

Incomplete Cross Approximation in the Mosaic-Skeleton Method

Eugene E. Tyrtysnikov *

ABSTRACT

The mosaic-skeleton method was bred in a simple observation that rather large blocks in very large matrices coming from integral formulations can be approximated accurately by a sum of just few rank-one matrices (skeletons). These blocks might correspond to a region where the kernel is smooth enough, and anyway it can be a region where the kernel is approximated by a short sum of separable functions (functional skeletons). Since the effect of approximations is like that of having small-rank matrices, we find it pertinent to say about *mosaic ranks* of a matrix which turn to be pretty small for many nonsingular matrices.

On the first stage, the method builds up an appropriate mosaic partitioning using the concept of a tree of clusters and some extra information rather than the matrix entries (related to the mesh). On the second stage, it approximates every (allowed) block by skeletons using the entries of some rather small cross which is chosen by an adaptive procedure. We focus chiefly on some aspects of practical implementation and numerical examples on which the approximation time was found to grow almost linearly in the matrix size.

*Institute of Numerical Mathematics, Russian Academy of Sciences, Gubkina 8, Moscow 117333, Russia. Supported by the RFBR Grant 97-01-00155 and Volkswagen-Stiftung.

1 Introduction

The mosaic-skeleton method [21, 22] was bred in a simple observation that rather large blocks in very large matrices coming from integral formulations can be approximated accurately by a sum of just few skeletons (some say dyads or rank-one matrices). These blocks might correspond to a region where the kernel is smooth enough, and anyway it can be a region where the kernel is approximated by a short sum of separable functions (in other words, functional skeletons). From a historical point of view, we may note that the earliest mention of low-rank blocks in dense matrices we are aware of was made in [26] (yet in the context entirely different from ours).

The mosaic-skeleton approximations are easy to result in fast approximate matrix-vector multiplication algorithms close by nature to those of multipoles [11, 15, 16, 19, 20], panel clustering [12], interpolation [2, 3, 17], and wavelet-based approaches [1, 13] (we apologize for not mentioning here many other important works related to these utterly topical fields). All the said techniques involve some hierarchy of interface regions and function approximants. What the mosaic-skeleton method differs in from others is a matrix analysis view on largely the same problem. Such a view can be very useful due to the generality of matrix theory approaches. Since the effect of approximations is like that of having small-rank matrices, we find it pertinent to say about *mosaic ranks* of a matrix which turn to be pretty small for many nonsingular matrices.

From a practical point of view, the mosaic-skeleton methods is the only one that works explicitly with the entries of a matrix. It is crucial that it works only with a small part of the entries. These are entries of some cross in every block allowing for a low-rank approximation. First of all, we need to find the list of such blocks (and other blocks as well). On this stage, we fall back to the concept of a tree of clusters due to Hackbusch and Novak [12]. Apart from the entries, however, we need some extra information related to the mesh. In this sense, I rather like to say that we discuss a *grey box* solver, in contrast to *black box* solvers, for large dense unstructured matrices (T. Chan told me that he already used this term, luckily in the same sense).

The first stage of the method is discussed in Section 3. Then, in Section 4, we get to the second stage where (allowed) blocks are to be approximated by skeletons. It is done using the entries of some rather small cross in the block. Since only a part of the whole block is involved, we call the approach an incomplete cross approximation. The cross is selected by an adaptive procedure involving more (yet not many) rows and columns step by step.

On the whole, the procedure is inspired by the concept of maximal volumes which is of tremendous value in the approximation theory and now adopted to matrix approximation problems. In Section 5, we present numerical examples showing that the solution time depends almost linearly on the matrix size.

In spite of our intention to focus here only on practical aspects, in the next section we begin still with a brief discussion why and when mosaic ranks appear to be small. We discuss estimates on mosaic ranks with a special stress on some important assumptions worthy to be put explicitly (though it is not the case in many papers on compression strategies). We also present some relations between mosaic ranks of matrices and their inverses.

