



A fast direct solver for boundary integral equations in two dimensions

P.G. Martinsson *, V. Rokhlin

Department of Mathematics, Yale University, 10 Hillhouse Avenue, New Haven, CT 06511, USA

Received 7 May 2004; accepted 5 October 2004
Available online 15 December 2004

Abstract

We describe an algorithm for the direct solution of systems of linear algebraic equations associated with the discretization of boundary integral equations with non-oscillatory kernels in two dimensions. The algorithm is “fast” in the sense that its asymptotic complexity is $O(n)$, where n is the number of nodes in the discretization. Unlike previous fast techniques based on iterative solvers, the present algorithm directly constructs a compressed factorization of the inverse of the matrix; thus it is suitable for problems involving relatively ill-conditioned matrices, and is particularly efficient in situations involving multiple right hand sides. The performance of the scheme is illustrated with several numerical examples.

© 2004 Elsevier Inc. All rights reserved.

1. Introduction

The boundary value problems of classical potential theory are ubiquitous in engineering and physics. Most such problems can be reduced to boundary integral equations which are, from a mathematical point of view, more tractable than the original differential equations. Although the mathematical benefits of such reformulations were realized and exploited in the 19th century, until recently boundary integral equations were rarely used as numerical tools, since most integral operators upon discretization turn into dense matrices. In the 1980s, the cost of applying dense matrices resulting from potential theory to arbitrary vectors was greatly reduced by the development of “fast” algorithms (Fast Multipole Methods, panel clustering, wavelets, etc.). Combining fast matrix-vector multiplication techniques with iterative schemes for the solution of large-scale systems of linear algebraic equations, it became possible to solve

* Corresponding author. Tel.: +1 203 432 1277.

E-mail address: per-gunnar.martinsson@yale.edu (P.G. Martinsson).

well-conditioned boundary integral equations of potential theory in $O(n)$ operations, where n is the number of unknowns. Today, such combinations are in many environments the fastest and most accurate numerical solution techniques available. Iterative linear solvers have certain drawbacks though; we briefly discuss these below.

- (1) The number of iterations required by an iterative solver is sensitive to the spectral properties of the matrix of the system to be solved; for sufficiently ill-conditioned problems, the number of iterations is proportional to n . Since each iteration (with FMM acceleration) requires $O(n)$ operations, the total operation count is then proportional to n^2 . While this is still better than the $O(n^3)$ estimate associated with direct solvers, in many situations $O(n^2)$ is not acceptable.
- (2) When one needs to solve a collection of problems involving a single matrix and multiple right-hand sides, the CPU time requirements of most iterative algorithms are close to the time required to solve one problem multiplied by the number of problems to be solved. With most direct solvers, the situation is different; once the matrix has been inverted (or factored), applying its inverse to each additional right-hand side is very inexpensive.
- (3) When a collection of linear systems has to be solved whose matrices are in some sense “close” to each other, iterative algorithms derive little (if any) advantage from the closeness of the matrices.
- (4) Most direct schemes for the solution of linear systems are closely related to efficient algorithms for the construction of their Singular Value Decompositions and certain other matrix factorizations (L-U, Q-R, etc.). The simplest such scheme is probably the inverse power method with shifts (see, for example, [6]), which converts *any* algorithm for the solution of a linear system into an algorithm for the determination of a prescribed singular value. Iterative techniques do not provide such a capability, except via the so-called Lanczos schemes, which tend to require a large number of iterations (see, for example, [14]).

The subject of this paper is a numerical technique that is intended to overcome these shortcomings by directly producing a compressed (“data-sparse”) factorization of the inverse of the matrix. When applied to contour integral equations of potential theory whose kernels are non-oscillatory, the asymptotic complexity of the solver is typically $O(n)$. When applied to problems involving oscillatory kernels, the asymptotic complexity deteriorates as the wavenumber increases but the scheme remains viable for objects up to a few hundred wavelengths in size. The factorization technique described in this paper is a multilevel extension of the compression technique described in [3]. The machinery underlying these techniques applies generally to matrices with rank-deficient off-diagonal submatrices; contour integral equations have been chosen by the authors simply as the most straightforward application.

It is not the purpose of this paper to provide an exhaustive survey of the literature on the subject we are addressing. A number of researchers have observed that matrices with rank-deficient off-diagonal blocks admit “fast” factorizations (see [8,9]); others have constructed “fast” algorithms in various environments (see [1,2,4,5,12]) where the operators in question possess rank-deficient off-diagonal blocks, without using this property explicitly. However, we observe that the algorithm of this paper is closely related to the scheme described in [13]. In fact, our algorithm could be viewed as a modification of the algorithm of [13] that replaces “elongated” objects in two or three dimensions with “curves”, extends the class of kernels addressed by [13], and introduces modifications in the scheme of [13] that are necessary for this extension to work.

The paper is organized as follows: In Section 2 we introduce our notation and list certain facts about compression of rank-deficient matrices. In Section 3 we demonstrate that the inverse of a matrix with rank-deficient off-diagonal blocks possesses a data-sparse hierarchical factorization. In Section 4 we present a generic numerical technique for constructing the factorization described in Section 3. In Section 5 we show how the generic numerical technique presented in Section 4 can be improved further when applied to contour integral equations. In Section 6 we illustrate through numerical examples the efficiency of the

technique presented in Section 5 when applied to a number of different kernels and contours. In Section 7 we summarize our findings and discuss possible extensions and generalizations.

2. Preliminaries

2.1. Notation

Throughout the paper, we use upper case letters for matrices and lower case letters for vectors and scalars. The canonical unit vectors in \mathbb{C}^n are denoted by e_j . Given a matrix $X \in \mathbb{C}^{m \times n}$, we let

X^* denote its adjoint (the complex conjugate transpose),

$\sigma_k(X)$ denote its k th singular value,

$\|X\|_2$ denote its L^2 operator norm,

$\|X\|_F$ denote its Frobenius norm, and

$x_j \in \mathbb{C}^{m \times 1}$ denote its j th column.

Given matrices A , B , C and D we let

$$[AB], \quad \begin{bmatrix} A \\ C \end{bmatrix}, \quad \text{and} \quad \begin{bmatrix} A & B \\ C & D \end{bmatrix}, \quad (2.1)$$

denote larger matrices obtained by stringing the blocks A , B , C and D together.

Definition 1. (Permutation vectors) Given a positive integer n , we define

$$\mathbb{J}_n = \text{the set of permutations of the integers } \{1, \dots, n\}. \quad (2.2)$$

Given two integers k and n such that $1 \leq k \leq n$, we define

$$\mathbb{J}_n^k = \text{the set of subsets of size } k \text{ of elements of } \mathbb{J}_n. \quad (2.3)$$

In other words, if $J \in \mathbb{J}_n^k$, then J is a vector of integers

$$J = [j_1, \dots, j_k], \quad (2.4)$$

where $1 \leq j_l \leq n$ and all j_l 's are different.

Definition 2. (Submatrix) When we use the term ‘‘submatrix’’ we do not insist that the submatrix must form a contiguous block. To be precise, we say that $B \in \mathbb{C}^{k \times l}$ is a submatrix of $A \in \mathbb{C}^{m \times n}$, if there exist permutations $I = [i_1, \dots, i_k] \in \mathbb{J}_m^k$ and $J = [j_1, \dots, j_l] \in \mathbb{J}_n^l$ such that

$$b_{pq} = a_{i_p j_q}, \quad \text{for } p = 1, \dots, k, \quad q = 1, \dots, l. \quad (2.5)$$

Definition 3. (Neutered rows and columns) Let A be a matrix consisting of $p \times p$ blocks,

$$A = \begin{bmatrix} A^{(1,1)} & \dots & A^{(1,p)} \\ \vdots & & \vdots \\ A^{(p,1)} & \dots & A^{(p,p)} \end{bmatrix}. \quad (2.6)$$

We refer to the submatrix formed by all blocks on the i th row except the diagonal one, i.e.

$$[A^{(i,1)} \dots A^{(i,i-1)} \quad A^{(i,i+1)} \dots A^{(i,p)}], \quad (2.7)$$

as the i th neutered row of blocks. A neutered column of blocks is defined analogously.

2.2. Compression of matrices

In this section we state a theorem on matrix compression that forms the foundation of the matrix factorization technique presented later in this paper. Roughly speaking, the theorem asserts that given a matrix A of rank k , it is possible to pick k of its columns in such a fashion that they form a well-conditioned basis for the remaining columns. It was first reported in slightly different form in [7].

Theorem 1. *Given an arbitrary matrix $A \in \mathbb{C}^{m \times n}$ and an integer k such that $1 \leq k < \min(m, n)$, there exists a (not necessarily unique) matrix $T \in \mathbb{C}^{k \times (n-k)}$ and a permutation $J = [j_1, \dots, j_n] \in \mathbb{J}_n$ such that*

$$\tilde{A}_2 = \tilde{A}_1 T + E. \quad (2.8)$$

Here, \tilde{A}_1 and \tilde{A}_2 are matrices formed by the columns of A ,

$$\begin{aligned} \tilde{A}_1 &= [a_{j_1}, \dots, a_{j_k}] \in \mathbb{C}^{m \times k}, \\ \tilde{A}_2 &= [a_{j_{k+1}}, \dots, a_{j_n}] \in \mathbb{C}^{m \times (n-k)}, \end{aligned} \quad (2.9)$$

the elements of the matrix $T \in \mathbb{C}^{k \times (n-k)}$ satisfy

$$|T_{ij}| \leq 1, \quad \text{for } 1 \leq i \leq k, \quad 1 \leq j \leq n - k, \quad (2.10)$$

and the matrix $E \in \mathbb{C}^{m \times (n-k)}$ satisfies the inequality

$$\|E\|_2 \leq \sigma_{k+1}(A) \sqrt{1 + k(n-k)}, \quad (2.11)$$

where $\sigma_{k+1}(A)$ is the $(k+1)$ th singular value of A .

