

Block Preconditioners for Linear Systems Arising from Multilevel RBF Collocation*

Patricio Farrell and Jennifer Pestana

November 5, 2015

Abstract

Symmetric multiscale collocation methods with radial basis functions allow approximation of the solution of a partial differential equation, even if the right-hand side is only known at scattered data points, without needing to generate a grid. However, the benefit of a guaranteed symmetric positive definite block system comes at a high computational cost. In particular, the condition number and sparsity deteriorate with the number of data points. Therefore, we study certain block diagonal and triangular preconditioners. We investigate ideal preconditioners and determine the spectra of the preconditioned matrices before proposing more practical preconditioners based on a restricted additive Schwarz method with coarse grid correction (ARASM). Numerical results verify the effectiveness of the preconditioners.

1 Introduction

Radial basis functions (RBFs) are a modern tool to flexibly approximate scattered data [22, 36, 3, 9, 30]. Although they were initially mainly used for interpolation problems, they have recently become an attractive alternative to solve PDEs—especially when the given data is arbitrarily scattered [13, 17, 11, 12, 37, 33, 16]. Their attractiveness stems from the fact that the RBF method dispenses with the expensive generation of a grid. Unlike a method yielding a nonsymmetric system proposed by Kansa [23], we would like to focus on *symmetric* RBF approximation, which has the advantage that it always yields a symmetric positive definite system.

The main problem with RBFs is that for a large number of data sites the condition number of the system one needs to solve becomes prohibitively large. Therefore, multiscale ideas in combination with compactly supported RBFs have been developed to reduce the computational cost [8, 10, 26, 38]. Here, we focus

*This is the pre-peer reviewed version of the following article: P. Farrell and J. Pestana. Block preconditioners for linear systems arising from multiscale collocation with compactly supported RBFs. *Numerical Linear Algebra with Applications* 22:731–747. 2015., which has been published in final form at <http://onlinelibrary.wiley.com/doi/10.1002/nla.1984/full>.

on a symmetric multiscale RBF collocation method for second-order elliptic PDEs on bounded domains [7, 8]. The method employs RBFs on a sequence of levels, and uses different numbers of data sites and RBF support radii on each level. Consequently, it is particularly suited to problems with multiple scales.

At each level of the multiscale collocation method a linear system of the form

$$\underbrace{\begin{bmatrix} A & B^T \\ B & C \end{bmatrix}}_{\mathcal{A}} \underbrace{\begin{bmatrix} \alpha^i \\ \alpha^b \end{bmatrix}}_{\alpha} = \underbrace{\begin{bmatrix} b^i \\ b^b \end{bmatrix}}_b, \quad (1)$$

must be solved, where \mathcal{A} is nonsingular and symmetric positive definite [7] and $B \in \mathbb{R}^{m \times n}$ has full rank. The positive definiteness of \mathcal{A} ensures that $A \in \mathbb{R}^{n \times n}$ and $C \in \mathbb{R}^{m \times m}$ are themselves positive definite. More detailed descriptions of these matrices will be provided in Section 2.2. We also assume that $n \geq m$, which means that more data sites are located in the interior of the domain than on its boundary.

Even though the multiscale approach is already much more efficient than the one-shot method, certain features of \mathcal{A} make (1) difficult to solve at later levels. First, the number of data sites, and thus the dimension of the matrix \mathcal{A} , increases at each level so that finer scales may be resolved. Second, the conditioning of the coefficient matrix deteriorates as the separation between data sites decreases, a fact we state more in carefully in Theorem 2. In particular, if one wants to ensure convergence, the density of nonzero elements increases [7, 8].

When the dimension of \mathcal{A} is large, iterative methods are more feasible than direct methods for solving (1). However, the ill-conditioning and density of nonzeros mean that fast convergence will typically only be achieved with suitable preconditioners. For interpolation problems solved by an analogous RBF multiscale method, convergence and constant condition numbers can be achieved at the same time [38]. However, applying the same approach to PDE problems is insufficient for level-independent convergence [7, 8] and a more sophisticated strategy is required.

Preconditioners for RBF matrices have previously been developed, with domain decomposition approaches among the most popular. Beatson *et al.* [2] employed a multiplicative Schwarz method to solve (rather than precondition) linear systems resulting from interpolation by polyharmonic splines. Yokota *et al.* [39] investigated restricted additive Schwarz (RAS) methods for interpolation by Gaussian RBFs, while Deng and Driscoll [6] used a two-level RAS method as a GMRES preconditioner for interpolation by multiquadrics. Additionally, Le Gia *et al.* [25] applied a two-level overlapping additive Schwarz method to the problem of solving PDEs by compactly-supported basis functions on spheres. Alternative preconditioners include those based on approximate cardinal functions [1, 29], which can be combined with domain decomposition [2, 28].

To construct preconditioners for the whole matrix \mathcal{A} , however, it seems sensible to exploit the block structure. Recently, Le Gia and Tran [27] examined effective block diagonal preconditioners for the RBF multiscale method for PDEs on the local spherical regions, with additive Schwarz preconditioners for

each block. In this complementary work we consider both block diagonal and block triangular preconditioners for the multiscale method for PDEs on general bounded domains. In contrast to Le Gia and Tran [27] we attempt to identify ideal preconditioners, from the point of view of fast convergence of the iterative solver of (1), and analytically determine the spectra of the preconditioned matrices. These ideal preconditioners then guide the development of more practical alternatives based on restricted additive Schwarz methods.

The rest of this paper is organised as follows. In Section 2 we describe in more detail the multiscale RBF method and give a brief overview of Krylov subspace methods and preconditioners for solving (1). We present our ideal block diagonal and block triangular preconditioners, and describe the spectra of the preconditioned matrices, in Section 3. We investigate the effect of replacing A by an additive Schwarz preconditioner in 4 and give numerical results in Section 5. Note that throughout, the transpose of a matrix A is represented by A^T and its nullspace by $\text{null}(A)$.

2 Background

In this section we present background material on the multiscale RBF method and on preconditioned Krylov subspace methods.

2.1 Second-order elliptic boundary value problems

Let $\Omega \subset \mathbb{R}^d$. We consider second-order elliptic boundary value problems of the form

$$\begin{aligned} \mathcal{L}u &= f && \text{in } \Omega, \\ u &= F && \text{on } \partial\Omega, \end{aligned} \tag{2}$$

where \mathcal{L} is a second-order elliptic linear differential operator defined by

$$\mathcal{L}u(x) = \sum_{i,j=1}^d a_{ij}(x) \partial_{ij} u(x) + \sum_{i=1}^d b_i(x) \partial_i u(x) + c(x)u(x),$$

which is strictly elliptic on Ω . That is, there exists a constant $c_E > 0$ such that

$$c_E \|\xi\|_2^2 \leq \sum_{i,j=1}^d a_{ij}(x) \xi_i \xi_j$$

for all $x \in \Omega \subset \mathbb{R}^d$ and $\xi = (\xi_i) \in \mathbb{R}^d$.

