Hierarchical Matrices for Convection-Dominated Problems

Sabine Le Borne

Tennessee Technological University, Mathematics
(http://math.tntech.edu/~sleborne/)

Summary. Hierarchical matrices provide a technique to efficiently compute and store explicit approximations to the inverses of stiffness matrices computed in the discretization of partial differential equations. In a previous paper, Le Borne [2003], it was shown how standard $\mathcal{H}$-matrices must be modified in order to obtain good approximations in the case of a convection dominant equation with a constant convection direction. This paper deals with a generalization to arbitrary (non-constant) convection directions. We will show how these $\mathcal{H}$-matrix approximations to the inverse can be used as preconditioners in iterative methods.

1 Introduction

Considerable advancements have been achieved in algebraic and geometric multigrid solvers, state-of-the art domain decomposition methods such as FETI (Farhat et al. [2001]), direct and approximate (inverse) factorization solvers (Grote and Huckle [1997], Chow and Saad [1998]) as well as custom strategies for coarsening, partitioning, ordering, pivoting, etc., which improve the effectiveness and robustness of these methods. However, many important challenges remain which in particular include the construction of a robust solver for convection-dominant systems of PDEs.

A completely new and powerful approach for the construction of efficient preconditioners and smoothing iterations has recently been introduced that involve so-called hierarchical matrices, or $\mathcal{H}$-matrices (see, e.g., Hackbusch [1999], Grasedyck and Hackbusch [2002], Le Borne [2003]). The $\mathcal{H}$-matrix technique is a generalization of the panel clustering method and permits the treatment of fully populated matrices while restricting the requirements for storage and arithmetics (approximate matrix-vector multiplication, matrix-matrix multiplication and matrix inversion) to nearly optimal complexity $\mathcal{O}(n \log_2^\alpha n)$ for some (small) constant α. Related methods are the multipole method and the mosaic skeleton method (Tyrtyshnikov [2000]). In this paper we use an $\mathcal{H}$-matrix as a preconditioner in an iterative method to solve a convection-dominant problem. The characteristic feature that distinguishes

$\mathcal{H}$-matrices from other sparse approximate inverse techniques (SPAI) (see, e.g., Grote and Huckle [1997], Chow and Saad [1998], Benzi and Tuma [1998]) is the particular storage format of an $\mathcal{H}$-matrix that will be further explained below. Whereas these sparse (SPAI) methods typically work well if the approximated matrix contains many very small entries, the $\mathcal{H}$-matrix techniques provides convergent approximations if (large) subblocks of the approximated matrix are smooth (but not necessarily have small entries).

The remainder of this paper is organized as follows: After the introduction of the model problem in Section 2.1 we review the construction of an $\mathcal{H}$-matrix in Section 2.2. In Sections 2.3 and 2.4, modifications to the standard $\mathcal{H}$-matrix are developed for the convection-dominant case with constant and non-constant convection directions, resp. In Section 3, we will provide the results of numerical tests where $\mathcal{H}$-matrices have been used in iterative methods.

2 Preliminaries

2.1 The model problem

In this paper we consider the two-dimensional convection-diffusion equation with Dirichlet boundary conditions

$$-\epsilon \Delta u + b \cdot \nabla u = f \text{ in } \Omega = (0,1)^2, \tag{1}$$

$$u = g \text{ on } \partial\Omega \tag{2}$$

for $0 < \epsilon \ll 1$ and an arbitrary convection $b : \mathbb{R}^2 \to \mathbb{R}^2$. An (upwind) finite element discretization leads to a linear system of equations $A_h x_h = f_h$ where the parameter h characterizes the grid width of the underlying mesh. The $\mathcal{H}$-matrix technique is applicable to matrices obtained by a wide range of discretizations since its theory is based upon the approximability of the underlying Green's function by a separable function (and not on a particular discretization technique). Even though the construction of an $\mathcal{H}$-matrix is based on some knowledge on the underlying Green's function, the Green's function need not be known explicitly.