Remark 4. (Computational complexity) While Theorem 1 asserts the theoretical existence of a matrix T and a permutation J with certain properties, it does not address the question of how to determine these numerically. In fact, the authors are not aware of any algorithm that finds these objects in polynomial time. However, in [7] an algorithm is presented that finds a matrix T and a permutation J such that all statements of Theorem 1 still hold, except that (2.10) and (2.11) are replaced by the weaker inequalities

$$|T_{ij}| \leq \sqrt{n}, \quad \text{for } 1 \leq i \leq k, \quad 1 \leq j \leq n - k \quad (2.12)$$

and

$$\|E\|_2 \leq \sigma_{k+1}(A) \sqrt{1 + nk(n-k)}. \quad (2.13)$$

When $m \geq n$, the computational complexity of this algorithm is typically $O(mnk)$, the same as for the pivoted QR -factorization. In rare cases, the computational complexity may be somewhat larger but it never exceeds $O(mn^2)$.

Observation 5. (Column compression) When applied to a matrix $A \in \mathbb{C}^{m \times n}$ of rank k , Theorem 1 asserts that there exists a well-conditioned column operation that leaves k of the columns of A unchanged while mapping the remaining $n - k$ columns to zero. More specifically, let us define

$$R = P_J \begin{bmatrix} I & -T \\ 0 & I \end{bmatrix} \in \mathbb{C}^{n \times n}, \quad (2.14)$$

where T and J are defined in Theorem 1 and the permutation matrix P_J is defined by

$$P_J = [e_{j_1}, \dots, e_{j_n}] \in \mathbb{C}^{n \times n}. \quad (2.15)$$

Then

$$AR = [A_{CS}0] \in \mathbb{C}^{m \times n}, \quad (2.16)$$

where the “column skeleton” A_{CS} , is formed by k of the columns of A ;

$$A_{CS} = [a_{j_1}, \dots, a_{j_k}] \in \mathbb{C}^{m \times k}. \quad (2.17)$$

Moreover, by virtue of (2.10) and the identity

$$R^{-1} = \begin{bmatrix} I & T \\ 0 & I \end{bmatrix} P_J^*, \quad (2.18)$$

it is clear that

$$\|R\|_F \leq \sqrt{n + k(n - k)}, \quad \text{and} \quad \|R^{-1}\|_F \leq \sqrt{n + k(n - k)}. \quad (2.19)$$

Observation 6. (Row compression) The argument of Observation 5 can equally well be applied to the rows of a matrix $A \in \mathbb{C}^{m \times n}$ of rank k . Thus, there exists a matrix $L \in \mathbb{R}^{m \times m}$ such that

$$LA = \begin{bmatrix} A_{RS} \\ 0 \end{bmatrix} \in \mathbb{C}^{m \times n}, \quad (2.20)$$

where the “row skeleton” $A_{RS} \in \mathbb{C}^{k \times n}$ is formed by k of the rows of A and

$$\|L\|_F \leq \sqrt{m + k(m - k)} \quad \text{and} \quad \|L^{-1}\|_F \leq \sqrt{m + k(m - k)}. \quad (2.21)$$

3. Analytical apparatus

Consider a $p \times p$ block matrix

$$A = \begin{bmatrix} A^{(11)} & \dots & A^{(1p)} \\ \vdots & & \vdots \\ A^{(p1)} & \dots & A^{(pp)} \end{bmatrix}, \quad (3.1)$$

such that any neutered row or column of blocks is rank-deficient. In this section we derive compressed factorizations of the inverse of such a matrix. Lemmas 2 and 3 provide factorizations for the case $p = 2$. Observation 8 extends the results of Lemma 3 to a general p . Observation 9 introduces hierarchical factorizations that further improve the efficiency.

Lemma 2 below asserts that for a given 2×2 block matrix with rank-deficient off-diagonal blocks, there exist well-conditioned row- and column-operations that (i) introduce zeros in the off-diagonal blocks and (ii) leave the remaining elements in the off-diagonal blocks untouched.

Lemma 2. *Let A be a non-singular matrix*

$$A = \begin{bmatrix} A^{(11)} & A^{(12)} \\ A^{(21)} & A^{(22)} \end{bmatrix} \in \mathbb{C}^{(n+m) \times (n+m)}, \quad (3.2)$$

where $A^{(11)} \in \mathbb{C}^{n \times n}$, $A^{(22)} \in \mathbb{C}^{m \times m}$ and the off-diagonal blocks $A^{(12)} \in \mathbb{C}^{n \times m}$, $A^{(21)} \in \mathbb{C}^{m \times n}$ have rank $k < \min(m, n)$. Then there exist matrices $R, L \in \mathbb{C}^{n \times n}$ such that

$$\begin{bmatrix} L & 0 \\ 0 & I \end{bmatrix} \begin{bmatrix} A^{(11)} & A^{(12)} \\ A^{(21)} & A^{(22)} \end{bmatrix} \begin{bmatrix} R & 0 \\ 0 & I \end{bmatrix} = \begin{bmatrix} X_{11} & X_{12} & A_{RS}^{(12)} \\ X_{21} & X_{22} & 0 \\ A_{CS}^{(21)} & 0 & A^{(22)} \end{bmatrix}. \quad (3.3)$$

Here, the matrix $A_{RS}^{(12)} \in \mathbb{C}^{k \times m}$ consists of k of the rows of $A^{(12)}$ and the matrix $A_{CS}^{(21)} \in \mathbb{C}^{m \times k}$ consists of k of the columns of $A^{(21)}$. Moreover, $X_{11} \in \mathbb{C}^{k \times k}$, $X_{12} \in \mathbb{R}^{k \times (n-k)}$, $X_{21} \in \mathbb{R}^{(n-k) \times k}$, $X_{22} \in \mathbb{R}^{(n-k) \times (n-k)}$, and the matrices R and L satisfy (2.19) and (2.21), respectively.

Proof. Due to Observations 5 and 6, there exist matrices $R, L \in \mathbb{C}^{n \times n}$ such that

$$LA^{(12)} = \begin{bmatrix} A_{RS}^{(12)} \\ 0 \end{bmatrix} \quad \text{and} \quad A^{(21)}R = \begin{bmatrix} A_{CS}^{(21)} & 0 \end{bmatrix}, \quad (3.4)$$

where $A_{RS}^{(12)}$ and $A_{CS}^{(21)}$ are submatrices of $A^{(12)}$ and $A^{(21)}$, respectively. The identity (3.3) now follows by partitioning

$$\begin{aligned} L &= \begin{bmatrix} L_1 \\ L_2 \end{bmatrix}, \quad \text{where } L_1 \in \mathbb{C}^{k \times n}, L_2 \in \mathbb{C}^{(n-k) \times n}, \\ R &= [R_1 \ R_2], \quad \text{where } R_1 \in \mathbb{C}^{n \times k}, R_2 \in \mathbb{C}^{n \times (n-k)}, \end{aligned} \quad (3.5)$$

and setting

$$\begin{aligned} X_{11} &= L_1 A^{(11)} R_1 \in \mathbb{C}^{k \times k}, \\ X_{12} &= L_1 A^{(11)} R_2 \in \mathbb{C}^{k \times (n-k)}, \\ X_{21} &= L_2 A^{(11)} R_1 \in \mathbb{C}^{(n-k) \times k}, \\ X_{22} &= L_2 A^{(11)} R_2 \in \mathbb{C}^{(n-k) \times (n-k)}. \quad \square \end{aligned} \quad (3.6)$$

The following lemma uses the results of Lemma 3 to reduce the problem of factoring the inverse of the matrix A in (3.2) to the problem of factoring the inverse of the smaller matrix \tilde{A} in (3.8).

Lemma 3. Let $A, X_{11}, X_{12}, X_{21}, X_{22}, A_{RS}^{(12)}$ and $A_{CS}^{(21)}$ be as in Lemma 2. Provided that the matrix X_{22} in (3.3) is non-singular, there exist matrices $B \in \mathbb{C}^{n \times k}$, $C \in \mathbb{C}^{k \times n}$ and $D \in \mathbb{C}^{n \times n}$ such that

$$A^{-1} = \begin{bmatrix} B & 0 \\ 0 & I \end{bmatrix} \tilde{A}^{-1} \begin{bmatrix} C & 0 \\ 0 & I \end{bmatrix} + \begin{bmatrix} D & 0 \\ 0 & 0 \end{bmatrix}, \quad (3.7)$$

where

$$\tilde{A} = \begin{bmatrix} \tilde{A}^{(11)} & A_{RS}^{(12)} \\ A_{CS}^{(21)} & A^{(22)} \end{bmatrix} \in \mathbb{C}^{(k+m) \times (k+m)} \quad (3.8)$$

and

$$\tilde{A}^{(11)} = X_{11} - X_{12} X_{22}^{-1} X_{21} \in \mathbb{C}^{k \times k}. \quad (3.9)$$

Proof. We let L_1, L_2, R_1 and R_2 be defined by (3.5). Inverting both sides of Eq. (3.3), we obtain the identity

$$A^{-1} = \begin{bmatrix} R_1 & R_2 & 0 \\ 0 & 0 & I \end{bmatrix} \begin{bmatrix} X_{11} & X_{12} & A_{RS}^{(12)} \\ X_{21} & X_{22} & 0 \\ A_{CS}^{(21)} & 0 & A^{(22)} \end{bmatrix}^{-1} \begin{bmatrix} L_1 & 0 \\ L_2 & 0 \\ 0 & I \end{bmatrix}. \quad (3.10)$$

