A note on the $O(n)$-storage implementation of the GKO algorithm

Federico Poloni*

Abstract

We propose a new $O(n)$-space implementation of the GKO-Cauchy algorithm for the solution of linear systems with Cauchy-like matrix. Despite its slightly higher computational cost, this new algorithm makes a more efficient use of the processor cache memory. Thus, for matrices of size larger than $n \approx 500 - 1000$, it outperforms the existing algorithms.

We present an applicative case of Cauchy-like matrices with non-reconstructible main diagonal. In this special instance, the $O(n)$ space algorithms can be adapted nicely to provide an efficient implementation of basic linear algebra operations in terms of the low displacement-rank generators.

1 Introduction

Several classes of algorithms for the numerical solution of Toeplitz-like linear systems exist in literature. We refer the reader to [16] for an extended introduction on this topic, with descriptions of each method and plenty of citations to the relevant papers, and only summarize them in the following table.

Name	Operations	Memory	Stability
Levinson	$O(n^2)$	$O(n)$	stable only for some symmetric matrices (positive reflection coefficients)
Schur-Bareiss	$O(n^2)$	$O(n^2)$	backward stable only for PSD matrices
GKO	$O(n^2)$	$O(n^2)$	stable in practice in most cases (like Gaussian elimination)
Superfast	$O(n \log^2 n)$	$O(n)$	leading constant may be large; approximate; may be unstable in the nonsymmetric case

Stabilization techniques such as *look-ahead* may raise the computational cost from $O(n^2)$ to $O(n^3)$ or from $O(n \log^2 n)$ to $O(n^2)$. Though superfast Toeplitz solvers have a lower computational cost, when the system matrix is nonsymmetric and ill-conditioned $O(n^2)$ algorithms such as the GKO [6] are still attractive.

In this paper we will deal with the GKO algorithm. It is composed of two steps: reduction of the Toeplitz matrix to a Cauchy-like matrix with displacement rank $r = 2$, which takes $O(n)$ space and $O(n \log n)$ ops, and actual solution

---

[*] Scuola Normale Superiore, Piazza dei Cavalieri,7; 56126 Pisa, Italy. E-mail `f.poloni@sns.it`

of the Cauchy-like system via a Schur-type algorithm, which takes $O(n^2)$ space and time.

In a 2006 paper by G. Rodriguez on Cauchy-like least squares problems, a strategy to implement the Cauchy-like system solution using only $O(n)$ memory locations (but $6rn^2$ instead of $4rn^2$ ops, where r is the displacement rank of the Cauchy-like matrix) first appeared. In the first part of the present paper, we will provide an alternative $O(n)$-space implementation of the Schur Cauchy-like system solution algorithm, having the same cost and memory requirements but some desirable computational properties.

Moreover, in some applications, a special kind of partially reconstructible Cauchy-like matrices appear, i.e., those in which the main diagonal is not reconstructible. We shall call them Trummer-like, as they are associated with Trummer's problem [5]. We will show how the $O(n)$-storage algorithms adapt nicely to this case, allowing one to develop an integrated algorithm for their fast inversion. In particular, one of the key steps in order to obtain a full representation of their inverse is the calculation of $\operatorname{diag}(T^{-1})$ for a given Trummer-like T, which is not so obvious to do with $O(n^2)$ ops, even using $O(n^2)$ space.

Structure of the paper In Section 2, we will recall the concept of displacement operators, Cauchy-like and Trummer-like matrices. In Sections 3 and 4 we will study respectively the original GKO algorithm and its first $O(n)$-space variant due to Rodriguez. In Section 5 we will introduce and analyze our new $O(n)$-space variant. In Section 6 we will deal with system solving and matrix inversion for Trummer-like matrices. Finally, Section 7 is dedicated to showing some numerical experiments that confirm the effectiveness of our approach, and Section 8 contains some conclusive remarks.

2 Basic definitions

Indexing and notation We will make use of some handy matrix notations taken from Matlab®. When M is a matrix, the symbol $M_{i:j,k:\ell}$ denotes the submatrix formed by rows i to j and columns k to ℓ of M, including extremes. The index i is a shorthand for $i:i$, and $:$ alone is a shorthand for $1:n$, where n is the maximum index allowed for that row/column. A similar notation is used for vectors. When $v \in \mathbb{C}^n$ is a vector, the symbol $\operatorname{diag}(v)$, or $\operatorname{diag}(v_1, \ldots, v_n)$ denotes the diagonal matrix $D \in \mathbb{C}^{n\times n}$ such that $D_{i,i} = v_i$. On the other hand, when $M \in \mathbb{C}^{n\times n}$ is a square matrix, the symbol $\operatorname{diag}(M)$ denotes the (column) vector with entries $M_{1,1}, M_{2,2}, \ldots, M_{n,n}$.

Displacement operators and Cauchy-like matrices Let $t, s \in \mathbb{C}^n$. We shall denote by $\nabla_{t,s}$ the operator $\mathbb{C}^{n\times n} \to \mathbb{C}^{n\times n}$ which maps M to

$$\nabla_{t,s}(M) = \operatorname{diag}(t)M - M\operatorname{diag}(s).$$

A matrix $C \in \mathbb{C}^{n\times n}$ is said *Cauchy-like* (with displacement rank r) if there are vectors s, t and matrices $G \in \mathbb{C}^{n\times r}, B \in \mathbb{C}^{r\times n}$ such that

$$\nabla_{s,t}(C) = GB. \tag{1}$$

Notice that if we allow $r = n$, then any matrix is Cauchy-like. In the applications, we are usually interested in cases in which $r \ll n$, since the computational costs of all the involved algorithms depend on r.

A Cauchy-like matrix is called a *quasi-Cauchy matrix* if $r = 1$, and *Cauchy matrix* if $G^* = B = (1, 1, \dots, 1)$. Usually, it is assumed that the operator $\nabla_{t,s}$ is nonsingular, or equivalently, $t_i \neq s_j$ for all pairs i, j. Under this assumption, the elements of C can be written explicitly as

$$C_{ij} = \frac{\sum_{l=1}^{r} G_{il} B_{lj}}{t_i - s_j}, \tag{2}$$

thus C can be fully recovered from G, B, t and s. Otherwise, the latter formula only holds for the entries C_{ij} such that $t_i \neq s_j$, and C is said to be *partially reconstructible*. The matrices G and B are called the *generators* of C, and the elements of t and s are called *nodes*.

Trummer-like matrices In Section 6, we will deal with the case in which $t = s$ is an injective sequence, that is, when the non-reconstructible elements are exactly the ones belonging to the main diagonal. We will use ∇_s as a shorthand for $\nabla_{s,s}$. If $\nabla_s(T) = GB$ has rank r, a matrix T will be called *Trummer-like* (with displacement rank r). Notice that a Trummer-like matrix can be fully recovered from G, B, s, and $d = \operatorname{diag}(T)$. Trummer-like matrices are related to interpolation problems [5], and may arise from the transformation of Toeplitz and similar displacement structure [11], or directly from the discretization of differential problems [2].

