressive coarsening* on some levels is applied to keep this property [22]. Unfortunately, to the best of the authors' knowledge, there is for now no theory on the decay of the number of non-zeros in the coarse grid matrices \mathbf{A}_j constructed by classical Ruge-Stüben AMG on multiple levels and for general M matrices \mathbf{A}_J . The operator complexity is therefore always used as empirical measure for coarsening quality.

2.2 Multilevel frames

Let us note here that the above algebraic construction naturally leads to *algebraic multilevel frames*, cf. [23], for the elliptic problem on Ω . That is, we can replace our original system of linear equations in (2) by the system

$$\mathbf{A}_J \mathbf{u}_J = \mathbf{f}_J \tag{5}$$

with

$$\mathbf{A}_J := \begin{pmatrix} \mathbf{A}_{11} & \cdots & \mathbf{A}_{1J} \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{J1} & \cdots & \mathbf{A}_{JJ} \end{pmatrix}, \quad \mathbf{u}_J := \begin{pmatrix} \mathbf{u}_0 \\ \vdots \\ \mathbf{u}_J \end{pmatrix}, \quad \mathbf{f}_J := \begin{pmatrix} \mathbf{f}_0 \\ \vdots \\ \mathbf{f}_J \end{pmatrix}$$

and set

$$\mathbf{A}_{j_1 j_2} = \mathbf{P}_{j_1}^J \mathbf{A}_J \mathbf{P}_J^{j_2}.$$

The diagonal matrices \mathbf{A}_{j_j} are the system matrices \mathbf{A}_j from the previous paragraph. Moreover, we have extended the prolongation / restriction to arbitrary levels. This is possible by concatenating the corresponding operators. Above, we further introduce the multi-index $\mathbf{j} = (j_1, j_2)$ allowing the abbreviated notation

$$\mathbf{A}_J = [\mathbf{A}_j]_{\|\mathbf{j}\|_{\ell^\infty} \leq J}, \quad \mathbf{u}_J = [\mathbf{u}_j]_{|\mathbf{j}| \leq J}, \quad \mathbf{f}_J = [\mathbf{f}_j]_{|\mathbf{j}| \leq J}.$$

As in multilevel frame discretizations based on geometric refinements / coarsening, cf. [14], the above system of linear equations now encodes the full information of the hierarchy of systems in (3). Especially, it is equivalent to the linear system of equations (2), if the BPX-preconditioner is applied, cf. [2, 5, 6, 16].

The system matrix in (5) has a large kernel, which can be ignored by using appropriate iterative linear solvers. Solutions \mathbf{u}_J can be projected back to single-level solutions \mathbf{u}_j by applying the operator

$$\mathcal{P}_J = [\mathbf{P}_0^J, \mathbf{P}_1^J, \dots, \mathbf{P}_J^J].$$

In [23], it has been shown by numerical experiments that the application of specific iterative solvers to (5) leads to problem-size independent convergence rates also in case of algebraically constructed multilevel frames.

3 Sparse algebraic tensor product approach

Next, we like to consider elliptic problems on tensor products $\Omega \times \Omega$ of the polygonally bounded domain Ω . That is, we consider problems of the form

$$\begin{aligned} (\Delta \otimes \Delta)u &= f \quad \text{on } \Omega \times \Omega, \\ u &= 0 \quad \text{on } \partial(\Omega \times \Omega). \end{aligned} \tag{6}$$

As in Section 2, we assume to have a discretization (e.g. by finite elements) for the problem on a level J resulting in the system of linear equations

$$(\mathbf{A}_J \otimes \mathbf{A}_J)\mathbf{U}_J = \mathbf{F}_J. \tag{7}$$

Here, $\mathbf{A}_J \in \mathbb{R}^{N_J \times N_J}$ is the system matrix from (2). The operator \otimes is the Kronecker product operator for matrices. For matrices $\mathbf{S} \in \mathbb{R}^{n_1 \times n_2}$, $\mathbf{T} \in \mathbb{R}^{m_1 \times m_2}$, it computes the Kronecker product

$$\mathbf{S} \otimes \mathbf{T} := \begin{pmatrix} s_{11}\mathbf{T} & \dots & s_{1n_2}\mathbf{T} \\ \vdots & \ddots & \vdots \\ s_{n_1 1}\mathbf{T} & \dots & s_{n_1 n_2}\mathbf{T} \end{pmatrix}.$$

$\widehat{\mathbf{A}}_{(0,3)}$	$\widehat{\mathbf{A}}_{(1,3)}$	$\widehat{\mathbf{A}}_{(2,3)}$	$\widehat{\mathbf{A}}_{(3,3)}$
$\widehat{\mathbf{A}}_{(0,2)}$	$\widehat{\mathbf{A}}_{(1,2)}$	$\widehat{\mathbf{A}}_{(2,2)}$	$\widehat{\mathbf{A}}_{(3,2)}$
$\widehat{\mathbf{A}}_{(0,1)}$	$\widehat{\mathbf{A}}_{(1,1)}$	$\widehat{\mathbf{A}}_{(2,1)}$	$\widehat{\mathbf{A}}_{(3,1)}$
$\widehat{\mathbf{A}}_{(0,0)}$	$\widehat{\mathbf{A}}_{(1,0)}$	$\widehat{\mathbf{A}}_{(2,0)}$	$\widehat{\mathbf{A}}_{(3,0)}$