2 Mosaic ranks

Consider a matrix A of size $m \times n$, and let A_k be a submatrix which is $m_k \times n_k$. Denote by $\Pi(A_k)$ a matrix of the same size as A with zeroes in the positions that are not occupied by A_k . A finite set of submatrices A_k is called a *covering* of A if $A = \sum_k \Pi(A_k)$, and a *mosaic partitioning* if the submatrices have no common entries. If A_k is of rank r_k , then A_k is a sum of k skeletons (rank-one matrices), and the memory for retaining A_k is $\text{mem}(A_k) = \min\{r_k(m_k + n_k), m_k n_k\}$. The mosaic rank associated with the given mosaic partitioning is defined as [22] $\text{mr}(A) = \sum_k \text{mem}(A_k)/(m + n)$. If $m = n$ and A is of rank r , then we can store A using only $2rn$ memory cells. The same holds true if A is a nonsingular matrix and r is its mosaic rank.

We are also interested in the case when A_k are approximated by matrices of rank r_k with accuracy ε . If the Frobenius norm is used, then, obviously, $\|A - \sum_k \Pi(A_k)\|_F \leq \varepsilon \|A\|_F$. In such cases, we consider *approximate mosaic ranks* (ε -ranks).

Very large nonsingular matrices coming from integral equations are usually dense and unstructured. Nevertheless, they may be approximated by matrices of low mosaic rank.

For example, consider two typical single layer potential equations related to the Dirichlet boundary value problems for the Laplace ($\Delta w = 0$) and Helmholtz ($(\Delta + k^2)w = 0$) equations in two dimensions. The first one is

$$-\frac{1}{2\pi} \int_{\partial\Omega} \log|x - y| U(y) ds(y) = F(x), \quad x \in \partial\Omega, \quad (1)$$

and the second is

$$\frac{i}{4} \int_{\partial\Omega} H_0^{(1)}(k|x-y|) U(y) ds(y) = F(x), \quad x \in \partial\Omega, \quad (2)$$

here $ds(y)$ is the arclength element (for simplicity, let $\partial\Omega$ be an infinitely smooth closed curve cutting the plane into two parts). The kernels for the equations are the fundamental solutions for the Laplace and Helmholtz equations in two dimensions. $H_0^{(1)}$ is the Hankel function of order 0 of the first kind. We assume that Ω is such that (1) and (2) have only trivial solution in case $F(x) = 0$.

Let $\partial\Omega$ be an ellipse with half-axes $a = 1$ and $b = 0.5$. We use the Galerkin method with piecewise constant functions. Let us see how approximate mosaic ranks behave as n increases. In Table 1, there are results computed for the equation (1) with $\varepsilon = 10^{-4}$. We also output the *compression factor* which is the total memory (that would have been needed to keep the original matrix) over the memory used to keep skeletons of the mosaic approximations. In Table 2, there are results computed for the equation (2) with $\varepsilon = 10^{-3}$; here $k = 1$ and the ellipse parameters are $a = 1$ and $b = 0.25$.

Matrix order	512	1024	2048	4096	8192	16384	32768
Mosaic rank	63.46	71.48	78.44	86.84	93.80	100.76	106.50
Compression factor	24.79%	13.96%	7.66%	4.24%	2.29%	1.23%	0.65%

Table 1. Mosaic ranks for the logarithmic-kernel equation.

Matrix order	256	512	1024	2048	4096
Mosaic rank	46.08	53.76	61.44	71.68	81.92
Compression factor	36%	21%	12%	7%	4%

Table 2. Mosaic ranks for the Hankel-kernel equation.

As is seen from Table 1, in matrix-vector multiplications we can work with the nonsingular matrix of order $n = 32768$ as if its classical rank were about 106. It is equal to say that the memory used is equal to 0.65% of n^2 .

An important observation from the tables is that the compression factor becomes about twice smaller as n increases twice. This means that the memory and arithmetic work for the matrix-vector multiplication manifest

about linear behavior in n . It is important to understand why and when this is the case. On the whole, it is sufficient to do this for matrices of the form

$$A_n = [f(x_{in}, y_{jn})], \quad 1 \leq i, j \leq n,$$

where $f(x, y)$ is a function of x and y from a bounded region S in the m -dimensional space, and x_{in} and y_{jn} are the nodes of some meshes.