3 Overview of the GKO Schur step

Derivation The fast LU factorization of a Cauchy-like matrix C is based on the following lemma.

Lemma 1. *[12] Let*

$$C = \begin{bmatrix} C_{1,1} & C_{1,2:n} \\ C_{2:n,1} & C_{2:n,2:n} \end{bmatrix}$$

satisfy the displacement equation (1) *with* $m = n$, *and suppose* $C_{1,1} \neq 0$. *Then its Schur complement* $C^{(2)} = C_{2:n,2:n} - C_{2:n,1} {C_{1,1}}^{-1} C_{1,2:n}$ *satisfies the displacement equation*

$$D_{t_{2:n}} S - S D_{s_{2:n}} = G^{(2)} B^{(2)},$$

with

$$G^{(2)} = G_{2:n,1:r} - C_{2:n,1} {C_{1,1}}^{-1} G_{1,1:r}, \; B^{(2)} = B_{1:r,2:n} - B_{1:r,1} {C_{1,1}}^{-1} C_{1,2:n}. \tag{3}$$

Using this lemma, we can construct the LU factorization of C with $O(n^2)$ floating point operations (ops). The algorithm goes on as follows. Given $G^{(1)} = G$, $B^{(1)} = B$, and the two vectors s and t, recover the pivot $C_{1,1}$, the first row $C_{1,2:n}$ and the first column $C_{2:n,1}$ of C using the formula (2). This allows to calculate easily the first row of U as $\begin{bmatrix} C_{1,1} & C_{1,2:n} \end{bmatrix}$ and the first column of L as $\begin{bmatrix} 1 & C_{2:n,1}^T {C_{1,1}}^{-1} \end{bmatrix}^T$. Then use equations (3) to obtain the generators $G^{(2)}$ and $B^{(2)}$ of the Schur complement $C^{(2)}$ of C. Repeat the algorithm setting

$G \leftarrow G^{(2)}$, $B \leftarrow B^{(2)}$, $s \leftarrow s_{2:n}$ and $t \leftarrow t_{2:n}$ to get the second row of U and the second column of L, and so on. A simple Matlab implementation is given in Algorithm 1.

```
function [L,U]=gko_lu(G,B,t,s)
  n=size(G,1);

  L=eye(n);
  U=zeros(n,n);

  for k=1:n
    U(k,k:n) = (G(k,:) * B(:,k:n))./ transpose(t(k)-s(k:n));
    L(k+1:n,k) = (G(k+1:n,:) * (B(:,k)/U(k,k))) ./ (t(k+1:n)-s(k));

    G(k+1:n,:) = G(k+1:n,:) - L(k+1:n,k) * (G(k,:));
    B(:,k+1:n) = B(:,k+1:n) - (B(:,k)/U(k,k)) * U(k,k+1:n);
  end
end
```

Algorithm 1: LU factorization of Cauchy-like matrices

When the LU factorization is only used for the solution of a linear system in the form $Cx = b$, with $b \in \mathbb{C}^{n\times m}$, it is a common technique to avoid constructing explicitly L, computing instead $L^{-1}b$ on-the-fly as the successive columns of L are computed. This is also possible with the GKO algorithm, as shown in Algorithm 2. This strategy effectively reduces the temporary storage needed by the algorithm from $2n^2 + o(n^2)$ to $n^2 + o(n^2)$ floating point numbers of memory space.

Comments Notice that Algorithm 2 includes partial pivoting. Its total cost is $(4r + 2m)n^2$ ops, when applied to a matrix C with displacement rank r and an $n \times m$ right-hand side. The algorithm works whenever C is a completely reconstructible Cauchy matrix; if it is not the case, when the number of non-reconstructible entries is small, the algorithm can be modified to store and update them separately, see e.g. Kailath and Olshevsky [11] or Section 6.

However, there is an important drawback in Algorithm 2: while the size of the input and output data is $O(n)$, $O(n^2)$ memory locations of temporary storage are needed along the algorithm to store U. Therefore, for large values of n the algorithm cannot be effectively implemented on a computer because it does not fit in the RAM.

Moreover, another important issue is caching. Roughly speaking, a personal computer has about 512 kb–8 Mb of cache memory, where the most recently accessed locations of RAM are copied. Accessing a non-cached memory location is an order of magnitude slower than a cached one. The real behavior of a modern processor is more complicated than this simple model, due to the presence of several different levels of cache, each with its own performance, and instruction pipelines [9]. Nevertheless, this should highlight that when the used data do not fit anymore into the cache, saving on memory could yield a greater speedup than saving on floating point operations.

```
function x=gko_solve(G,B,t,s,b)
%solves Cx=b, where diag(t)*C-C*diag(s)=G*B
  n=size(G,1);

  l=zeros(n,1);
  U=zeros(n,n);

  x=b;
  for k=1:n
    %finds pos=pivot row
    l(k:n) = (G(k:n,:)*B(:,k))./(t(k:n)-s(k));
    [abspiv pos]=max(abs(l(k:n)));
    if(abspiv==0) error('singular matrix');end
    %pivots
    pos=pos+k-1;
    l([k pos])=l([pos k]);
    x([k pos],:)=x([pos k],:);
    G([k pos],:)=G([pos k],:);
    t([k pos])=t([pos k]);
    %LU step
    U(k,k+1:n)=(G(k,:)*B(:,k+1:n))./ transpose(t(k)-s(k+1:n));
    pivot=l(k);
    U(k,k)=pivot;
    x(k+1:n,:) = x(k+1:n,:) - l(k+1:n)* (x(k,:)/pivot);
    G(k+1:n,:) = G(k+1:n,:) - l(k+1:n) * (G(k,:)/pivot);
    B(:,k+1:n) = B(:,k+1:n) - (B(:,k)/pivot) * U(k,k+1:n);
  end
  %performs back-substitution
  for k=n:-1:1
    x(k,:)= (x(k,:) - U(k,k+1:n)*x(k+1:n,:)) / U(k,k);
  end
end
```

Algorithm 2: Solving a system $Cx = b$ with implicit L factor and pivoting

4 Low-storage version of GKO: the extended matrix approach

Derivation Up to our knowledge, the first algorithm to solve the high storage issue in GKO was proposed by Rodriguez [14] in 2006, while dealing with least squares Cauchy-like problems. More recently, a deeper analysis and a complete ready-to-use Matlab implementation were provided by Aricò and Rodriguez [1].