$\widehat{\mathbf{A}}_{(0,3)}$	$\widehat{\mathbf{A}}_{(1,3)}$	$\widehat{\mathbf{A}}_{(2,3)}$	$\widehat{\mathbf{A}}_{(3,3)}$
$\widehat{\mathbf{A}}_{(0,2)}$	$\widehat{\mathbf{A}}_{(1,2)}$	$\widehat{\mathbf{A}}_{(2,2)}$	$\widehat{\mathbf{A}}_{(3,2)}$
$\widehat{\mathbf{A}}_{(0,1)}$	$\widehat{\mathbf{A}}_{(1,1)}$	$\widehat{\mathbf{A}}_{(2,1)}$	$\widehat{\mathbf{A}}_{(3,1)}$
$\widehat{\mathbf{A}}_{(0,0)}$	$\widehat{\mathbf{A}}_{(1,0)}$	$\widehat{\mathbf{A}}_{(2,0)}$	$\widehat{\mathbf{A}}_{(3,0)}$

$\widehat{\mathbf{A}}_{(0,3)}$	$\widehat{\mathbf{A}}_{(1,3)}$	$\widehat{\mathbf{A}}_{(2,3)}$	$\widehat{\mathbf{A}}_{(3,3)}$
$\widehat{\mathbf{A}}_{(0,2)}$	$\widehat{\mathbf{A}}_{(1,2)}$	$\widehat{\mathbf{A}}_{(2,2)}$	$\widehat{\mathbf{A}}_{(3,2)}$
$\widehat{\mathbf{A}}_{(0,1)}$	$\widehat{\mathbf{A}}_{(1,1)}$	$\widehat{\mathbf{A}}_{(2,1)}$	$\widehat{\mathbf{A}}_{(3,1)}$
$\widehat{\mathbf{A}}_{(0,0)}$	$\widehat{\mathbf{A}}_{(1,0)}$	$\widehat{\mathbf{A}}_{(2,0)}$	$\widehat{\mathbf{A}}_{(3,0)}$

Figure 1: For discretization level $J = 3$, multilevel frames on the full tensor product space require a very densely populated system matrix $\widehat{\mathbf{A}}_{\mathbf{J}}$ (*left*), while sparse approximation leads to the system matrix $\widehat{\mathbf{A}}_{\mathbf{J}}$ (*center*) with smaller size due to fewer active (i.e. gray) matrix subblocks. The sparse grid combination technique (*right*) leads to the most efficient approximation.

Consequently, $(\mathbf{A}_J \otimes \mathbf{A}_J)$ becomes a matrix of size $N_J^2 \times N_J^2$. Moreover, $\mathbf{U}_J, \mathbf{F}_J \in \mathbb{R}^{N_J \cdot N_J}$ are the solution and the right-hand side, respectively.

By assuming an underlying d -dimensional finite element discretization with mesh width h and a multigrid-type linear solver, solving the linear system in (7) would require at least $O(h^{-2d})$ operations, in contrast to $O(h^{-d})$ for the problem given by (2). This amount of computational work is prohibitively large, especially for larger d . Therefore, we shall find a way to reduce the amount of work to solve this problem. Before we do that, we change the problem discretization to a multilevel discretization, which is the basis for the subsequent sparse approaches.

3.1 Multilevel frames for tensor product constructions

As in Section 2.1, we can for $\mathbf{j} = (j, j')$ introduce multiple levels of systems of linear equations

$$(\mathbf{A}_j \otimes \mathbf{A}_{j'})\mathbf{U}_j = \mathbf{F}_j \quad (8)$$

with

$$\mathbf{A}_j = \mathbf{P}_J^j \mathbf{A}_J \mathbf{P}_j^J \quad \text{and} \quad \mathbf{F}_j = \left(\mathbf{P}_J^j \otimes \mathbf{P}_J^{j'} \right) \mathbf{F}_J$$

by applying coarsening and the transfer operators of algebraic multigrid. Thereby, we obtain

$$(\mathbf{A}_j \otimes \mathbf{A}_{j'}) \in \mathbb{R}^{N_j N_{j'} \times N_j N_{j'}} \quad \text{and} \quad \mathbf{U}_j, \mathbf{F}_j \in \mathbb{R}^{N_j N_{j'}}.$$

$\mathbf{P}_j^{j'}$ is the prolongation / restriction matrix introduced in Section 2.1.