Since X_{22} is non-singular,

$$\begin{aligned} \begin{bmatrix} X_{11} & X_{12} & A_{RS}^{(12)} \\ X_{21} & X_{22} & 0 \\ A_{CS}^{(21)} & 0 & A^{(22)} \end{bmatrix}^{-1} &= \begin{bmatrix} I_k & 0 \\ -X_{22}^{-1}X_{21} & 0 \\ 0 & I_m \end{bmatrix} \begin{bmatrix} X_{11} - X_{12}X_{22}^{-1}X_{21} & A_{RS}^{(12)} \\ A_{CS}^{(21)} & A^{(22)} \end{bmatrix}^{-1} \\ &= \begin{bmatrix} I_k & -X_{12}X_{22}^{-1} & 0 \\ 0 & 0 & I_m \end{bmatrix} + \begin{bmatrix} 0 & 0 & 0 \\ 0 & X_{22}^{-1} & 0 \\ 0 & 0 & 0 \end{bmatrix}. \end{aligned} \quad (3.11)$$

Now we obtain (3.7) by combining (3.10) and (3.11) and setting

$$\begin{aligned} B &= R_1 - R_2X_{22}^{-1}X_{21} \in \mathbb{C}^{n \times k}, \\ C &= L_1 - X_{12}X_{22}^{-1}L_2 \in \mathbb{C}^{k \times n}, \\ D &= R_2X_{22}^{-1}L_2 \in \mathbb{C}^{n \times n}. \quad \square \end{aligned} \quad (3.12)$$

Remark 7. (Symmetric factorizations) It is possible to force the factorization (3.7) to be symmetric in the sense that $R = L^*$ (which does not imply that $C = B^*$ unless A itself is Hermitian). To this end, we define L and J_R as the matrix and index vector that compress the rows of the matrix $[A^{(12)}A^{(21)*}] \in \mathbb{R}^{n \times 2m}$ (rather than the rows of $A^{(12)}$ alone), and set $R = L^*$ and $J_C = J_R$. This modification typically results in a poorer compression ratio but may dramatically improve the conditioning of the transformation matrices, as discussed in Section 4.4.

Observation 8. (One-level compression of a block matrix) Consider a matrix

$$A = \begin{bmatrix} A^{(11)} & \dots & A^{(1p)} \\ \vdots & & \vdots \\ A^{(p1)} & \dots & A^{(pp)} \end{bmatrix} \in \mathbb{C}^{pm \times pn}, \quad (3.13)$$

where $A^{(ij)} \in \mathbb{C}^{n \times n}$ for $i, j = 1, \dots, p$. We assume that any neutered row or column of blocks has rank at most k . Through p applications of Lemma 3, it is possible to reduce the problem of inverting A to the problem of inverting the smaller matrix

$$\tilde{A} = \begin{bmatrix} \tilde{A}^{(11)} & \dots & \tilde{A}^{(1p)} \\ \vdots & & \vdots \\ \tilde{A}^{(p1)} & \dots & \tilde{A}^{(pp)} \end{bmatrix} \in \mathbb{C}^{pk \times pk}, \quad (3.14)$$

where $\tilde{A}^{(ij)} \in \mathbb{C}^{k \times k}$ for $i, j = 1, \dots, p$, and $\tilde{A}^{(ij)}$ is a submatrix of $A^{(ij)}$ whenever $i \neq j$.

More specifically, applying Lemma 3 to each of the p diagonal blocks of A , we obtain the factorization

$$A^{-1} = \begin{bmatrix} B_1 & 0 & \cdots & 0 \\ 0 & B_2 & & 0 \\ \vdots & & & \vdots \\ 0 & 0 & \cdots & B_p \end{bmatrix} \tilde{A}^{-1} \begin{bmatrix} C_1 & 0 & \cdots & 0 \\ 0 & C_2 & & 0 \\ \vdots & & & \vdots \\ 0 & 0 & \cdots & C_p \end{bmatrix} + \begin{bmatrix} D_1 & 0 & \cdots & 0 \\ 0 & D_2 & & 0 \\ \vdots & & & \vdots \\ 0 & 0 & \cdots & D_p \end{bmatrix}, \tag{3.15}$$

where $B_i \in \mathbb{C}^{n \times k}$, $C_i \in \mathbb{C}^{k \times n}$ and $D_i \in \mathbb{C}^{n \times n}$, for $i = 1, \dots, p$.

The single-level matrix compression is illustrated graphically in Fig. 1.

Observation 9. (Hierarchical compression of a block matrix) Observation 8 reduces the problem of inversion of a block matrix with rank-deficient neutered rows and columns to the problem of inversion of a block matrix with smaller blocks. If by clustering these smaller blocks, we can create a matrix with off-diagonal rank-deficiencies, then the process can be repeated recursively to further improve the compression.

More specifically, let us change notation so that the objects labeled A , \tilde{A} and k in Observation 8 are now labeled $A^{(1)}$, $\tilde{A}^{(1)}$ and k_1 , respectively. Eq. (3.15) then reads

$$(A^{(1)})^{-1} = B^{(1)} \left(\tilde{A}^{(1)} \right)^{-1} C^{(1)} + D^{(1)}, \tag{3.16}$$

where $B^{(1)}$, $C^{(1)}$, $D^{(1)}$ are block diagonal matrices whose p diagonal blocks are of sizes $n \times k_1$, $k_1 \times n$, $n \times n$, respectively. We then cluster the blocks of the matrix $\tilde{A}^{(1)}$ to form a matrix $A^{(2)}$ with $(p/2) \times (p/2)$ blocks of size $2k_1 \times 2k_1$ and apply the factorization (3.16) to it, thus obtaining a telescoped factorization

$$(A^{(1)})^{-1} = B^{(1)} \left[B^{(2)} \left(\tilde{A}^{(2)} \right)^{-1} C^{(2)} + D^{(2)} \right] C^{(1)} + D^{(1)}. \tag{3.17}$$

Here, $A^{(2)}$, $B^{(2)}$, $C^{(2)}$, $D^{(2)}$ are all block matrices with $(p/2) \times (p/2)$ blocks. Letting k_2 denote the rank of the neutered rows and columns of $A^{(2)}$, the blocks of $\tilde{A}^{(2)}$ have size $k_2 \times k_2$, while $B^{(2)}$, $C^{(2)}$, $D^{(2)}$ are diagonal block matrices with diagonal blocks of sizes $2k_1 \times k_2$, $k_2 \times 2k_1$ and $2k_1 \times 2k_1$, respectively. This process can be continued until no further clustering is advantageous.

The multi-level matrix compression is illustrated graphically in Fig. 2.

Remark 10. (Adjoint of the inverse) Obviously, the factorizations (3.15) and (3.17) provide a mechanism for the accelerated application of both A^{-1} and $[A^{-1}]^*$.

Remark 11. (Block sizes) In Observations 8 and 9, it was assumed that all blocks within one of the matrices A , \tilde{A} , $A^{(1)}$, $A^{(2)}$, \dots , have the same size. This assumption was made for notational convenience only and is in no way essential to the results.

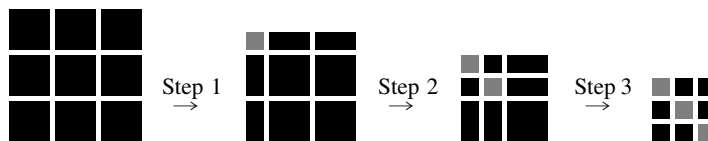


Fig. 1. A 3×3 matrix $[A^{(ij)}]_{i,j=1}^3$ is compressed in three steps, cf. Observation 8. In step $j = 1, 2, 3$, the single-block compression of Lemma 3 is applied to compress the interaction between $A^{(ij)}$ and the rest of the matrix. *Black* blocks represents entries that have not been changed beyond row and column permutations and *gray* represents entries that have been updated but are not (necessarily) zero.

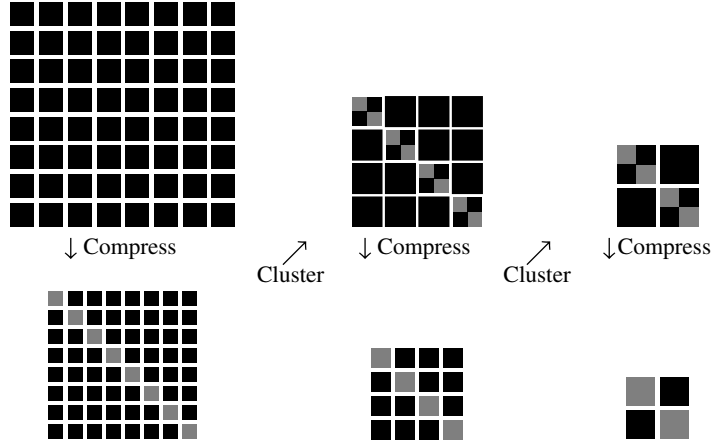


Fig. 2. An 8×8 block matrix is compressed through a three-level compression scheme in the vein of Observation 9. The gray scale coding is the same as in Fig. 1.

4. An algorithm for the computation of a compressed inverse

In Section 3 we demonstrate the existence of a compact factorization of the inverse of any block matrix whose neutered rows and columns of blocks are rank-deficient. In this section, we describe a numerical scheme for the construction of such factorizations, and estimate its efficiency.

Remark 12. The inversion scheme presented in this section is fairly generic, depending only on the ranks of off-diagonal blocks of the matrix to be inverted. In situations where the structure of the matrix is known, further improvements are possible. For instance, when applied to a dense $n \times n$ matrix resulting from the discretization of a contour integral operator, the generic algorithm of this section requires $O(n^2)$ arithmetic operations to construct its inverse, while the customized technique presented in Section 5 requires $O(n \log^2 n)$ operations or less, depending on the integral operator.

4.1. Single block compression

Lemmas 2 and 3 assert that the inverse of a 2×2 block matrix of the form (3.2) can be factored in the compressed form (3.7). The quantities $\tilde{A}_{(11)}$, R , L , $A_{RS}^{(12)}$ and $A_{CS}^{(21)}$ that appear in (3.7) can be determined by taking the following steps:

- (1) Determine a matrix $L \in \mathbb{C}^{n \times n}$ and a permutation $J_R \in \mathbb{J}_n^k$ such that

$$LA^{(12)} = \begin{bmatrix} A_{RS}^{(12)} \\ 0 \end{bmatrix},$$

where $A_{RS}^{(12)}$ is formed by the k rows of $A^{(12)}$ specified by J_R , as described in Observation 6.