If we assume that the right-hand sides f and F are chosen such that the solution u lies in the Sobolev space $H^\sigma(\Omega)$ with $\sigma > d/2+2$, then the differential operator $\mathcal{L}u$ is in fact well-defined since we know by the Sobolev embedding theorem that $H^\sigma(\Omega) \subseteq C^2(\Omega)$.

To ensure that \mathcal{L} is a bounded operator from $H^\sigma(\Omega)$ to $H^{\sigma-2}(\Omega)$, we impose some restrictions on the coefficients. For $k := \lfloor \sigma \rfloor > 2 + d/2$ we demand that

a_{ij}, b_i, c lie in $W_\infty^{k-1}(\Omega)$. Due to our previous assumption on σ , we see that $(k-1) - 1 > d/2$, which implies by the Sobolev embedding theorem that the coefficients are continuously differentiable, see [17] for details.

2.2 Multiscale RBF collocation

In order to solve the boundary value problem (2), we will construct a numerical approximation from a linear combination of translated radial basis functions. These basis functions are particularly useful in the context of scattered data approximation. Therefore, we introduce two measures that help us to describe scattered data points $X = \{x_1, \dots, x_N\}$ in $\Omega \subset \mathbb{R}^d$. The mesh norm

$$h_{X,\Omega} = \sup_{x \in \Omega} \min_{x_j \in X} \|x - x_j\|_2$$

is the radius of the largest data-free hole that is contained in the domain of interest Ω . On the other hand, the separation distance

$$q_X = \min_{j \neq k} \|x_j - x_k\|_2$$

is the shortest distance between any two data points in X .

Definition 1 (Radial basis function). *A continuous function $\Phi: \mathbb{R}^d \rightarrow \mathbb{R}$ is called positive definite on \mathbb{R}^d if for any d -dimensional data set $X = \{x_1, \dots, x_N\}$ of pairwise distinct points the matrix*

$$A_{\Phi,X} := (\Phi(x_j - x_k))_{1 \leq j, k \leq N}$$

is positive definite. We refer to Φ a radial basis function if it is a radial positive definite function.

Thus, radial basis functions lead naturally to symmetric and positive definite matrices.

There are many different examples of radial basis functions. Among them Gaussians, (inverse) multiquadrics and polyharmonic splines have been popular. Here, we are interested in compactly supported radial basis functions since in this case the matrix $A_{\Phi,X}$ is sparse. The Fourier transform of the compactly supported radial basis function shall satisfy

$$c_1(1 + \|\omega\|_2^2)^{-\sigma} \leq \widehat{\Phi}(\omega) \leq c_2(1 + \|\omega\|_2^2)^{-\sigma} \quad (3)$$

for $0 < c_1 \leq c_2$. For $\sigma > \frac{d+1}{2}$ it is indeed possible to find such a function, see [5].

The most prominent examples of such compactly supported RBFs were given by Wendland [35]. For any given integer smoothness degree and dimension Wendland was able to construct radial basis functions that are polynomials within the unit ball and vanish outside of it. They have the special property

that for given smoothness and dimension these polynomials are of the smallest degree such that the compactly supported RBF is still positive definite.

The numerical solution we would like to consider is of the form

$$s(x) = \sum_{j=1}^n \alpha_j^i \mathcal{L}^{(2)} \Phi(x - x_j) + \sum_{j=1}^m \alpha_j^b \Phi(x - y_j), \quad (4)$$

where $X = \{x_1, \dots, x_n\} \subset \Omega$ and $Y = \{y_1, \dots, y_m\} \subset \partial\Omega$ denote scattered data points in the interior and on the boundary, respectively, for which the right-hand sides are known. The superscript next to the differential operator in the first sum means that the operator is applied to the second argument of RBF and then evaluated at data site x_j . The real interior and boundary coefficients α_j^i and α_j^b are determined by applying the boundary value problem (2) to the numerical approximation s at these scattered data points. The system one needs to solve is then given by (1), where the different parts are given by

$$\begin{aligned} A &= (\mathcal{L}^{(1)} \mathcal{L}^{(2)} \Phi(x_i - x_j))_{1 \leq i, j \leq n}, \\ B &= (\mathcal{L}^{(2)} \Phi(x_i - x_j))_{\substack{1 \leq i \leq m, \\ 1 \leq j \leq n}}, & C &= (\Phi(x_i - x_j))_{1 \leq i, j \leq m}, \\ \alpha^i &= (\alpha_j^i)_{1 \leq j \leq n}, & \alpha^b &= (\alpha_j^b)_{1 \leq j \leq m}, \\ b^i &= (f(x_j))_{1 \leq j \leq n}, & b^b &= (F(y_j))_{1 \leq j \leq m}. \end{aligned}$$

In the context of generalised interpolation [36], it can be shown that the matrices \mathcal{A} and A are symmetric and positive definite. The matrix C , on the other hand, is symmetric and positive by Definition 1. This guarantee that the block matrix \mathcal{A} is always symmetric and positive definite is the main reason we choose our numerical approximation as in (4). However, the above approach also differs fundamentally from other methods commonly used to approximate PDEs such as finite difference, finite element and finite volume methods. Unlike for these methods, the matrix system (1) is not a discrete approximation of the PDE. Hence, classical preconditioning theory for PDE problems does not directly apply.

Though it would be feasible to use just this one-shot solution as a numerical approximation to the solution of the PDE, it is not very efficient to do so since the system suffers from severe ill-conditioning. One way around this is to employ the following multiscale strategy.

We will choose a sequence of denser data sets as well as smaller support radii. We denote a sequence of point sets in the bounded domain Ω by X_1, X_2, X_3, \dots and sequence of point sets on the domain's boundary $\partial\Omega$ by Y_1, Y_2, Y_3, \dots . For support radii $\delta_j > 0$ and a compactly supported basis function Φ we define

$$\Phi_j(x - y) = \Phi_{\delta_j}(x - y) = \Phi\left(\frac{x - y}{\delta_j}\right).$$

So if Φ has unit support, the Φ_j indeed have support radii δ_j . Instead of using just the previously introduced one-shot approximation, we will construct several

ones, coming from the spaces

$$\mathcal{L}V_{X_j} + V_{Y_j} = \text{span}\{\mathcal{L}\Phi_j(\cdot - x) \mid x \in X_j\} + \text{span}\{\Phi_j(\cdot - y) \mid y \in Y_j\}.$$

Each approximation from these spaces shall resolve the current residual so that the sum of all those approximations yields a numerical solution to the original PDE. This is achieved quite naturally by the following multiscale RBF collocation algorithm.