By extending the solution approach in Section 2.2 to tensor product problems, we finally obtain the multilevel frame linear system for the tensor product problem as

$$\widehat{\mathbf{A}}_{\mathbf{J}} \mathbf{U}_{\mathbf{J}} = \mathbf{F}_{\mathbf{J}}. \quad (9)$$

Here, we have

$$\widehat{\mathbf{A}}_{\mathbf{J}} = [\mathbf{A}_j \otimes \mathbf{A}_{j'}]_{j,j' \leq J} =: [\widehat{\mathbf{A}}_{\mathbf{j}}]_{\|\mathbf{j}\|_{\ell^\infty} \leq J},$$

and

$$\mathbf{U}_{\mathbf{J}} = [\mathbf{U}_{\mathbf{j}}]_{\|\mathbf{j}\|_{\ell^\infty} \leq J}, \quad \mathbf{F}_{\mathbf{J}} = [\mathbf{F}_{\mathbf{j}}]_{\|\mathbf{j}\|_{\ell^\infty} \leq J}.$$

Note that we indeed construct frames over the tensor product problems instead of constructing a tensor product of frame discretizations, cf. [14].

In order to characterize the computational complexity for the solution of (9), we recall that we assume to have a constant operator complexity for the sequence of matrices \mathbf{A}_j , i.e. $\sum \eta(\mathbf{A}_j) \leq c\eta(\mathbf{A}_J)$. Moreover, by definition of the Kronecker product, we have the number of non-zeros in each block of $\widehat{\mathbf{A}}_{\mathbf{J}}$ given by

$$\eta(\mathbf{A}_j \otimes \mathbf{A}_{j'}) = \eta(\mathbf{A}_j)\eta(\mathbf{A}_{j'}).$$

From that, we can estimate the total number of non-zeros in $\widehat{\mathbf{A}}_{\mathbf{J}}$ by

$$\eta(\widehat{\mathbf{A}}_{\mathbf{J}}) = \sum_{j,j' \leq J} \eta(\mathbf{A}_j)\eta(\mathbf{A}_{j'}) \leq c^2\eta(\mathbf{A}_J)^2.$$

This means that the computational work to solve (9) is asymptotically identical to a solve of (7).

3.2 Sparse tensor product construction

Solving (7) or (9) would be prohibitively expensive, cf. Figure 1. As in the geometric multilevel case, we assume that the solution of the elliptic problem (1) on Ω is H^s regular. Therefore, the solution of the tensor product problem (6) becomes H_{mix}^s -regular, see [18]. This allows to follow, for example, the lines of [14] to introduce a sparse, however now algebraically constructed, version of the discretized problem. Instead of using all sub-problems for multi-indices $\|\mathbf{j}\|_{\ell^\infty} \leq J$, the sparse approximation is reduced to multi-indices $\|\mathbf{j}\|_{\ell^1} \leq J$. Thereby, we obtain a new system of linear equations

$$\widetilde{\mathbf{A}}_{\mathbf{J}} \widetilde{\mathbf{U}}_{\mathbf{J}} = \widetilde{\mathbf{F}}_{\mathbf{J}}$$

with

$$\widetilde{\mathbf{A}}_{\mathbf{J}} := [\widehat{\mathbf{A}}_{\mathbf{j}}]_{\|\mathbf{j}\|_{\ell^1} \leq J}, \quad \mathbf{U}_{\mathbf{J}} = [\mathbf{U}_{\mathbf{j}}]_{\|\mathbf{j}\|_{\ell^1} \leq J}, \quad \mathbf{F}_{\mathbf{J}} = [\mathbf{F}_{\mathbf{j}}]_{\|\mathbf{j}\|_{\ell^1} \leq J}.$$

Figure 1 compares both choices in the plots on the left-hand side and the center.

As discussed before, there is not much theory on the size of the levels in the algebraic multilevel construction. The only available information is the assumed bound on the operator complexity. However, this does not give

enough information to discuss the expected improvement in performance due to the sparse construction. Nevertheless, the bound implies a similar scaling of the non-zeros with level j as in the geometric multilevel construction. Therefore, we here briefly discuss the number of non-zeros in $\widehat{\mathbf{A}}_J$ for the geometric construction to give a hint towards the possible performance improvement by the algebraic sparse construction.

With this in mind, we follow the previous example of (linear) finite elements on a mesh with mesh width h . The number of non-zero entries for matrix \mathbf{A}_j is proportional to the number of elements and therefore

$$\eta(\mathbf{A}_j) = O(2^{dj}).$$

Moreover, we have $J = O(|\log h|)$. By evaluating

$$\eta(\widehat{\mathbf{A}}_J) = \sum_{\|\mathbf{j}\|_{\ell^1} \leq J} \eta(\widehat{\mathbf{A}}_{\mathbf{j}}),$$

one can easily verify that the number of non-zeros in the system matrix in $\widehat{\mathbf{A}}_J$ is asymptotically

$$\eta(\widehat{\mathbf{A}}_J) = O(|\log h| h^{-d}).$$

That is, in case a BPX-type preconditioner [2, 5, 6, 16] and an optimal approach for the construction of the sub-problem matrices $\widehat{\mathbf{A}}_{\mathbf{j}}$ [1, 3, 24] is used, the computational complexity of the problem on the tensor product domain $\Omega \times \Omega$ is (up to a logarithmic factor) reduced to the computational complexity of the problem on domain Ω .

