# A SUPERFAST STRUCTURED SOLVER FOR TOEPLITZ LINEAR SYSTEMS VIA RANDOMIZED SAMPLING* 

JIANLIN XIA ${ }^{\dagger}$, YUANZHE XI ${ }^{\dagger}$, AND MING $\mathrm{GU}^{\ddagger}$


#### Abstract

We propose a superfast solver for Toeplitz linear systems based on rank structured matrix methods and randomized sampling. The solver uses displacement equations to transform a Toeplitz matrix $T$ into a Cauchy-like matrix $\mathcal{C}$, which is known to have low-numerical-rank offdiagonal blocks. Thus, we design a fast scheme for constructing a hierarchically semiseparable (HSS) matrix approximation to $\mathcal{C}$, where the HSS generators have internal structures. Unlike classical HSS methods, our solver employs randomized sampling techniques together with fast Toeplitz matrixvector multiplications, and thus converts the direct compression of the off-diagonal blocks of $\mathcal{C}$ into the compression of much smaller blocks. A strong rank-revealing QR factorization method is used to generate/preserve certain special structures, and also to ensure stability. A fast ULV HSS factorization scheme is provided to take advantage of the special structures. We also propose a precomputation procedure for the HSS construction so as to further improve the efficiency. The complexity of these methods is significantly lower than some similar Toeplitz solvers for large matrix size $n$. Detailed flop counts are given, with the aid of a rank relaxation technique. The total cost of our methods includes $O(n)$ flops for HSS operations and $O\left(n \log ^{2} n\right)$ flops for matrix multiplications via FFTs, where $n$ is the order of $T$. Various numerical tests on classical examples, including ill-conditioned ones, demonstrate the efficiency, and also indicate that the methods are stable in practice. This work shows a practical way of using randomized sampling in the development of fast rank structured methods.


Key words. superfast Toeplitz solver, hierarchically semiseparable matrix, randomized sampling, structure-preserving rank-revealing factorization, precomputation

AMS subject classifications. 65F05, 65F30, 15A06
DOI. 10.1137/110831982

1. Introduction. In this work, we consider the solution of Toeplitz linear systems. Toeplitz systems arise in many numerical and engineering applications, such as PDE and integral equation solutions, image and signal processing, time series analysis, orthogonal polynomials, etc. A Toeplitz system looks like

$$
\begin{equation*}
T x=b, \tag{1.1}
\end{equation*}
$$

where $T$ is a Toeplitz matrix in the following form:

$$
T=\left(\begin{array}{ccccc}
t_{0} & t_{-1} & t_{-2} & \cdots & t_{-(n-1)}  \tag{1.2}\\
t_{1} & t_{0} & t_{-1} & \ddots & \vdots \\
t_{2} & t_{1} & t_{0} & \ddots & t_{-2} \\
\vdots & \ddots & \ddots & \ddots & t_{-1} \\
t_{n-1} & \cdots & t_{2} & t_{1} & t_{0}
\end{array}\right) .
$$

[^0]That is, the entries along every diagonal of $T$ are constant. $T$ is defined by a vector of $2 n-1$ parameters $t_{-(n-1): n-1}$, called the Toeplitz vector, that generates $T$.

The system can be solved quickly. There are fast solvers that need about $O\left(n^{2}\right)$ flops, such as Schur-type and Levinson-type methods [28]. Methods requiring less than $O\left(n^{2}\right)$ costs (or roughly $O(n)$ costs) are called superfast solvers. However, some fast and superfast algorithms are unstable. See [4, 6, 14, 41] for some discussions. Stable generalized Schur algorithms [13, 27] and look-ahead algorithms [8] are proposed to enhance the numerical stability. Superfast solvers with stabilization also exist [44]. Preconditioning techniques are first studied in $[35,39]$. See [7, 34] and a survey paper [9], as well as the references therein. Fast displacement-equation type algorithms in $[15,17,21,22,23]$ are numerically or analytically shown to be stable. Some of these methods transform $T$ into a Cauchy-like matrix $\mathcal{C}$ using displacement structures [26] and then solve $\mathcal{C}$.
1.1. Displacement structures and a low-rank property. Displacement structures are first introduced in [26]. Selected detailed studies can be found in $[15,21,24,36,38]$. It is straightforward to verify that the following Sylvester-type displacement equation for $T$ holds:

$$
\begin{equation*}
Z_{1} T-T Z_{-1}=G H^{T} \tag{1.3}
\end{equation*}
$$

where $G \in \mathbb{R}^{n \times d}, H \in \mathbb{R}^{n \times d}$, and $Z_{t}=\left(\begin{array}{cc}0 & t \\ I_{n-1} & 0\end{array}\right)$, with $I_{n-1}$ the size $n-1$ identity matrix. Here, $d=2$ is called the displacement rank. Since $Z_{1}$ is a circulant matrix, it can be diagonalized by the normalized inverse discrete Fourier transform matrix

$$
\begin{equation*}
\mathcal{F}_{n}=\frac{1}{\sqrt{n}}\left(\omega^{2(i-1)(j-1)}\right)_{1 \leq i, j \leq n}, \quad \omega=e^{\frac{\pi \mathbf{i}}{n}}, \quad \mathbf{i}=\sqrt{-1} \tag{1.4}
\end{equation*}
$$

This leads to another displacement equation [21],

$$
\begin{equation*}
\mathcal{D}_{1} \mathcal{C}-\mathcal{C} \mathcal{D}_{-1}=\hat{G} \hat{H}^{T} \tag{1.5}
\end{equation*}
$$

where

$$
\begin{gather*}
\mathcal{C}=\mathcal{F} T \mathcal{D}_{0}^{H} \mathcal{F}^{H}, \quad \hat{G}=\mathcal{F} G, \quad \hat{H}=\mathcal{F} \mathcal{D}_{0} H  \tag{1.6}\\
\mathcal{D}_{0}=\operatorname{diag}\left(1, \omega, \ldots, \omega^{n-1}\right), \quad \mathcal{D}_{1}=\operatorname{diag}\left(1, \omega^{2}, \ldots, \omega^{2(n-1)}\right), \quad \mathcal{D}_{-1}=\omega \mathcal{D}_{1} .
\end{gather*}
$$

Equation (1.5) indicates that $\mathcal{C}$ is a Cauchy-like matrix in the following form:

$$
\begin{equation*}
\mathcal{C}=\mathcal{F}_{n} T \mathcal{D}_{0}^{H} \mathcal{F}_{n}^{H}=\left(\frac{\hat{G}_{i} \hat{H}_{j}^{T}}{\omega^{2(i-1)}-\omega^{2 j-1}}\right)_{1 \leq i, j \leq n} \tag{1.7}
\end{equation*}
$$

where $\hat{G}_{i}$ is the $i$ th row of $\hat{G}$, and $\hat{H}_{j}$ is the $j$ th row of $\hat{H}$. Thus with FFTs, the system (1.1) can be converted into a Cauchy-like system

$$
\begin{equation*}
\mathcal{C} \mathrm{x}=\mathbf{b} \tag{1.8}
\end{equation*}
$$

where $\mathbf{x}=\mathcal{F}_{n} \mathcal{D}_{0} x$ and $\mathbf{b}=\mathcal{F}_{n} b$. Such a strategy for Toeplitz solutions is first used in [21].

The matrix $\mathcal{C}$ and similar representations have a low-rank property [12, 33]. That is, its off-diagonal blocks have small numerical ranks for a given tolerance, which are bounded by $O(\log n)$. Based on this property, two superfast and numerically stable

Toeplitz solvers are developed in [12, 33]. The solvers compress the off-diagonal blocks of $\mathcal{C}$ or its variations and approximate $\mathcal{C}$ by rank structured matrices. (Here by compression, we mean the computation of a truncated SVD or a rank-revealing factorization.) The method in [12] first uses $O\left(n^{2} \log n\right)$ flops to approximate the following Cauchy matrix in a precomputation stage:

$$
\begin{equation*}
\mathbf{C}=\left(\frac{1}{\omega^{2(i-1)}-\omega^{2 j-1}}\right)_{1 \leq i, j \leq n} \tag{1.9}
\end{equation*}
$$

Then a rank structured approximation to $\mathcal{C}$, called a sequentially semiseparable (SSS) matrix [11], can be constructed with about $O\left(n \log ^{2} n\right)$ cost. After this, (1.1) can be solved quickly with about $O(n \log n)$ cost. The method in [33] has similar complexity.
1.2. Main results. The methods in [12, 33] need $O\left(n^{2} \log n\right) \sim O\left(n^{2} \log ^{2} n\right)$ flops for precomputations, and then $O\left(n \log ^{2} n\right)$ flops for the structure construction and solution for (1.8). They use only the low-rank property of $\mathcal{C}$ and ignore its Cauchylike structure. Here, we consider further taking advantage of this special Cauchy-like structure (in fast matrix-vector multiplications) as well as special matrix forms (e.g., (3.10)-(3.11)), so as to reduce the costs of rank structured operations to $O(n)$ after an $O\left(n \log ^{2} n\right)$ cost for FFTs. We achieve this from several aspects.