Algorithm 1 (Multilevel RBF collocation algorithm). *Given right-hand sides f and F do:*

1. Set $u_0 = 0, f_0 = f, F_0 = F$

2. For $j = 1, 2, 3 \dots$ do

(a) Determine the correction $s_j \in \mathcal{L}V_{X_j} + V_{Y_j}$ to residuals f_{j-1} and F_{j-1} from the equations

$$\begin{aligned} \mathcal{L}s_j(x) &= f_{j-1}(x), & x \in X_j, \\ s_j(y) &= F_{j-1}(y), & y \in Y_j. \end{aligned}$$

(b) Update the final approximation and the residuals

$$\begin{aligned} u_j &= u_{j-1} + s_j, \\ f_j &= f_{j-1} - \mathcal{L}(s_j|\Omega), \\ F_j &= F_{j-1} - s_j|_{\partial\Omega}. \end{aligned}$$

A variant of this algorithm for pure interpolation problems was shown to converge [38] if the support radii are chosen proportionally to the mesh norms. This in turn implied that for quasi-uniform data sets (i.e. when the separation distance is comparable to the mesh norm) the condition numbers of the interpolation matrices could be bounded independently of the current level. Thus, the conjugate gradient method would converge in a fixed number of steps, regardless of the size of the problem.

Unfortunately, the same does not hold for the multiscale collocation algorithm. We state the two main results from [7] concerning convergence and stability, omitting some technical details in favour of readability. The main point here is that in order to guarantee convergence, one has to cope with ill-conditioning issues.

Theorem 1 (Convergence). *Let Φ satisfy (3) for $\sigma > 2 + d/2$, and h denote the maximum of the boundary and interior mesh norms. Then, the multiscale collocation algorithm for elliptic boundary value problems converges if the support radius δ is chosen proportional to $h^{1-2/\sigma}$.*

Theorem 2 (Stability). *Under the same assumptions of Theorem 1, the condition number of the block matrices $\mathcal{A} = \mathcal{A}(\delta)$ arising in Algorithm 1 can be bounded by*

$$\text{cond}(\mathcal{A}) \leq Cq_{X \cup Y}^{-8+8/\sigma}.$$

Employing the diagonal preconditioner \mathcal{M} with entries

$$m_{ij} = \begin{cases} 0, & i \neq j, \\ \delta^2, & 1 \leq i \leq n, \\ 1, & n \leq i \leq n + m, \end{cases}$$

the condition number can be bounded by

$$\text{cond}(\mathcal{M}^{-1}\mathcal{A}\mathcal{M}^{-1}) \leq Cq_{X \cup Y}^{-4}.$$

The preconditioner \mathcal{M} , first suggested by Fasshauer [8], aims to mitigate different scaling of the blocks of \mathcal{A} in (1), since due to the chain rule the block A scales like $O(\delta^{-4})$, B like $O(\delta^{-2})$ and C like $O(1)$.

The preconditioned system is

$$\underbrace{\mathcal{M}^{-1}\mathcal{A}\mathcal{M}^{-1}}_{\tilde{\mathcal{A}}}y = \mathcal{M}^{-1}b, \quad \mathcal{M}^{-1}y = x, \quad \tilde{\mathcal{A}} = \begin{bmatrix} \tilde{A} & \tilde{B}^T \\ \tilde{B} & \tilde{C} \end{bmatrix}, \quad (5)$$

Due to the obvious improvement in the condition number the diagonal preconditioner achieves, we assume that this preconditioner has been applied to (1) and solve (5) in the following sections, dropping the tildes for convenience. Lastly, we point out that employing $\sqrt{\text{diag}(\tilde{\mathcal{A}})}$ yields better condition numbers than \mathcal{M} but the result in Theorem 2 still applies as the diagonals only differ by constants.

2.3 Preconditioned Krylov subspace methods

When the system (5) is large, iterative methods are often used to obtain its solution. Krylov subspace methods are iterative solvers for linear systems with large, sparse coefficient matrices. When the coefficient matrix is symmetric positive definite, as in (5), we can apply the conjugate gradient method (CG) [19]. For nonsymmetric systems, iterative methods such as GMRES [32], QMR [14] or Bi-CGSTAB [34] are required.

Whichever Krylov method is employed, the rate of convergence can be sensitive to the conditioning of the matrix \mathcal{A} . In particular, the convergence rate of Krylov methods for symmetric positive definite matrices typically decreases when the condition number of \mathcal{A} increases, i.e., when the eigenvalues are less clustered. Moreover, for any of the mentioned Krylov methods small eigenvalues can cause slow convergence. The condition numbers of the RBF multiscale matrices \mathcal{A} increase at each level and thus to achieve fast convergence of the Krylov subspace method it is necessary to precondition (5). In right preconditioning $\mathcal{A}\mathcal{P}^{-1}y = b$, $\alpha = \mathcal{P}^{-1}y$, is solved in place of (5). Symmetry can be preserved when the preconditioner is symmetric positive definite, (see, for example, Greenbaum [18, Chapter 8]) and CG can be applied to the preconditioned system.

2.4 Test problem

The next sections are devoted to overcoming the ill-conditioning described by Theorem 2. To illustrate the improvements made, as well as certain features of \mathcal{A} , we test our preconditioners using the following example problem.

We have implemented a Poisson problem on the unit square, which comes from [9], namely

$$\Delta u = -\frac{5}{4} \sin(\pi x) \cos(\pi y/2) \quad \text{on } (0, 1)^2,$$

$$u = \begin{cases} \sin(\pi x) & \text{on } 0 \leq x \leq 1, y = 0, \\ 0 & \text{otherwise.} \end{cases}$$

We use Wendland's compactly supported radial basis function, $\phi_{2,3}(r) = (1 - r)_+^8 (32r^3 + 25r^2 + 8r + 1) \in C^6(\mathbb{R}^2)$, which satisfies (3) with $\sigma = 4.5$ as well as support radii of the form

$$\delta = \nu (h/\mu)^{1-2/4.5}$$

with $\mu = 0.5$ and $\nu = 2.4$. Again, h denotes the maximum of the boundary and interior mesh norms on uniform grids. By Theorem 1, we have ensured convergence of the multiscale collocation algorithm. Theorem 2 states that the condition number of the collocation matrix in (1) should behave like $q_{X \cup Y}^{-6.2}$ and the condition number of the preconditioned collocation matrix in (5) should behave like $q_{X \cup Y}^{-4}$. These theoretical results were verified in [7].