The approach is based on an idea first appeared in Kailath and Chun [10]. Let us suppose that C is a completely reconstructible Cauchy matrix and that s is an injective sequence, i.e., $s_i \neq s_j$ for all i, j. The solution of the linear system $Cx = b$ can be expressed as the Schur complement of C in the rectangular matrix

$$\widetilde{C} = \begin{bmatrix} C & b \\ -I & 0 \end{bmatrix}.$$

Thus, we can compute x by doing n steps of Gaussian elimination on $\widetilde{C}$. Moreover, the first block column of $\widetilde{C}$ is is partially reconstructible Cauchy-like with respect to $\widetilde{s} = \begin{bmatrix} s^T & t^T \end{bmatrix}^T$ and t; therefore, while performing the Gaussian elimination algorithm, the entries of this block can be stored and updated in terms of the generators, as in Algorithm 1. Unlike the previous algorithms, we may discard the rows of U and columns of L as soon as they are computed, keeping only the generators. Instead, the entries in the second block column are computed with customary Gaussian elimination and stored along all the algorithm.

The following observations, which will be needed later, should make clearer what is going on with this approach.

Lemma 2. *Suppose for simplicity that no pivoting is performed; let L and U be the LU factors of C, x be the solution to the linear system $Cx = b$, y be the solution to $L^{-1}y = b$, and $W = U^{-1}$. Let k denote the step of Gaussian elimination being performed, with e.g. $k = 1$ being the step that zeroes out all the elements of the first column but the first. During the algorithm,*

1. *The (i, j) entry of the $(1, 1)$ block is updated at all steps k with $k < j$. After the last update, it contains $U_{i,j}$*
2. *The (i, j) entry of the $(1, 2)$ block is updated at all steps k with $k < j$. After the last update, it contains $y_{i,j}$.*
3. *The (i, j) entry of the $(2, 2)$ block is updated at all steps k with $k \geq j$. In particular, the last step ($k = n$) updates all entries, and after that the $(2, 2)$ block contains $x_{i,j}$.*
4. *The (i, j) entry of the $(2, 1)$ block is updated at all steps k with $i \leq k \leq j$. After the last update, it contains 0. Immediately before that, i.e., just after step $j - 1$, it contains $-W_{i,j}U_{j,j}$.*

Proof. From the structure of Gaussian elimination, it can easily be verified that the entries are only updated during the abovementioned steps. In particular, for the condition on updates to the $(2, 1)$ block, it is essential that the initial $(2, 1)$ block initially contains a diagonal matrix. Regarding which values appear finally in each position,

1. is obvious: in fact, if we ignore all the other blocks, we are doing Gaussian elimination on C.
2. is easily proved: since the row operations we perform transform $C = LU$ to U, they must be equivalent to left multiplication by L^{-1}.
3. is a consequence of the well-known fact that after n steps of Gaussian elimination we get the Schur complement of the initial matrix in the trailing diagonal block.
4. is less obvious. For all $i \geq j$, let us call $Z_{i,j}$ the value of the (i, j) entry right after step $j - 1$, and consider how the entries of the $(2, 2)$ block are updated along the algorithm. They are initially zero, and at the kth step the one in place (i, j) is incremented by $-(Z_{i,k}/U_{k,k})y_{k,j}$, so its final value is
$$x_{i,j} = -\sum_k (Z_{i,k}/U_{k,k})y_{k,j}.$$

Since for each choice of b (and thus of $y = L^{-1}b$) $Z_{i,k}$ and $W_{i,k}$ are unchanged, as they only depend on C, and it holds that

$$x_{i,j} = (U^{-1}y)_{i,j} = \sum_k W_{i,k} y_{k,j},$$

the only possibility is that $W_{i,k} = -Z_{i,k}/U_{k,k}$ for each i, k.

$\square$

We report here the resulting Algorithm 3.

```
function x=gko_solve_em(G,B,t,s,b)
%solves Cx=b, where diag(t)*C-C*diag(s)=G*B
%G,B are generators of the 1st block col, x is the 2nd block col
%G,l,x have 2n rows, but they are stored ”modulo n”
  n=size(G,1);
  l=zeros(n,1);
  u=zeros(1,n);
  x=b;
  for k=1:n
    %finds pos=pivot row (among k: n)
    l = (G*B(:,k))./([s(1:k-1);t(k:n)]-s(k));
    [abspiv pos]=max(abs(l(k:n)));
    if(abspiv==0) error('singular matrix');end
    %pivots
    pos=pos+k-1;
    l([k pos])=l([pos k]);
    x([k pos],:)=x([pos k],:);
    G([k pos],:)=G([pos k],:);
    t([k pos])=t([pos k]);
    %LU step
    u(k+1:n)=(G(k,:)*B(:,k+1:n))./ transpose(t(k)-s(k+1:n));
    pivot = l(k);
    l(k) = 0;
    x(k,:) = x(k,:)/pivot;
    x = x - l*x(k,:);
    G(k,:) = G(k,:)/pivot;
    G = G - l*G(k,:);
    B(:,k+1:n) = B(:,k+1:n) - (B(:,k)/pivot) * u(k+1:n);
  end
end
```

Algorithm 3: Solving a system $Cx = b$ with the extended matrix algorithm

Comments It is worth mentioning that several nice properties notably simplify the implementation.

- The partial reconstructibility of $\widetilde{C}$ is not an issue. If the original matrix C is fully reconstructible and s is injective, then the non-reconstructible entries of $\widetilde{C}$ are the ones in the form $C(n+k,k)$ for $k = 1,\dots,n$, that is, the ones in which the -1 entries of the $-I$ block initially lie. It is readily shown that whenever the computation of such entries is required, their value is the initial one of -1.

- At each step of the algorithm, the storage of only a limited number of rows of G and of the right block column x is required: at step k, we only need the rows with indices from k to $n+k-1$ (as the ones below are still untouched by the algorithm, and the ones above are not needed anymore). It is therefore possible to store the rows modulo n, thus halving the storage space needed for some of the matrices.
- Pivoting can be easily included without destroying the block structure by acting only on the rows belonging to the first block row of $\widetilde{C}$.

Algorithm 3 uses $(6r+2m)n^2+o(n^2)$ floating point operations, and it can be implemented so that the input variables G, B, t, b are overwritten during the algorithm, with x overwriting b, so that it only requires $2n$ memory locations of extra storage (to keep l and u).

As we stated above, for the algorithm to work we need the additional assumption that s is an injective sequence, i.e., $s_i \neq s_j$ for all j. This is not restrictive when working with Cauchy-like matrices deriving from Toeplitz matrices or from other displacement structure; in fact, in this case the entries s_i are the n complex nth roots of a fixed complex number, thus not only are they different, but their differences $s_i - s_j$ can be easily bounded from below, which is important in order to assure that the algorithm is numerically stable. This is a common assumption when dealing with Cauchy matrices, since a Cauchy (or quasi-Cauchy) matrix is nonsingular if and only if x and y are injective sequences. For Cauchy-like matrices this does not hold, but the injectivity of the two sequences is still related to the nonsingularity of the matrix: for instance, we have the following result.