One aspect is that we use randomized sampling together with fast Toeplitz matrixvector multiplications. The ideas of randomized sampling for constructing rank structured approximations have been extensively studied, and are shown to be very useful in dealing with rank deficiency $[29,30,31,37]$ (and others). As one example, to compress a low-rank block $\Phi$, multiply a random matrix $X$ to $\Phi$. The product $\Phi X$ is then compressed with rank-revealing factorizations, which gives the compression information of $\Phi[29]$. For a given matrix with small off-diagonal ranks, the method in [31] uses $O(1)$ such matrix multiplications to approximate the matrix by a rank structured form called hierarchically semiseparable (HSS) matrix [10, 46]. An HSS matrix is a data-sparse form given by a sequence of internal small dense matrices (called generators). The method in [30] uses $O(\log n)$ multiplications without knowing the entries of the original matrix. Here, it is well known that Toeplitz matrix-vector products can be quickly computed [16]. Together with (1.7), we can compute the product of the Cauchy-like matrix $\mathcal{C}$ with an $n \times O(\log n)$ Gaussian random matrix $X$ as needed here in $O\left(n \log ^{2} n\right)$ flops via FFTs. Then, we use a procedure similar to the one in [31] to approximate $\mathcal{C}$ by an HSS matrix.

Another aspect is that we use a special HSS form by incorporating additional structures into the HSS generators. The generators obtained from the off-diagonal block compression via randomized sampling has certain special structure (e.g., (3.5)). Unlike the method in [31] which uses an extra step to reorthonormalize an HSS form, we avoid this step to not only save the cost but also preserve the special structures. In fact, the special structures of the generators can be used to improve the HSS solution efficiency.

By taking advantage of both the Cauchy-like and additional structures, we gain significant improvements to classical HSS algorithms. Unlike some inversion-based algorithms, we propose special ULV-type [10, 46] HSS factorization and solution methods which are often more efficient. The complexity of all these algorithms is $O(n)$, and is significantly lower than both standard HSS methods [45] and the structured operations in the Toeplitz solvers in [12, 33]. We give detailed flop counts, which are not available in the Toeplitz solvers in $[12,33]$ or the randomized sampling methods in $[30,31]$. We also use an idea of rank relaxation [45] so that even if the off-diagonal
rank bound of $\mathcal{C}$ is $O(\log n)$, we can still achieve truly $O(n)$ complexity in the HSS operations, in practice, instead of $O(n \log n)$ or more. The usage of strong RRQR factorizations leads to a stable process, which is verified by numerical experiments. A variety of Toeplitz examples, especially ill-conditioned ones, demonstrate both the efficiency and the accuracy of our methods.

We also discuss a precomputation procedure to first compress the off-diagonal blocks of $\mathbf{C}$, which depends on $n$ only, instead of the entries of $T$. The information is then used in the HSS construction for $\mathcal{C}$ so as to further reduce or avoid the cost of multiplying $\mathcal{C}$ and the random matrix $X$. The total precomputation cost is only $O(n)$ (after FFTs), in contrast with $O\left(n^{2} \log n\right)$ or $O\left(n^{2} \log ^{2} n\right)$ in [12, 33]. This version is especially useful when there are multiple $T$ of the same large size.
1.3. Outline. The remaining sections are organized as follows. Section 2 provides brief discussions of HSS representations and the low-rank property of $\mathcal{C}$. The HSS construction, factorization, and solution procedures for $\mathcal{C}$ are presented in section 3. Section 4 then briefly describes an HSS construction method with precomputations. A variety of numerical examples are shown in section 5, and we draw some concluding remarks in section 6. The following notation is used in the presentation:

- Let $\mathbf{I}_{i}$ be a subset of $\{1: n\} \equiv\{1,2, \ldots, n\}$ with contiguous indices. We use $\mathbf{I}_{i}^{c}$ and $\mathbf{I}_{i}^{r}$ to denote the subsets of contiguous indices in $\{1: n\}$ that are smaller and larger than those in $\mathbf{I}_{i}$, respectively.
- $\left.A\right|_{\mathbf{I}_{i}}$ is the submatrix formed by all rows of $A$ with row index set $\mathbf{I}_{i}$. Also, $\left.A\right|_{\mathbf{I}_{i} \times \mathbf{I}_{j}}$ is the submatrix of $A$ with row index set $\mathbf{I}_{i}$ and column index set $\mathbf{I}_{j}$.
- $\operatorname{diag}\left(D_{1}, \ldots, D_{k}\right)$ is a diagonal matrix with diagonal blocks $D_{1}, \ldots, D_{k}$, and $\operatorname{diag}\left(v_{1: k}\right)$ is a diagonal matrix with diagonal entries $v_{1}, v_{2}, \ldots, v_{k}$.


## 2. Hierarchically semiseparable matrices and the low-rank property.

2.1. Review of hierarchically semiseparable representations. HSS representations enable us to conveniently take advantage of the low-rank property of the Cauchy-like matrix in (1.7). An $n \times n$ HSS matrix $A$ with postordering notation is defined hierarchically in terms of three components [46]:
(1) There is a full binary tree $\mathcal{T}$ with nodes labeled as $i=1,2, \ldots$, which are postordered. That is, any nonleaf node $i$ has a left child $c_{1}$ and a right child $c_{2}$, ordered as $c_{1}<c_{2}<i$.
(2) There is an index set $\mathbf{I}_{i}$ (with contiguous indices) associated with each node $i$ of $\mathcal{T}$. The index sets are defined so that $\mathbf{I}_{i}=\{1,2, \ldots, n\}$ for the root node $i=\operatorname{root}(\mathcal{T})$, and $\mathbf{I}_{c_{1}} \cup \mathbf{I}_{c_{2}}=\mathbf{I}_{i}, \mathbf{I}_{c_{1}} \cap \mathbf{I}_{c_{2}}=\emptyset$ for the children $c_{1}$ and $c_{2}$ of each nonleaf node $i$.
(3) There are a sequence of matrices $D_{i}, U_{i}, V_{i}, R_{i}, W_{i}, B_{i}$ (called HSS generators) associated with the nodes $i=1,2, \ldots$, so that for any nonleaf node $i$ and its children $c_{1}$ and $c_{2}$,

$$
\begin{gather*}
\left.D_{i} \equiv A\right|_{\mathbf{I}_{i} \times \mathbf{I}_{i}}=\left(\begin{array}{cc}
D_{c_{1}} & U_{c_{1}} B_{c_{1}} V_{c_{2}}^{T} \\
U_{c_{2}} B_{c_{2}} V_{c_{1}}^{T} & D_{c_{2}}
\end{array}\right)  \tag{2.1}\\
U_{i}=\left(\begin{array}{cc}
U_{c_{1}} & \\
& U_{c_{2}}
\end{array}\right)\binom{R_{c_{1}}}{R_{c_{2}}}, \quad V_{i}=\left(\begin{array}{cc}
V_{c_{1}} & \\
& V_{c_{2}}
\end{array}\right)\binom{W_{c_{1}}}{W_{c_{2}}} . \tag{2.2}
\end{gather*}
$$

Initially, the $U, V$ generators are defined for leaf nodes, and $R, W$ generators are defined for all nodes. $R, W$ are called translation operators and are used to recursively reconstruct the $U, V$ generators associated with $i$ from those
associated with $c_{1}$ and $c_{2}$ [10]. If $i=\operatorname{root}(\mathcal{T})$, then $D_{i} \equiv A$, and $U_{i}, V_{i}$, $R_{i}, W_{i}, B_{i}$ are set to be empty matrices (since $\operatorname{root}(\mathcal{T})$ is associated with a diagonal block which is the entire $A$, instead of any off-diagonal block).
Then $A$ is said to be in an $H S S$ form with the corresponding $H S S$ tree $\mathcal{T}$. See Figure 2.1 for an example. Note that we use matrix transposes $V^{T}$ in (2.1) even if $A$ is complex. This is convenient for our operations below.


Fig. 2.1. An HSS matrix and its HSS tree.


Fig. 2.2. The numerical ranks $r_{l}$ (versus the row sizes $N_{l}$ ) of the $H S S$ block rows $A_{i}^{-}$at level $l$ of the HSS partition for order $n$ Cauchy matrices $\mathbf{C}$ in (1.9), where the relative tolerance for computing the ranks is $\tau=10^{-8}$, the bottom level HSS block row size is $N_{l_{\max }}=20$, and for comparison purposes, we also plot shifted $\log N_{l}$ and $\log \log N_{l}$ curves.