3 Ideal block preconditioners

To achieve better conditioning as the levels increase we need effective preconditioners. In this section we introduce ideal block diagonal and block triangular preconditioners for the linear system (5) that depend on the blocks and on the Schur complement $S = C - BA^{-1}B^T$. We additionally examine the effect of the preconditioning blocks on the spectrum and conditioning of the preconditioned matrix. Although the preconditioners we obtain may be too expensive to apply in practice, since they all involve linear solves with A , in subsequent sections we show that A can be replaced by a restricted additive Schwarz preconditioner to obtain an efficient alternative.

Our block diagonal preconditioners are positive definite and the preconditioned system can be solved by the conjugate gradient method. On the other hand, block triangular preconditioners are nonsymmetric and can only be used in conjunction with a Krylov method for nonsymmetric matrices. However, if the speed of convergence is significantly faster with the block triangular preconditioner this may outweigh any extra cost associated with using a nonsymmetric solver; this is the case here.

Note that since $\mathcal{P}^{-1}\mathcal{A}$, $\mathcal{P}^{-\frac{1}{2}}\mathcal{A}\mathcal{P}^{-\frac{1}{2}}$ and $\mathcal{A}\mathcal{P}^{-1}$ are similar they all have the same eigenvalues, although their eigenvectors may differ. Thus, in the following we determine only the eigenvalues of $\mathcal{P}^{-1}\mathcal{A}$.

3.1 Block diagonal preconditioners

For saddle point matrices, for which $C = 0$ in (5), it is well known that if the preconditioner

$$\mathcal{P}_0 = \begin{bmatrix} A & 0 \\ 0 & BA^{-1}B^T \end{bmatrix},$$

is applied then $\mathcal{P}_0^{-1}\mathcal{A}$ has the three eigenvalues 1, $(1 \pm \sqrt{5})/2$ [21, 24, 31]. Krylov methods applied to this preconditioned system usually converge rapidly; for example, CG converges in at most three iterations. When C is small in norm, we might expect the preconditioner \mathcal{P}_0 to continue to perform well but when the norm of C increases a different choice for the (2,2) block might be preferable. Accordingly, we consider the block diagonal preconditioner

$$\mathcal{P}_D = \begin{bmatrix} A & 0 \\ 0 & \widehat{S} \end{bmatrix}, \quad (6)$$

where $\widehat{S} \in \mathbb{R}^{m \times m}$ is symmetric positive definite, and investigate the spectrum of $\mathcal{P}_D^{-1}\mathcal{A}$ for different choices of \widehat{S} . The following lemma describes the spectrum of $\mathcal{P}_D^{-1}\mathcal{A}$.

Lemma 3. *The matrix $\mathcal{P}_D^{-1}\mathcal{A}$, with \mathcal{P}_D and \mathcal{A} defined by (6) and (5), respectively, has the eigenvalue 1 with multiplicity $n - m$. Each of the remaining $2m$ eigenvalues λ , with corresponding eigenvector $[u^T v^T]^T$, satisfy*

$$\lambda = \frac{1}{2} \left(1 + \frac{v^T C v}{v^T \widehat{S} v} \right) \pm \sqrt{\frac{1}{4} \left(1 - \frac{v^T C v}{v^T \widehat{S} v} \right)^2 + \frac{v^T B A^{-1} B^T v}{v^T \widehat{S} v}}. \quad (7)$$

Proof. The eigenvalues λ of $\mathcal{P}_D^{-1}\mathcal{A}$ satisfy

$$\begin{aligned} Au + B^T v &= \lambda Au, \\ Bu + Cv &= \lambda \widehat{S} v, \end{aligned} \quad (8)$$

where $u \in \mathbb{R}^n$ and $v \in \mathbb{R}^m$ are not simultaneously 0. Both \mathcal{A} and \mathcal{P}_D are nonsingular and so $\lambda \neq 0$. If $\lambda = 1$, then $B^T v = 0$ which, since B has full rank and $m \leq n$, implies that $v = 0$. Then (8) shows that $u \in \text{null}(B)$. We can find $n - m$ linearly independent vectors $u \in \text{null}(B)$ and so $\lambda = 1$ is an eigenvalue of $\mathcal{P}_D^{-1}\mathcal{A}$ with multiplicity $n - m$.

If $\lambda \neq 1$, then $u = 1/(\lambda - 1)A^{-1}B^T v$, from which we see that $v \neq 0$. Substituting for u in (8), dividing by $v^T \widehat{S} v$ and simplifying gives (7). \square

The $2m$ non-unit eigenvalues (7) lie in two intervals on the real line which should be small if the rate of convergence of the Krylov subspace method is to be fast. That is, the boundary introduces problematic eigenvalues. This makes sense when comparing the matrix to that obtained by a similar algorithm for problems on the sphere [26].

Since (7) contains the terms $v^T C v$ and $v^T B A^{-1} B^T v$, we consider the choices C , S , and $B A^{-1} B^T$ for \widehat{S} and examine their effect on the non-unit eigenvalues. To investigate all three choices simultaneously, we let

$$\widehat{S} = \beta C + \gamma B A^{-1} B^T, \quad (9)$$

where β and γ are either 0 or 1. Note that since A^{-1} may be dense even when A is sparse, the choices S and $B A^{-1} B^T$ are not necessarily practical. However, they provide insight into the best theoretical choice and how different choices affect the quality of \mathcal{P}_D .