In Le Borne [2003], the case of a constant convection b was analysed. The numerical results showed better results in the case where the convection b aligned with the grid compared to a general, non-aligning convection direction. This can be explained by the numerical diffusion produced by the discretization scheme in the case of a non-aligning convection direction. Therefore, the construction of the $\mathcal{H}$-matrix should not only depend on the continuous problem but also on the amount of numerical diffusion, especially since in the case of an arbitrary, non-constant convection we typically cannot expect the grid to align with the convection.

2.2 $\mathcal{H}$-matrices

We will briefly review the the definition and standard construction of an $\mathcal{H}$-matrix in order to later derive modifications for the convection-dominant case. More details on $\mathcal{H}$-matrix approximations can be found in, e.g. Hackbusch [1999], Hackbusch and Khoromskij [2000b,a], Hackbusch et al. [2003], Grasedyck and Hackbusch [2002], and the references therein. An $\mathcal{H}$-matrix approximation to a given (fully populated) matrix $A \in \mathbb{R}^{I\times I}$ for a finite index set I is obtained by first constructing a certain block partitioning of the matrix index set $I \times I$, and then replacing each subblock $b = b_1 \times b_2 \subset I \times I$ of this partitioning that is larger than a certain threshold by a matrix of low rank $k(b)$. If this rank $k(b)$ is small compared to the number of indices contained in b_1 and b_2, then such a low rank matrix has much lower storage requirements than the approximated full matrix.

Definition 1 (R(k)-matrix representation). *Let $k, n, m \in \mathbb{N}_0$, and let $M \in \mathbb{R}^{n\times m}$ be a matrix of at most rank k. A representation of M in factorised form $M = AB^T,\ A \in \mathbb{R}^{n\times k}, B \in \mathbb{R}^{m\times k}$, with A and B stored as full matrices, is called an R(k)-matrix representation of M, or, in short, we call M an R(k)-matrix.*

Remark 1. The storage requirement $N_{R,St}(n, m, k)$ and the costs $N_{R\cdot v}(n, m, k)$ for the matrix-vector product with a matrix $M \in \mathbb{R}^{n\times m}$ in R(k)-matrix representation are $N_{R,St}(n, m, k) = k(n+m)$ and $N_{R\cdot v}(n, m, k) = 2k(n+m)-n-k$.

Compared to the respective complexities for full matrices, $\mathcal{O}(nm)$, we have significant savings for the R(k)-matrix if the rank k is small compared to the size of the matrix.

Definition 2 ($\mathcal{H}$-matrix). *Let $n_{\min} \in \mathbb{N}_0$. Let $\mathcal{P}$ be a partition of the block index set $I \times I$. Let $k : \mathcal{P} \to \mathbb{N}_0$ be a mapping that assigns a rank $k(b)$ to each block $b = s \times t \in \mathcal{P}$. The set of $\mathcal{H}$-matrices induced by the partition $\mathcal{P}$ and with minimum block size $n_{\min}$ is defined by*

$$\mathcal{H}(\mathcal{P}, k) := \{M \in \mathbb{R}^{I\times I}\ |\forall s \times t \in \mathcal{P} : \operatorname{rank}(M|_{s\times t}) \leq k(s \times t) \text{ or } \min\{\#s, \#t\} \leq n_{\min}\}.$$

A matrix $M \in \mathcal{H}(\mathcal{P}, k)$ is said to be given in $\mathcal{H}$-matrix representation if the blocks $M|_{s\times t}$ with rank$(M|_{s\times t}) \leq k(s \times t)$ are stored in R(k)-matrix representation and the remaining blocks with $\min\{\#s, \#t\} \leq n_{\min}$ *as full matrices.*