An HSS form can be constructed by compressing certain off-diagonal blocks, called HSS blocks of $A$ associated with each node $i$ of $\mathcal{T}$ and defined as

$$
A_{i}^{-}=\left(\begin{array}{ll}
\left.A\right|_{\mathbf{I}_{i} \times \mathbf{I}_{i}^{c}} & \left.A\right|_{\mathbf{I}_{i} \times \mathbf{I}_{i}^{r}} \tag{2.3}
\end{array}\right), \quad A_{i}^{\prime}=\binom{\left.A\right|_{\mathbf{I}_{i}^{c} \times \mathbf{I}_{i}}}{\left.A\right|_{\mathbf{I}_{i}^{r} \times \mathbf{I}_{i}}} .
$$

Clearly, the columns of $U_{i}$ form a basis for the column space of the HSS block row $A_{i}^{-}$, and the rows of $V_{i}^{T}$ form a basis for the row space of the HSS block column $A_{i}^{1}$. We call the maximum (numerical) rank of all HSS blocks the HSS rank of $A$.
2.2. Low-rank property of C. According to [12] and section 4, the numerical rank of an HSS block of $\mathcal{C}$ in (1.7) is at most twice that of the corresponding HSS
block of $\mathbf{C}$ in (1.9). Thus to examine the low-rank property of $\mathcal{C}$, it is sufficient to look at $\mathbf{C}$. The HSS rank of $\mathbf{C}$ can be shown to be at most $O(\log n)$ [12, 33]. In Figure 2.2, we plot the numerical ranks of the HSS block rows of two Cauchy matrices C. Each matrix is hierarchically partitioned into $l_{\text {max }}$ levels of HSS blocks. At level $l=0,1, \ldots, l_{\text {max }}$, the HSS block has row size $N_{l}=n / 2^{l}$ and numerical rank $r_{l}$. Clearly, when $l$ decreases, $N_{l}$ doubles, but $r_{l}$ only increases slightly. In fact, Figure 2.2 indicates that $r_{l}$ increases much slower than $O\left(\log N_{l}\right)$ (it is roughly in the pattern of $O\left(\log \log N_{l}\right)$ in this computation). Although not yet analytically justified, this is observed to hold numerically, in practice. This observation is useful for the derivation of the truly linear complexity of our HSS construction and solution as in section 3.6.
3. Fast HSS construction and solution for the Cauchy-like system. In this section, we study the fast construction of an HSS approximation to $\mathcal{C}$ in (1.7). Let $r$ be the HSS rank of $\mathcal{C}$. Then $r=O(\log n)$. Unlike the HSS construction algorithm for a general dense matrix (which costs $O\left(r n^{2}\right)$ flops or more), here we need only $O\left(n \log ^{2} n\right)$ flops for $\mathcal{C}$. The main framework of our method is listed in Table 3.1.

TABLE 3.1
Overview of our superfast Toeplitz solver.

| Stage | Step | Section |
| :--- | :--- | :---: |
| $T$ to $\mathcal{C}$ | Conversion of $T$ to $\mathcal{C}$ | 1.1 |
| $\mathcal{C}$ to HSS approximate | Computation of $\mathcal{C} X$ and $\mathcal{C}^{T} X$ with random $X$ | 3.1 |
|  | Compression via randomized sampling | 3.2 |
|  | HSS construction for $\mathcal{C}$ with $\mathcal{C} X$ and $\mathcal{C}^{T} X$ | 3.3 |
| Factorization/solution | HSS factorization | 3.4 |
|  | HSS solution for $\mathcal{C} \mathbf{x}=\mathbf{b}$ | 3.5 |
|  | Recovery of the solution to $T x=b\left(x=\mathcal{D}_{0}^{-1} \mathcal{F}_{n}^{*} \mathbf{x}\right)$ | 1.1 |

Some major ideas used are as follows:
(1) A structure-preserving rank-revealing factorization via randomized sampling.
(2) Fast Toeplitz matrix-vector multiplication via circulant matrix-vector multiplication and FFT.
(3) HSS construction via hierarchical application of the compression and randomized sampling.
(4) Rank relaxation in the detailed analysis of the complexity.
3.1. Fast matrix-vector multiplication for $\mathcal{C}$. Our HSS construction for $\mathcal{C}$ needs the multiplication of $\mathcal{C}$ with an $n \times(r+\mu)$ Gaussian random matrix $X$, where $r$ is the HSS rank of $\mathcal{C}$, and $\mu$ is a small integer. (The values of $r+\mu$ in some practical examples are shown in section 5.) According to (1.7), we have

$$
\begin{equation*}
\mathcal{C} X=\mathcal{F}_{n} T\left(\mathcal{D}_{0}^{H} \mathcal{F}_{n}^{H} X\right) \tag{3.1}
\end{equation*}
$$

The product $\tilde{X}=\mathcal{D}_{0}^{H} \mathcal{F}_{n}^{H} X$ can be obtained efficiently and stably via FFT and diagonal scaling. Thus, we need to compute $T \tilde{X}$. This can be done quickly by extending the Toeplitz matrix to a circulant matrix (see, e.g., [16])

$$
\mathbf{T}=\left(\begin{array}{cc}
T & S \\
S & T
\end{array}\right)
$$

where $S$ is a Toeplitz matrix generated by the Toeplitz vector $\left(t_{1: n-1}, 0, t_{-(n-1):-1}\right)^{T}$. Let $v=\left(t_{0: n-1}, 0, t_{-(n-1):-1}\right)^{T}$. Then $\mathbf{T}$ can be diagonalized by the normalized inverse

DFT matrix $\mathcal{F}_{2 n}($ defined in (1.4)):

$$
\mathbf{T}=\mathcal{F}_{2 n}^{H} \operatorname{diag}\left(\mathcal{F}_{2 n} v\right) \mathcal{F}_{2 n}
$$

Thus,

$$
\begin{equation*}
\mathbf{T}\binom{\tilde{X}}{0}=\binom{T \tilde{X}}{S \tilde{X}}=\mathcal{F}_{2 n}^{H} \operatorname{diag}\left(\mathcal{F}_{2 n} t\right) \mathcal{F}_{2 n}\binom{\tilde{X}}{0} \tag{3.2}
\end{equation*}
$$

The last product can be quickly computed via FFTs and diagonal scaling, and its first half is $T \tilde{X}$. The total cost for computing $\mathcal{C} X$ with (3.1) and (3.2) is about $40(r+\mu) n \log n$ flops for FFTs. Note that such a fast multiplication scheme works since $\mathcal{C}$ is obtained from $T$, and is generally not applicable to an arbitrary Cauchy-like matrix.
3.2. Compression method: Structure-preserving rank-revealing factorization via randomized sampling. The construction of an HSS form usually involves hierarchical compression of the HSS blocks. Without loss of generality, assume $\Phi$ is an $M \times N$ block with numerical rank $\alpha$. Compute a strong rank-revealing RQ factorization [18]

$$
\begin{equation*}
\underset{M \times N}{\Phi} \approx \underset{M \times \alpha}{P} \quad \underset{\alpha \times N}{\Omega} \tag{3.3}
\end{equation*}
$$

where $P$ is a permutation matrix. Let

$$
R=\binom{R_{1}}{R_{2}} \begin{gathered}
\alpha \\
M-\alpha
\end{gathered}, \quad E=R_{2} R_{1}^{-1}
$$

To ensure stability, the permutations are done by increasing $\operatorname{det}\left(R_{1}\right)$ so that the entries of $|E|$ are small enough [18]. The increase of the determinant can be quickly detected.

Rewrite (3.3) as

$$
\begin{equation*}
\Phi \approx P\binom{R_{1}}{R_{2}} \Omega=P\binom{I}{E}\left(R_{1} \Omega\right)=\left.P\binom{I}{E} \Phi\right|_{\hat{\mathbf{I}}} \tag{3.4}
\end{equation*}
$$

where $\left.\Phi\right|_{\hat{\mathbf{I}}}$ is a submatrix of $\Phi$ with row indices $\hat{\mathbf{I}}$. This holds due to the permutation and identity matrices. To emphasize the fact that $\left.\Phi\right|_{\hat{\mathbf{I}}}$ preserves part of the structure of $\Phi$, we call (3.4) a structure-preserving rank-revealing (SPRR) factorization of $\Phi$. This factorization is closely related to skeleton approximations [43], and is also called interpolative decomposition [29, 31] or subset selection decomposition [25]. A similar decomposition called representative row/column selection is also studied in [19] based on strong rank-revealing LU factorizations.

To enhance the efficiency of the compression, especially when the size of $\Phi$ is large, (3.4) can be combined with random sampling [29,31]. Generate an $N \times(\alpha+\mu)$ Gaussian random matrix $X_{1}$, where $\mu$ is a small integer. Next, apply the SPRR factorization to the product of $\Phi$ and $X_{1}$, which is a much smaller matrix:

$$
Y_{1}=\left.\Phi X_{1} \approx P_{1}\binom{I}{E_{1}} Y_{1}\right|_{\hat{\mathbf{I}}_{1}}
$$

Then an approximate factorization of $A$ can be obtained as [29]

$$
\begin{equation*}
\left.\Phi \approx U_{1} \Phi\right|_{\hat{\mathbf{I}}_{1}} \text { with } U_{1}=P_{1}\binom{I}{E_{1}} \tag{3.5}
\end{equation*}
$$

It is shown in $[20,29,32]$ that, under very mild assumptions for $\mu$, the probability for the approximation error to satisfy the following relationship is $1-6 \mu^{-\mu}$ :

$$
\begin{equation*}
\left\|\Phi-\left.U_{1} \Phi\right|_{\hat{\mathbf{I}}_{1}}\right\|_{2} \leq(1+11 \sqrt{\alpha+\mu} \sqrt{\min (M, N)}) \sigma_{\alpha+1} \tag{3.6}
\end{equation*}
$$

where $\sigma_{\alpha+1}$ is the $(\alpha+1)$ st singular value of $\Phi$. For example, if $\mu=20$, then the probability of failure is less than $10^{-17}$. If, in addition, $M=N=1000, \alpha=100$, and $\sigma_{\alpha+1}=10^{-15}$, then the right-hand side in (3.6) is about $3.8 \times 10^{-12}$. In practice, the bound often overestimates the error. In our tests below, we use $\mu=10$ as in [31] and a relative tolerance $10^{-15}$ (for $\sigma_{1} / \sigma_{\alpha+1}$ ).