Lemma 3. *Let s have $r+1$ repeated elements, that is, $s_{i_1} = s_{i_2} = \cdots = s_{i_{r+1}} = s$. Then the Cauchy-like matrix* (2) *is singular.*

Proof. Consider the submatrix C' formed by the $r+1$ columns of C with indices $i_1, \dots, i_{r+1}$. It is the product of the two matrices $G' \in \mathbb{C}^{n\times r}$ and $B' \in \mathbb{C}^{r\times r+1}$, with

$$(G')_{ij} = \frac{G_{ij}}{t_i - s},\ (B')_{ij} = B_{is_j}.$$

Therefore C' (and thus C) cannot have full rank. ☐

5 Low-storage version of GKO: the downdating approach

Derivation In this section, we shall describe a different algorithm to solve Cauchy-like system using only $O(n)$ locations of memory. Our plan is to perform the first **for** loop in Algorithm 2 unchanged, thus getting $y = L^{-1}b$, but discarding the computed entries of U which would take $O(n^2)$ memory locations, and then to recover them via additional computations on the generators.

For the upper triangular system $s = U^{-1}(L^{-1}b)$ to be solved incrementally by back-substitution, we need the entry of the matrix U to be available one row at a time, starting from the last one, and *after* the temporary value $y = L^{-1}b$ has been computed, that is, after the whole LU factorization has been performed.

The key idea is trying to undo the transformations performed on B step by step, trying to recover $B^{(k)}$ from $B^{(k+1)}$. Because of the way in which the generators are updated in Algorithms 1 and 2, the first row of $G^{(k)}$ and the first column of $B^{(k)}$ are kept in memory untouched by iterations $k+1, \dots, n$ of the GKO algorithm. Thus we can use them in trying to undo the kth step of Gaussian elimination.

Let us suppose we know $B^{(k+1)}$, i.e., the contents of the second generator B after the $(k+1)$st step of Gaussian elimination, and the values of $G^{(k)}_{k,:}$ and $B^{(k)}_{:,k}$, which are written in G and B by the kth step of Gaussian elimination and afterwards unmodified (since the subsequent steps of 2 do not use those memory locations anymore).

We start from the second equation of (3) and (2) for the kth row of U, written using the colon notation for indices.

$$\begin{aligned} B^{(k+1)}_{:,\ell} &= B^{(k)}_{:,\ell} - B^{(k)}_{:,k} {U_{k,k}}^{-1} U_{k,\ell}, \\ U_{k,\ell} &= \frac{G^{(k)}_{k,:} B^{(k)}_{:,\ell}}{t_k - s_\ell}, \end{aligned}$$

both valid for all $\ell \geq k$. Substituting $B^{(k)}_{:,\ell}$ from the first into the second, and using the $k = \ell$ case pf the formula above to deal with $U_{k,k}$, we get

$$U_{k,\ell} = \frac{G^{(k)}_{k,:} B^{(k+1)}_{:,\ell}}{t_k - s_\ell} + \frac{G^{(k)}_{k,:} B^{(k)}_{:,k} {U_{k,k}}^{-1}}{t_k - s_\ell} U_{k\ell} = \frac{G^{(k)}_{k,:} B^{(k+1)}_{:,\ell}}{t_k - s_\ell} + \frac{t_k - s_k}{t_k - s_\ell} U_{k\ell}$$

and thus

$$U_{k,\ell} = \frac{G^{(k)}_{k,:} B^{(k+1)}_{:,\ell}}{s_k - s_\ell}, \tag{4}$$

$$B^{(k)}_{:,\ell} = B^{(k+1)}_{:,\ell} + B^{(k)}_{:,k} {U_{k,k}}^{-1} U_{k,\ell}. \tag{5}$$

The above equations allow one to recover the value of $B^{(k)}_{:,\ell}$ for all $\ell > k$ using only $B^{(k+1)}_{:,\ell}$, $G^{(k)}_{k,:}$ and $B^{(k)}_{:,k}$ as requested.

If we apply the method just described for $k = n-1, n-2, \dots, 1$, we are able to recover the different values of $B^{(k)}_{:,\ell}$ after the Gaussian elimination step, that is, "downdate" B to its previous value. In the process, we get at each iteration the kth row of U. In this way, the entries of U are computed in the correct order to solve the system $Ux = y$ incrementally by back-substitution. We report here the Matlab source for the resulting Algorithm 4.

Comments Notice that pivoting only affects the first phase of the algorithm, since the whole reconstruction stage can be performed on the pivoted version of C without additional row exchanges.

This algorithm has the same computational cost, $(6r+2m)n^2$, and needs the same number of memory locations, $2n$, as the extended matrix approach. Moreover, they both need the additional property that s be an injective sequence, as an $s_k - s_\ell$ denominator appears in (4). These facts may lead one to suspect

```
function x=gko_solve_dd(G,B,t,s,b)
%solves Cx=b, where diag(t)*C-C*diag(s)=G*B
n=size(G,1);
l=zeros(n,1);
u=zeros(1,n);
x=b;
for k=1:n
    %finds pos=pivot row
    l(k:n) = (G(k:n,:)*B(:,k))./(t(k:n)-s(k));
    [abspiv pos]=max(abs(l(k:n)));
    if(abspiv==0) error('singular matrix');end
    %pivots
    pos=pos+k-1;
    l([k pos])=l([pos k]);
    x([k pos],:)=x([pos k],:);
    G([k pos],:)=G([pos k],:);
    t([k pos])=t([pos k]);
    %LU step
    u(k+1:n)=(G(k,:)*B(:,k+1:n))./ transpose(t(k)-s(k+1:n));
    pivot=l(k);
    u(k)=pivot;
    x(k+1:n,:) = x(k+1:n,:) - l(k+1:n)* (x(k,:)/pivot);
    G(k+1:n,:) = G(k+1:n,:) - l(k+1:n) * (G(k,:)/pivot);
    B(:,k+1:n) = B(:,k+1:n) - (B(:,k)/pivot) * u(k+1:n);
end
%now x contains y=L\b, downdating phase
for k=n:-1:1
    u(k+1:n)=(G(k,:)*B(:,k+1:n))./transpose(s(k)-s(k+1:n));
    B(:,k+1:n) = B(:,k+1:n) + (B(:,k) / u(k)) * u(k+1:n);
    %back-substitution
    x(k,:)=(x(k,:)-u(k+1:n)*x(k+1:n,:)) / u(k);
end
end
```

Algorithm 4: Solving a system $Cx = b$ with the downdating algorithm

that they are indeed the same algorithm. However, it is to be noted the two algorithms notably differ in the way in which the system $Ux = y$ is solved. In fact, in the extended matrix approach we solve this system by accumulating the explicit multiplication $U^{-1}y$, while in the downdating approach we solve it by back-substitution.

Moreover, there are several small favorable details that suggest adopting the latter algorithm.