The accuracy of an $\mathcal{H}$-matrix approximation depends on how well the individual blocks in the partition can be approximated by low rank matrices, which in turn depends on the approximability of the underlying Green's function by separable functions as well as the ordering of the unknowns. To obtain a suitable block partition, we construct a hierarchy of partitionings from which we choose the "coarsest" one that satisfies a certain admissibility condition which shall ensure the approximability by a low rank matrix. The construction of a hierarchy of partitionings of an index set is shown in Figure 1. The hierarchical index set partition of Figure 1 does not state how to divide an index set into two subsets. Typically, the indices are ordered in a certain

Let $I = I_{0,0}$ be a finite index set. If the jth subset on level ℓ, $I_{\ell,j} \subset I$, contains more than one index, we subdivide it into two disjoint successor index sets $I_{\ell+1,j\ell-1}$ and $I_{\ell+1,j\ell}$ of approximately the same size on the next level $\ell+1$ that satisfy $I_{\ell,j} = I_{\ell+1,j\ell-1} \cup I_{\ell+1,j\ell}$.

Fig. 1. Hierarchical index set partitioning

way based upon the geometric information associated with the indices, and then this ordered list of indices is bisected into two sets of approximately the same size. In the case of uniformly elliptic differential operators, it has been shown in Bebendorf and Hackbusch [2003] that a partitioning into subsets with small diameters (with respect to the Euclidean norm) will lead to a convergent $\mathcal{H}$-matrix approximation. Such a partition is obtained if the indices within each index set $I_{\ell,j}$ are ordered as follows:

$$
\begin{aligned}
&if \max_{v,w\in I_{\ell,j}} |x_v - x_w| > \max_{v,w\in I_{\ell,j}} |y_v - y_w| \text{ then} \\
&\qquad n(v) < n(w) \text{ if } x_v < x_w \text{ or } (x_v = x_w \text{ and } y_v < y_w) \\
&\text{else} \quad n(v) < n(w) \text{ if } y_v < y_w \text{ or } (y_v = y_w \text{ and } x_v < x_w).
\end{aligned}
$$

Here, (x_v, y_v) describes the geometric location associated with an index v, and $n(v) \in \{1, \cdots, \#I_{\ell,j}\}$ assigns the index number. We will refer to this type of bisection as the *standard partition* or *geometric bisection.*

In order to define an admissibility condition, let $B_{\ell,j} := B_{I_{\ell,j}}$ be an axially parallel bounding box that contains the union of the supports of the basis functions corresponding to the indices in $I_{\ell,j}$. Then, the standard *admissibility condition* is given by: $I_{\ell,j} \times I_{\ell,k}$ is admissible if

$$
\min\{\operatorname{diam}(B_{\ell,j}), \operatorname{diam}(B_{\ell,k})\} \leq \eta \operatorname{dist}(B_{\ell,j}, B_{\ell,k}) \tag{3}
$$

for some parameter $\eta > 0$.

Given a hierarchical index set partitioning, a hierarchy of partitionings of the block index set $I \times I$ is obtained in a canonical way as shown in Figure 2.

Let a hierarchical index set partitioning be given. We define a hierarchy of block partitionings by defining $I \times I = I_{0,0} \times I_{0,0}$, and a block $b := I_{\ell,j_1} \times I_{\ell,j_2}$ satisfies exactly one of the following three conditions:

(i) b satisfies an admissibility condition (3),
(ii) $\min\{\#I_{\ell,j_1}, \#I_{\ell,j_2}\} \leq n_{\min}$,
(iii) b has (four) successors $I_{\ell+1,k_1} \times I_{\ell+1,k_2}$ where $I_{\ell+1,k_1}$ and $I_{\ell+1,k_2}$ are successors of I_{ℓ,j_1} and I_{ℓ,j_2}, resp..

Fig. 2. Hierarchical block index set partitioning

In terms of the respective matrix blocks, the three cases (i) - (iii) correspond to (i) the approximation of a block that satisfies the admissibility condition (3) by an $R(k)$-matrix, (ii) the representation of small blocks as full matrices, and, (iii) the subdivision of blocks that have successors in the hierarchical block index set partition.