- (2) Determine a matrix $R \in \mathbb{C}^{n \times n}$ and a permutation $J_C \in \mathbb{J}_n^k$ such that

$$A^{(21)}R = \begin{bmatrix} A_{CS}^{(21)} & 0 \end{bmatrix},$$

where $A_{CS}^{(21)}$ is formed by the columns of $A^{(21)}$ specified by J_C , as described in Observation 5.

- (3) Partition R and L as specified in (3.5) and form the blocks X_{ij} as in (3.6).
 (4) Compute $\tilde{A}_{(11)}$, B , C and D using the formulas (3.9) and (3.12).

Steps (1) and (2) require $O(mnk)$ floating point operations while steps (3) and (4) require $O(n^3)$ operations. The total cost is thus $O(mnk + n^3)$.

4.2. Single-level compression

Let A denote a matrix consisting of $p \times p$ blocks, each of size $n \times n$, in which every neutered row or column has rank k such that $k < n$. Observation 8 states that such a matrix can be factored in the sparse form (3.15). This factorization contains the entities B_i , C_i , D_i , $\tilde{A}_{(ij)}$ for $i, j = 1, \dots, p$, which can be computed through p applications of the single-block compression technique of Section 4.1 – one application for each diagonal block. Each one of the p steps requires $O(pkn^2 + n^3)$ floating point operations resulting in a total computational cost of $O(p^2kn^2 + pn^3)$.

Remark 13. The off-diagonal blocks of the compressed matrix \tilde{A} are never explicitly computed. Instead, the block $\tilde{A}^{(ij)} \in \mathbb{C}^{k \times k}$ is specified by giving the index vectors $J_R^{(i)}, J_C^{(j)} \in \mathbb{J}_n^k$ that define the rows and columns of $A^{(ij)} \in \mathbb{C}^{k \times k}$, whose intersections form $\tilde{A}^{(ij)}$. (Here $J_R^{(i)}$ is the index vector obtained when compressing the i th row of blocks and $J_C^{(j)}$ is the index vector obtained when compressing the j th column of blocks.)

4.3. Multi-level compression

The single-level technique compresses a block matrix A to form another block matrix \tilde{A} with smaller blocks. Now, if by clustering blocks, we can create rank-deficiencies in the neutered rows and columns of \tilde{A} , then the single-level technique can be applied recursively. The algorithmic implementation entirely follows the description in Observation 9.

When estimating the computational cost for the multi-level technique we use $r = 1, \dots, R$ as an index for the levels (with $r = 1$ being the finest level), we let p_r denote the number of blocks on level r , n_r the average block size and k_r the average rank. The cost for step r is then

$$t_r \sim k_r p_r^2 n_r^2 + p_r n_r^3. \quad (4.1)$$

We assume that $p_r k_r \geq n_r$ so that the second term is dominated by the first. Using that $p_r k_r = p_{r+1} n_{r+1}$, we then find that the total cost for all R steps is

$$T \sim \sum_{r=1}^R t_r \sim \sum_{r=1}^R p_{r+1} p_r n_{r+1} n_r^2. \quad (4.2)$$

At each level, the number of blocks is cut in half, so

$$p_r = \frac{p_1}{2^{r-1}}. \quad (4.3)$$

We let $\gamma_r = n_{r+1}/2n_r$ denote the compression ratio so that

$$n_r = (2\gamma_{r-1}) \cdots (2\gamma_1) n_1. \quad (4.4)$$

Assuming that there exists a constant γ such that $\gamma_r \leq \gamma$, we obtain the bound

$$n_r \leq (2\gamma)^{r-1} n_1. \quad (4.5)$$

Combining (4.2), (4.3) and (4.5), we find that the total cost is

$$T \sim \sum_{r=1}^R \frac{p_1}{2^r} \frac{p_1}{2^{r-1}} (2\gamma)^r n_1 (2\gamma)^{2r-2} n_1^2 \sim p_1^2 n_1^3 \sum_{r=1}^R (2\gamma^3)^r. \tag{4.6}$$

We assume that $\gamma < \sqrt[3]{4} = 0.7937\dots$ so that the sum is bounded by $(1 - 2\gamma^3)^{-1}$. Letting N denote the size of the matrix we find that $N = p_1 n_1$ and thus

$$T \sim N^2 n_1. \tag{4.7}$$

The assumption that (4.5) holds for some $\gamma < 0.7939\dots$ is valid in many environments relating to discretization of contour integral equations. We will return to this point in Section 6.

4.4. Conditioning

All factorizations computed in this section are variations of (3.15). For this formula to be of practical use, the matrices B_i , C_i and D_i must not be excessively large (in say the ℓ^2 operator norm) and the condition number of \tilde{A} has to be similar to that of A . The formulas (3.12) imply that this is true if $\|X_{22}^{-1}\|_2$ is of moderate size (since (2.19) and (2.21) assert that R and L are well-conditioned). Under the assumptions of this section (that the global matrix be non-singular and the off-diagonal blocks have low rank) it is not possible to prove any such bound.

However, in the context of contour integral equations, the problem can largely be avoided by enforcing that the compression be symmetric in the sense of Remark 7. The reason is that the diagonal blocks of the original matrix tend to have the form

$$A^{(11)} = D + E, \tag{4.8}$$

where D is a positive definite Hermitian matrix and E is “small” compared to D in operator norm. Since $R_2 = L_2^*$ when symmetry is enforced, we find that, cf. (3.6),

$$X_{22} = L_2(D + E)L_2^* = (L_2 D^{1/2})(L_2 D^{1/2})^* + L_2 E L_2^*. \tag{4.9}$$

Here, the first term is well-conditioned, and the second has at most a few non-small singular values. Thus, it is very unlikely that the sum of the two matrices should have any small singular values. Furthermore, should such a coincidence happen, the algorithm detects it and avoids the problem by locally re-partitioning the matrix.

4.5. Error estimation

Given a prescribed accuracy ε , the numerical scheme presented in this section solves the equation

$$Au = f \tag{4.10}$$

by constructing an approximation A_ε that satisfies

$$\|A - A_\varepsilon\|_2 \leq \varepsilon \tag{4.11}$$

and is such that the approximate solution $u_\varepsilon = A_\varepsilon^{-1}f$ can be computed fast. The error in u satisfies

$$u - u_\varepsilon = (A^{-1} - A_\varepsilon^{-1})f = A_\varepsilon^{-1}(A_\varepsilon - A)A^{-1}f = A_\varepsilon^{-1}(A_\varepsilon - A)u. \tag{4.12}$$

The relative error is therefore bounded as follows:

$$\frac{\|u - u_\varepsilon\|}{\|u\|} \leq \|A_\varepsilon^{-1}(A_\varepsilon - A)\|_2 \leq \varepsilon \|A_\varepsilon^{-1}\|_2. \tag{4.13}$$

While the algorithm cannot possibly control $\|A_\varepsilon^{-1}\|_2$, this quantity can be computed cheaply using power iteration, see Remark 10. Thus, an assured bound for the relative error can be computed à posteriori.

5. An accelerated algorithm applicable to contour integral equations

The bulk of the computational cost of the matrix compression technique presented in Section 4 consists of the cost of determining index vectors and transformation matrices that compress the neutered rows and columns. When the matrix under consideration is a discrete approximation of a contour integral operator, it is possible to determine these quantities through an entirely local operation whose cost only depends on the size of the diagonal block to be compressed (i.e., not on the size of the rest of the matrix). This is possible since the column and row operations employed in the present matrix compression technique do not update the elements of the off-diagonal blocks, as discussed in Remark 13.

This section is structured as follows: In Section 5.1 we describe a single-block compression technique for the boundary integral equations associated with Laplace’s equation in two dimensions that is faster than the generic single-block technique of Section 4.1. In Section 5.2 we describe single and multi-level techniques for contour integral equations obtained by repeated application of the single-block compression technique of Section 5.1. Section 5.3 discusses generalizations of the technique to other equations of potential theory.

Remark 14. (Numerically rank-deficient matrices) In this section, we say that a matrix has rank k provided that it has only k singular values that are larger than some preset accuracy. In other words, we do not distinguish between what is sometimes called “numerical rank” and actual rank.

5.1. Single-block compression

The following observation summarizes the principle finding of this section:

Observation 15. Let the matrix A in (3.2) represent the discretization of the integral operator

$$\int_{\Gamma} K(x, y)u(y)ds(y), \quad \text{for } x \in \Gamma, \quad (5.1)$$

where $\Gamma = \Gamma_1 + \Gamma_2$ is a contour (Fig. 3 shows one example), the block structure of A corresponds to the partitioning of Γ (so that, e.g., $A^{(12)}$ represents evaluation on Γ_1 of the potential generated by a charge distribution on Γ_2), and K is the kernel of a single and/or double layer potential for the Laplace operator. Then under very mild assumptions on the contour Γ , the factorization (3.3) can be computed using $O(n^3)$ floating point operations, where n is the number of points used in the discretization of Γ_1 .

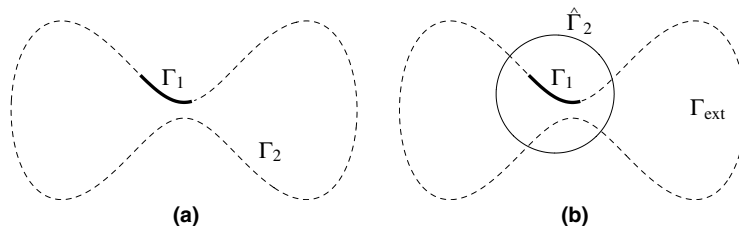


Fig. 3. A contour Γ . In figure (a), the partitioning $\Gamma = \Gamma_1 + \Gamma_2$ is shown with Γ_1 drawn with a bold line. In figure (b) the contour $\hat{\Gamma}_2$ is drawn with a thin solid line and Γ_{ext} with a dashed line.