This SPRR factorization has more advantages other than structure preservation. It uses only matrix-vector multiplications without requiring the explicit entries of $\Phi$. If $\alpha \ll N$, the compression of $Y_{1}$ is significantly faster than that of $\Phi$.
3.3. HSS construction for $\mathcal{C}$. Our HSS construction algorithm has a scheme similar to the one in [31], but with additional structures fully considered and without using any extra reorthonormalization. Also, we use postordered traversal of a given HSS tree $\mathcal{T}$. Assume $r$ is the HSS rank of $\mathcal{C}$, and the HSS block row $i$ have row dimension $m_{i}$ and starting row index $l_{i}$. That is,

$$
l_{i}=\left\{\begin{array}{cl}
\sum_{j: \text { leaf, } j<i} m_{j}+1 & \text { if } i \text { is a leaf, }  \tag{3.7}\\
l_{c_{1}} & \text { otherwise, if } c_{1} \text { is the left child of } i
\end{array}\right.
$$

First, use the method in section 3.1 to compute

$$
\begin{equation*}
Y=\mathcal{C} X, \quad Z=\mathcal{C}^{T} X \tag{3.8}
\end{equation*}
$$

(In [31], two Gaussian random matrices are generated. Here, one is sufficient.)
Next, we show the HSS construction process following the traversal of the nodes $i=1,2, \ldots$ of $\mathcal{T}$. We compress the HSS blocks hierarchically to get the HSS generators $U_{i}, V_{i}, R_{i}, W_{i}$, and $B_{i}$.

If $i$ is a leaf, let $\left.X_{i} \equiv X\right|_{\mathbf{I}_{i}},\left.Y_{i} \equiv Y\right|_{\mathbf{I}_{i}}$, and $\left.Z_{i} \equiv Z\right|_{\mathbf{I}_{i}}$. Set $\left.D_{i} \equiv \mathcal{C}\right|_{\mathbf{I}_{i} \times \mathbf{I}_{i}}$. Then we compress $\mathcal{C}_{i}^{-}$and $\mathcal{C}_{i}^{l}$ with SPRR factorizations and randomized sampling. We need to form the products $\left.\Phi_{i} \equiv \mathcal{C}_{i}^{-} X\right|_{\left(\mathbf{I}_{i}^{c} \cup \mathbf{I}_{i}^{r}\right)}$ and $\left.\Theta_{i} \equiv\left(\mathcal{C}_{i}^{1}\right)^{T} X\right|_{\left(\mathbf{I}_{i}^{c} \cup \mathbf{I}_{i}^{r}\right)}$ for compression. The $i$ th block rows of $\mathcal{C} X$ and $\mathcal{C}^{T} X$ are

$$
D_{i} X_{i}+\left.\mathcal{C}_{i}^{-} X\right|_{\left(\mathbf{I}_{i}^{c} \cup \mathbf{I}_{i}^{r}\right)}=Y_{i} \quad \text { and } \quad D_{i}^{T} X_{i}+\left.\left(\mathcal{C}_{i}^{\mid}\right)^{T} X\right|_{\left(\mathbf{I}_{i}^{c} \cup \mathbf{I}_{i}^{r}\right)}=Z_{i}
$$

respectively. This leads to

$$
\begin{equation*}
\Phi_{i}=Y_{i}-D_{i} X_{i}, \quad \Theta_{i}=Z_{i}-D_{i}^{T} X_{i} \tag{3.9}
\end{equation*}
$$

See Figure 3.1(i).
Thus, to compress $\mathcal{C}_{i}^{-}$and $\mathcal{C}_{i}^{\mid}$, we compute SPRR factorizations

$$
\begin{align*}
& \left.\Phi_{i} \approx U_{i} \Phi_{i}\right|_{\hat{\mathbf{I}}_{i}} \text { with } U_{i}=P_{i}\binom{I}{E_{i}}  \tag{3.10}\\
& \Theta_{i}-r  \tag{3.11}\\
& \left.\Theta_{i} \approx V_{i} \Theta_{i}\right|_{\hat{\mathbf{J}}_{i}} \text { with } V_{i}=Q_{i}\binom{I}{F_{i}} \begin{array}{c}
r \\
m_{i}-r
\end{array}
\end{align*}
$$

(To simplify the notation, $r$ is used to denote the numerical ranks of all HSS blocks.) Then

$$
\begin{equation*}
\left.\mathcal{C}_{i}^{-} \approx U_{i} \mathcal{C}_{i}^{-}\right|_{\hat{\mathbf{I}}_{i}},\left.\quad\left(\mathcal{C}_{i}^{\mid}\right)^{T} \approx V_{i}\left(\mathcal{C}_{i}^{\mid}\right)^{T}\right|_{\hat{\mathbf{J}}_{i}} \tag{3.12}
\end{equation*}
$$



Fig. 3.1. How $\Phi_{i}=\left.\mathcal{C}_{i}^{-} X\right|_{\left(\mathbf{I}_{i}^{c} \cup \mathbf{I}_{i}^{r}\right)}$ is obtained when node $i$ is a leaf node or a nonleaf node, where the index sets satisfy $\mathbf{I}_{i}^{c} \cup \mathbf{I}_{i} \cup \mathbf{I}_{i}^{r}=\{1,2, \ldots, n\}$.

Also, let

$$
\hat{Y}_{i}=V_{i}^{T} X_{i}=\left(\begin{array}{cc}
I & F_{i}^{T}
\end{array}\right) Q_{i}^{T} X_{i}, \quad \hat{Z}_{i}=U_{i}^{T} X_{i}=\left(\begin{array}{cc}
I & E_{i}^{T} \tag{3.13}
\end{array}\right) P_{i}^{T} X_{i}
$$

where the identity matrices (as additional structures) are used to save the multiplication cost. Since $\left.\mathcal{C}_{i}^{-}\right|_{\hat{\mathbf{I}}_{i}}$ and $\left.\left(\mathcal{C}_{i}^{\mid}\right)^{H}\right|_{\hat{\mathbf{J}}_{i}}$ in (3.12) are blocks in $\mathcal{C}$, we identify their global indices in $\mathcal{C}$ as

$$
\begin{equation*}
\tilde{\mathbf{I}}_{i}=l_{i}+\hat{\mathbf{I}}_{i}-1, \quad \tilde{\mathbf{J}}_{i}=l_{i}+\hat{\mathbf{J}}_{i}-1 \tag{3.14}
\end{equation*}
$$

If $i$ is a nonleaf node, assume $c_{1}$ and $c_{2}$ are the left and right children of $i$, respectively, which have been visited before. Let

$$
\begin{equation*}
B_{c_{1}}=\left.\mathcal{C}\right|_{\tilde{\mathbf{I}}_{c_{1}} \times \tilde{\mathbf{J}}_{c_{2}}}, \quad B_{c_{2}}=\left.\mathcal{C}\right|_{\tilde{\mathbf{I}}_{c_{2}} \times \tilde{\mathbf{J}}_{c_{1}}} . \tag{3.15}
\end{equation*}
$$

Due to recursion, all generators associated with $c_{1}$ and $c_{2}$ are available, and

$$
\begin{equation*}
Y_{c_{1}}-D_{c_{1}} X_{c_{1}}=\left.U_{c_{1}} \Phi_{c_{1}}\right|_{\mathbf{I}_{1}}, \quad Y_{c_{2}}-D_{c_{2}} X_{c_{2}}=\left.U_{c_{2}} \Phi_{c_{2}}\right|_{\mathbf{I}_{c_{2}}} \tag{3.16}
\end{equation*}
$$

For example, the first formula in (3.16) is a direct result of the steps associated with $c_{1},(3.10)$, and (3.18), together with (2.2). Therefore, we can ignore existing $U, V$ generators in our further compression as follows:

$$
\left.\begin{array}{rl} 
& \left.\mathcal{C}_{i}^{-} X\right|_{\left(\mathbf{I}_{i}^{c} \cup \mathbf{I}_{i}^{r}\right)}=Y_{i}-D_{i} X_{i} \\
\approx & \binom{Y_{c_{1}}}{Y_{c_{2}}}-\left(\begin{array}{cc}
D_{c_{1}} & U_{c_{1}} B_{c_{1}} V_{c_{2}}^{T} \\
U_{c_{2}} B_{c_{2}} V_{c_{1}}^{T} & D_{c_{2}}
\end{array}\right)\binom{X_{c_{1}}}{X_{c_{2}}} \\
= & \binom{\left(Y_{c_{1}}-D_{c_{1}} X_{c_{1}}\right)-U_{c_{1}} B_{c_{1}} V_{c_{2}}^{T} X_{c_{2}}}{\left(Y_{c_{2}}-D_{c_{2}} X_{c_{2}}\right)-U_{c_{2}} B_{c_{2}} V_{c_{1}}^{T} X_{c_{1}}} \approx\binom{\left.U_{c_{1}} \Phi_{c_{1}}\right|_{\mathbf{I}_{c_{1}}}-U_{c_{1}} B_{c_{1}} V_{c_{2}}^{T} X_{c_{2}}}{U_{c_{2}} \Phi_{c_{2}} \mid \mathbf{I}_{c_{2}}} U_{c_{2}} B_{c_{2}} V_{c_{1}}^{T} X_{c_{1}}
\end{array}\right) .
$$