3.3 Sparse grid combination technique

It has been shown [13] that the previous sparse approximation is equivalent to the so-called sparse grid combination technique. The latter one requires to solve a set of decoupled problems

$$\widehat{\mathbf{A}}_{\mathbf{j}} \mathbf{U}_{\mathbf{j}} = \mathbf{F}_{\mathbf{j}}, \quad \text{where } \|\mathbf{j}\|_{\ell^1} \in \{J, J-1\}. \quad (10)$$

These are afterwards combined to a solution

$$\widehat{\mathbf{U}}_J = \sum_{\|\mathbf{j}\|_{\ell^1} = J} (\mathbf{P}_{\mathbf{j}}^J \otimes \mathbf{P}_{\mathbf{j}'}^J) \mathbf{U}_{\mathbf{j}} - \sum_{\|\mathbf{j}\|_{\ell^1} = J-1} (\mathbf{P}_{\mathbf{j}}^J \otimes \mathbf{P}_{\mathbf{j}'}^J) \mathbf{U}_{\mathbf{j}}. \quad (11)$$

On the right-hand side of Figure 1, the sub-matrices $\widehat{\mathbf{A}}_{\mathbf{j}}$ used in this approximation have been marked gray. As before, one can easily verify that the total number of non-zeros of the matrices in (10) is asymptotically $O(|\log h| h^{-d})$ for the case of linear finite elements on a tetrahedral mesh with mesh width h

in d dimensions and a geometrically constructed multilevel structure. However, Figure (1) easily clarifies that the pre-asymptotic number of non-zeros in the matrices involved in the combination technique is much smaller than the non-zeros in the sparse approximation discussed before.

In terms of computational complexity of the combination technique, let us remind that the (approximate) solution of each sub-problem in (10) can be realized by an iterative linear solver with matrix-vector products. To be more specific, tensor product versions of standard iterative solvers can be constructed, by reshaping a given iterate $\mathbf{U}_{j=(j,j')} \in \mathbb{R}^{N_j \cdot N_{j'}}$ (and the appropriate right-hand side) to a matrix of size $N_j \times N_{j'}$. Then, the action of one step of an iterative solver for matrix $\widehat{\mathbf{A}}_j = \mathbf{A}_j \otimes \mathbf{A}_{j'}$ is done by first applying the iterative solver step for \mathbf{A}_j to all $N_{j'}$ columns of the reshaped matrix and by second applying the iterative solver step for $\mathbf{A}_{j'}$ to all N_j rows of the reshaped matrix. Since we have all prolongation and restriction operators from AMG at our disposal, we can construct, in the above way, a tensor product version of algebraic multigrid. Given this solver, we obtain roughly problem-size independent convergence for each sub-problem in (10), i.e. we need $O(N_j N_{j'})$ operations for each sub-problem. In the geometric setting, we would again have the relation $N_j = O(2^{dj})$ and thereby $O(2^{d(j+j')})$ operations per sub-problem. Since it holds $\|\mathbf{j}\|_{\ell^1} \in \{J, J-1\}$ and the number of sub-problems is $O(J)$, we would finally end up with a computational complexity of $O(J2^{dJ})$ or $O(N_J \log N_J)$.

4 Implementation

In our numerical results, we approximate solutions for tensor product finite element discretizations of elliptic problems based on the combination technique. To this end, we assemble system matrices for a given problem, construct the multilevel hierarchies, solve the decoupled, anisotropic problems in (10) and combine the solutions following the combination rule (11).

Assembly of system matrices. The discretization by the finite element method is done with the *Matlab PDE Toolbox* of *Matlab 2017a*. We use linear finite elements and construct meshes with maximum element size $\mathbf{Hmax} = 2^{-J}$. Furthermore, we use the option `Jiggle` to optimize the mesh in quality. The stiffness matrix (incorporating boundary conditions) is constructed by using the *Matlab* command `assembleFEMatrices` with option `nullspace`. In a similar way, we extract the mass matrix. Afterwards, both matrices and the mesh node coordinates are stored to files.

Construction of the multilevel hierarchy. From within *Matlab* we call a newly implemented code that uses the parallel linear solver library *hypra* in version 2.11.1. This library contains the implementation *BoomerAMG*

of classical Ruge-Stüben AMG. The code reads the matrix from file and creates the AMG multilevel hierarchy by using *hypr*. In addition to *standard coarsening* with strength measure $\epsilon_{str} = 0.25$ and *standard interpolation*, we use two passes of *Jacobi* interpolation [21] with a truncation of the Jacobi interpolation with a threshold of 0.001 for the two-dimensional problems and 0.01 for the three-dimensional problem. All other parameters are kept as the defaults of *BoomerAMG*. After having created the multigrid hierarchy, the program stores the prolongation matrices of all created levels to files. These are read by *Matlab*.