The idea behind the construction alluded to in Observation 15 is simple: Instead of compressing the interaction between Γ_1 and Γ_2 , it is sufficient to compress the interaction between Γ_1 and a small contour $\hat{\Gamma}_2$, formed by the union of an artificial circular contour enclosing Γ_1 and the part of Γ_2 that is inside this circle (as shown in Fig. 3(b)). The reason is that by virtue of Green’s theorem, any potential field generated by charges on Γ_2 can equally well be generated by charges on $\hat{\Gamma}_2$. Finally we note that if Γ_1 is discretized using n nodes, then typically $\hat{\Gamma}_2$ can be discretized using $O(n)$ nodes, yielding a total cost for the procedure of $O(n^3)$.

The remainder of this subsection is devoted to substantiating Observation 15. We start by introducing some notation; let Γ_{circ} denote the circle in Fig. 3(b) and let Γ_{ext} denote the part of Γ_2 outside of Γ_{circ} . Furthermore, let $\mathbf{S}_{\Gamma_2 \rightarrow \Gamma_1}$ denote the integral operator that evaluates a potential on Γ_1 caused by a charge distribution on Γ_2 . In other words, $\mathbf{S}_{\Gamma_2 \rightarrow \Gamma_1}$ acts on a charge distribution u as follows:

$$[\mathbf{S}_{\Gamma_2 \rightarrow \Gamma_1} u](x) = \int_{\Gamma_2} K(x, y) u(y) ds(y), \quad \text{for } x \in \Gamma_1. \quad (5.2)$$

Observation 15 rests on the following claim:

Lemma 4. *Let $H \in \mathbb{C}^{n \times n'}$ denote the matrix discretizing $\mathbf{S}_{\Gamma_{\text{circ}} \rightarrow \Gamma_1}$, and let the index vector $J_R \in \mathbb{J}_n^k$ and the transformation matrix L be such that they compress H in the sense of Observation 6. Then J_R and L also compress the matrix $B \in \mathbb{C}^{n \times m}$ that approximates the operator $\mathbf{S}_{\Gamma_{\text{ext}} \rightarrow \Gamma_1}$.*

Sketch of proof. It is sufficient to prove that there exists a matrix $W \in \mathbb{C}^{n' \times m}$ with moderate l^2 operator norm such that

$$B = HW. \quad (5.3)$$

(The matrix W is the matrix that maps a charge distribution on Γ_{ext} to an equivalent charge distribution on Γ_{circ} .) Now, Eq. (5.3) is the discrete approximation of the operator relation

$$\mathbf{S}_{\Gamma_{\text{ext}} \rightarrow \Gamma_1} = \mathbf{S}_{\Gamma_{\text{circ}} \rightarrow \Gamma_1} \left[(\mathbf{S}_{\Gamma_{\text{circ}} \rightarrow \Gamma_{\text{circ}}})^{-1} \mathbf{S}_{\Gamma_{\text{ext}} \rightarrow \Gamma_{\text{circ}}} \right]. \quad (5.4)$$

The matrix W in (5.3) corresponds to the operator in square brackets in (5.4). That this operator is bounded is a consequence of Green’s theorem. \square

5.2. Single- and multi-level compression

The generic single- and multi-level compression techniques of Sections 4.2 and 4.3 were obtained by repeated application of the single-block technique described in Section 4.1. Single- and multi-level techniques for contour integral equations are analogously obtained by repeated application of the single-block technique of Section 5.1.

It remains to estimate the computational cost of the accelerated compression technique. By virtue of Observation 15, the cost for a single level compression at level $r = 1, \dots, R$ is now

$$t_r \sim p_r n_r^3, \quad (5.5)$$

where p_r denotes the number of clusters on level r and n_r is the (average) cluster size (the corresponding estimate for the non-accelerated compression technique is given in (4.1)). We next use that n_r is proportional to the rank of interaction k_r between a cluster on level r and the rest of the world (to be precise, $n_r = 2k_{r-1} \leq 2k_r$). For the integral equations associated with Laplace’s equation, k_r is for a typical contour bounded by the logarithm of the number of points in the original discretization of a cluster on level r so that

$$k_r \leq C \log(2^{r-1} n_1) \leq C'(r + \log n_1), \quad (5.6)$$

where n_1 is the number of points inside a cluster on the finest level, and C and C' are numbers that do not depend on r . Using that $p_r = p_1/2^{r-1}$, we then find that the total cost for all R steps is given by

$$T \sim \sum_{r=1}^R p_r n_r^3 \sim \sum_{r=1}^R p_r k_r^3 \sim \sum_{r=1}^R \frac{p_1}{2^r} (r + \log n_1)^3 \sim p_1 (\log n_1)^3. \quad (5.7)$$

On the finest level, we partition the contour so that there are a small, fixed, number of points in each cluster. (In practise, using n_1 between 30 and 50, depending on the desired accuracy, works well.) We then find that $p_1 \sim N$, where N denotes the total number of points in the discretization. Thus,

$$T \sim N. \quad (5.8)$$

Remark 16. The single-block compression technique described in Observation 15 requires the algorithm to determine which of the nodes of Γ_2 lie inside the artificial circle Γ_{circ} . If this search would be done by brute force, the computational cost for a single level solve would include a term $p_r^2 n_r^2$, cf. (5.5). Even though the constant in front of this term is small, it would dominate the computation for large problems (in our implementation, this would happen for $N \geq 25000$). One solution to this problem is to perform the search via a hierarchical search tree; the estimate (5.5) then remains valid.

5.3. Generalizations

The technique presented in Section 5.1 for Laplace's equation is readily applicable to other equations of classical potential theory; Helmholtz, Yukawa, Shrödinger, Maxwell, Stokes, elasticity, etc. The only complication occurs when working with equations that may have resonances. In such cases, it is possible that the operator of self-interaction for the artificial circle (the operator $\mathcal{S}_{\Gamma_{\text{circ}} \rightarrow \Gamma_{\text{circ}}}$ in (5.4)) has a non-trivial nullspace. This complication can be rectified by letting the artificial charges on Γ_{circ} consist of both monopoles and dipoles. Alternatively, it is possible to consider only one type of charges but placing them on two concentric circles instead of a single one.

When applied to oscillatory problems such as Helmholtz' and Maxwell's equations, the efficiency of the technique deteriorates when the wave number increases since then the compression rate deteriorates as the blocks grow larger (in other words, the assumption (4.5) no longer holds). In practice, it appears that the method experiences very few problems for objects smaller than about 50 wavelengths. After that, the computational complexity increases superlinearly with the problem size although the technique remains viable for equations set on contours a few hundred wavelengths in size. This effect will be illustrated in the numerical examples in Section 6.1.

Finally we remark that the scheme has $O(n \log^k n)$ complexity when applied to integral equations defined on one-dimensional curves in *any* dimension. The fact that we have so far only discussed equations embedded in two space dimensions is simply that contour integral equations associated with boundary value problems in two dimensions is the most common source of such equations.

6. Numerical examples

In this section we present the results of a number of numerical experiments performed to assess the efficiency of the numerical scheme presented in Sections 4 and 5. In every experiment, we compute a compressed factorization of the inverse of the matrix resulting from Nyström discretization of one of the following three integral equations:

$$\pm \frac{1}{2}u(x) + \frac{1}{2\pi} \int_{\Gamma} [n(y) \cdot \nabla_y \log |x - y|] u(y) ds(y) = f(x), \quad x \in \Gamma, \quad (6.1)$$

$$\int_{\Gamma} [\log |x - y|] u(y) ds(y) = f(x), \quad x \in \Gamma, \quad (6.2)$$

$$\mp 2iu(x) + \int_{\Gamma} [(n(y) \cdot \nabla_y + ik)H_0(k|x - y|)] u(y) ds(y) = f(x), \quad x \in \Gamma, \quad (6.3)$$

where $n(y)$ is the outward pointing unit normal of Γ at y and $H_0(x) = J_0(x) + iY_0(x)$ is the Hankel function of zeroth order. Eqs. (6.1) and (6.2) are the double and single layer equations associated with Laplace Dirichlet problems, and (6.3) is an equation associated with the Helmholtz Dirichlet problem with wave number k . In Eqs. (6.1) and (6.3), the top sign in front of the first term refers to exterior problems and the lower sign refers to interior problems.

The kernel in (6.1) is smooth and the equation was discretized using the trapezoidal rule (which is exponentially convergent on a smooth contour). The Eqs. (6.2) and (6.3) involve log-singular kernels that were discretized using the modified trapezoidal quadrature rules of [10] of orders 6 and 10, respectively. The algorithm was implemented in Fortran 77 and the experiments were run on a 2.8 GHz Pentium 4 desktop with 512Mb of RAM memory.

When presenting the numerical results, we use the following notation:

R	the number of levels in the multi-level solver,
N_{start}	the size of the discrete problem at the start,
N_{final}	the size of the compressed problem,
t_{tot}	the total CPU time (in seconds),
t_{solve}	the CPU time required to apply the factorized inverse (in seconds),
c_{top}	the condition number of the compressed matrix,
σ_{min}	the smallest singular value of the original matrix,
M	the amount of memory used (in MB),
E_{actual}	the relative error in u , $E_{\text{actual}} = \ u_\varepsilon - u\ /\ u\ $,
E_{res}	the relative residual error, $E_{\text{res}} = \ Au_\varepsilon - f\ /\ f\ $,

In each experiment, the right hand side f was the Dirichlet data corresponding to a potential field generated by a few randomly placed point charges. Since the exact potential field was known, we could compare the potential field generated by the numerical solution to the exact one. We did this at J random points on a circle enclosing Γ and separated from Γ by one quarter of its radius. Letting $\{v^{(j)}\}_{j=1}^J$ denote the exact potential and $\{v_\varepsilon^{(j)}\}_{j=1}^J$ denote the potential generated by u_ε , we define the relative l^2 -norm error in the potential as $E_{\text{pot}} = \|v - v_\varepsilon\|/\|v\|$.