Assumption 1. Let f be *asymptotically smooth* in the sense that there exist $c, d > 0$ and a real number g such that

$$|\partial^p f(x, y)| \leq c_p |x - y|^{g-p} \quad (3)$$

where

$$c_p \leq c d^p p!. \quad (4)$$

Here, ∂^p is any p -order derivative in $y = (y_1, \dots, y_m)$:

$$\partial^p = \left(\frac{\partial}{\partial y_1} \right)^{i_1} \cdots \left(\frac{\partial}{\partial y_m} \right)^{i_m}, \quad i_1 + \dots + i_m = p.$$

Assumption 2. Let the meshes be *quasi-uniform* in the sense that there are positive constants c_1 and c_2 such that

$$c_1 \frac{\text{mes } S'}{\text{mes } S} \leq \mu(S') \leq c_2 \frac{\text{mes } S'}{\text{mes } S},$$

where S' is any subregion of S and μ counts how many nodes of the n th mesh fall into S' .

Theorem 1. [22] *Under Assumptions 1 and 2, for any $\delta > 0$ there are splittings $A_n = T_n + R_n$ where*

$$\text{mr } T_n = O(\log^{m+1} n), \quad \|R_n\|_F = O(n^{-\delta}). \quad (5)$$

Note that the concept of asymptotical smoothness in the sense of (3) was introduced by Brandt [2]. However, any proof for compression strategies would be incomplete if we do not say how c_p may grow in p . Thus, (4) should be regarded as an essential complement to (3). (It appeared explicitly, likely first, in [22].)

Theorem 1 can be generalized in several ways. First of all, we can substitute Assumption 1 with the following.

Assumption 1'. Assume that for any y_0 there are $\alpha(p)$ functions $u_i(x, y_0)$ and $v_i(y, y_0)$ such that, whenever $\|y - y_0\| < \|x - y_0\|$,

$$f(x, y) = \sum_{i=0}^{\alpha(p)} u_i(x, y_0)v_i(y, y_0) + E_p, \quad (6)$$

where

$$|E_p| \leq \beta(p) \left(\frac{\|y - y_0\|}{\|x - y_0\|} \right)^p. \quad (7)$$

Moreover, for some positive c , d , and γ ,

$$\alpha(p) \leq c p^\gamma, \quad \beta(p) \leq c d^p. \quad (8)$$

Theorem 2. Under Assumptions 1' and 2, for any $\delta > 0$ there are splittings $A_n = T_n + R_n$ where

$$\text{mr } T_n = O(\log^{\gamma+1} n), \quad \|R_n\|_F = O(n^{-\delta}). \quad (9)$$

We omit the proof because, on the whole, it would repeat part of the proof from [22].

In comparison with Theorem 1, the new formulation has at least two advantages. First, with some special expansions other than the Taylor series, it might be proved, sometimes, that $\gamma < m$, and this leads to a finer estimate. Second, some oscillatory kernels as that of the equation (2) are not asymptotically smooth. It is still possible to prove that these kernels fulfil Assumption 2 [7].

We can not do without any assumption on meshes. However, Assumption 2 seemed to be essential chiefly for the algebraic technique proposed and used in [22], and we always thought that it could be weakened. Now it is done in [7]. More precisely, the lower estimate on $\mu(S')$ is no longer in need. Consequently, S could be a closed curve on the plane or a surface in the 3D case [7].

Besides the above theorems, there are other general classes of low-mosaic-rank matrices. For example, such are the inverses to banded matrices. To

present this result, we introduce the concept of weakly semiseparable matrices.

Assume that an $n \times n$ matrix A is such that the rank of any submatrix located in its upper triangular (and lower triangular) part is at most p (and q). Then A is said to be a (p, q) *weakly semiseparable* matrix.