Solution of the anisotropic tensor product problems. Based on the prolongation matrices and the system matrix \mathbf{A}_J on the finest levels, the decoupled problems in (10) can be set up. As discussed before, a tensor product version of AMG is used to solve the systems of linear equations. In our implementation, we construct the sub-problem operators in (10) by individually multiplying the transfer operators between two consecutive levels.

Our tensor product AMG is iterated until the convergence criterion

$$\|\mathbf{R}_j^{it}\|_{\ell^2} / \|\mathbf{F}_j\|_{\ell^2} \leq \epsilon_{tol}$$

is fulfilled, where \mathbf{R}_j^{it} is the residual of the current iterate \mathbf{U}_j^{it} in the solver. Since the problems in (10) completely decouple, we can easily parallelize their solution process by a `parfor` loop in *Matlab*. In case an individual problem becomes very expensive, we further implemented a distributed memory parallelization for the tensor product AMG based on *Matlab*'s `distributed` function. Thereby, we overcome the limitation of a non-existing multi-core parallelization for sparse matrix-vector products in *Matlab*.

Combination of the solutions. In the combination phase, we avoid to prolongate the full partial solutions to the finest level J . Instead, we randomly chose N_{eval} nodes on the product of the finest meshes on $\Omega \times \Omega$. On these points, we evaluate the combination formula (11) and compute the empirical error measure

$$e(\mathbf{U}_{approx}) = \|\mathbf{U}_{approx} - \mathbf{U}_{ref}\|_{\ell^2} / \|\mathbf{U}_{ref}\|_{\ell^2},$$

where \mathbf{U}_{approx} is the approximated solution and \mathbf{U}_{ref} is an appropriately evaluated reference solution. Note that we do not multiply the tensor product of the prolongation with the solution. However, we follow the ideas from Section 3.3 for the construction of the tensor product AMG and apply the prolongations direction-wise. The prolongation for each sub-problem is also parallelized by a `parfor` loop.

5 Numerical results

In our empirical studies, we consider the numerical solution of the problem

$$\begin{aligned} (\Delta \otimes \Delta)u &= f \quad \text{on } \Omega \times \Omega, \\ u &= 0 \quad \text{on } \partial(\Omega \times \Omega). \end{aligned} \tag{12}$$

by means of the combination technique based on the algebraic multilevel hierarchy. Different choices will be made for the domain Ω and the right-hand side f .

5.1 Analytic example on a disk

The first study is done on a disk domain Ω with center $(0, 0)^\top$ and radius 0.5. We set

$$f(\mathbf{x}, \mathbf{y}) = 1.$$

The exact solution of the resulting problem is

$$u(\mathbf{x}, \mathbf{y}) = \frac{1}{16} (x_1^2 + x_2^2 - 0.5^2) (y_1^2 + y_2^2 - 0.5^2).$$

To approximate the solution u by the combination technique, we follow the methodology discussed in Section 4. As part of this, we triangulate the geometry with a maximum element width of 2^{-J} . Figure 2 shows on the left-hand side the resulting mesh for $J = 5$. It is obvious that the resulting mesh is unstructured. Therefore, classical geometric constructions for the sparse grid combination technique would not be feasible on that mesh. In contrast, our new algebraic approach can solve this problem.

This is shown on the right-hand side of Figure 2, where we compare the numerically approximated solution against the above exact solution. Convergence results for the choices $J = 3, \dots, 8$ are given. From literature, compare e.g. [13], we know that the error of the geometrically constructed sparse grid combination technique scales for the problem under consideration like $J4^{-J}$. As we can see from the convergence results in Figure 2, the algebraically constructed combination technique shows the same convergence behavior, while being applicable to unstructured grids.

Figure 3 shows on the left-hand side computing times for growing problem size N_J of the univariate discretization of Ω . We compare the time required for the solution of the combination technique sub-problems with the time required to solve the full tensor-product problem (7) by our tensor-product AMG implementation. Note that we use the coarse grid hierarchies reported in Table 1 for both the combination technique and the full tensor-product approach. All measurements were done on a compute server with dual 20-core Intel Xeon E5-2698 v4 CPU at 2.2 GHz and 768 GB RAM. It becomes evident that our algebraically constructed combination technique




Figure 2: The combination technique based on our algebraic multilevel hierarchy and applied to the tensor product of a disk geometry with an unstructured mesh (*left, triangulated with $J = 5$*) shows the same convergence as the geometrically constructed combination technique (*right*).

approach beats the full tensor-product approach in both, computational complexity and effective runtime. However, both results do not show the predicted computational complexity of $O(N_J \log N_J)$ and $O(N_J^2)$. There are several reasons for this behavior.

- First, algebraic multigrid often shows a small, roughly logarithmic, growth in the number of iterations for larger problem sizes, resulting in a slow-down by a logarithmic factor.
- Second, we observe a certain fill-in in the system matrices for coarser problems in the algebraic construction due to our choice of an additional (truncated) Jacobi interpolation. However, this should be pre-asymptotic behavior.
- Third, as can be seen in Table 1, the AMG coarsening approach chosen in our implementation does not show the exact same decay rate in the number of levels as we expect from the geometric construction. In fact, this leads to a problem-size dependent growth of the coarsest grid. While this growth does not affect the error decay, it shows up in the computational complexity.