- With the extended matrix approach, we do not get any entry of x before the last step. On the other side, with the downdating approach, as soon as the first `for` cycle is completed, we get x_1, and then after one step of the downdating part we get x_2, and so on, getting one new component of the solution at each step. This is useful because in the typical use of this algorithm on Toeplitz matrices, x is the Fourier transform of a "meaningful" vector, such as one representing a signal, or an image, or the solution to an equation. Using the correct ordering, the first entries of a Fourier transform can be used to reconstruct a lower-sampled preview of the original data, with no additional computational overhead, see e.g. Walker[17]. Thus with this approach we can provide an approximate solution after only the first part of the algorithm is completed.

- In the extended matrix version, each step of the algorithm updates $O(nr)$ memory locations. Instead, in the downdating version, for each k, the $(n-k)$th and $(n+k)$th step work on $O(kr)$ memory locations. Therefore, the "innermost" iterations take only a small amount of memory and thus fit better into the processor cache. This is a desirable behavior similar to the one of *cache-oblivious algorithms* [4].
- In exact arithmetic, at the end of the algorithm the second generator B of the matrix C is reconstructed as it was before the algorithm. In floating point arithmetic, this can be used as an *a posteriori* accuracy test: if one or more entries of the final values of B are not close to their initial value, then there was a noticeable algorithmic error.

6 Computations with Trummer-like matrices

Theoretical results The following results, which are readily proved by expanding the definition of ∇_s on both sides, are simply the adaptation of classical results on displacement ranks (see e.g. Heinig and Rost [8]) to the Trummer-like case. Notice the formal similarity with the derivative operator.

Theorem 4. *Let* $A, B \in \mathbb{C}^{n\times n}$*, and let* $r(X) = \operatorname{rk} \nabla_s(A)$*.*

1. $\nabla_s(A+B) = \nabla_s(A) + \nabla_s(B)$*, so* $r(A+B) \leq r(A) + r(B)$*.*
2. $\nabla_s(AB) = \nabla_s(A)B + A\,\nabla_s(B)$*, so* $r(AB) \leq r(A) + r(B)$*.*
3. $\nabla_s(A^{-1}) = -A^{-1}\,\nabla_s(A)A^{-1}$*, so* $r(A^{-1}) = r(A)$*.*

As we saw in Section 2, a Trummer-like matrix can be completely reconstructed by knowing only the node vector s, the generators G and B, and its diagonal $d = \operatorname{diag}(T)$. In this section, we are interested in implementing fast—i.e., using $O(n^2)$ ops—and space-efficient—i.e., using $O(n)$ memory locations—matrix-vector and matrix-matrix operations involving Trummer-like matrices stored in this form.

First operations For the matrix-vector product, all we have to do is reconstructing one row at a time of the matrix T and then computing the customary matrix-vector product via the usual formula $(Tv)_i = \sum_j T_{ij}v_j$. The matrix product between two Trummer-like matrices T and S is easy to implement: let G_T and B_T (resp. G_S and B_S) be the generators of T (resp. S); then, by Theorem 4, the generators of TS are

$$\begin{bmatrix} TG_S & G_T \end{bmatrix}, \quad \begin{bmatrix} B_S \\ B_T S \end{bmatrix},$$

while $\operatorname{diag}(TS)$ can be computed in $O(n^2)$ by recovering at each step one row of T and one column of S and computing their dot product. Sums can be done exactly in the same way: the generators of $S+T$ are

$$\begin{bmatrix} G_S & G_T \end{bmatrix}, \quad \begin{bmatrix} B_S \\ B_T \end{bmatrix},$$

and the diagonal is $d_S + d_T$

Linear systems Linear system solving is less obvious. In Kailath and Olshevsky [11], the following algorithm was suggested: the GKO Gaussian elimination algorithm is performed, but at the same time the computed row $U_{k,k:n}$ and column $L_{k:n,k}$ are used to update the diagonal d to the diagonal of the Schur complement, with the customary Gaussian elimination formula

$$T^{(k+1)}_{i,i} = T^{(k)}_{i,i} - L_{i,k}(T^{(k)}_{k,k})^{-1}U_{k,i}. \tag{6}$$

It is easy to see that this strategy can be adapted to work with both the extended matrix and the downdating version of the algorithm, thus allowing one to implement GKO with $O(n)$ storage also for this class of matrices.

However, a more delicate issue is pivoting. Kailath and Olshevsky do not deal with the general case, since they work with symmetric matrices and with a symmetric kind of pivoting that preserves the diagonal or off-diagonal position of the entries. Let us consider the pivoting operation before the kth step of Gaussian elimination, which consists in choosing an appropriate row p and exchanging the kth and pth rows. The main issue here is that the two non-reconstructible entries that were in position T_{kk} and T_{pp}, now are in positions T_{pk} and T_{kp}. This requires special handling in the construction of the kth row in the Gaussian elimination step, but luckily it does not affect the successive steps of the algorithm, since the kth column and row are not used from step $k+1$ onwards. On the other hand, the entry T_{pp}, which used to be non-reconstructible before pivoting, now is reconstructible. We may simply ignore this fact, store it in d and update it again with the formula (6) as if it were not reconstructible.

The complete Algorithm 5 is reported here.

Matrix inversion Matrix inversion poses an interesting problem too. The generators of T^{-1} can be easily computed as $-T^{-1}G$ and BT^{-1} by resorting to Algorithm 5. However, whether we try to compute the representation of T^{-1} or directly that of $T^{-1}S$ for another Trummer-like matrix S, we are faced with the problem of computing $\operatorname{diag}(T^{-1})$ given a representation of T. There appears to be no simple direct algorithm to extract it in time $O(n^2)$ from the LU factors of T.

We shall present here a solution based on the observations of Lemma 2. Let us ignore pivoting in this first stage of the discussion. Notice that the last part of Lemma 2 shows us a way to compute $(U^{-1})_{1:k,k}$ at the kth step of the extended matrix algorithm. Our plan is to find a similar way to get $(L^{-1})_{k,1:k}$ at the same step, so that we can compute the sums

$$(T^{-1})_{i,i} = \sum_k (U^{-1})_{i,k}(L^{-1})_{k,i}, \quad i = 1, \dots, n, \tag{7}$$

one summand at each step, accumulating the result in a temporary vector.

The following result holds.