See Figure 3.1(ii). Similarly,

$$
\left.\left(\mathcal{C}_{i}^{\mid}\right)^{T} X\right|_{\left(\mathbf{I}_{i}^{c} \cup \mathbf{I}_{i}^{r}\right)} \approx\left(\begin{array}{cc}
V_{c_{1}} & \\
& V_{c_{2}}
\end{array}\right)\binom{\Theta_{c_{1}} \mid \mathbf{I}_{c_{1}}-B_{c_{2}}^{T} \hat{Z}_{c_{1}}}{\Theta_{c_{2}} \mid \mathbf{I}_{c_{2}}-B_{c_{1}}^{T} \hat{Z}_{c_{2}}} .
$$

This means we need only compress

$$
\begin{equation*}
\Phi_{i} \equiv\binom{\Phi_{c_{1}} \mid \mathbf{I}_{c_{1}}-B_{c_{1}} \hat{Y}_{c_{2}}}{\Phi_{c_{2}} \mid \mathbf{I}_{c_{2}}-B_{c_{2}} \hat{Y}_{c_{1}}}, \quad \Theta_{i} \equiv\binom{\Theta_{c_{1}} \mid \mathbf{I}_{c_{1}}-B_{c_{2}}^{T} \hat{Z}_{c_{1}}}{\Theta_{c_{2}} \mid \mathbf{I}_{c_{2}}-B_{c_{1}}^{T} \hat{Z}_{c_{2}}} \tag{3.17}
\end{equation*}
$$

Compute SPRR factorizations

$$
\begin{align*}
\Phi_{i} & \left.\approx\binom{R_{c_{1}}}{R_{c_{2}}} \Phi\right|_{\hat{\mathbf{I}}_{i}} \text { with }\binom{R_{c_{1}}}{R_{c_{2}}}=P_{i}\binom{I}{E_{i}}  \tag{3.18}\\
\Theta_{i} & \left.\approx\binom{W_{c_{1}}}{W_{c_{2}}} \Theta\right|_{\hat{\mathbf{J}}_{i}} \text { with }\binom{W_{c_{1}}}{W_{c_{2}}}=Q_{i}\binom{I}{F_{i}} \tag{3.19}
\end{align*}
$$

Also, let

$$
\begin{align*}
& \hat{Y}_{i}=\left(\begin{array}{ll}
W_{c_{1}}^{T} & W_{c_{2}}^{T}
\end{array}\right)\binom{\hat{X}_{c_{1}}}{\hat{X}_{c_{2}}}=\left(\begin{array}{ll}
I & F_{i}^{T}
\end{array}\right) Q_{i}^{T}\binom{\hat{X}_{c_{1}}}{\hat{X}_{c_{2}}},  \tag{3.20}\\
& \hat{Z}_{i}=\left(\begin{array}{ll}
R_{c_{1}}^{T} & R_{c_{2}}^{T}
\end{array}\right)\binom{\hat{Z}_{c_{1}}}{\hat{Z}_{c_{2}}}=\left(\begin{array}{ll}
I & E_{i}^{T}
\end{array}\right) P_{i}^{T}\binom{\hat{Z}_{c_{1}}}{\hat{Z}_{c_{2}}},
\end{align*}
$$

where, again, the identity matrices are useful in saving the multiplication cost.
Similarly, use (3.14) to update the indices. Then the process repeats for the nodes in $\mathcal{T}$. The algorithm is summarized as follows.

Algorithm 1 (HSS construction for $\mathcal{C}$ ).

- Compute the products (3.8)
- for nodes $i=1,2, \ldots$ of $\mathcal{T}$
if $i$ is a leaf
- Set $\left.D_{i} \equiv \mathcal{C}\right|_{\mathbf{I}_{i} \times \mathbf{I}_{i}}$
- Compute $\Phi_{i}$ and $\Theta_{i}$ in (3.9)
- Compute $U_{i}$ in (3.10) and $V_{i}$ as in (3.11)
else
- Obtain $B_{c_{1}}, B_{c_{2}}$ as in (3.15)
- Compute $\Phi_{i}$ and $\Theta_{i}(3.17)$
- Compute $R_{c_{1}}, R_{c_{2}}$ in (3.18) and $W_{c_{1}}, W_{c_{2}}$ as in (3.19)
end
Remark 1. Unlike the method in [31], here we do not further reorthonormalize the HSS form. That is, we maintain the forms of $U_{i}, V_{i}$ in (3.10)-(3.11) and $R_{i}, W_{i}$ in (3.18)-(3.19), not only to avoid the extra cost, but also to preserve the internal structures (permutations and identity matrices).
3.4. Fast ULV HSS factorization. For notational convenience, we do not distinguish between $\mathcal{C}$ and its HSS approximation. Here, we directly use the special forms of $U_{i}, V_{i}, R_{i}, W_{i}$ to compute a ULV-type factorization. (ULV represents a sequence of orthonormal and triangular factors. Here, the factors may be special structured matrices.) Standard ULV HSS factorization [10, 46] has three main steps:
(1) Introduce zeros into HSS block row $i$ by introducing zeros into $U_{i}$ or $R_{i}$.
(2) Partially factorize $D_{i}$.
(3) Merge remaining blocks and repeat.

The first step is often expensive. For example, it is usually done via a full QR factorization $U_{i}=\Omega_{i}\left(\stackrel{\sim}{U}_{i}\right)$. Then $\Omega_{i}^{H}$ is multiplied to the HSS block row $\mathcal{C}_{i}^{-}$so that
zeros are introduced into $\mathcal{C}_{i}^{-}$, or

$$
\begin{equation*}
\Omega_{i}^{H} U_{i}=\binom{0}{\tilde{U}_{i}} \tag{3.21}
\end{equation*}
$$

This also needs a dense matrix multiplication $\tilde{D}_{i}=\Omega_{i}^{H} D_{i}$. Here, since $U_{i}$ and $R_{i}$ have special forms as in (3.10) and (3.18), respectively, this step can be significantly improved. That is, we can avoid the QR step and also reduce the matrix multiplication cost. Our ULV factorization has a framework similar to those in [10, 46] (where pictorial representations are available), but is much faster with internal structured operations. Again, we traverse the tree $\mathcal{T}$ for nodes $i=1,2, \ldots$.

If $i$ is a leaf, according to (3.10), the following formula automatically introduces zeros into $U_{i}$ without any actual cost:

$$
\left[\left(\begin{array}{cc}
-E_{i} & I  \tag{3.22}\\
I & 0
\end{array}\right) P_{i}^{T}\right] U_{i}=\binom{0}{I}
$$

That is, in (3.21), we simply set

$$
\Omega_{i}^{H}=\left(\begin{array}{cc}
-E_{i} & I \\
I & 0
\end{array}\right) P_{i}^{T}, \quad \tilde{U}_{i}=I
$$

Then we also update the diagonal block $D_{i}$ as

$$
\tilde{D}_{i}=\left(\begin{array}{cc}
-E_{i} & I \\
I & 0
\end{array}\right) P_{i}^{T} D_{i}
$$

Note that this requires only the multiplication of $E_{i}$ with the pivot block of $P_{i}^{T} D_{i}$, which reduces the $\Omega_{i}^{H} D_{i}$ dense multiplication cost from $2 m_{i}^{3}$ to $2 r^{2}\left(m_{i}-r\right)$. Therefore, zeros are introduced into the $\operatorname{HSS}$ block row $\mathcal{C}_{i}^{-}$, and (3.12) is updated to

$$
\left.\left(\begin{array}{cc}
-E_{i} & I \\
I & 0
\end{array}\right) P_{i}^{T} \mathcal{C}_{i}^{-} \approx\binom{0}{I} \mathcal{C}_{i}^{-}\right|_{\hat{\mathbf{I}}_{i}}
$$

Then partition $\tilde{D}_{i}$ as $\tilde{D}_{i}=\left(\begin{array}{cc}\tilde{D}_{i ; 1,1} & \tilde{D}_{i ; 1,2} \\ \tilde{D}_{i ; 2,1} & \tilde{D}_{i ; 2,2}\end{array}\right)_{r}^{m_{i}-r}$ and compute a QR factorization

$$
\binom{\tilde{D}_{i ; 1,1}^{T}}{\tilde{D}_{i ; 1,2}^{T}}=\tilde{Q}_{i}\binom{\hat{D}_{i ; 1,1}^{T}}{0}
$$

Update $\tilde{D}_{i}$ and the HSS block column $\mathcal{C}_{i}^{\mid}$by computing

$$
\hat{D}_{i}=\tilde{D}_{i} \tilde{Q}_{i}=\left(\begin{array}{cc}
L_{i} & 0 \\
\hat{D}_{i ; 2,1} & \hat{D}_{i ; 2,2}
\end{array}\right), \quad \tilde{V}_{i}=\tilde{Q}_{i}^{T} V_{i} \equiv\binom{\tilde{V}_{i ; 1}}{\tilde{V}_{i ; 2}}
$$

where $\tilde{V}_{i}$ is partitioned conformably. Therefore, $\hat{D}_{i ; 1,1}$ can be eliminated. Similarly, the blocks associated with the sibling node of $i$ can also be partially eliminated.