Since the fractions $v^T C v / v^T \widehat{S} v$ and $v^T B A^{-1} B^T v / v^T \widehat{S} v$ appear in (7), for our choice of \widehat{S} we shall find the generalised Rayleigh quotient

$$\mu(v) = \frac{v^T B A^{-1} B^T v}{v^T C v}$$

useful for assessing the quality of our preconditioners. This ratio is bounded, for any $v \in \mathbb{R}^m$, $v \neq 0$, by [20, Theorems 4.2.11 and 7.7.6]

$$0 < \lambda_{\min}(C^{-1} B A^{-1} B^T) \leq \mu(v) \leq \lambda_{\max}(C^{-1} B A^{-1} B^T) < 1. \quad (10)$$

Substituting (9) for \widehat{S} in (7) and simplifying gives

$$\lambda_{1,2} = \frac{1}{2} \left(1 + \frac{1}{\beta + \mu(v)\gamma} \right) \pm \sqrt{\frac{1}{4} \left(1 - \frac{1}{\beta + \mu(v)\gamma} \right)^2 + \frac{\mu(v)}{\beta + \mu(v)\gamma}}.$$

Table 1 shows the non-unit eigenvalues we expect for each choice of \widehat{S} . If $\mu(v) \ll 1$, both $\widehat{S} = C$ and $\widehat{S} = S$ are good choices, since the eigenvalues of $\mathcal{P}_D^{-1} \mathcal{A}$ are clustered near 1. However, $\widehat{S} = B A^{-1} B^T$ is not a good choice, since the eigenvalues of $\mathcal{P}_D^{-1} \mathcal{A}$ may be spread out. The limiting case $\mu(v) = 1$ shows that if $\mu(v)$ is close to 1 then $\mathcal{P}_D^{-1} \mathcal{A}$ will have very small eigenvalues and will, therefore, be ill-conditioned for all three choices. This suggests scaling \widehat{S} , i.e., choosing $\omega \widehat{S}$ for some positive scalar ω , to ensure that $\mu(v)$ is not too close to 1.

Table 1: Approximate values of non-unit eigenvalues of $\mathcal{P}_D^{-1} \mathcal{A}$ for different choices of $\mu(v)$. Note that S is singular when $\mu(v) = 1$.

\widehat{S}	$\lambda_{1,2}$	$\mu(v) \ll 1$	$\mu(v) = 1$
C	$1 \pm \sqrt{\mu(v)}$	$1 \pm \sqrt{\mu(v)}$	0, 2
$B A^{-1} B^T$	$\frac{1}{2} \left(1 + \frac{1}{\mu(v)} \right) \pm \sqrt{\frac{1}{4} \left(1 - \frac{1}{\mu(v)} \right)^2 + 1}$	$1, \frac{1}{\mu(v)}$	0, 2
S	$\frac{1}{2} \left(1 + \frac{1}{1-\mu(v)} \right) \pm \sqrt{\frac{1}{4} \left(1 - \frac{1}{1-\mu(v)} \right)^2 + \frac{\mu(v)}{1-\mu(v)}}$	$1, \frac{1}{1-\mu(v)}$	—

Because the optimal choice of \widehat{S} depends on $\mu(v)$, we compute the extreme values of $\mu(v)$, i.e., the extreme eigenvalues of $C^{-1} B A^{-1} B^T$, for the test problem

from Section 2.4, with the results given in Table 2. We see that the smallest value of $\mu(v)$ for this problem is on the order of 10^{-3} and that the largest is near 1. The spread of eigenvalues means that an optimal choice is not straightforward to determine, but since the smallest eigenvalue is closer to zero than the largest eigenvalue is to 1, the best choice appears to be $\hat{S} = C$ (since $BA^{-1}B^T$ is expensive to apply). Note that this gives an analogous preconditioner to that proposed by Le Gia and Tran [27]. Our numerical results in Section 5.1 confirm this. However, we note that if eigenvalues of $\mathcal{P}_D^{-1}\mathcal{A}$ become too small, the convergence rate of the conjugate gradient method may decrease.

Table 2: Extreme eigenvalues of $C^{-1}BA^{-1}B^T$ for different problem sizes.

$n + m$	289	1 089	4 225	16 641
$\lambda_{\min}(C^{-1}BA^{-1}B^T)$	0.0011	0.0030	0.0065	0.012
$\lambda_{\max}(C^{-1}BA^{-1}B^T)$	0.77	0.88	0.94	0.97

3.2 Block triangular preconditioners

We now consider block triangular preconditioners for (5). Since \mathcal{A} has the block decomposition

$$\begin{bmatrix} A & B^T \\ B & C \end{bmatrix} = \begin{bmatrix} A & 0 \\ B & S \end{bmatrix} \begin{bmatrix} I & A^{-1}B^T \\ 0 & I \end{bmatrix},$$

we choose the preconditioner

$$\mathcal{P}_T = \begin{bmatrix} A & 0 \\ B & \hat{S} \end{bmatrix}. \quad (11)$$

It is easy to see that

$$\mathcal{P}_T^{-1}\mathcal{A} = \begin{bmatrix} I & A^{-1}B^T \\ 0 & \hat{S}^{-1}S \end{bmatrix}$$

and so the ideal choice is $\hat{S} = S$ since then all eigenvalues of $\mathcal{P}_T^{-1}\mathcal{A}$ are 1 and GMRES converges in two steps [21]. However, the Schur complement is prohibitively expensive to apply and we are unaware of a spectrally equivalent approximation. If $\hat{S} \neq S$ then $\mathcal{P}_T^{-1}\mathcal{A}$ has n eigenvalues at 1 and the remainder are the eigenvalues of $\hat{S}^{-1}S$. We again consider the choices $\hat{S} = C$ and $\hat{S} = BA^{-1}B^T$, which correspond to each term in the ideal choice S .

When $\hat{S} = C$ the eigenvalues of $\hat{S}^{-1}S$ are $1 - \eta$, where η is an eigenvalue of $C^{-1}BA^{-1}B^T$, and (10) shows the eigenvalues of $\hat{S}^{-1}S$ lie in $(0,1)$. Thus, the eigenvalues of $\mathcal{P}_T^{-1}\mathcal{A}$ are contained in an interval but $\eta \approx 1$ will result in small eigenvalues, which may cause slow convergence rates. When $\hat{S} = BA^{-1}B^T$ the eigenvalues of $\hat{S}^{-1}S$ are given by $1/\eta - 1$. If η is small the eigenvalues of $\mathcal{P}_T^{-1}\mathcal{A}$ may be spread out while if $\eta \approx 1$ some eigenvalues of $\mathcal{P}_T^{-1}\mathcal{A}$ may be very small; both situations can negatively affect the speed of convergence. Our numerical experiments in Section 5.1 verify that C is a better choice for \hat{S} than $BA^{-1}B^T$ for our problem.