Lemma 5. *Let T be Trummer-like with generators G and B, nodes s and diagonal d, $T = LU$ be the its LU factorization, and $D = \operatorname{diag}(p)$, where $p_i = U_{i,i}$ are the pivots.*

1. *The LU factorization of T^*, the transpose conjugate of T, is $(U^* D^{-1})(DL^*)$.*
2. *The matrix T^* is Trummer-like with nodes s, diagonal d and generators B^* and G^*.*
3. *Let $G^{(k)}$ and $B^{(k)}$ be the content of the variables G and B after the kth step of the GKO algorithm on T, and $\bar{G}^{(k)}$ and $\bar{B}^{(k)}$ be the content of the same variables after the same step of the GKO algorithm on T^*. Then, $\bar{G}^{(k)} = (B^{(k)})^*$ and $\bar{B}^{(k)} = (G^{(k)})^*$*

Proof. The matrices $U^* D^{-1}$ and DL^* are respectively unit lower triangular and upper triangular. Thus the first part holds by the uniqueness of the LU factorization. The second part is clear, and the last one follows by writing down the formula (3) for T and T^*. □

Therefore, there is much in common between the GKO algorithm on T and T^*, and the two can be carried on simultaneously saving a great part of the computations involved. Moreover, in the same way as we obtain $(U^{-1})_{1:k,k}$, we may also get at the kth step its equivalent for T^*, i.e., $((DL^*)^{-1})_{1:,k} = p_k (L^{-1})_{k,1:k}$. Since p_k, the kth pivot, is also known, this allows us to recover $(L^{-1})_{k,1:k}$.

Thus we have shown a way to recover both $(U^{-1})_{1:k,k}$ and $(L^{-1})_{k,1:k}$ at the kth step of the extended matrix algorithm, and this allows us to compute the kth summand of (7) for each i.

Pivoting How does pivoting affect this scheme for the computation of $\operatorname{diag}(T^{-1})$? If $T = PLU$, formula (7) becomes

$$(T^{-1})_{i,i} = \sum_k (U^{-1})_{i,k} (L^{-1} P^{-1})_{k,i}, \quad i = 1, \ldots, n. \tag{8}$$

The permutation matrix P, of which we already have to keep track during the algorithm, acts on L^{-1} by scrambling the column indices i, so this does not affect our ability to reconstruct the diagonal, as we still have all the entries needed to compute the kth summand at each step k. We only need to take care of the order in which the elements of $(U^{-1})_{1:k,k}$ and $(L^{-1})_{k,1:k}$ are paired in (8).

The complete algorithm, which includes pivoting, is reported here as Algorithm 6

Comments It is worth noting with the same run of the GKO algorithm we can both compute $\operatorname{diag}(T^{-1})$ and solve linear systems with matrices T and T^*, as the two algorithms share much of their computations. In particular, the solution of the two systems giving $T^{-1}G$ and BT^{-1}, which are the generators of T^{-1}, are computed by the algorithm with no additional effort: the transformations on G and B needed to solve them are exactly the ones that are already performed by the factorization algorithm. Also observe that, since the computation of $(T^{-1})_{ii}$ spans steps i to n, while d_i is needed from step 1 to step i, we may reuse the vector d to store the diagonal of the inverse. The resulting algorithm has a total computational cost of $(12r+2m_1+2m_2+3)n^2$ ops, if we are solving at the same time a system $Tx = b$ with $b \in \mathbb{C}^{n\times m_1}$ and a system $yT = c$ with $c \in \mathbb{C}^{m_2\times n}$. The only extra storage space needed is that used for u and l, i.e., $2n$ memory locations.

Another observation is that we did not actually make use of the fact that the diagonal of T is non-reconstructible: in principle, this approach works even if C is a Cauchy matrix with respect to two different node vectors t and s. This might be useful in cases in which we would rather not compute explicitly the diagonal elements, e.g. because $t_i - s_i$ is very small, and thus would lead to ill-conditioning.

7 Numerical experiments

Speed measurements Due to the computational overhead caused by its being interpreted, Matlab® is not the best tool to measure subtle speed improvements such as the ones due to caching issues. Instead, for the speed experiments we converted the Matlab® source into Fortran 90. The compiler used was the lf95 Fortran compiler version 6.20c, with command-line options −o2 −tp4 −lblasmtp4. The experiments took place on two different computers:

C1 a machine equipped with four Intel® Xeon™ 2.80Ghz CPUs, each equipped with 512kb of L2 cache, and 6 GB of RAM. Since we did not develop a parallel implementation, only one of the processor was actually used for the computations.

C2 a machine equipped with one Intel® Pentium® 4 2.80Ghz CPU with 1024kb of L2 cache, and 512Mb of RAM.

The results are are shown in Table 1.

Comments It is clear from the table that two different behaviors arise for different sizes of the input. For small values of n, the clear winner is the traditional $O(n^2)$-space algorithm, due to its lower net computational cost of $(4r+2m)n^2$ instead of $(6r+2m)n^2$ (for these tests, $r = 2$, $m = 1$). As the dimension of the problem increases, cache efficiency starts to matter, and the traditional algorithm becomes slower than its counterparts. This happens starting from $n \approx 256 - 512$. Quick calculations show that the memory occupation of the full $n \times n$ matrix is 512kb for $n = 256$ and 2Mb for $n = 512$, so the transition takes indeed place when the $O(n^2)$ algorithm starts to suffer from cache misses.

Accuracy measurements For the accuracy experiments, we chose three test problems, the first two taken from Boros, Kailath and Olshevsky [3] and the third from Gohberg, Kailath and Olshevsky [6].

P1 is a Cauchy-like matrix with $r = 2$, and nodes $t_i = a + ib$, $s_j = jb$ for $a = 1$ and $b = 2$. It is an example of well-conditioned Cauchy-like matrix; in fact, for $n = 512$, its condition number (estimated with the Matlab® function **condest**) is 4E+02, and for $n = 4096$ it is 1.3E+03.

P2 is the same matrix but with $a = 1$ and $b = -0.3$. It is an ill-conditioned Cauchy like matrix; in fact, the condition number estimate is 5E+20 for $n = 4096$ and 1E+17 for $n = 512$.

	n	$O(n^2)$-space GKO	Extended matrix	Downdating
C1	128	1.5072e-03	2.2062e-03	2.2340e-03
	256	7.9456e-03	7.6737e-03	7.9024e-03
	512	5.4915e-02	2.9391e-02	3.0240e-02
	1024	2.2892e-01	1.1660e-01	1.1922e-01
	2048	9.0878e-01	4.6298e-01	4.6911e-01
	4096	4.5026e+00	1.9070e+00	1.9039e+00
	8192	3.0390e+01	8.5737e+00	8.1583e+00
	16384	Out of memory	4.0430e+01	3.6570e+01
	32768	Out of memory	1.8971e+02	1.6531e+02
	65536	Out of memory	7.8461e+02	7.0309e+02

	n	$O(n^2)$-space GKO	Extended matrix	Downdating
C2	128	1.5786e-03	2.1707e-03	2.2363e-03
	256	6.2504e-03	8.1438e-03	8.1638e-03
	512	5.6890e-02	3.1951e-02	3.1486e-02
	1024	2.7656e-01	1.2755e-01	1.2406e-01
	2048	1.0801e+00	5.1083e-01	5.0123e-01
	4096	4.3176e+00	2.0409e+00	1.9748e+00
	8192	Out of memory	8.2725e+00	7.9705e+00
	16384	Out of memory	3.7706e+01	3.3402e+01
	32768	Out of memory	1.8448e+02	1.5066e+02
	65536	Out of memory	8.9376e+02	7.3051e+02