If $i$ is a nonleaf node, its children $c_{1}$ and $c_{2}$ can be removed from $\mathcal{T}$, and $i$ becomes a leaf with new generators

$$
D_{i}=\left(\begin{array}{cc}
\hat{D}_{c_{1} ; 2,2} & B_{c_{1}} \tilde{V}_{c_{2} ; 2}^{T} \\
B_{c_{2}} \tilde{V}_{c_{1} ; 2}^{T} & \hat{D}_{c_{2} ; 2,2}
\end{array}\right), \quad U_{i}=\binom{R_{c_{1}}}{R_{c_{2}}}, \quad V_{i}=\binom{\tilde{V}_{c_{1} ; 2} W_{c_{1}}}{\tilde{V}_{c_{2} ; 2} W_{c_{2}}}
$$

(Here, $D_{i}, U_{i}, V_{i}$ are used for notational convenience, and they are not the original generators associated with $i$.) This is the merge process. Note that, due to (3.18)(3.19), we have

$$
U_{i}=P_{i}\binom{I}{E_{i}}, \quad V_{i}=\left(\begin{array}{cc}
\tilde{V}_{c_{1} ; 2} & \\
& \tilde{V}_{c_{2} ; 2}
\end{array}\right) Q_{i}\binom{I}{F_{i}}
$$

The formation of $V_{i}$ can also take advantage of the identity matrix. Here again, due to the form $U_{i}$, we can apply the same procedure to node $i$ as in the step for a leaf. Thus, the fast elimination steps repeat, until we reach the root node where a direct LU factorization is used.
3.5. Fast ULV HSS solution. Similarly, the solution for (1.8) (with the ULV factors of the HSS approximation to $\mathcal{C}$ ) can also be quickly computed. The framework follows those in $[10,46]$. Initially, partition $\mathbf{b}$ into individual pieces $\mathbf{b}_{i}$ following all the leaf level $D_{i}$ sizes. The solution consists of two stages.

The first stage is a forward substitution process. For a leaf $i$, let

$$
\tilde{\mathbf{b}}_{i}=\left(\begin{array}{cc}
-E_{i} & I  \tag{3.23}\\
I & 0
\end{array}\right) P_{i}^{T} \mathbf{b}_{i} \equiv\binom{\tilde{\mathbf{b}}_{i, 1}}{\tilde{\mathbf{b}}_{i, 2}} \begin{aligned}
& m-r \\
& r
\end{aligned}, \quad \mathbf{y}_{i} \equiv L_{i}^{-1} \tilde{\mathbf{b}}_{i, 1}
$$

The multiplication for $\tilde{\mathbf{b}}_{i}$ can be quickly conducted. For a nonleaf node $i$ with children $c_{1}$ and $c_{2}$, let $\mathbf{b}_{i}=\binom{\tilde{\mathbf{b}}_{c_{1}, 2}}{\tilde{\mathbf{b}}_{c_{2}, 2}}$, and the above process (3.23) repeats for $i$. When $i=\operatorname{root}(\mathcal{T})$ is reached, we compute $\mathbf{y}_{i} \equiv D_{i}^{-1} \tilde{\mathbf{b}}_{i, 1}$ using the LU factorization of $D_{i}$.

The second stage is a backward substitution process. Let $\mathbf{x}_{\operatorname{root}(\mathcal{T})}=\mathbf{y}_{\operatorname{root}(\mathcal{T})}$. For any nonleaf node $i$ and its children $c_{1}$ and $c_{2}$, partition $\mathbf{x}_{i}$ as $\mathbf{x}_{i}=\binom{\mathbf{x}_{i, 1}}{\mathbf{x}_{i, 2}}_{r}^{r}$. Then compute

$$
\mathbf{x}_{c_{1}}=\tilde{Q}_{c_{1}}\binom{\mathbf{y}_{c_{1}}}{\mathbf{x}_{i, 1}}, \quad \mathbf{x}_{c_{2}}=\tilde{Q}_{c_{2}}\binom{\mathbf{y}_{c_{2}}}{\mathbf{x}_{i, 2}}
$$

This process then continues in a top-down traversal of $\mathcal{T}$. When it finishes, we just need to merge the $\mathbf{x}_{i}$ pieces associated with all leaves $i$ to form the solution $\mathbf{x}$.
3.6. Complexity, rank relaxation, and accuracy. To estimate the complexity, we list the flop counts of some basic matrix operations in Table 3.2. Here, we need only count the cost of the HSS block row compression, since the block column compression cost is the same. For simplicity, assume all HSS blocks have numerical ranks $r$ and $m_{i}=m=O(r)$. Also let $\tilde{r}=r+\mu$, which is close to $r$.

TABLE 3.2
Flop counts of some basic matrix operations, which can be found in, say, [16] or can be derived based on those. The RRQR factorization in our work is based on the modified Gram-Schmidt process with column pivoting [16]. The $S P R R$ factorization uses the $R R Q R$ factorization.

| Operation | Flops |
| :--- | :--- |
| Product of an $m \times q$ matrix and an $q \times r$ matrix | $2 m q r$ |
| RRQR factorization of an $m \times q$ matrix with rank $r$ | $4 m q r-2 r^{2}(m+q)+\frac{4}{3} r^{3}$ |
| SPRR factorization of an $m \times q$ matrix with rank $r$ | $4 m q r-r^{2}(m+2 q)+\frac{1}{3} r^{3}$ |
| Full QR factorization of an $m \times q$ tall matrix $(m>q)$ | $2 q^{2}\left(m-\frac{q}{3}\right)$ |
| Product of the $Q$ factor and an $m \times r$ matrix | $2 r q(2 m-q)$ |
| Solution of an order $m$ triangular system $L x=b$ | $k m^{2}$ |

The flop counts for the main steps of the HSS construction for $\mathcal{C}$ are listed as follows, with low-order terms dropped:

- Matrix multiplications:
- Computing the product $\mathcal{C} X: 40 \tilde{r} n \log n$.
- For each leaf node $i$,
* Computing $\Phi_{i}$ in (3.9): $2 m^{2} \tilde{r}$.
* Computing $\hat{Y}_{i}$ in (3.13): $2 r(m-r) \tilde{r}$.
- For each nonleaf node $i$,
* Computing $\Phi_{i}$ in (3.17): $4 r^{2} \tilde{r}$.
* Computing $\hat{Y}_{i}$ in (3.20): $2 r^{2} \tilde{r}$.
- SPRR factorizations:
- (3.10) for each leaf node $i$ : $4 m \tilde{r} r-r^{2}(m+2 \tilde{r})+\frac{1}{3} r^{3}$.
- (3.18) for each nonleaf node $i: 8 \tilde{r} r^{2}-r^{2}(2 r+2 \tilde{r})+\frac{1}{3} r^{3}$.

The total matrix multiplication cost (for both row and column operations) is thus

$$
\begin{equation*}
\xi_{0} \equiv \xi_{0,1}+\xi_{0,2} \tag{3.24}
\end{equation*}
$$

where $\xi_{0,1}$ is for computing $\mathcal{C} X$, and $\xi_{0,2}$ is for updating the products in the construction process, and they are given by

$$
\begin{aligned}
& \xi_{0,1}=80 \tilde{r} n \log n \\
& \xi_{0,2}=2\left[\sum_{i: \text { leaf }}\left(2 m^{2} \tilde{r}+2 r(m-r) \tilde{r}\right)+\sum_{i: \text { nonleaf }}\left(4 r^{2} \tilde{r}+2 r^{2} \tilde{r}\right)\right]=4 \tilde{r}\left(m+r+2 \frac{r^{2}}{m}\right) n
\end{aligned}
$$

The total HSS construction cost (for the compression and excluding $\xi_{0}$ ) is

$$
\begin{align*}
\xi_{1} & =2\left[\sum_{i: \text { leaf }}\left(4 m \tilde{r} r-r^{2}(m+2 \tilde{r})+\frac{1}{3} r^{3}\right)+\sum_{i: \text { nonleaf }}\left(8 \tilde{r} r^{2}-r^{2}(2 r+2 \tilde{r})+\frac{1}{3} r^{3}\right)\right]  \tag{3.25}\\
& \approx \frac{2}{3} r\left(12 \mu+9 r+8 \frac{r^{2}}{m}+12 \frac{r \mu}{m}\right) n .
\end{align*}
$$

Similarly, the ULV factorization and the solution costs are

$$
\xi_{2} \approx 2\left(\frac{4}{3} m^{2}+4 m r-2 r^{2}+16 \frac{r^{3}}{m}\right) n \quad \text { and } \quad \xi_{3} \approx 2\left(4 m+4 r+8 \frac{r^{2}}{m}\right) n
$$

respectively.
We then compare the performance of our methods with some similar ones from several aspects. First, if $m$ is chosen to be $2 r$ as often used in HSS methods [45], we get the costs as illustrated in Table 3.3. The methods are much more efficient than classical HSS methods.

$$
\text { Table } 3.3
$$

Flop counts of our new HSS algorithms for $\mathcal{C}$ when $m=2 r$, as compared with classical HSS algorithms, where the terms with $\mu$ are dropped since $\mu$ is a small constant.

|  | HSS construction | ULV factorization | ULV solution |
| :--- | :---: | :---: | :---: |
| Classical HSS | $6 r n^{2}$ | $42 r^{2} n$ | $37 r n$ |
| New HSS | $\left(80 r \log n+\frac{74}{3} r^{2}\right) n$ | $\frac{58}{3} r^{2} n$ | $16 r n$ |

Next, unlike [31], we do not need an extra step to reorthonormalize the HSS form. Such a step may cost as much as $413 r^{2} n$ flops when $m=2 r$ [45], and it even destroys the special structures in the generators.