4 Domain decomposition preconditioner

We see that preconditioners built from blocks A , B and C , and the Schur complement S , are effective at reducing the number of Krylov subspace iterations required to solve (5). However, these preconditioners are costly to apply. For larger problems A is typically much bigger than C , since there are usually many more interior than boundary points. Consequently, we examine a two-level restricted additive Schwarz (RAS) domain decomposition preconditioner [4] for A , although we note that the same procedure could easily be applied to C as well. Although the RAS preconditioner is nonsymmetric it requires less communication than the additive Schwarz (AS) method, making it better suited to parallel implementations. (Note that the block triangular discussed in the previous section is already nonsymmetric.) Moreover, the rate of convergence of RAS methods is typically similar to, or better than, that of AS methods [4, 15].

In the traditional overlapping AS method, we divide the domain $V = \Omega$ into a set of disjoint subdomains, so that $V = V_{1,0} + V_{2,0} + \dots + V_{k,0}$. These subdomains are then extended by θ to give $V = V_{1,\theta} + V_{2,\theta} + \dots + V_{k,\theta}$, where θ is the number of nodes in each direction by which the overlapping domain is larger than the non-overlapping domain (see Figure 1a). The restriction of V to the i th domain $V_{i,\theta}$ is associated with the operator $R_{i,\theta}$ while the corresponding prolongation operator is $R_{i,\theta}^T$. The restriction of A to the i th domain is $A_{i,\theta} = R_{i,\theta} A R_{i,\theta}^T$ so that $A_{i,\theta}$ contains the rows and columns of A associated with the data sites in subdomain $V_{i,\theta}$. Then, the additive Schwarz (AS) preconditioner M of A is

$$M_{AS}^{-1} = \sum_{i=1}^k R_{i,\theta}^T A_{i,\theta}^{-1} R_{i,\theta}.$$

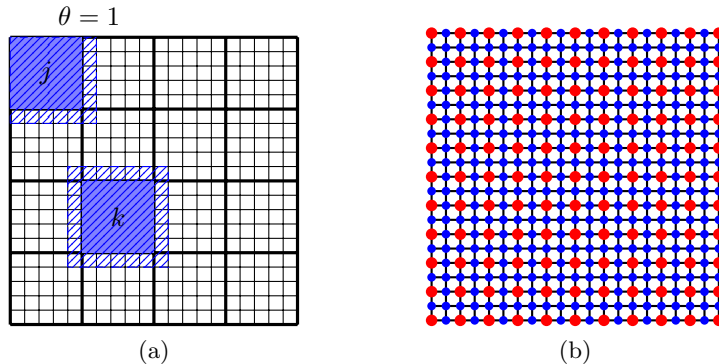
The restricted additive Schwarz preconditioner is a slight modification of the AS preconditioner which uses the overlap for computations but which only projects information from the non-overlapping domain. In this case the prolongation operator becomes $R_{i,0}$, and the RAS preconditioner is

$$M_{RAS}^{-1} = \sum_{i=1}^k R_{i,0}^T A_{i,\theta}^{-1} R_{i,\theta}.$$

Note that the subdomain solves in the AS and RAS methods could be replaced by inexact solves, and all subdomain solves can be computed in parallel, but we do not consider this here.

As the number of subdomains increases, the speed of convergence of the RAS scheme can be slow because information takes longer to propagate through all subdomains. However, this slow convergence can be remedied by incorporating a coarse-grid correction to give a two-level, or augmented, RAS method (ARASM). We consider only additive corrections and let $V_0 \subset V$ be the set of coarse grid RBF centres and R_0 and R_0^T be the corresponding restriction and prolongation operators. The coarse grid representation of A is then $A_0 =$

Figure 1: (a) Additive Schwarz subdomains on the boundary (j) and in the interior (k). The overlapping domain is shown with diagonal lines and the non-overlapping domain is shaded. The overlap is $\theta = 1$ because the overlapping domain is obtained from the non-overlapping domain by extending by one node in each direction. (b) Coarse and fine grid points with coarse grid points denoted by larger red circles.



$R_0 A R_0^T$ and the two-level preconditioner is

$$M_{ARASM}^{-1} = R_0^T A_0^{-1} R_0 + \sum_{i=1}^k R_{i,0}^T A_{i,\theta}^{-1} R_{i,\theta}.$$

In our numerical results the coarse grid contains every second node of the fine grid (see Figure 1b).

To demonstrate the effectiveness of the ARASM preconditioner, and its sensitivity to the size of the problem, number of domains and overlap, we apply right-preconditioned GMRES with the ARASM preconditioner to the matrix A in the saddle point system (5) for the test problem defined in Section 2.4. The initial vector for GMRES is the zero vector and we stop when the relative residual $\|r_k\|_2 / \|r_0\|_2$ falls below 10^{-6} , where $r_k = b^i - A\alpha_k^i$ is the residual at the k th iteration.

Table 3 shows the GMRES iteration counts. Without a preconditioner the iteration number grows rapidly with the problem size. In contrast, with the ARASM preconditioner the iteration number has a much more modest dependence on n . Moreover, this dependence is weaker when the overlap is larger and an overlap of $\theta = 4$ appears to be sufficient for n -independent iteration counts. We also note that for larger problems it seems advantageous to have more subdomains—a possible explanation is that the conditioning of the subproblems improves as the subdomains get smaller.

In summary, it seems that the ARASM preconditioner is a good approximation of A in the block diagonal and block triangular preconditioners. We verify this in the following section.

Table 3: GMRES iteration counts for the RAS preconditioner applied to A in (5). The number of subdomains is k and the overlap is θ . We denote by ‘-’ problems for which the dimension of A is too small for the number of subdomains and overlap.

n	GMRES	$k = 4$			$k = 8$			$k = 16$		
		$\theta = 1$	$\theta = 2$	$\theta = 4$	$\theta = 1$	$\theta = 2$	$\theta = 4$	$\theta = 1$	$\theta = 2$	$\theta = 4$
961	44	17	14	11	19	16	-	24	-	-
3 969	92	21	18	16	23	20	19	23	19	-
16 129	192	25	22	20	24	22	22	26	21	23
65 025	402	26	23	23	25	23	23	26	22	22
261 121	849	28	25	23	27	24	23	28	23	21

5 Numerical results

In this section we apply our preconditioners to the test problem described in Section 2.4. The problem sizes we consider are $n + m = 289, \dots, 64049$. Note that without preconditioning, conjugate gradients terminates in $n + m$ steps (see Table 4). The condition numbers for the unpreconditioned coefficient matrix are given in Table 5. Throughout, we terminate computations when the relative residual $\|r_k\|_2 / \|r_0\|_2$ falls below 10^{-8} .

5.1 Ideal preconditioner

The ideal preconditioners \mathcal{P}_D and \mathcal{P}_T described in Section 3 are examined first. We apply conjugate gradients when the block diagonal preconditioner \mathcal{P}_D is used and right-preconditioned GMRES when the block triangular preconditioner \mathcal{P}_T is used.