Theorem 3. *For any (p, q) weakly semiseparable matrix A_n of order $n > n_0 > 0$,*

$$\text{mr } A_n \leq c(n_0) \max\{p, q\} \log_2 2n,$$

where

$$c(n_0) = \frac{1}{\log_2 \frac{2}{\frac{1}{4n_0^2} + 1}}.$$

If n is a power of two, we subdivide A_n into square blocks of order $n/2$ and, then, proceed by induction:

$$\frac{np + nq + 2n \max\{p, q\} \log n}{2n} = \max\{p, q\} \left(\frac{1}{2} + \log n \right) \leq \max\{p, q\} \log 2n.$$

Thus, we obtained the estimate even sharper than what we are after. However, the case of an arbitrary n seems to be more tricky (the embedding might not work because the mosaic rank of a submatrix is not necessarily smaller than that of the matrix), and that is where we luckily fall back to the algebraic approach proposed in [22]. We can interpret the result also as follows: for any $\varepsilon > 0$, it holds

$$\text{mr } A_n \leq (1 + \varepsilon) \max\{p, q\} \log_2 2n$$

for all n sufficiently large.

Now, the estimation of mosaic ranks for the inverses to banded matrices reduces to the same problem for weakly semiseparable matrices because the inverse to a nonsingular (p, q) banded matrix is a (p, q) weakly semiseparable matrix. This is due to the following proposition.

Proposition. *A nonsingular matrix is (p, q) weakly semiseparable if and only if its inverse possesses the same property.*

Proof. It is sufficient to prove the following assertion. Let

$$A = \begin{bmatrix} A_{11} & A_{12} \\ A_{21} & A_{22} \end{bmatrix}, \quad A^{-1} = \begin{bmatrix} B_{11} & B_{12} \\ B_{21} & B_{22} \end{bmatrix},$$

where the diagonal blocks are square. Then $\text{rank } B_{12} = \text{rank } A_{12}$.

If A_{11} is nonsingular, then this emanates from the Frobenius formulas. Otherwise, consider matrices $A(t) = A + tI$ with a real parameter t . For all sufficiently small $0 < t < \varepsilon$ the block $A_{11}(t)$ is nonsingular (it follows that $A_{22}(t)$ is also nonsingular). Consequently, $\text{rank } B_{12}(t) \leq \text{rank } A_{12}(t)$ for any $0 < t < \varepsilon$. Obviously, $A(t) \rightarrow A$ as $t \rightarrow 0$, and, in the limit, we come to $\text{rank } B_{12} \leq \text{rank } A_{12}$. (It might be worthy to remark that we can not pass to the limit in the Frobenius formulas.) \square

Now we finish the discussion of theoretical estimates (which is not the main purpose in this paper) and get to practical issues.

3 Mosaic partitionings

On the first stage of the mosaic-skeleton method, we are to choose a suitable mosaic partitioning. To do this, we do not compute any entries of the given matrix. Instead of this, we rely on some (rather weak indeed) geometrical information.

We require that every entry is associated with two points, x_i and y_j , in the m -dimensional space. We think it is sufficient to have one mesh instead of two ($x_i = y_i$), though we could keep the two if necessary.

We call a *cluster* any subset of nodes x_1, \dots, x_n furnished with a finite (independent of the number of nodes) set of attributes. In the present algorithms, we use only two attributes: the center c , defined as the mean radius-vector for the nodes, and radius r , defined as the maximal distance between c and the nodes. Prior to forming the list of blocks we construct a *tree of clusters* suggested in [12].

The root of this tree is the cluster containing all the nodes. Apart from the nodes, there are two input parameters in our algorithm: the maximal number of levels L_{\max} and a *separator* for any given cluster (the procedure that subdivides a cluster into several subclusters, if possible). The tree of clusters is described by a list of clusters \mathcal{T} and an integer array \mathcal{P} containing a permutation of $\{1, 2, \dots, n\}$. Each cluster is identified by its index in the list \mathcal{T} . Every item in \mathcal{T} has the following components:

- Level where the cluster belongs.
- Index of the parent cluster.
- Number of kid-clusters.

- Index of a cluster that is followed contiguously by the kid-clusters.
- Number of nodes in the cluster.
- Reference to the position in the permutation array \mathcal{P} which is followed contiguously by indices of the nodes forming the cluster.

Algorithm 1. Let $N(\mathcal{T})$ denote the number of items in \mathcal{T} and l the maximal level in the subtree already constructed.