2.3 Modifications for constant convection directions

In Le Borne [2003], modifications to the standard $\mathcal{H}$-matrix have been developed at first for the pure convection case $\epsilon = 0$ and then been generalized for arbitrary $\epsilon > 0$ to an ϵ- and $\mathbf{b}$-dependent partitioning and admissibility condition which produce a gradual transition from the standard partitioning and admissibility condition to their modified counterparts as $\epsilon \to 0$. In order to generate such a gradual transition for a *constant* convection direction $\mathbf{b}$, the (Euclidean) norm that was used for the calculation of the diameter and distance of clusters in the admissibility (3) has been replaced by the norm

$$\|\mathbf{x}\|_{\alpha,\mathbf{b}} := \sqrt{\alpha(\mathbf{b}\cdot\mathbf{x})^2 + (\mathbf{b}^\perp \cdot \mathbf{x})^2} \qquad \text{for} \qquad \mathbf{x} = (x_1, x_2)^T \in \mathbb{R}^2$$

where $\mathbf{b}$ is the convection vector in the convection-diffusion equation, $\mathbf{b}^\perp$ is its orthogonal complement, and $\alpha \in \mathbb{R}^+$ is a parameter that depends on the convection dominance given by ϵ, the mesh width h, and the numerical viscosity induced by the discretization.

In the index partitioning algorithm we will now use bounding boxes that are parallel to the convection $\mathbf{b}$ and its orthogonal complement $\mathbf{b}^\perp$, and the objective is no longer to produce subsets with small diameters but rather to produce subsets *stretched in convection direction.* The modified partition is obtained as follows:

$$\text{if} \quad \big(\max_{v,w\in B_{\ell,j}} \{\alpha|\mathbf{b}\cdot(v-w)|\} > \max_{v,w\in B_{\ell,j}} |\mathbf{b}^\perp \cdot (v-w)|\big) \text{ then}$$

partition cluster $I_{\ell,j}$ along $\mathbf{b}^\perp$ (orthogonal complement of $\mathbf{b}$)

else partition cluster $I_{\ell,j}$ along $\mathbf{b}$ (convection vector);

If we set $\alpha = 1$ and $\mathbf{b} = (1,0)^T$ we obtain the standard partition. Given such a hierarchical index partitioning, we will then construct the hierarchy of block partitioning in the canonical way described in Figure 2.

2.4 Modifications for non-constant convection directions

In the case of a non-constant convection $\mathbf{b}$, we begin our consideration with an example where the convection aligns perfectly with the underlying grid as shown in Figure 3. We will later generalize our strategy to the more realistic case of a convection $\mathbf{b}$ that does not necessarily align with the grid. In the

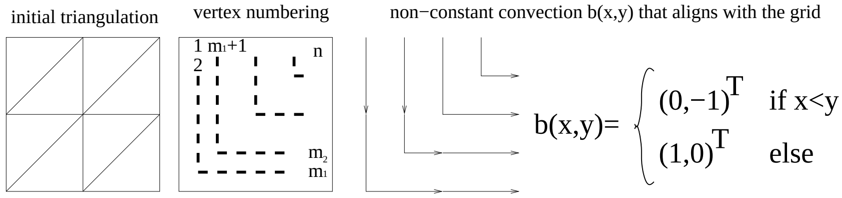

Fig. 3. Non-constant convection that aligns with the grid

given example, we will order the n unknowns with respect to the convection as indicated in Figure 3.

We now construct an $\mathcal{H}$-matrix structure using the index and block index partitioning as described in Figures 1 and 2. Using the weak admissibility condition: $s \times t$ is admissible if $s \neq t$ (i.e., all off-diagonal blocks are admissible), we can represent the exact inverse in the case of $\epsilon = 0$ (pure convection problem) as an $\mathcal{H}$-matrix with local ranks $k(b) = 1$. The storage costs for this $\mathcal{H}$-matrix structure amount to $\mathcal{O}(n \log_2 n)$ as proven in [Hackbusch, 1999, Lemma 3.1]. The fact that we indeed represent the exact inverse results from the particular ordering which guarantees that off-diagonal blocks have at most rank 1.