From Table 4 we see that, as predicted in Section 3.1, $\hat{S} = C$ gives the fastest convergence rate for the block diagonal preconditioner \mathcal{P}_D . Although the convergence speed is mesh-dependent, the iteration growth is quite modest. On the other hand, CG requires more iterations when $BA^{-1}B^T$ and S are used and the iteration growth is more rapid. Additional computations, which for brevity are not reported, show that GMRES achieves similar iteration counts to CG when $\hat{S} = C$. The CG iteration counts are reflected in the condition numbers of $\mathcal{P}_D^{-1}\mathcal{A}$ (see Table 5) which are smallest when $\hat{S} = C$. In addition, $\hat{S} = C$ has smaller condition number growth than that for the exact Schur complement S .

We consider now the ideal block triangular preconditioner \mathcal{P}_T in (11). As predicted in Section 3.2, GMRES converges in 2 steps when $\hat{S} = S$ (see Table 6). However, forming the Schur complement is prohibitively expensive for large problems. The choice $\hat{S} = C$ is a reasonable alternative and performs better than $BA^{-1}B^T$, as we might expect from the analysis in Section 3.2. An added bonus is that when $\hat{S} = C$ the growth in the iteration count appears to slow as the problem size increases. The condition number of $\mathcal{P}_T^{-1}\mathcal{A}$ is not as relevant to convergence rates as that of $\mathcal{P}_D^{-1}\mathcal{A}$, because $\mathcal{P}_T^{-1}\mathcal{A}$ is nonsymmetric, and we

Table 4: Iteration counts for PCG to reach a tolerance of 10^{-8} with the ideal block diagonal preconditioner.

\widehat{S}	289	1 089	4 225	16 641	64 049
Unpreconditioned	289	1 089	4 225	16 641	64 049
C	27	38	46	56	66
$BA^{-1}B^T$	80	114	129	149	168
S	32	51	74	108	149

Table 5: Condition numbers of $\mathcal{P}_D^{-1}\mathcal{A}$. The asterisk indicates that condition numbers are estimated for dimension 16 641 using the Matlab function `condst`.

\widehat{S}	289	1 089	4 225	16 641*
Unpreconditioned	3.3×10^4	5.5×10^5	9.0×10^6	3.4×10^8
C	2.9×10^2	2.6×10^3	2.9×10^4	3.4×10^5
$BA^{-1}B^T$	4.1×10^5	1.7×10^6	9.2×10^6	3.6×10^7
S	1.3×10^3	2.5×10^4	5.4×10^5	7.3×10^6

do not report this information. However, for all problem sizes the condition numbers of $\mathcal{P}_T^{-1}\mathcal{A}$ are of similar magnitudes to those of $\mathcal{P}_D^{-1}\mathcal{A}$ for $\widehat{S} = C$, $BA^{-1}B^T$ and are less than 100 for the Schur complement S .

5.2 Additive Schwarz preconditioners

Although Tables 4 and 6 show that we can develop preconditioners that significantly reduce the number of GMRES iterations, our ideal preconditioners require a linear solve with A at each iteration. As the level increases, both the dimension of A and the density of nonzeros grow, making this solve costly. Consequently, we see how the preconditioners are effected by replacing A by the ARASM preconditioner described in Section 4. The ARASM preconditioner is nonsymmetric and so for both the block diagonal preconditioner and block triangular preconditioner we apply right-preconditioned GMRES and set $\widehat{S} = C$. If $n < 2000$ we use 4 subdomains and an overlap of $\theta = 2$ data sites. Otherwise, we use 8 subdomains and an overlap of $\theta = 4$. When $n + m = 289$ the problem is too small to apply the ARASM preconditioner with the choice of subdomains and overlap. Our results are reported in Table 7.

Table 6: Iteration counts for GMRES to reach a tolerance of 10^{-8} with the ideal block triangular preconditioner.

\widehat{S}	289	1 089	4 225	16 641	64 049
C	14	20	23	28	29
$BA^{-1}B^T$	33	48	62	76	79
S	2	2	2	2	2

Table 7: Iteration counts for GMRES to reach a tolerance of 10^{-8} with the additive Schwarz preconditioners.

	1 089	4 225	16 641	64 049
Block diagonal	41	53	57	66
Block triangular	28	34	46	51

The ARASM block diagonal preconditioner gives higher iteration counts than the corresponding ideal preconditioner (see Table 4) for small problems but both perform similarly for larger problems. We caution that conjugate gradients is used with the ideal preconditioner and so the results are not directly comparable with those in Table 4 but, as noted in Section 5.1, GMRES behaves similarly to CG when the ideal block diagonal preconditioner (6) is applied with $\widehat{S} = C$. The convergence speed of the ARASM block triangular preconditioner is slower than for the ideal preconditioner, but the growth with dimension is still modest and iteration counts are lower than for the ARASM block diagonal preconditioner.

6 Conclusions

We have developed block diagonal and block triangular preconditioners for the symmetric positive definite systems that arise in the RBF multiscale collocation method for PDEs and have described the spectra of the preconditioned matrices for different choices of (2,2) block \widehat{S} . We find, analytically and experimentally, that the block diagonal preconditioner with $\widehat{S} = C$ has the most favourable spectrum for fast CG convergence. On the other hand, choosing $\widehat{S} = S = C - BA^{-1}B^T$ in the block triangular preconditioner guarantees convergence of GMRES in two iterations. However, the block triangular preconditioner is costly to form when $\widehat{S} = S$ and so a reasonable alternative is to choose $\widehat{S} = C$ in the block triangular preconditioner. Practical preconditioners that replace A by a restricted additive Schwarz method augmented with a coarse-grid correction (ARASM) were introduced and numerical experiments show that they are effective. Further speed-ups could be achieved by using inexact solves within the ARASM preconditioner and by exploiting parallelism. Additionally, a spectrally equivalent approximation to \widehat{S} should improve the rate of convergence of the block triangular preconditioned system.

Acknowledgements This publication was based on work supported in part by Award No KUK-C1-013-04, made by King Abdullah University of Science and Technology (KAUST).