- Set $N_{\text{beg}} = 0$, $N_{\text{end}} = 1$, and form Item no. 1 in \mathcal{T} (the number of kid-clusters is 0).
- For every i from $N_{\text{beg}} + 1$ to N_{end} , check if the i th cluster is a leaf (has no kid-clusters). If so, try to split it into subclusters using the given separator. Every subcluster, if any, is successively added to the current end of \mathcal{T} . Also, the contiguous indices of \mathcal{P} corresponding to the i th cluster are permuted to make the indices associated with every subcluster run contiguously. Correct the kid-information in Item no. i of \mathcal{T} .
- Increase l by 1 and quit if it is equal to L_{max} . If not, set $N_{\text{beg}} = N_{\text{end}}$, $N_{\text{end}} = N(\mathcal{T})$ and go to one step back.

The complexity of Algorithm 1 depends on the separator used. If the separation time is linear in the number of nodes of a cluster under separation and the number of subclusters on output is upper bounded uniformly in n , then the working time for Algorithm 1 is linear in n .

At present, we tested two methods of separation. The first is strictly motivated by the asymptotical-smoothness property. The second is somewhat heuristic. As we found, both lead to about equal results in practice.

1. Find a (minimal) parallelepiped containing all the nodes of a cluster, subdivide it into 2^m parallelepipeds, and make up subclusters of the nodes fallen into each of them.

2. Let a cluster consist of the nodes x_1, \dots, x_k with a center c . Find a hyperplane passing through c so that the sum of squared projections of $x_i - c$ onto this hyperplane is maximal. If h is a unit vector normal to the hyperplane in question, then we need to minimize

$$\Phi \equiv \sum_{i=1}^k |(x_i - c, h)|^2 = h^T M h, \quad M = \sum_{i=1}^k (x_i - c)(x_i - c)^T.$$

It is easy to see that h is the eigenvector of M for its minimal eigenvalue. The hyperplane subdivides the space into two subspaces which accumulate the nodes of two possible subclusters.

When we proceed to preparing the list of blocks, we confine ourselves to considering only those blocks that are associated with pairs of clusters in the tree of clusters (or two of them, in case there are two meshes). We permute the original rows and columns in line with \mathcal{P} ; then the tree-of-clusters blocks contain contiguous rows and columns.

It is possible to produce many different lists of blocks based on the same tree of clusters. Our final goal is a tree-of-cluster mosaic partitioning minimizing approximate mosaic rank. However, it might be a difficult discrete optimization problem. We use the following heuristic approach: low-rank blocks should be of maximal possible size. We also assume that, for every pair of clusters, there is an easily computable tag showing if the block is allowed to be compressed or not; for brevity, call the former *allowed* blocks.

Algorithm 2. Let \mathcal{M} be the target list of blocks and $\mathcal{M}_1, \mathcal{M}_2$ auxiliary lists.

- Put in \mathcal{M}_1 the root-cluster block (original matrix).
- Take up successively the blocks from \mathcal{M}_1 . If the block is allowed, relegate it immediately to \mathcal{M} . If not, consider the clusters a and b defining this block. Add to \mathcal{M}_2 all the blocks associated with the subclusters of a and b . If there are no subclusters, move this block to \mathcal{M} .
- Quit if \mathcal{M}_2 is empty. Otherwise, substitute \mathcal{M}_1 with \mathcal{M}_2 , empty \mathcal{M}_2 , and go to one step back.

In practice, the algorithms of this section consume rather negligible part of the whole time for the mosaic-skeleton method.

4 Incomplete cross approximation

On the second stage of the mosaic-skeleton method, we browse in the list of blocks and try to *compress* (approximate by skeletons) every allowed block. It can be done by the Lanczos bidiagonalization method. Although we eventually compute all the entries and the compression time does not behave any

close to linear in n , the multiplication time is about linear in n and it might be still useful for some practical problems [10]. Here we propose an entirely different approach.

Let A denote an allowed block of size $m \times n$. If $\text{rank } A = r$ then we can obtain skeletons using any r rows and columns for which the intersection block is nonsingular. The problem is that A is of rank r only up to some small perturbation. We rely on the following (nontrivial) result.