In the case of a non-zero parameter ϵ or a non-aligning convection direction, the discrete system will contain some artificial diffusion. In the case of a constant convection $\mathbf{b}$, we introduced a parameter α to let the amount of diffusion control the partition and admissibility. In the case of non-constant convection, the general idea to proceed is as follows: For the first p refinement steps, we use a precomputed downwind ordering of the unknowns (along the non-constant convection direction) to partition the index set. Suitable downwind ordering strategies can be found in, e.g, Le Borne [2000]. For any further refinement steps, we use the standard partitioning (trying to obtain subsets with small diameters).

3 $\mathcal{H}$-matrices in iterative methods and numerical results

The $\mathcal{H}$-matrix technique allows to compute a data-sparse approximation $A^{-\mathcal{H}}$ to a (typically fully populated) matrix A^{-1} in nearly optimal complexity. Such an approximation can be used

- in a linear iteration $x_{i+1} = x_i - A^{-\mathcal{H}}(Ax_i - b)$,
- as a preconditioner in a Krylov subspace method (e.g., BiCG-stab, GMRES, etc.), or
- as a smoother in a multigrid iteration,
- for the computation of Schur complements and their inverses, etc.

Here we provide numerical results for the first two applications. The convection-diffusion equation (1) serves as a test problem for various values of ϵ and convection directions $\mathbf{b} = (1, 0)^T$ (Table 1) and $\mathbf{b}(x, y) = (0.5-y, x-0.5)^T$

(Table 2). In Tables 1 and 2 we provide the number of iteration steps that are necessary to reduce the Euclidean norm of the residual $\|b - Ax_i\|_2$ to an accuracy of 10^{-8} for $n = 16129$ unknowns (with a maximum number of iterations of 100 and initial iterate $x_0 = 100(1, \cdots, 1)^T$). The time per iteration step is recorded in the second column of Table 1 (computed on a DELL Precision Workstation, 2.4GHz, compare with 0.012s for a classical Gauß-Seidel step).

Table 1. Iteration steps for $\mathbf{b} = (0,1)^T$, modified $\mathcal{H}$-matrix

	time per	basic iteration				bicg-stab			
k(b)/ϵ	step (s)	1	1e-2	1e-4	1e-6	1	1e-2	1e-4	1e-6
1	0.117	100	19	4	2	100	9	2	1
2	0.128	61	7	3	2	11	4	2	1
3	0.143	8	5	3	2	4	3	2	1
4	0.155	6	4	3	2	3	2	2	1
5	0.165	4	3	2	2	2	2	1	1
6	0.183	3	3	2	2	2	2	1	1

As expected, the number of necessary steps decreases considerably as we increase the local rank of the $\mathcal{H}$-matrix. For the numerical tests reported in Table 2 we used the standard $\mathcal{H}$-matrix. For the convection-dominant case $\epsilon = 10^{-6}$, we also provide in parentheses the results for the modified partition. Here, in the first two index partitions the indices have been ordered with respect to their distance to the circle origin $(0.5, 0.5)$. All further partitions were performed in the standard way. We observe slight improvements.