6.1. Example: a smooth contour

In this subsection we present results pertaining to the smooth contour shown in Fig. 4. The contour was discretized using between 800 and 102 400 points and the integral equations associated with exterior Dirichlet problems were solved. Tables 1–3 present the results for the kernels (6.1)–(6.3), respectively. As a reference, we give in Table 4 the timings for highly optimized implementations of the LU-factorization, direct matrix-vector multiplication and FMM-accelerated matrix-vector multiplication.

For the two Laplace problems considered, we see that both the computational cost and the memory requirement scale more or less linearly with the problem size, as expected. We recall that this expectation was based on the postulate that for Laplace problems, the interaction rank between adjacent clusters depend only very weakly (logarithmically) on their size. Fig. 5 illustrates this point; it shows that after two

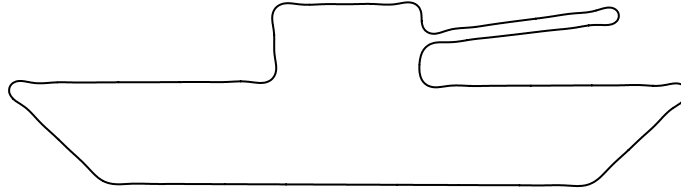


Fig. 4. A smooth contour. The length of the contour is roughly 5.1 and its horizontal width is 2.

Table 1

Computational results for the double layer potential (6.1) associated with an exterior Laplace Dirichlet problem on the contour shown in Fig. 4

N_{start}	N_{final}	t_{tot}	t_{solve}	E_{actual}	E_{res}	E_{pot}	c_{top}	σ_{min}	M
400	301	5.3e-01	2.9e-03	2.3e-10	4.7e-10	3.0e-06	4.3e+00	1.3e-02	4.2e+00
800	351	9.6e-01	4.1e-03	2.5e-10	2.2e-10	6.3e-10	9.1e+00	1.2e-02	6.5e+00
1600	391	1.6e+00	6.3e-03	1.4e-10	1.3e-10	1.6e-10	1.6e+01	1.2e-02	9.2e+00
3200	391	1.8e+00	8.5e-03	–	6.6e-11	3.7e-10	3.2e+01	1.2e-02	1.1e+01
6400	391	2.2e+00	1.2e-02	–	5.9e-11	8.9e-11	7.7e+01	1.2e-02	1.4e+01
12 800	390	2.6e+00	1.9e-02	–	3.6e-11	5.9e-11	1.4e+02	1.2e-02	2.1e+01
25 600	391	3.9e+00	3.4e-02	–	2.7e-11	4.7e-10	2.1e+02	–	3.5e+01
51 200	393	6.5e+00	6.5e-02	–	2.5e-11	5.3e-11	2.0e+02	–	6.3e+01
102 400	402	1.3e+01	1.2e-01	–	2.0e-11	–	1.1e+03	–	1.2e+02

Table 2

Computational results for the single layer potential (6.2) associated with an exterior Laplace Dirichlet problem on the contour shown in Fig. 4

N_{start}	N_{final}	t_{tot}	t_{solve}	E_{actual}	E_{res}	E_{pot}	c_{top}	σ_{min}	M
400	253	4.1e-01	1.9e-03	4.6e-09	2.7e-09	1.6e-04	2.2e+01	3.5e-02	3.1e+00
800	306	8.2e-01	3.3e-03	7.5e-09	9.9e-09	2.4e-06	1.6e+02	2.9e-04	5.4e+00
1600	353	1.6e+00	6.2e-03	4.9e-09	6.3e-09	1.6e-09	1.5e+02	1.4e-04	8.6e+00
3200	369	2.3e+00	9.7e-03	–	2.5e-07	1.2e-10	2.1e+02	4.2e-05	1.2e+01
6400	379	3.2e+00	1.6e-02	–	1.3e-08	6.8e-12	2.6e+02	2.1e-05	1.8e+01
12 800	395	4.8e+00	2.7e-02	–	1.7e-08	3.4e-12	2.8e+02	2.7e-06	2.9e+01
25 600	409	7.7e+00	4.8e-02	–	3.6e-08	1.4e-11	3.5e+02	2.7e-07	5.0e+01
51 200	419	1.4e+01	9.0e-02	–	2.7e-07	–	3.7e+02	3.5e-07	9.1e+01
102 400	429	3.6e+01	1.7e-01	–	1.6e-08	–	5.2e+02	–	1.7e+02

Table 3

Computational results for the kernel (6.3) associated with an exterior Helmholtz Dirichlet problem on the contour shown in Fig. 4

k	N_{start}	N_{final}	t_{tot}	t_{solve}	E_{actual}	E_{res}	E_{pot}	c_{top}	σ_{min}	M
21	800	435	1.5e+01	3.3e-02	2.7e-07	9.7e-08	7.1e-07	4.1e+03	6.5e-01	1.3e+01
40	1600	550	3.0e+01	6.7e-02	1.6e-07	6.2e-08	4.0e-08	6.1e+03	8.0e-01	2.5e+01
79	3200	683	5.3e+01	1.2e-01	–	5.3e-08	3.8e-08	2.1e+04	3.4e-01	4.5e+01
158	6400	870	9.2e+01	2.0e-01	–	3.9e-08	2.9e-08	4.0e+04	3.4e-01	8.2e+01
316	12 800	1179	1.8e+02	3.9e-01	–	2.3e-08	2.0e-08	4.2e+04	3.4e-01	1.6e+02
632	25 600	1753	4.3e+02	7.5e+00	–	1.7e-08	1.4e-08	9.0e+04	3.3e-01	3.5e+02
1264	51 200	2864	(1.5e+03)	(2.3e+02)	–	9.5e-09	–	–	–	8.3e+02

The Helmholtz parameter was chosen to keep the number of discretization points per wavelength constant at roughly 45 points per wavelength (resulting in a quadrature error about 10^{-12}). The times in parenthesis refer to experiments that did not fit in RAM.

Table 4

Timings (in seconds) for highly optimized implementations of the LU-factorization, matrix-vector multiplication and FMM accelerated matrix-vector multiplication

N	400	800	1600	3200	6400	12 800	25 600	51 200	102 400
t_{LU}	2.8e−02	2.0e−01	1.6e+00	1.3e+01	(1.0e+02)	(8.3e+02)	(6.7e+03)	(5.3e+04)	(4.3e+05)
t_{mult}	7.5e−04	2.9e−03	1.2e−02	4.8e−02	(1.9e−02)	(7.7e−01)	(3.1e+00)	(1.2e+01)	(4.9e+01)
t_{FMM}	3.8e−03	8.0e−03	1.6e−02	3.0e−02	6.0e−02	1.2e−01	2.4e−01	4.8e−01	9.6e−01

The FMM was run at a relative accuracy of 10^{-10} with the same kernel as in the Eq. (6.2). The numbers in parenthesis are extrapolated.

rounds of compression, almost the only nodes that have survived are the ones near the border to the neighboring clusters. The figure also illustrates that the algorithm detects the need to keep more nodes in the interior of those clusters that run close to other clusters. (For an example of a situation where the Eq. (6.1) needs to be discretized using a large number of nodes in spite of the fact that the contour is uncomplicated, see [11].)

Since the scheme presented in this paper relies on rank-considerations only, it works for oscillatory problems with low wave numbers but it eventually fails as the wavenumber is increased. Table 5 illustrates this point by showing how the compression ratios deteriorate as the wavenumber k in Eq. (6.3) is increased from 1 to 1200. However, the authors were surprised to find that the method remains viable up to objects about 200 wavelengths across, as indicated in Table 3.

Remark 17. (Comparison with the fast multipole method) From Tables 1 and 4, we see that a single FMM matrix-vector multiply is about 15–20 times faster than a matrix inversion. Thus, if an iterative solver requires less than 15–20 iterations to solve Eq. (6.1), this would beat the direct method for a single solve. However, once the inverse has been computed, it can be applied to additional right hand sides in about one tenth of the time required for a single FMM accelerated matrix-vector multiply.

6.2. A rippled contour that almost self-intersects

In this subsection we present results pertaining to the rippled contour shown in Fig. 6. The contour was discretized using between 800 and 102 400 points and integral equations associated with exterior Dirichlet problems were solved. The number of ripples in the experiments increase with the number of discretization nodes in such a fashion that there are roughly 80 nodes for each wavelength. Tables 6–8 present the results for the kernels (6.1)–(6.3), respectively.

We see that the asymptotic complexity of the algorithm remains essentially the same as for the smooth contour shown in Fig. 4. However, the constants involved are larger since more degrees of freedom are required to resolve the contour at the finest levels.

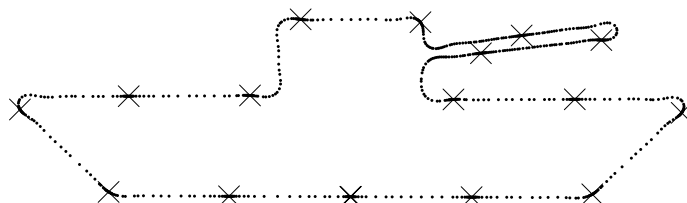


Fig. 5. The points left after two rounds of compression of the contour shown in Fig. 4. The crosses mark the boundary points between adjacent clusters.

Table 5
Deterioration of compression rates for large wavenumber Helmholtz problems

k	γ_1	γ_2	γ_3	γ_4	γ_5	γ_6	γ_7	γ_8	N_{final}	M
1	0.68	0.58	0.54	0.55	0.58	0.64	0.64	0.72	512	167
100	0.72	0.56	0.55	0.56	0.60	0.68	0.72	0.82	777	197
500	0.72	0.58	0.58	0.62	0.68	0.76	0.84	0.91	1522	303

The table shows the compression ratio γ_j , see (4.4), at each of the levels $j = 1, \dots, 8$ for the Helmholtz kernel (6.3) on the smooth contour in Fig. 4, discretized with $N = 25\,600$ points. The three rows correspond to wave numbers $k = 1, 100, 500$. The second to last column shows the number of degrees of freedom left on the finest level and the last column shows the total memory requirement (in MB).