Theorem 4. [8, 9] *Let A and F are $m \times n$, $\text{rank}(A+F) \leq r$, and $\|F\|_2 \leq \varepsilon$. Then there exists a cross in A with columns C and rows R , for which, for some $r \times r$ matrix G ,*

$$\|A - CGR\|_2 \leq \varepsilon \left(1 + \left(\sqrt{t(r, m)} + \sqrt{t(r, n)} \right)^2 \right)$$

with

$$t(r, n) \equiv \max_U \min_{P \in \mathcal{M}(U)} \sigma_{\min}^{-1}(P),$$

where U means any r columns of a unitary matrix of order n , $\mathcal{M}(U)$ is the set of all $r \times r$ submatrices in U , and σ_{\min} designates the minimal singular value.

It was proved also [9] that

$$t(r, n) \leq t_V(r, n) \equiv \sqrt{r(n-r) + 1},$$

though we believe that $t(r, n) \leq \sqrt{n}$ (not proved so far). The idea behind the proof of Theorem 4 was the splitting of the singular value decomposition of $A = U_1 \Sigma_1 V_1 + U_2 \Sigma_2 V_2$, where the first term corresponds to r senior singular triplets, and, then, choosing those submatrices in U_1 and V_1 that have the reciprocal to the minimal singular value majorized by $t(r, m)$ and $t(r, n)$, respectively. The rows chosen in U_1 and columns in V_1 determine the cross at issue.

The choice of the above submatrices in U_1 and V_1 can be performed constructively: it is sufficient to find the submatrices of maximal determinant in modulus (we call this quantity a *volume*) [9].

Unfortunately, the singular value decomposition of A is too heavy a tool to be practical for our purposes. Instead, we would fall back to the concept of maximal volumes.

To give more motivation, recall the role of maximal volumes in the interpolation theory. Assume that Ω is a compact domain in the m -dimensional

space. Assume that $f(x)$ is to be interpolated by $\phi(x) = \sum_{l=1}^k \alpha_l \phi_l(x)$ in the nodes $x_1, \dots, x_k \in \Omega$. Then, as is readily verified,

$$\phi(x) = \sum_{l=1}^k f(x_l) \frac{M_l}{M},$$

where $M = \det\{\phi_i(x_j)\}_{i,j=1}^k$ and M_l is obtained from M by replacing the l th column by $[\phi_1(x_l), \dots, \phi_k(x_l)]^T$. If we are allowed to choose the nodes, which is a reasonable choice? A classical result (see [6]) is that a good idea is to maximize $|\det M|$ over $x_1, \dots, x_k \in \Omega$.

Theorem 5. *Let M maximize $|\det M|$ and $\|f\|_C \equiv \max_{x \in \Omega} |f(x)|$. Then $\|f - \phi\|_C \leq (1 + k) E_{\text{best}}$, where $E_{\text{best}} = \inf_{\beta_1, \dots, \beta_k} \|f - \sum_{l=1}^k \beta_l \phi_l\|_C$.*

Luckily, this theorem is as profound as elementary allowing for a one-line proof:

$$|f - \phi| \leq |f - \phi_{\text{best}}| + |\phi - \phi_{\text{best}}| \leq E_{\text{best}} + \sum_{l=1}^k |f(x_l) - \phi_{\text{best}}(x_l)| \left| \frac{M_l}{M} \right|,$$

where ϕ_{best} corresponds to the optimal choice of β_l on which the best uniform bound is attained. A matrix analogue of Theorem 5 is the following.