Table 1: Speed experiments: CPU times for the solution of a linear system with the different algorithms

	n	$O(n^2)$-space GKO	Extended matrix	Downdating
P1	128	1.262497e-15	1.160020e-15	1.062489e-15
	256	1.520695e-15	1.812447e-15	1.463218e-15
	512	2.979162e-15	3.063677e-15	3.091645e-15
	1024	2.790466e-15	3.429299e-15	3.068041e-15
	2048	4.568803e-15	5.921849e-15	5.044874e-15
	4096	5.231503e-15	7.448194e-15	5.461259e-15
	8192	7.491095e-15	1.250913e-14	7.287788e-15
	16384	Out of memory	1.648221e-14	1.154215e-14
	32768	Out of memory	2.624266e-14	1.757211e-14
	65536	Out of memory	3.929339e-14	2.209921e-14

	n	$O(n^2)$-space GKO	Extended matrix	Downdating
P2	128	4.226744e-05	4.226745e-05	4.226745e-05
	256	2.498321e-03	2.498321e-03	2.498321e-03
	512	1.307574e-01	1.307574e-01	1.307574e-01
	1024	1.634538e+01	1.634538e+01	1.634538e+01
	2048	4.367616e+02	4.367616e+02	4.367616e+02
	4096	2.311074e+04	2.311075e+04	2.311075e+04

	n	$O(n^2)$-space GKO	Extended matrix	Downdating
P3	128	5.691225e-08	3.039562e-08	4.267730e-08
	256	4.777684e-07	2.600382e-07	2.036880e-07
	512	1.176901e-07	9.031807e-08	1.163921e-07
	1024	1.247047e-07	1.213764e-07	1.266929e-07
	2048	1.854518e-07	1.817114e-07	1.873034e-07
	4096	2.799968e-07	2.750396e-07	2.733141e-07
	8192	4.180990e-07	4.028210e-07	3.834197e-07

Table 2: Relative forward errors

P3 is the Cauchy-like matrix arising from the (complex) Fourier transform of the Gaussian Toeplitz matrix [6] with $a = 0.9$, i.e., the Cauchy-like system to be solved when applying the GKO algorithm to said Toeplitz matrix. It is an interesting test case, since it is a matrix for which the Levinson-based Toeplitz solvers are unstable. It is a symmetric, moderately ill-conditioned Cauchy-like matrix with $r = 2$. Its condition number estimate is 7E+10 for $n = 512$ and 1E+11 for $n = 4096$.

For each experiment, we chose several different values of n, and generated the corresponding matrix C. We then computed $v = Ce$, with $e = [11\ldots1]^T$, and applied the old and new GKO algorithms to solve the system $Cx = v$. We computed the relative error as

$$err = \frac{\|x - e\|}{\|e\|},$$

where $\|\cdot\|$ is the Euclidean norm. The results are shown in Table 2

Comments There are no significant differences in the accuracy of the three algorithms. This shows that, at least in our examples, despite the larger number of operations needed, the space-efficient algorithms are as stable as the original GKO algorithm.

We point out that a formal stability proof of the GKO algorithm cannot be established, since it is ultimately based on Gaussian elimination with partial pivoting, for which counterexamples to stability exist, and since in some limit cases there are other issues such as generators growth [15]. However, both computational practice and theoretical analysis suggest that the GKO algorithm is in practice a reliable algorithm [13]. Several pivoting strategies, such as the one proposed by Gu [7], exist in order to avoid generator growth, and they can be applied to both the original GKO algorithm and its space-efficient versions.

8 Conclusions

In this paper, we proposed a new $O(n)$-space version of the GKO algorithm for the solution of Cauchy-like linear systems. Despite the larger number of operations needed, this algorithms succeeds in making a better use of the internal cache memory of the processor, thus providing an improvement with respect to both the customary GKO algorithm and a similar $O(n)$-space algorithm proposed by Rodriguez [14, 1]. Starting from $n \approx 500 - 1000$, the algorithm outperforms the existing ones. When applying this algorithm to special cases such as the Trummer-like matrices, several small optimizations reduce the number of steps needed to make an integrated algorithm to perform the inversion of a Trummer-like matrix represented in terms of its generators. We provided a ready-to-use Matlab® implementation of all the proposed algorithms.

Acknowledgements The author would like to thank Antonio Aricò, Dario Bini and Vadim Olshevsky for helping him with many useful discussions on the subject of the paper.

References

[1] Antonio Aricò and Giuseppe Rodriguez. A fast solver for linear systems with displacement structure. in preparation, 2009.

[2] Dario A. Bini, Bruno Iannazzo, and Federico Poloni. A fast Newton's method for a nonsymmetric algebraic Riccati equation. *SIAM J. Matrix Anal. Appl.*, 30(1):276–290, 2008.

[3] Tibor Boros, Thomas Kailath, and Vadim Olshevsky. Pivoting and backward stability of fast algorithms for solving Cauchy linear equations. *Linear Algebra Appl.*, 343/344:63–99, 2002. Special issue on structured and infinite systems of linear equations.

[4] Matteo Frigo, Charles E. Leiserson, Harald Prokop, and Sridhar Ramachandran. Cache-oblivious algorithms. *Foundations of Computer Science, Annual IEEE Symposium on*, 0:285, 1999.

[5] A. Gerasoulis. A fast algorithm for the multiplication of generalized Hilbert matrices with vectors. *Math. Comp.*, 50(181):179–188, 1988.

[6] I. Gohberg, T. Kailath, and V. Olshevsky. Fast Gaussian elimination with partial pivoting for matrices with displacement structure. *Math. Comp.*, 64(212):1557–1576, 1995.

[7] Ming Gu. Stable and efficient algorithms for structured systems of linear equations. *SIAM J. Matrix Anal. Appl.*, 19(2):279–306 (electronic), 1998.

[8] Georg Heinig and Karla Rost. *Algebraic methods for Toeplitz-like matrices and operators*, volume 13 of *Operator Theory: Advances and Applications*. Birkhäuser Verlag, Basel, 1984.

[9] John L. Hennessy and David A. Patterson. *Computer Architecture: A Quantitative Approach (The Morgan Kaufmann Series in Computer Architecture and Design)*. Morgan Kaufmann, May 2002.

[10] T. Kailath and J. Chun. Generalized displacement structure for block-Toeplitz, Toeplitz-block, and Toeplitz-derived matrices. *SIAM J. Matrix Anal. Appl.*, 15(1):114–128, 1994.

[11] Thomas Kailath and Vadim Olshevsky. Diagonal pivoting for partially reconstructible Cauchy-like matrices, with applications to Toeplitz-like linear equations and to boundary rational matrix interpolation problems. In *Proceedings of the Fifth Conference of the International Linear Algebra Society (Atlanta, GA, 1995)*, volume 254, pages 251–302, 1997.