Furthermore, even our bounds above can be further improved. Assume, for the HSS block rows corresponding to the nodes at level $l$ of the HSS tree $\mathcal{T}$, the row size is $N_{l}=N / 2^{l}$ and the maximum numerical rank is $r_{l}$. (Here, the root is at level 0 and the leaves are at the largest level $l_{\text {max }}$.) According to the idea of rank relaxation in [45], when $N_{l}$ and $r_{l}$ increase (as $l$ decreases along the tree levels) following certain patterns, HSS algorithms may achieve similar orders of complexity as when $r_{l}=O(1)$. Here as observed in section $2, r_{l} \leq O\left(\log N_{l}\right)=O\left(l_{\max }-l\right)$ in practice. Thus, let $m=O(1)$ so that $r_{l_{\max }}=O(1)$. The count (3.25) is essentially improved to

$$
\xi_{1}=\sum_{i: \text { leaf }} O(1)+\sum_{i: \text { nonleaf }} O\left(\left(l_{\max }-l\right)^{2}\left[\frac{13}{3}\left(l_{\max }-l\right)+6 \mu\right]\right)=O(n) .
$$

Therefore, the actual HSS construction cost is truly linear in $n$, in practice. This also holds for the ULV factorization and solution. Thus, the costs of our methods are

$$
\begin{equation*}
\xi_{0,1}=O\left(n \log ^{2} n\right), \quad \xi_{0,2}=O(n), \quad \xi_{1}=O(n), \quad \xi_{2}=O(n), \quad \xi_{3}=O(n) . \tag{3.26}
\end{equation*}
$$

Remark 2. The accuracy of the method can be roughly discussed as follows:
(1) The HSS blocks of $\mathcal{C}$ have fast decaying singular values, and $\mathcal{C}$ can be accurately approximated by our HSS construction scheme. The compression of each HSS block introduces an error that satisfies the bound in (3.6) with a high probability. (The bound is related to the tolerance.) Such a bound is then magnified by an appropriate factor in the hierarchical compression. In general, this often overestimates the error, as observed in other HSS methods [10, 46] and randomized HSS methods or similar [30, 31].
(2) The original ULV HSS factorization methods in $[10,46]$ produce a sequence of intermediate QR factorizations. This converts the HSS matrix into the products of smaller orthogonal and triangular matrices at the hierarchical levels, which can be used to solve HSS systems stably. Here, these QR factorizations are replaced by a stable process (3.22). Also, the pivot growth factor is generally much smaller than the worst case bound in the classical LU factorization [19]. This helps ULV HSS schemes achieve good accuracies. Our numerical tests below indicate that our method works well also for ill-conditioned problems, although a mathematical justification is not yet available. A comprehensive error and stability analysis will be conducted in future work. A closely related study can be found in [3].
4. Precomputations. The previous HSS construction algorithm requires the SPRR factorizations of $\Phi_{i}$ and $\Theta_{i}$ for each node $i$. We can further improve the efficiency of the HSS construction and randomized sampling by using a precomputation stage. As pointed out in [12], we can first compress the HSS blocks of the Cauchy matrix $\mathbf{C}$ in (1.9). The compression steps in section 3.3 can be significantly improved so that the SPRR factorizations take only $O\left(\log ^{3} n\right)$ flops. Also, the multiplication with $X$ is needed only in this precomputation. This is briefly explained as follows.

Due to the special structure of $\mathbf{C}$, its compression can be performed with high efficiency. That is, at each level of the HSS tree, only one block needs to be compressed. Then a shifting strategy similar to the one in [12] can be used to get the compressed forms for other blocks at the same level. Similarly, the column compression stage can also take advantage of the compression information from the row compression stage.

The compression information from $\mathbf{C}$ can then be used to quickly compute an HSS approximation to $\mathcal{C}$. Without loss of generality, we consider only a leaf node $i$.

For convenience, we use similar notation as in the previous section, but with different fonts. That is, $\mathcal{C}_{i}^{-}, U_{i}$, etc., are replaced by $\mathbf{C}_{i}^{-}, \mathbf{U}_{i}$, etc., respectively. Once we have an approximate factorization $\left.\mathbf{C}_{i}^{-} \approx \mathbf{U}_{i} \mathbf{C}_{i}^{-}\right|_{\hat{\mathbf{I}}_{1}}$, similar to the method in [12], the corresponding HSS block $\mathcal{C}_{i}^{-}$of $\mathcal{C}$ in (1.7) can be approximated by

$$
\begin{aligned}
& \mathcal{C}_{i}^{-} \approx \hat{U}_{i} \hat{\Psi}_{i} \text { with } \\
& \hat{U}_{i}=\left(\operatorname{diag}\left(\left.\hat{G}\right|_{\left(l_{i}: l_{i}+m-1\right), 1}\right) \mathbf{U}_{i} \quad \operatorname{diag}\left(\left.\hat{G}\right|_{\left(l_{i}: l_{i}+m-1\right), 2}\right) \mathbf{U}_{i}\right), \\
& \hat{\Psi}_{i}=\binom{\left.\mathbf{C}_{i}^{-}\right|_{\hat{\mathbf{I}}_{1}} \operatorname{diag}\left(\left.\hat{H}\right|_{\left(l_{i}: l_{i}+m-1\right), 1}\right)}{\left.\mathbf{C}_{i}^{-}\right|_{\hat{\mathbf{I}}_{1}} \operatorname{diag}\left(\left.\hat{H}\right|_{\left(l_{i}: l_{i}+m-1\right), 2}\right)},
\end{aligned}
$$

where $\hat{G}$ and $\hat{H}$ are given in (1.6) and $\left.\hat{G}\right|_{\left(l_{i}: l_{i}+m_{i}-1\right), 1}$ is the vector given by rows $l_{i}$ to $l_{i}+m-1$ of the first column of $\hat{G}$.

Then to preserve the structure of $\mathcal{C}_{i}^{-}$in the factorization, we further compute an SPRR factorization

$$
\begin{equation*}
\left.\hat{U}_{i} \approx U_{i} \hat{U}_{i}\right|_{\hat{\mathbf{I}}_{i}} \text { with } U_{i}=P_{i}\binom{I}{E_{i}} \tag{4.1}
\end{equation*}
$$

Thus, due to the special structure of $U_{i}$ (permutation and identity matrix), we have

$$
\begin{equation*}
\mathcal{C}_{i}^{-} \approx U_{i} \hat{U}_{i| |_{\mathbf{I}_{i}}} \hat{\Psi}_{i}=\left.U_{i} \mathcal{C}_{i}^{-}\right|_{\hat{\mathbf{I}}_{i}} \tag{4.2}
\end{equation*}
$$

Therefore, we get a factorization (4.2) that is similar to (3.12) and still preserves the Cauchy-like structure of $\mathcal{C}$ in $\left.\mathcal{C}_{i}^{-}\right|_{\hat{\mathbf{I}}_{i}}$.

This idea can be modified to handle upper level HSS blocks similarly. After the precomputation stage, the HSS construction for $\mathcal{C}$ costs only $O(n)$ flops.

Remark 3. The cost of the SPRR factorization (4.1) is comparable to that of (3.10). However, we save the cost of $\xi_{0}=O\left(n \log ^{2} n\right)$ which is needed only in the precomputation. In practice, $\xi_{0}$ is insignificant if $n$ is not very large. Hence, the scheme with precomputations is suitable for large $n$. This section shows the potential of this precomputation stage for multiple $T$ of the same large $n$, and the details will appear in future work. Our current implementation and numerical experiments in the next section focus on the previous direct HSS construction for $\mathcal{C}$ without precomputations. This direct version gives satisfactory results and is sufficient in general.
5. Numerical experiments. In this section, we test our algorithm (denoted NEW) in section 3, which is implemented in MATLAB on a Unix machine. For convenience, we use the following notation in the experiments:

- $m$ is the leaf level HSS block row size.
- $\tilde{r}=r+\mu$ is the column size of $X$ as in section 3.3.
- $\xi_{\text {constr }}(\mathrm{NEW})=\xi_{0}+\xi_{1}$ and $\xi_{\text {sol }}(\mathrm{NEW})=\xi_{2}+\xi_{3}$ as in (3.26) are the flop counts of the HSS construction and solution, respectively, for $\mathcal{C}$ by NEW.
- $e_{2}=\frac{\|x-\tilde{x}\|_{2}}{\|x\|_{2}}$ is the relative error, where $x$ is the exact solution to $T x=b$ and $\tilde{x}$ is the numerical one.
- $\gamma_{2}=\frac{\|T \tilde{x}-b\|_{2}}{\||\tilde{x}|+|b|\|_{2}}$ is the relative residual for the numerical solution.
- $\sigma$ is the storage size, such as the total number of nonzeros in the factors.