Theorem 6. *Let A be $m \times n$ and C consist of the first k columns of A . Assume that $\text{rank } C = k$ and let B be the maximal-volume submatrix in C . Denote by R the rows of A defined by B . Then*

$$\|A - CB^{-1}R\|_2 \leq (1 + t_V(k, m)) \inf_W \|A - CW\|_2. \quad (10)$$

At the same time, if $\mu(A)$ is the maximal entry of A in modulus, then

$$\mu(A - CB^{-1}R) \leq (1 + k) \inf_W \mu(A - CW). \quad (11)$$

Proof. Let $A = [C, A_2]$ and, similarly, $W = [W_1, W_2]$. Then

$$A - CB^{-1}R = (A_2 - CW_2) + CB^{-1}(BW_2 - R).$$

Since B is of maximal volume in C , all the entries of CB^{-1} are not greater than 1 in modulus. Indeed, the postmultiplication of C by any nonsingular

matrix does not change the ratio of volumes for any two submatrices. With no loss of generality, let

$$H \equiv CB^{-1} = \begin{bmatrix} & I & \\ h_{k+11} & \cdots & h_{k+1k} \\ \cdots & \cdots & \cdots \end{bmatrix}.$$

If $|h_{k+l j}| > 1$, then H would have a submatrix of volume greater than 1. This submatrix would be I with row j substituted with row $k+l$. Now, (10) are (11) are evident. \square

The above proof incorporates an algorithm which we use to find maximal-volume submatrices. On input, we have C and some $\gamma \geq 1$. On output, we obtain a submatrix B whose volume is greater than or equal to the maximal volume over γ . With such a B , we modify the estimate (11) to the form

$$\mu(A - CB^{-1}R) \leq (1 + \gamma k) \inf_W \mu(A - CW). \quad (12)$$

Algorithm 3.

- Reduce C to H by column transformations using a complete pivoting.
- Find the maximal in modulus entry below the k th row of H . Let it have indices $k+l, j$. Quit if it does not exceed γ .
- Interchange row j and $k+l$ and eliminate nondiagonal zeroes in the upper part of H using column transformations. Go to one step back.

We are ready to formulate the incomplete cross approximation algorithm. Presently it is a succession of the prescribed number of cross-refinement cycles.

Cross-Refinement Cycle. Given a cross with row indices I and column indices J , find the maximal-volume submatrices in the columns and the rows. Let the former has column indices \hat{J} and the latter has row indices \hat{I} . Then update the cross changing I onto $I \cup \hat{I}$ and J onto $J \cup \hat{J}$.

Thus, we form the cross adaptively adding to it new rows and columns. Using a sufficiently large cross, we provide an accurate skeleton approximation for the whole block. At the same time, we hope that we finish with a cross small enough.

5 Time versus size

An important issue is how we get an initial cross. Sometimes, as in our experiments with the equation (1), it could be almost arbitrary (even of size 1×1 or 2×2). This is not entirely clear yet from the point of theory. However, below we present a practical justification.

Consider more information related to Table 1. On Fig. 1 we can see that the approximation time grows almost linearly in n . This is due to the incomplete cross approximation approach. Note also that we used the conjugate gradients with the circulant preconditioner [4, 5, 24, 25]. We had only 3 iteration to reduce the residual to a factor of 10^{-4} .




Fig 1. $\log(\text{TIME})$ versus $\log(\text{SIZE})$.

The maximal size we tested was $n = 1\,048\,576$. The computed mosaic rank was 159.72 and the corresponding compression factor was 0.03%. We used about 9 Gb of the disc space and 27 283 seconds on the Silicon Graphic workstation in the University of Saarland. Special thanks for the computing facilities go to Prof. S. Rjasanow.