[12] Thomas Kailath and Ali H. Sayed. Displacement structure: theory and applications. *SIAM Rev.*, 37(3):297–386, 1995.

[13] Vadim Olshevsky. Pivoting for structured matrices and rational tangential interpolation. In *Fast algorithms for structured matrices: theory and applications (South Hadley, MA, 2001)*, volume 323 of *Contemp. Math.*, pages 1–73. Amer. Math. Soc., Providence, RI, 2003.

[14] G. Rodriguez. Fast solution of Toeplitz- and Cauchy-like least-squares problems. *SIAM J. Matrix Anal. Appl.*, 28(3):724–748 (electronic), 2006.

[15] D. R. Sweet. Error analysis of a fast partial pivoting method for structured matrices. In *Proceedings SPIE, Volume 2563, Advanced Signal Processing Algorithms SPIE*, pages 266–280, 1995.

[16] Marc Van Barel, Georg Heinig, and Peter Kravanja. A stabilized superfast solver for nonsymmetric Toeplitz systems. *SIAM J. Matrix Anal. Appl.*, 23(2):494–510 (electronic), 2001.

[17] James S. Walker. *Fast Fourier transforms*. Studies in Advanced Mathematics. CRC Press, Boca Raton, FL, second edition, 1996. With 1 IBM-PC floppy disk (3.5 inch; HD).

```
function x=tr_sv(G,B,s,d,b)
%solves Tx=b, where diag(s)*T-T*diag(s)=G*B and diag(T)=b
n=size(G,1);
l=zeros(n,1);
u=zeros(1,n);
p=[1:n]'; %to keep track of the pivoting of s
x=b;
for k=1:n
    %finds pos=pivot row
    l(k)=d(k);
    l(k+1:n) = (G(k+1:n,:)*B(:,k))./(s(p(k+1:n))-s(k));
    [abspiv pos]=max(abs(l(k:n)));
    %pivots
    pos=pos+k-1;
    l([k pos])=l([pos k]);
    x([k pos],:)=x([pos k],:);
    G([k pos],:)=G([pos k],:);
    p([k pos])=p([pos k]);
    pivot=l(k);
    %LU step
    %there may be an off-diagonal non-reconstructible entry in u
    u(k+1:pos-1)=(G(k,:)*B(:,k+1:pos-1))./transpose(s(p(k))-s(k+1:pos-1));
    u(pos)=d(pos);
    u(pos+1:n)=(G(k,:)*B(:,pos+1:n))./transpose(s(p(k))-s(pos+1:n));
    x(k+1:n,:) = x(k+1:n,:) - l(k+1:n) * (x(k,:)/pivot);
    G(k+1:n,:) = G(k+1:n,:) - l(k+1:n) * (G(k,:)/pivot);
    B(:,k+1:n) = B(:,k+1:n) - (B(:,k)/pivot) * u(k+1:n);
    %Gaussian elimination on the diagonal
    d(k+1:n)=d(k+1:n)-l(k+1:n).*(transpose(u(k+1:n))/pivot);
    d(k)=pivot;
    %d(pos) may be reconstructible now -- but we store it anyway
    if (pos~=k)
        d(pos)=(G(pos,:)*B(:,pos))./(s(p(pos))-s(pos));
    end
end
%now x contains L\b, downdating phase
for k=n:-1:1
    pivot=d(k);
    u(k+1:n)=(G(k,:)*B(:,k+1:n))./transpose(s(k)-s(k+1:n));
    B(:,k+1:n) = B(:,k+1:n) + (B(:,k) / pivot) * u(k+1:n);
    %back-substitution
    x(k,:)=(x(k,:)-u(k+1:n)*x(k+1:n,:)) / pivot;
end
end
```

Algorithm 5: Solving a system $Tx = b$ with the downdating algorithm

```
function [G,B,s,d,x,y]=tr_inv(G,B,s,d,b,c)
%given the representation of T, i.e., G,B,s,d such that
%diag(s)*T-T*diag(s)=G*B, diag(T)=d,
%computes the representation of inv(T).
%If b and c are given, also solves Tx=b and yT=c
n=size(G,1);l=zeros(n,1);u=zeros(1,n);
if(nargin<6) c=zeros(0,n);end
if(nargin<5) b=zeros(n,0);end
x=b; y=c; p=[1:n]';
for k=1:n
    l(1:k-1)=(G(1:k-1,:)*B(:,k)) ./ (s(1:k-1)-s(k));
    l(k)=d(k);
    l(k+1:n)=(G(k+1:n,:)*B(:,k)) ./ (s(p(k+1:n))-s(k));
    %finds pos=pivot row (among k:n)
    [abspiv, pos]=max(abs(l(k:n)));
    if(abspiv==0) error('singular matrix');end
    %pivots
    pos=pos+(k-1);
    G([k pos],:)=G([pos k],:);
    x([k pos],:)=x([pos k],:);
    p([k pos])=p([pos k]);
    pivot=l(pos); l(pos)=l(k);
    %LU step
    %there may be an off-diagonal non-reconstructible entry in u
    u(1:k-1)=(G(k,:)*B(:,1:k-1)) ./ transpose(s(p(k))-s(p(1:k-1)));
    u(k+1:pos-1)=(G(k,:)*B(:,k+1:pos-1))./transpose(s(p(k))-s(k+1:pos-1));
    u(pos)=d(pos);
    u(pos+1:n)=(G(k,:)*B(:,pos+1:n))./transpose(s(p(k))-s(pos+1:n));

    l(k) = 0; u(k) = 0;
    G(k,:) = G(k,:)/pivot;
    G = G-l*G(k,:);
    x(k,:) = x(k,:)/pivot;
    x = x-l*x(k,:);
    B(:,k) = B(:,k)/pivot;
    B=B-B(:,k)*u;
    y(:,k) = y(:,k)/pivot;
    y = y-y(:,k)*u;
    %Gaussian elimination on the diagonal
    d(k+1:n)=d(k+1:n)-l(k+1:n).*u(1,k+1:n).'/pivot;
    %d(pos) may be reconstructible now -- but we store it anyway
    if (pos~=k)
        d(pos)=(G(pos,:)*B(:,pos))./(s(p(pos))-s(pos));
    end
    %overwrites d with the diagonal of the inverse
    l(k)=-1;u(k)=-1;d(k)=0;
    %u(1:k) contains -inv(L)(k,p(:)), l(1:k) contains -inv(U)(:,k)*pivot
    u(k+1:n)=0;
    u(p)=u;
    d(1:k)=d(1:k) + (l(1:k)/pivot) .* u(1:k).';
end
%as we pivoted on rows, now instead of y we have y*P
y(:,p)=y(:,:); B(:,p)=-B(:,:);
end
```

Algorithm 6: Computing the representation of T^{-1} (and solving systems with matrix T and T^*)