References

- [1] R. K. Beatson, J. B. Cherrie, and C. T. Mouat. Fast fitting of radial basis functions: Methods based on preconditioned GMRES iteration. *Advances in Computational Mathematics*, 11:253–270, 1999.
- [2] R. K. Beatson, W. A. Light, and S. Billings. Fast solution of the radial basis function interpolation equations: Domain decomposition methods. *SIAM Journal on Scientific Computing*, 22:1717–1740, 2000.
- [3] M. D. Buhmann. *Radial Basis Functions: Theory and Implementations*, volume 12 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2003.
- [4] X.-C. Cai and M. Sarkis. A restricted additive Schwarz preconditioner for general sparse linear systems. *SIAM Journal on Scientific Computing*, 21:792–797, 1999.
- [5] A. Chernih and S. Hubbert. Closed form representations and properties of the generalised Wendland functions. *Journal of Approximation Theory*, 177:17–33, 2014.
- [6] Q. Deng and T. A. Driscoll. A fast treecode with multiquadric interpolation with varying shape parameters. *SIAM Journal on Scientific Computing*, 34:A1126–A1140, 2012.
- [7] P. Farrell and H. Wendland. RBF multiscale collocation for second order elliptic boundary value problems. *SIAM Journal on Numerical Analysis*, 51(4):2403–2425, 2013.
- [8] G. E. Fasshauer. Solving differential equations with radial basis functions: Multilevel methods and smoothing. *Advances in Computational Mathematics*, 11:139–159, 1999.
- [9] G. E. Fasshauer. *Meshfree Approximation Methods with MATLAB*, volume 6 of *Interdisciplinary Mathematical Sciences*. World Scientific Publishing, Hackensack, NJ, 2007.
- [10] M. S. Floater and A. Iske. Multistep scattered data interpolation using compactly supported radial basis functions. *Journal of Computational and Applied Mathematics*, 73:65–78, 1996.
- [11] N. Flyer and G. B. Wright. Transport schemes on a sphere using radial basis functions. *Journal of Computational Physics*, 226:1059–1084, 2007.
- [12] B. Fornberg and E. Lehto. Stabilization of RBF-generated finite difference methods for convective PDEs. *Journal of Computational Physics*, 230(6):2270–2285, 2011.

- [13] C. Franke and R. Schaback. Convergence order estimates of meshless collocation methods using radial basis functions. *Advances in Computational Mathematics*, 8:381–399, 1998.
- [14] R. W. Freund and N. M. Nachtigal. QMR: A quasi-minimal residual method for non-Hermitian linear systems. *Numer. Math.*, 60:315–339, 1991.
- [15] A. Frommer and D. B. Szyld. An algebraic convergence theory for restricted additive Schwarz methods using weighted max norms. *SIAM Journal on Numerical Analysis*, 39:463–479, 2001.
- [16] E. J. Fuselier. Sobolev-type approximation rates for divergence-free and curl-free RBF interpolants. *Mathematics of Computation*, 77(263):1407–1423, 2008.
- [17] P. Giesl and H. Wendland. Meshless collocation: Error estimates with application to dynamical systems. *SIAM Journal on Numerical Analysis*, 45(4):1723–1741, 2007.
- [18] A. Greenbaum. *Iterative Methods for Solving Linear Systems*. SIAM, Philadelphia, PA, 1997.
- [19] M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49:409–436, 1952.
- [20] R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, 1990.
- [21] I. C. F. Ipsen. A note on preconditioning nonsymmetric matrices. *SIAM Journal on Scientific Computing*, 23:1050–1051, 2001.
- [22] A. Iske. *Multiresolution Methods in Scattered Data Modelling*, volume 37 of *Lecture Notes in Computational Science and Engineering*. Springer-Verlag, Berlin, 2004.
- [23] E. J. Kansa. Multiquadrics—A scattered data approximation scheme with applications to computational fluid-dynamics—I Surface approximations and partial derivative estimates. *Computers & Mathematics with Applications*, 19(8–9):127–145, 1990.
- [24] Y. A. Kuznetsov. Efficient iterative solvers for elliptic finite element problems on nonmatching grids. *Russian Journal of Numerical Analysis and Mathematical Modelling*, 10:187–211, 1995.
- [25] Q. T. Le Gia, I. H. Sloan, and T. Tran. Overlapping additive Schwarz preconditioners for elliptic PDEs on the unit sphere. *Mathematics of Computation*, 78:79–101, 2009.
- [26] Q. T. Le Gia, I. H. Sloan, and H. Wendland. Multiscale RBF collocation for solving PDEs on spheres. *Numerische Mathematik*, 121(1):99–125, 2012.

- [27] Q. T. Le Gia and T. Tran. Fast iterative solvers for boundary value problems on a local spherical region. *ANZIAM Journal*, 54:C492–C507, 2013.
- [28] L. Ling and E. J. Kansa. Preconditioning for radial basis functions with domain decomposition methods. *Mathematical and Computer Modelling*, 40:1413–1427, 2004.
- [29] L. Ling and E. J. Kansa. A least-squares preconditioner for radial basis functions collocation methods. *Advances in Computational Mathematics*, 23:31–54, 2005.
- [30] W. R. Madych and S. A. Nelson. Multivariate interpolation and conditionally positive definite functions. II. *Mathematics of Computation*, 54:211–230, 1990.
- [31] M. F. Murphy, G. H. Golub, and A. J. Wathen. A note on preconditioning for indefinite linear systems. *SIAM Journal on Scientific Computing*, 21:1969–1972, 2000.
- [32] Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7:856–869, 1986.
- [33] D. Schröder and H. Wendland. A high-order, analytically divergence-free discretization method for Darcy’s problem. *Math. Comput.*, 80(273):263–277, 2011.
- [34] H. A. van der Vorst. Bi-CGSTAB: A fast and smoothly converging variant of BiCG for the solution of nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 13:631–644, 1992.
- [35] H. Wendland. Piecewise polynomial, positive definite and compactly supported radial functions of minimal degree. *Advances in Computational Mathematics*, 4(4):389–396, 1995.
- [36] H. Wendland. *Scattered Data Approximation*, volume 17 of *Cambridge Monographs on Applied and Computational Mathematics*. Cambridge University Press, Cambridge, 2005.
- [37] H. Wendland. Divergence-free kernel methods for approximating the Stokes problem. *SIAM Journal on Numerical Analysis*, 47(4):3158–3179, 2009.
- [38] H. Wendland. Multiscale analysis in Sobolev spaces on bounded domains. *Numerische Mathematik*, 116(3):493–517, 2010.
- [39] R. Yokota, L. Barba, and M. G. Knepley. PetRBF – A parallel $O(N)$ algorithm for radial basis function interpolation with Gaussians. *Computer Methods in Applied Mechanics and Engineering*, 199:1793–1804, 2010.