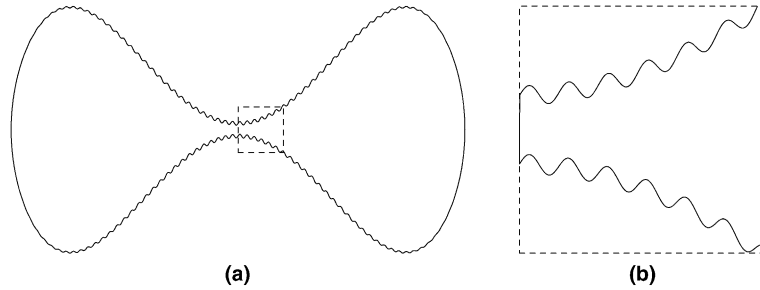


Fig. 6. (a) A ruffled contour. (b) A close-up of the area marked by a dashed rectangle in (a). The horizontal axis of the contour has length 1 and the number of ripples change between the different experiments to keep a constant ratio of 80 discretization nodes per wavelength.

Table 6
Computational results for the double layer potential (6.1) associated with an exterior Laplace Dirichlet problem on the ruffled contour shown in Fig. 6

N_{start}	N_{final}	t_{tot}	t_{solve}	E_{actual}	E_{res}	E_{pot}	c_{top}	σ_{min}	M
400	171	2.3e-01	1.0e-03	1.5e-10	1.1e-10	1.3e-07	7.4e+00	1.1e-01	1.5e+00
800	228	3.5e-01	1.0e-02	1.9e-10	1.3e-10	3.8e-08	9.7e+00	7.6e-02	3.0e+00
1600	306	7.3e-01	5.8e-03	1.3e-10	1.6e-10	5.5e-08	1.6e+01	5.2e-02	6.2e+00
3200	386	2.2e+00	8.5e-03	–	1.4e-10	7.5e-08	3.1e+01	3.9e-02	1.2e+01
6400	460	4.4e+00	1.7e-02	–	7.2e-11	8.2e-08	7.0e+01	3.3e-02	2.1e+01
12 800	536	9.6e+00	3.5e-02	–	5.9e-11	3.7e-08	1.4e+02	2.9e-02	4.0e+01
25 600	597	2.0e+01	7.6e-02	–	2.0e-11	1.4e-09	2.2e+02	–	7.6e+01
51 200	641	4.0e+01	1.5e-01	–	2.9e-11	–	6.2e+02	–	1.5e+02
102 400	688	(1.8e+01)	3.9e-01	–	1.2e-11	–	7.8e+02	–	2.9e+02

Table 7
Computational results for the single layer potential (6.2) associated with an exterior Laplace Dirichlet problem on the ruffled contour shown in Fig. 6

N_{start}	N_{final}	t_{tot}	t_{solve}	E_{actual}	E_{res}	E_{pot}	c_{top}	σ_{min}	M
400	176	2.4e-01	9.2e-04	2.1e-09	1.7e-09	2.4e-05	1.6e+02	5.5e-04	1.6e+00
800	220	3.9e-01	3.8e-03	1.6e-08	3.0e-08	8.0e-06	1.1e+03	1.0e-05	3.1e+00
1600	256	6.9e-01	5.3e-03	5.2e-09	7.0e-09	9.8e-08	2.8e+02	1.6e-05	5.3e+00
3200	286	1.3e+00	7.6e-03	–	7.0e-09	1.6e-08	3.3e+02	1.2e-05	9.1e+00
6400	314	2.5e+00	1.4e-02	–	1.5e-07	2.3e-09	7.5e+02	2.1e-06	1.6e+01
12 800	342	4.6e+00	2.8e-02	–	2.4e-08	1.5e-09	4.7e+02	1.7e-07	2.9e+01
25 600	362	8.8e+00	6.2e-02	–	2.3e-08	2.2e-09	1.1e+03	9.7e-08	5.5e+01
51 200	374	1.7e+01	1.2e-01	–	2.1e-08	–	1.8e+03	3.1e-08	1.1e+02
102 400	386	(8.1e+0)	2.3e-01	–	1.5e-07	–	3.1e+03	–	2.1e+02

Table 8

Computational results for the kernel (6.3) associated with an exterior Helmholtz Dirichlet problem on the rippled contour shown in Fig. 6

k	N_{start}	N_{final}	t_{tot}	t_{solve}	E_{actual}	E_{res}	E_{pot}	c_{top}	σ_{min}	M
7	400	224	2.9e+00	9.0e−03	1.4e−07	6.9e−08	9.4e−07	1.2e+04	7.9e−01	3.2e+00
15	800	320	7.7e+00	1.9e−02	1.6e−07	7.4e−08	1.2e−07	3.9e+03	7.9e−01	8.2e+00
29	1600	470	2.1e+01	4.6e−02	–	6.7e−08	8.1e−08	7.4e+03	7.8e−01	2.0e+01
58	3200	704	6.1e+01	1.1e−01	–	5.2e−08	6.4e−08	1.2e+04	8.0e−01	5.0e+01
115	6400	1122	1.4e+02	2.9e−01	–	4.8e−08	7.5e−08	1.4e+04	8.0e−01	1.3e+02
230	12 800	1900	(4.7e+02)	(2.5e+01)	–	5.5e−08	7.5e−08	8.8e+04	8.0e−01	3.4e+02
461	25 600	3398	–	–	–	–	–	–	–	9.8e+02

The Helmholtz parameter k was chosen to keep the number of discretization points per wavelength constant at roughly 55 points per wavelength (resulting in a quadrature error about 10^{-12}). The times in parenthesis refer to experiments that did not fit in RAM.

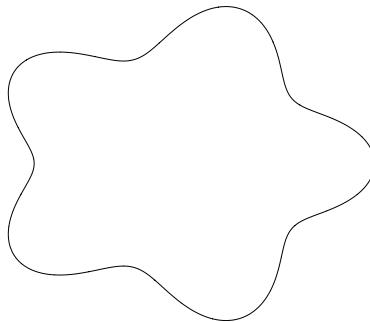


Fig. 7. A contour the shape of a smooth pentagram. Its diameter is 2.5 and its length is roughly 8.3.

6.3. An interior problem close to a resonance

In this section we present results pertaining to *interior* Dirichlet problem on the contour shown in Fig. 7. While interior and exterior Laplace Dirichlet problems are quite similar in nature, the corresponding Helmholtz Dirichlet problems are fundamentally different in that the interior problem possesses resonances while the exterior does not. We will therefore focus exclusively on interior Helmholtz problems.

We present the results of two computational experiments, both relating to the Helmholtz kernel (6.3). In the first experiment, we scan a range of wave numbers k between 99.9 and 100.1. For each wave number, we computed the smallest singular value σ_{min} of the integral operator using the iteration technique described in Section 4.5. The resulting graph of σ_{min} versus k , shown in Fig. 8, clearly indicates the location of each resonance in this interval. The second experiment consists of factoring the inverse of the matrix corresponding to $k = 100.0110276\dots$ for which $\sigma_{\text{min}} = 0.00001366\dots$. The results, shown in Table 9, illustrate that the method does not experience any difficulty in factoring the inverse of a reasonably ill-conditioned matrix. In particular, the table shows that the factorization matrices $B^{(j)}$, $C^{(j)}$ and $D^{(j)}$, see (3.17), are well-conditioned.

6.4. A contour resembling an area integral

The final numerical experiment that we present is included to demonstrate that the efficiency of the factorization scheme deteriorates when it is applied to a curve for which the physical distance between two

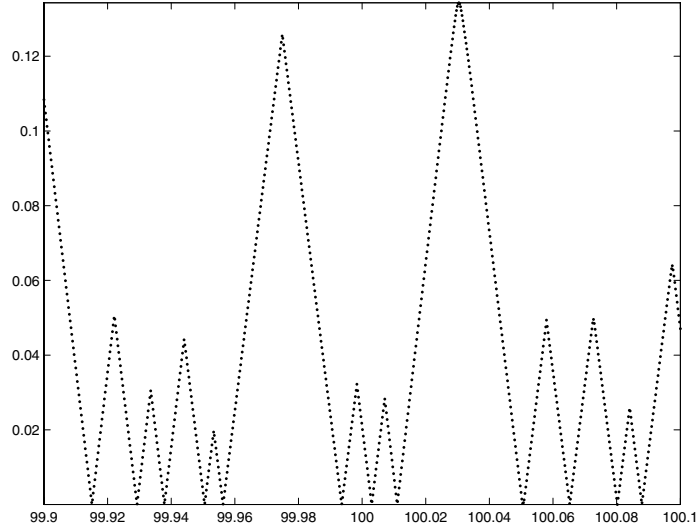


Fig. 8. Plot of σ_{\min} versus k for an interior Helmholtz problem on the contour shown in Fig. 7. The values shown were computed using the iteration technique of Section 4.5 applied to a matrix of size $N = 6400$. Each point in the graph required about 60 s of CPU time.