Meanwhile, as stated before, we are able to beat the solution approach based on the full tensor-product approach in terms of computational complexity. Even more, in terms of runtime, we are by more than two orders of magnitude faster.




Figure 3: We compare the runtime of the new combination technique approach (CT) with the runtime of the traditional the full tensor-product approach (full TP) for the solution of the tensor product elliptic problems on the disk geometry (*left*) and the plate geometry (*right*).

Ω	$J \setminus j$	# dofs on algebraically coarsened level j									
		0	1	2	3	4	5	6	7	8	9
disk	3	3	11	28	71						
	4	8	21	52	119	320					
	5	12	31	84	207	495	1292				
	6	20	51	139	348	852	2009	5234			
	7	35	93	244	606	1473	3510	8415	22118		
	8	46	130	366	978	2469	5983	14480	34081	89097	
plate	3	5	20	61							
	4	16	36	90	230						
	5	28	68	168	414	1072					
	6	46	116	297	745	1813	4703				
	7	63	184	515	1272	3117	7491	19611			
	8	103	302	815	2124	5301	12822	30639	80146		
spanner	3	4	10	19	50	117	247				
	4	11	22	59	147	326	689	1454			
	5	40	114	300	689	1516	3216	6484	13939		
	6	210	548	1364	3123	6708	14109	29103	57438	125223	
	7	1386	3120	6627	14016	29533	61150	124921	253291	496614	1082581

Table 1: For a given problem on level J , the algebraic multilevel construction on our example domains Ω constructs coarser levels with a decrease of the number of unknowns roughly similar to geometric multilevel constructions. Above, only those levels j are reported that are used in the convergence study.




Figure 4: Even for a covariance load on a complex geometry (*left*, triangulated for $J = 5$), the algebraic construction shows the appropriate convergence rate after a short pre-asymptotic phase (*right*).

5.2 Example on complex geometry with covariance load

The next numerical study is concerned with the solution of the problem (12) with the load

$$f(\mathbf{x}, \mathbf{y}) = \exp\left(\frac{-\|\mathbf{x} - \mathbf{y}\|^2}{\ell}\right)$$

that corresponds to an (unscaled) Gaussian covariance kernel with correlation length ℓ . This is a prototype version of the tensor product elliptic problem on $\Omega \times \Omega$ showing up in the computation of the output covariance of an elliptic problem on Ω with random input, cf. [14].

In addition to the more complicated right-hand side, we solve the problem for a rather complex geometry Ω . We choose the geometry of a square plate on $[0, 1]^2$ with circular wholes of radius 0.15 which are centered at the points

$$\{(0.25, 0.25), (0.25, 0.75), (0.75, 0.25), (0.75, 0.75)\}.$$

Figure 4 shows its triangulation for $J = 5$ on the left-hand side. Note that it would be almost impossible to solve a problem on such a geometry with the geometrical construction for the sparse grid combination technique. However, with the algebraic construction, a coarsening to very few degrees of freedom becomes easily possible, compare Table 1.

To be able to compare the above problem against a numerically computed reference solution, we replace the (sampled) covariance kernel for $\ell = 1$ by its low-rank approximation computed with the pivoted Cholesky factorization [12], truncated for a trace norm of 10^{-8} . In this case, depending on the problem size, the truncation results in roughly twenty low-rank terms.




Figure 5: In our large-scale real world example, we solve an elliptic problem on the tensor product of the three-dimensional geometry of a spanner. For a discretization level of $J = 7$, the discretization of Ω has more than a million unknowns. This would lead to 10^{12} , that is a *trillion*, unknowns in the full tensor product discretization.

On the right-hand side of Figure 4, we show the convergence results with errors computed against the numerically approximated exact solution by use of the low-rank approximation. After a pre-asymptotic phase, we are able to attain an error that scales like $J4^{-J}$ as in the geometric construction.

The problem size dependent runtime to compute the subspace solutions for the plate geometry is given in Figure 3 on the right-hand side. We observe similar computational complexities and similar runtimes as in the previous example on the disk.

5.3 Large-scale real-world example

Our last numerical study treats a large-scale problem with a complex real-world geometry Ω . We again aim at solving (12) for $f(\mathbf{x}, \mathbf{y}) = 1$. However, we choose the *three-dimensional* spanner geometry found in Figure 5. In contrast to the previous examples, we set the maximum mesh width to 2^{5-J} , since the geometry is contained in the rather large bounding box $[-5, 5] \times [-12.2, 112] \times [-15.7, 15.7]$. Note that the triangulation of Ω results for level $J = 7$ in a discretization with 1,082,581 unknowns. That is, if we would want to solve the full tensor product problem on $\Omega \times \Omega$, cf. (6), then we would have to solve a problem with about 10^{12} , that is a *trillion*, unknowns. This




Figure 6: Our algebraic multilevel construction for the sparse grid combination technique on the large-scale three-dimensional spanner geometry gradually approaches the optimal convergence rate of $J2^{-dJ}$.

would be clearly out of scope even for large parallel clusters. In contrast, the combination technique allows to solve this problem. Nevertheless, we still have to solve, e.g. for level $J = 7$ and the system matrix $\widehat{\mathbf{A}}_{(0,J)}$ a problem with $1,082,581 \times 1,386$ unknowns, compare Table 1.