References

- [1] G. Beylkin, R. Coifman and V. Rokhlin, Fast wavelet transform and numerical algorithms. I. *Comm. Pure Appl. Math.* 44: 141–183, 1991.
- [2] A. Brandt, Multilevel computations of integral transforms and particle interactions with oscillatory kernels, *Computer Physics Communications* 65: 24–38, 1991.
- [3] A. Brandt and A. A. Lubrecht, Multilevel matrix multiplication and fast solution of integral equations, *J. Comput. Phys.* 90: 348–370, 1990.
- [4] R. Chan, H. Sun, and W. Ng, Circulant preconditioners for ill-conditioned boundary integral equations from potential equations, to appear.
- [5] T. Chan, An optimal circulant preconditioner for Toeplitz systems, *SIAM J. Sci. Stat. Comput.* 9: 766–771 (1988).
- [6] D. Gaier, *Vorlesungen über Approximation im Complexen*, Birkhäuser Verlag, Basel-Boston-Berlin, 1980.
- [7] S. A. Goreinov, Mosaic-skeleton approximations of matrices generated by asymptotically smooth and oscillatory kernels, *Matrix Methods and Computations*, INM RAS, Moscow, pp. 41–76, 1999. (In Russian.)
- [8] S. A. Goreinov, E. E. Tyrtysnikov, N. L. Zamarashkin, Pseudo-Skeleton Approximations of Matrices, *Reports of the Russian Academy of Sciences*, 343(2): 151–152, 1995. (In Russian.)
- [9] S. A. Goreinov, E. E. Tyrtysnikov, N. L. Zamarashkin, A Theory of Pseudo-Skeleton Approximations, *Linear Algebra Appl.* 261: 1–21, 1997.
- [10] S. A. Goreinov, E. E. Tyrtysnikov, A. Yu. Yeregin, Matrix-Free Iterative Solution Strategies for Large Dense Linear Systems, *Numerical Linear Algebra with Applications*, 4(4): 273–294, 1997.
- [11] L. Greengard and V. Rokhlin, A fast algorithm for particle simulations. *J. Comput. Physics* 73: 325–348, 1987.
- [12] W. Hackbusch and Z. P. Novak, On the fast matrix multiplication in the boundary element method by panel clustering, *Numer. Math.* 54(4): 463–491, 1989.

- [13] A. Harten, *Multiresolution representation and numerical algorithms: a brief review*, ICASE Report 94-59, October 1994.
- [14] M. S. Martynov, Usage of fast multiplication methods for solving integral equation of potential theory, *Matrix Methods and Computations*, INM RAS, Moscow, pp. 77–130, 1999. (In Russian.)
- [15] S. V. Myagchilov and E. E. Tyrtysnikov, A fast matrix-vector multiplier in discrete vortex method, *Russian J. Numer. Anal. Math. Modelling* 7(4): 325–342, 1992.
- [16] K. Nabors, F. T. Korsmeyer, F. T. Leighton, and J. White, *Preconditioned, Adaptive, Multipole-Accelerated Iterative Methods for Three-Dimensional Potential Integral Equations of the First Kind*, Dept. of Electrical Eng. and Computer Science, Massachusetts Institute of Technology, 1994.
- [17] Yu. M. Nechepurenko, *Fast numerically stable algorithms for a wide class of discrete linear transforms*, Preprint 92, OVM RAN, 1985. (In Russian.)
- [18] I. Yu. Nikolsky, Interpolation method for fast approximate multiplication for a matrix generated by a function on a contour, *Matrix Methods and Computations*, INM RAS, Moscow, pp. 131–145, 1999. (In Russian.)
- [19] V. Rokhlin, Rapid solution of integral equations of classical potential theory, *J. Comput. Physics* 60: 187–207, 1985.
- [20] V. Rokhlin, Rapid solution of integral equations of scattering theory in two dimensions, *J. Comput. Physics* 86: 414–439, 1990.
- [21] E. E. Tyrtysnikov, Mosaic ranks and skeletons, in *L. Vulkov et al. (eds.), Lecture Notes in Computer Science 1196: Numerical Analysis and Its Applications. Proceedings of WNAA-96*. Springer-Verlag, 1996. pp. 505–516.
- [22] E. E. Tyrtysnikov, Mosaic–skeleton approximations, *Calcolo*, 33(1-2): 47–57, 1996.
- [23] E. Tyrtysnikov, Methods for fast multiplication and solution of equations, *Matrix Methods and Computations*, INM RAS, Moscow, pp. 4–40, 1999. (In Russian.)

- [24] E. E. Tyrtyshnikov, Optimal and superoptimal circulant preconditioners, *SIAM J. Matrix Anal. Appl.* 12(2): 459–473 (1992).
- [25] E. Tyrtyshnikov, *A Brief Introduction to Numerical Analysis*, Birkhauser, Boston, 1997.
- [26] V. V. Voevodin, On an order-reduction method for matrices arising in the solution of integral equations, *Numerical Analysis on FORTRAN*, Moscow State University Press, 21–26, 1979. (In Russian.)