Table 9

Details of the computation for the Helmholtz kernel (6.3) associated with an interior Dirichlet problem on the smooth pentagram shown in Fig. 7 for the case $N = 6400$ and $k = 100.011027569 \dots$

j	p_j	n_j	γ_j	t_j	$\ C^{(j)}\ _{\infty}$	$\ B^{(j)}\ _{\infty}$	$\ D^{(j)}\ _{\infty}$
1	128	50.00	0.76	15.50	1.12e+00	1.12e+00	4.20e-02
2	64	76.00	0.59	14.32	3.27e+01	3.27e+01	1.75e+00
3	32	89.72	0.60	8.94	1.63e+01	1.62e+01	9.28e-01
4	16	107.00	0.64	6.27	9.09e+00	9.17e+00	2.41e+00
5	8	138.00	0.72	5.97	7.32e+00	7.31e+00	3.64e+00
6	4	199.50	0.80	7.76	3.22e+00	3.23e+00	3.86e+00

For each level j , the table shows the number of clusters p_j on that level, the average size of a cluster n_j , the compression ratio γ_j , the time required for the factorization t_j and the size of the matrices $B^{(j)}$, $C^{(j)}$ and $D^{(j)}$ (see (3.17)) in the maximum norm. For this computation, $E_{\text{res}} = 2.8 \times 10^{-10}$, $E_{\text{pot}} = 3.3 \times 10^{-5}$ and $\sigma_{\min} = 1.4 \times 10^{-5}$.

random points on the contour is not well predicted by their physical separation. One example of such a curve is the star-fish lattice illustrated in Fig. 9. Focusing on the double layer Laplace problem (6.1), we apply the factorization scheme to a matrix of size $N = 25\,600$ and compare the performance to that for the rippled dumb-bell shown in Fig. 6. Table 10 shows that the factorization of the matrix related to the starfish lattice took almost five times as long and resulted in a compressed matrix of over twice the size.

To understand the difference in performance between the different contours, we need to consider how the interaction rank of a cluster depends on its size. For the contours shown in Figs. 4, 6, and 7, we have seen that the rank of the interaction between a cluster of size m and the rest of the contour is effectively bounded by $\log m$. However, for the contour shown in Fig. 9 the corresponding bound is \sqrt{m} . Figs. 5 and 10 illustrate the difference. Thus, the asymptotic complexity of the scheme when applied to a contour similar to the star-fish lattice is $O(n^{3/2})$ rather than $O(n \log n)$.

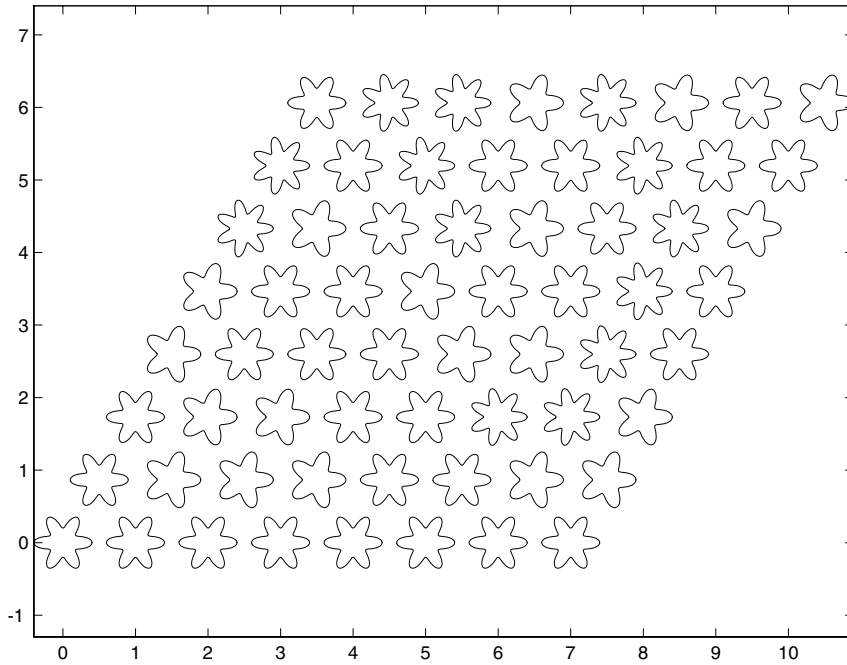


Fig. 9. The star-fish lattice contour; the physical distance between two random points on the contour is not well predicted by their distance along the contour.

Table 10

Test results for two experiments concerning the matrix obtained by discretizing the double layer Laplace problem (6.1)

Contour:	t_{tot}	N_{start}	N_{final}	M
Rippled dumb-bell	37s	25 600	559	86Mb
Star-fish lattice	172s	25 600	1202	210Mb

The table illustrates the difference in performance of the algorithm when applied to, on the one hand, the contour shown in Fig. 6 (top line), and on the other hand, the contour shown in Fig. 9 (lower line).

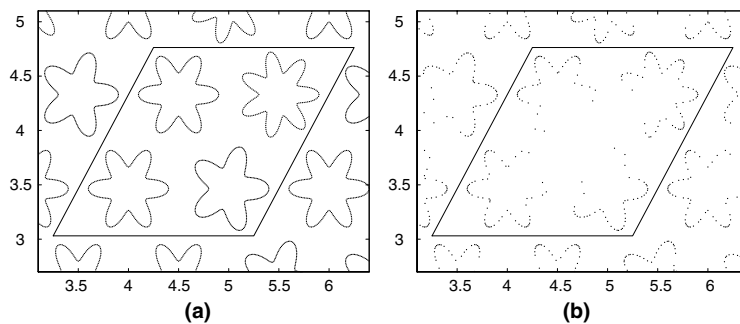


Fig. 10. Figure (a) shows a close-up of the star-fish lattice of Fig. 9. Figure (b) shows the nodes remaining after the interaction between the cluster formed by the points inside the parallelogram and the remainder of the contour has been compressed, cf. Fig. 5.

7. Generalizations and conclusions

We have presented a numerical scheme that constructs a compressed factorization of the inverse of a matrix. The scheme is applicable to generic matrices whose off-diagonal blocks have rank-deficiencies but is most efficient when applied to matrices arising from the discretization of integral equations defined on one-dimensional contours. (Although such integral equations frequently arise in the analysis of boundary value problems in two dimensions, the dimension of the underlying space is of little relevance to the algorithm.) For equations with non-oscillatory kernels, the computational complexity of the algorithm is $O(n)$, where n is the number of nodes in the discretization of the contour.

Comparing our implementations of the direct factorization scheme on the one hand and the FMM matrix-vector multiplication scheme on the other, we observed (i) that in a typical environment, the cost of constructing a factorization of the inverse is 15–20 times larger than the cost of a single FMM matrix-vector multiply, and (ii) that once the factorization of the inverse has been computed, the cost to apply it to a vector is 5–10 times smaller than the cost of a single FMM matrix-vector multiply. Thus, if an iterative solver requires less than 20 steps to converge, the iterative solver outperforms the direct solver for a single solve. However, if multiple right-hand sides are involved, the direct solver has a clear advantage. This observation is the foundation for [11].

Since the scheme is based on rank considerations only, it cannot work for boundary integral equations involving highly oscillatory kernels. However, since the interaction ranks are determined dynamically, the oscillation must be quite significant before the scheme becomes impracticable. Empirically, it was found that the scheme remains efficient for contours a couple of hundred wavelengths in size.

Another limitation of the scheme is that it does not achieve optimal efficiency when applied to a boundary integral equation set on either a contour similar to the one shown in Fig. 9, or on a two-dimensional surface. In either case, its computational complexity is $O(n^{3/2})$. Overcoming this limitation is a subject of on-going research.

Finally, we mention that the matrix factorization scheme presented in this paper can be modified to construct certain standard matrix factorizations (such as the singular value decomposition). This modification will be reported at a later date.

Acknowledgement

The first author was supported in part by the Office of Naval Research under contract #N00014-01-1-0364. The second author was supported in part by the Defense Advanced Research Projects Agency under contract # MDA972-00-1-0033.

References

- [1] F.X. Canning, K. Rogovin, Fast direct solution of standard moment-method matrices, *IEEE Antennas Propagation Mag.* 40 (1998) 15–26.
- [2] Yu. Chen, A fast direct algorithm for the Lippmann–Schwinger integral equation in two dimensions, *Adv. Comput. Math.* 16 (2–3) (2002) 175–190, Yu Chen, Modeling and computation in optics and electromagnetics. MR 2003b:65139.
- [3] H. Cheng, Z. Gimbutas, P.G. Martinsson, and V. Rokhlin, On the compression of low rank matrices, Tech. report, Yale University, Department of Computer Science, 2003.
- [4] W.C. Chew, An n^2 algorithm for the multiple scattering problem of n scatterers, *Micro. Optical Tech. Lett.* 2 (1989) 380–383.
- [5] D. Gines, G. Beylkin, J. Dunn, LU factorization of non-standard forms and direct multiresolution solvers, *Appl. Comput. Harmon. Anal.* 5 (2) (1998) 156–201, MR99c:65087.
- [6] Gene H. Golub, Charles F. Van Loan, Matrix computations, third ed., Johns Hopkins Studies in the Mathematical Sciences, Johns Hopkins University Press, Baltimore, MD, 1996, MR 97g:65006.

- [7] Ming Gu, Stanley C. Eisenstat, Efficient algorithms for computing a strong rank-revealing QR factorization, *SIAM J. Sci. Comput.* 17 (4) (1996) 848–869, MR 97h:65053.
- [8] W. Hackbusch, A sparse matrix arithmetic based on H-matrices. I. Introduction to H-matrices, *Computing* 62 (2) (1999) 89–108, MR 2000c:65039.
- [9] W. Hackbusch, S. Börm, Data-sparse approximation by adaptive H^2 -matrices, *Computing* 69 (1) (2002) 1–35, MR 1 954 142.
- [10] S. Kapur, V. Rokhlin, High-order corrected trapezoidal quadrature rules for singular functions, *SIAM J. Numer. Anal.* 34 (4) (1997) 1331–1356, MR 98k:65011.
- [11] P.G. Martinsson, Fast evaluation of electro-static interactions in two phase dielectric media, Tech. report, Yale University, Dept. of Computer Science, 2004.
- [12] E. Michielssen, A. Boag, A multilevel matrix decomposition algorithm for analysing scattering from large structures, *IEEE Trans. Antennas and Propagation* 44 (8) (1996) 1086–1093.
- [13] E. Michielssen, A. Boag, W.C. Chew, Scattering from elongated objects: direct solution in $O(n \log^2 n)$ operations, *IEEE Proc. H* 143 (1996) 277–283.
- [14] D. Scott, Analysis of the symmetric lanczos process, Tech. report, University of California at Berkeley, 1978.