In Figure 6, we show the convergence results for this large scale problem relative to a numerical approximation of the solution. Due to the high dimensionality and complexity of the domain Ω , the convergence results in Figure 6 are only gradually approaching the optimal scaling of $J2^{-dJ}$. Nevertheless, we are able to solve this problem up to a certain accuracy. This shows that even very complex problems of large scale can be solved by the proposed approach.

6 Conclusions

In this work, we have introduced an algebraic construction method for the sparse approximation of tensor product elliptic problems by means of the combination technique. While previous approaches were tight to geometric hierarchies of mesh refinements to build the underlying multilevel discretization, we were able to solve the given type of problems on complex geometries and for unstructured grids by an algebraic multilevel hierarchy based on AMG. We could show that our approach has the same convergence rates as the geometric construction. Measurements of the computational complexity were in the linear range with poly-logarithmic factors. Overall, we are now able to apply sparse approximation for elliptic tensor product problems in a black-box fashion.

References

- [1] R. Balder and C. Zenger. The solution of multidimensional real Helmholtz equations on sparse grids. *SIAM Journal on Scientific Computing*, 17(3):631–646, 1996.
- [2] J. Bramble, J. Pasciak, and J. Xu. Parallel multilevel preconditioners. *Mathematics of Computation*, 55:1–22, 1990.
- [3] H.-J. Bungartz. A multigrid algorithm for higher order finite elements on sparse grids. *ETNA. Electronic Transactions on Numerical Analysis*, 6:63–77, 1997.
- [4] H.-J. Bungartz and M. Griebel. Sparse grids. *Acta Numerica*, 13:1–123, 2004.
- [5] W. Dahmen. Wavelet and multiscale methods for operator equations. *Acta Numerica*, 6:55–228, 1997.
- [6] M. Griebel. *Multilevelmethoden als Iterationsverfahren über Erzeugendensystemen*. Teubner Skripten zur Numerik. B.G. Teubner, Stuttgart, 1993.
- [7] M. Griebel. Multilevel algorithms considered as iterative methods on semidefinite systems. *SIAM International Journal Scientific Statistical Computing*, 15(3):547–565, 1994.
- [8] M. Griebel and H. Harbrecht. On the convergence of the combination technique. In J. Garcke and D. Pflüger, editors, *Sparse grids and Applications – Stuttgart 2014*, volume 97 of *Lecture Notes in Computational Science and Engineering*, pages 55–74. Springer, 2014.
- [9] M. Griebel and P. Oswald. Greedy and randomized versions of the multiplicative Schwarz method. *Linear Algebra and its Applications*, 7:1596–1610, 2012.
- [10] M. Griebel, M. Schneider, and C. Zenger. A combination technique for the solution of sparse grid problems. In P. de Groen and R. Beauwens, editors, *Iterative Methods in Linear Algebra*, pages 263–281. IMACS, Elsevier, North Holland, 1992.
- [11] H. Harbrecht. A finite element method for elliptic problems with stochastic input data. *Applied Numerical Mathematics*, 60(3):227–244, Mar. 2010.
- [12] H. Harbrecht, M. Peters, and R. Schneider. On the low-rank approximation by the pivoted Cholesky decomposition. *Applied Numerical Mathematics*, 62(4):428–440, 2012.

- [13] H. Harbrecht, M. Peters, and M. Siebenmorgen. Combination technique based k -th moment analysis of elliptic problems with random diffusion. *Journal of Computational Physics*, 252(C):128–141, Nov. 2013.
- [14] H. Harbrecht, R. Schneider, and C. Schwab. Multilevel frames for sparse tensor product spaces. *Numerische Mathematik*, 110(2):199–220, July 2008.
- [15] M. Hegland, J. Garcke, and V. Challis. The combination technique and some generalisations. *Linear Algebra and its Applications*, 420(2):249–275, 2007.
- [16] P. Oswald. *Multilevel finite element approximation. Theory and applications*. Teubner Skripten zur Numerik. B.G. Teubner, Stuttgart, 1994.
- [17] J. Ruge and K. Stüben. Algebraic multigrid (AMG). In S. McCormick, editor, *Multigrid Methods, Frontiers in Applied Mathematics*, volume 5. SIAM, Philadelphia, 1986.
- [18] C. Schwab and R. A. Todor. Sparse finite elements for elliptic problems with stochastic loading. *Numerische Mathematik*, 95(4):707–734, 2003.
- [19] C. Schwab and R. A. Todor. Sparse finite elements for stochastic elliptic problems: Higher order moments. *Computing*, 71(1):43–63, Sept. 2003.
- [20] K. Stüben. A review of algebraic multigrid. *Journal of Computational and Applied Mathematics*, 128(12):281–309, 2001. Numerical Analysis 2000. Vol. VII: Partial Differential Equations.
- [21] U. Trottenberg and A. Schuller. *Multigrid*. Academic Press, Inc., Orlando, FL, USA, 2001.
- [22] U. M. Yang. On long-range interpolation operators for aggressive coarsening. *Numerical Linear Algebra with Applications*, 17(2-3):453–472, 2010.
- [23] P. Zaspel. Subspace correction methods in algebraic multi-level frames. *Linear Algebra and its Applications*, 488:505–521, 2016.
- [24] A. Zeiser. Fast matrix-vector multiplication in the sparse-grid Galerkin method. *SIAM Journal of Scientific Computing*, 47(3):328–346, 2011.