Table 2. Iteration steps for $\mathbf{b}$ =circle, standard $\mathcal{H}$-matrix

	basic iteration				bicg-stab			
k(b)/ϵ	1	1e-2	1e-4	1e-6	1	1e-2	1e-4	1e-6
1	100	100	100	100 (100)	100	100	63	69 (90)
2	57	26	100	100 (100)	11	9	26	40 (29)
3	8	6	39	55 (72)	4	3	11	13 (11)
4	6	4	20	34 (23)	3	2	7	9 (6)
5	4	3	9	14 (10)	2	2	4	5 (4)
6	3	3	8	11 (8)	2	2	4	4 (3)

The $\mathcal{H}$-matrix approximations $A^{-\mathcal{H}}$ have been computed using a block Gauß elimination process and are therefore not necessarily the best possible approximations. When evaluating a preconditioner, the costs for the construction of the preconditioner have to be taken into account. In this case, the construction of the $\mathcal{H}$-matrix $A^{-\mathcal{H}}$ is of nearly optimal complexity $\mathcal{O}(n \log_2^2 n)$, however, with a relatively high constant, see Grasedyck and Hackbusch [2002] (taking 73s ($k = 1$) up to 166s ($k = 6$) for $\epsilon = 1.0$ and 40s ($k = 1$) up to 76s ($k = 6$) for $\epsilon = 1e - 6$). $A^{-\mathcal{H}}$ can, however, be computed more efficiently via a (parallelizable) domain decomposition algorithm, see Hackbusch [2002].

These positive results encourage the further study of $\mathcal{H}$-matrix preconditioners in harder problems involving non-constant, cyclic convection directions in systems of PDEs in three spatial dimensions where there is still a need for efficient iteration methods.

References

M. Bebendorf and W. Hackbusch. Existence of $\mathcal{H}$-matrix Approximants to the inverse FE-Matrix of Elliptic Operators with L^∞-Coefficients. *Numerische Mathematik*, 95:1–28, 2003. available electronically at http://www.mis.mpg.de/preprints/2002/prepr2002_21.html.

M. Benzi and M. Tuma. A sparse approximate inverse preconditioner for nonsymmetric linear systems. *SIAM J. Sci. Comput.*, 19:968–994, 1998.

E. Chow and Y. Saad. Approximate inverse preconditioners via sparse-sparse iterations. *SIAM J. Sci. Comp.*, 19:995–1023, 1998.

C. Farhat, M. Lesoinne, P. L. Tallec, K. Pierson, and D. Rixen. FETI-DP: A dual-primal unified FETI method - part I: A faster alternative to the two-level FETI method. *J. Numer. Methods. Eng.*, 50:1523–1544, 2001.

L. Grasedyck and W. Hackbusch. Construction and arithmetics of $\mathcal{H}$-matrices. *Computing*, 70:295–334, 2002.

M. Grote and T. Huckle. Parallel preconditioning with sparse approximate inverses. *SIAM J. Sci. Comput.*, 18:838–853, 1997.

W. Hackbusch. A sparse matrix arithmetic based on $\mathcal{H}$-matrices. Part I: Introduction to $\mathcal{H}$-matrices. *Computing*, 62:89–108, 1999.

W. Hackbusch. Direct domain decomposition using the hierarchical matrix technique. In *Proceedings of the 14th International Conference on Domain Decomposition Methods in Cocoyoc, Mexico*, pages 39–50. ddm.org, 2002.

W. Hackbusch, L. Grasedyck, and S. Börm. Introduction to hierarchical matrices with applications. *Engineering Analysis with Boundary Elements*, 27: 405–422, 2003.

W. Hackbusch and B. Khoromskij. A sparse $\mathcal{H}$-matrix arithmetic: General complexity estimates. *J. Comp. Appl. Math.*, 125:479–501, 2000a.

W. Hackbusch and B. Khoromskij. A sparse $\mathcal{H}$-matrix arithmetic. Part II: Application to multi-dimensional problems. *Computing*, 64:21–47, 2000b.

S. Le Borne. Ordering techniques for two- and three-dimensional convection-dominated elliptic boundary value problems. *Computing*, 64:123–155, 2000.

S. Le Borne. $\mathcal{H}$-matrices for convection-diffusion problems with constant convection. *Computing*, 70:261–274, 2003.

E. E. Tyrtyshnikov. Incomplete cross approximation in the mosaic-skeleton method. *Computing*, 64:367–380, 2000.