LATEST PREPRINTS

- | No. | Author: Title |
|---------|--|
| 2016-30 | W. D. Brownawell, D. W. Masser
<i>Unlikely intersections for curves in additive groups over positive characteristic</i> |
| 2016-31 | M. Dambrine, H. Harbrecht, M. D. Peters, B. Puig
<i>On Bernoulli's free boundary problem with a random boundary</i> |
| 2016-32 | H. Harbrecht, J. Tausch
<i>A fast sparse grid based space-time boundary element method for the nonstationary heat equation</i> |
| 2016-33 | S. Iula
<i>A note on the Moser-Trudinger inequality in Sobolev-Slobodeckij spaces in dimension one</i> |
| 2016-34 | C. Bürli, H. Harbrecht, P. Odermatt, S. Sayasone, N. Chitnis
<i>Mathematical analysis of the transmission dynamics of the liver fluke, <i>Opisthorchis viverrini</i></i> |
| 2017-01 | J. Dölz and T. Gerig, M. Lüthi, H. Harbrecht and T. Vetter
<i>Efficient computation of low-rank Gaussian process models for surface and image registration</i> |
| 2017-02 | M. J. Grote, M. Mehlin, S. A. Sauter
<i>Convergence analysis of energy conserving explicit local time-stepping methods for the wave equation</i> |
| 2017-03 | Y. Bilu, F. Luca, D. Masser
<i>Collinear CM-points</i> |
| 2017-04 | P. Zaspel
<i>Ensemble Kalman filters for reliability estimation in perfusion inference</i> |
| 2017-05 | J. Dölz and H. Harbrecht
<i>Hierarchical Matrix Approximation for the Uncertainty Quantification of Potentials on Random Domains</i> |
| 2017-06 | P. Zaspel
<i>Analysis and parallelization strategies for Ruge-Stüben AMG on many-core processor</i> |
| 2017-07 | H. Harbrecht and M. Schmidlin
<i>Multilevel Methods for Uncertainty Quantification of Elliptic PDEs with Random Anisotropic Diffusion</i> |

LATEST PREPRINTS

- | No. | Author: <i>Title</i> |
|------------|---|
| 2017-08 | M. Griebel and H. Harbrecht
<i>Singular value decomposition versus sparse grids: Refined complexity Estimates</i> |
| 2017-09 | J. Garcke and I. Kalmykov
<i>Efficient Higher Order Time Discretization Schemes for Hamilton-Jacobi-Bellman Equations Based on Diagonally Implicit Symplectic Runge-Kutta Methods</i> |
| 2017-10 | M. J. Grote and U. Nahum
<i>Adaptive Eigenspace Regularization For Inverse Scattering Problems</i> |
| 2017-11 | J. Dölz, H. Harbrecht, S. Kurz, S. Schöps and F. Wolf
<i>A Fast Isogeometric BEM for the Three Dimensional Laplace- and Helmholtz Problems</i> |
| 2017-12 | P. Zaspel
<i>Algorithmic patterns for H-matrices on many-core processors</i> |
| 2017-13 | R. Brügger, R. Croce and H. Harbrecht
<i>Solving a free boundary problem with non-constant coefficients</i> |
| 2017-14 | M. Dambrine, H. Harbrecht and B. Puig
<i>Incorporating knowledge on the measurement noise in electrical impedance tomography</i> |
| 2017-15 | C. Bürli, H. Harbrecht, P. Odermatt, S. Sayasone and N. Chitnis
<i>Analysis of Interventions against the Liver Fluke, <i>Opisthorchis viverrini</i></i> |
| 2017-16 | D. W. Masser
<i>Abcological anecdotes</i> |
| 2017-17 | P. Corvaja, D. W. Masser and U. Zannier
<i>Torsion hypersurfaces on abelian schemes and Betti coordinates</i> |
| 2017-18 | F. Caubet, M. Dambrine and H. Harbrecht
<i>A Newton method for the data completion problem and application to obstacle detection in Electrical Impedance Tomography</i> |
| 2018-01 | H. Harbrecht and P. Zaspel
<i>On the algebraic construction of sparse multilevel approximations of elliptic tensor product problems</i> |