Example 1. Matrix A: The KMS Toeplitz matrix [42] defined by the Toeplitz vector $t$ with

$$
\begin{equation*}
t_{k}=\varphi^{|k|}, \quad 0<\varphi<1 \tag{5.1}
\end{equation*}
$$



(iii) Storage/number of nonzeros

(iv) Relative residual $\gamma_{2}$

(v) Relative error $e_{2}$

Fig. 5.1. Example 1: Numerical results for the well-conditioned Matrix $A$ with $\varphi=0.5$, where $m=80$ and $\tilde{r}=40$ in NEW, and NEW(0) (or NEW(2)) is for the accuracies of NEW without (or with two steps of) iterative refinement.

Our method NEW is compared with four other methods (also in MATLAB):

- SSS: The superfast solver in [11] based on SSS structures. SSS uses an $O\left(n^{2}\right)$ complexity precomputation procedure, whose cost is excluded here. (Even if so, NEW is still much faster than SSS in general, as illustrated below.)
- DIS: The fast solver in [2] with displacement structures.
- SCH1: The superfast Schur algorithm in [40] with improved stability.
- SCH2: The superfast Schur algorithm in [1].

For the case when $\varphi=0.5$ in (5.1), the Toeplitz matrix $T$ is well conditioned. For example, when $n=10^{3}$, the 2 -norm condition number $\kappa_{2} \approx 9$. The flops, timings (in MATLAB), and accuracies of the algorithms are shown in Figure 5.1. The relative tolerance in RRQR factorizations is set to be $10^{-15}$ in all our tests, although sometimes a larger tolerance may save some costs. We also show the solution costs of NEW and SSS after the factorizations in Figure 5.2, where $b$ is generated with $T x$, and $x$ is random. For $\varphi=1-10^{-12}$ in (5.1), $T$ is ill conditioned, and the results are given in Figure 5.3. We make the following observations:

- The total costs, timings, and factorization costs of NEW are roughly linear in $n$. In Figures 5.1(i)-(iii), 5.2, and 5.3(i)-(iii), we also included a line for $O(n)$. Clearly, this line matches well each performance line for NEW.
- Each flop/timing line for NEW has smaller slopes than those for other methods. In fact, in Figures 5.1(i), NEW starts to be faster than SSS, DIS, and SCH1 around $n=2^{9}, 2^{13}$, and $2^{15}$, respectively. In Figure 5.3(i), such break-even points are even smaller.
- The storage of NEW is also nearly $O(n)$. When $n$ is around $2^{15}$, NEW requires less storage than all the other methods except DIS (which is much slower).
- NEW performs better for the ill-conditioned case, both in the costs and the accuracies. When $n$ gets large, the other methods either require too much memory, are too slow, or break down. For example, in Figure 5.3, SCH1 and SCH2 produce NaN results in MATLAB. For the well-conditioned case, NEW gives relatively lower accuracies. However, high accuracies can be quickly reached with few steps of iterative refinement.
- Note that NEW is also very flexible, and works for general $n$ and (nonsymmetric/indefinite) $T$. However, SCH1 and SCH2 works only for symmetric positive definite $T$. Currently, SCH1 also requires $n$ to be a power of 2 .


FIG. 5.2. Example 1: Solution flops and timings for NEW corresponding to Figure 5.1.
Example 2. Then we test the algorithm NEW on various matrices as follows:

- Matrix B: The Toeplitz vector $t$ that defines the matrix $T$ is

$$
t_{-(n-1):-1}=1+\operatorname{randn}(n-1,1) / n, \quad t_{0: n-1}=1+\operatorname{randn}(n, 1) / n,
$$

where $\operatorname{randn}(n, 1)$ generates $n$ normally distributed random numbers. $T$ is mildly ill conditioned.


Fig. 5.3. Example 1: Numerical results for the ill-conditioned Matrix A with $\varphi=1-10^{-12}$, where $m=80$ and $\tilde{r}=40$ in NEW.

- Matrix C: The Prolate Toeplitz matrix defined by

$$
t_{0}=2 \omega, \quad t_{k}=\frac{\sin (2 \omega k \pi)}{k \pi}, \quad k \neq 0, \quad \omega \in[0,1 / 2]
$$

This is a classical test matrix which is very ill conditioned [15, 17]. Here, we use $\omega=1 / 4$.

- Matrix D: The Toeplitz matrix defined by

$$
t_{k}=\sqrt{\epsilon^{2} k^{2}+1}
$$

The matrix has a 2 -norm condition number $\kappa_{2} \approx 0.363 n^{2} e^{3.07 / \epsilon}[5]$. Here, we choose $\epsilon=\frac{1}{8}$, and then $\kappa_{2} \approx 1.7 n^{2} \times 10^{10}$.

- Matrix E: The Toeplitz matrix defined following a Gaussian radial basis function as

$$
t_{k}=e^{-\epsilon^{2} k^{2}}
$$

The matrix has a condition number $\kappa_{2} \approx \frac{1}{2} e^{\pi^{2} /\left(4 \epsilon^{2}\right)}$ and is highly ill conditioned for small $\epsilon[5]$. Here, we choose $\epsilon=\frac{1}{6}$, and then $\kappa_{2} \approx 1.9 \times 10^{38}$.

The results are shown in Table 5.1. We see that the construction and solution costs and storage still follow the patterns of nearly linear complexity. NEW gives small residuals for all the matrices. In some cases, the solution errors are relatively larger for large $n$ and $\kappa_{2}$. We may similarly use iterative refinement to improve the accuracies.

TABLE 5.1
Example 2: HSS construction cost ( $\xi_{\text {constr }}$ ), ULV solution cost $\left(\xi_{\text {sol }}\right)$, and relative residual $e_{2}$ of NEW for the Matrices $B-E$, where $m=80$ and $\tilde{r}=40$.

| $n$ | $n$ | $2^{5} \times 10$ | $2^{6} \times 10$ | $2^{7} \times 10$ | $2^{8} \times 10$ | $2^{9} \times 10$ | $2^{10} \times 10$ | $2^{11} \times 10$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| $\left\|\begin{array}{c} \text { Matrix } \\ \text { B } \end{array}\right\|$ | $\xi_{\text {constr }}$ | 3.33 e 7 | 7.95 e 7 | $1.69 e 8$ | 3.68 e8 | 7.37 e 8 | $1.53 e 9$ | $3.05 e 9$ |
|  | $\xi_{\text {sol }}$ | $1.16 e 7$ | $2.75 e 7$ | $5.92 e 7$ | $1.23 e 8$ | 2.50 e8 | $5.04 e 8$ | 1.01 e9 |
|  | Time | $3.40 e-1$ | $4.40 e-1$ | $9.20 e-1$ | $1.96 e 0$ | $3.91 e 0$ | $8.44 e 0$ | $1.61 e 1$ |
|  | $\sigma$ | $1.52 e 5$ | $3.3 e 5$ | $6.96 e 5$ | $1.42 e 6$ | $2.87 e 6$ | 5.78 e 6 | $1.56 e 7$ |
|  | ${ }_{2}$ | $6.76 e-14$ | $1.52 e-12$ | $6.35 e-12$ | $4.03 e-11$ | $8.50 e-11$ | 5.92e-10 | $1.85 e-9$ |
|  | $e_{2}$ | $3.40 e-9$ | $4.48 e-7$ | $3.15 e-6$ | $5.74 e-5$ | $6.25 e-4$ | $7.16 e-2$ | $2.34 e-1$ |
| $\left\|\begin{array}{c} \text { Matrix } \\ \mathrm{C} \end{array}\right\|$ | $\xi_{\text {constr }}$ | 3.63 e 7 | 8.69 e 7 | $1.90 e 8$ | $3.97 e 8$ | $8.02 e 8$ | $1.62 e 9$ | $3.30 e 9$ |
|  |  | $9.96 e 6$ | 2.58 e7 | $5.95 e 7$ | $1.23 e 8$ | $2.50 e 8$ | 5.05 e 8 | $1.01 e 9$ |
|  | Time | $4.00 e-1$ | $5.20 e-1$ | $1.19 e 0$ | $2.47 e 0$ | $5.04 e 0$ | 1.05 e 1 | 2.05 e 1 |
|  | $\sigma$ | $1.44 e 5$ | $3.26 e 5$ | $6.97 e 5$ | 1.43 e 6 | 2.88 e 6 | 5.78 e6 | $1.16 e 7$ |
|  | $\gamma_{2}$ | 1.44-15 | $2.47 e-14$ | $4.96 e-12$ | $4.78 e-12$ | $7.82 e-11$ | $8.52 e-10$ | $8.00 e-10$ |
|  | $e_{2}$ | $1.66 e 2$ | $2.85 e 2$ | $7.68 e 1$ | $3.32 e 2$ | $8.52 e 2$ | $1.21 e 4$ | $1.14 e 6$ |
| $\left.\begin{gathered} \text { Matrix } \\ D \end{gathered} \right\rvert\,$ | nstr | $3.42 e 7$ | 8.08 e 7 | $1.80 e 8$ | $3.73 e 8$ | $7.87 e 8$ | 1.55 e 9 | 9 |
|  | $\xi_{\text {sol }}$ | $7.47 e 6$ | $2.11 e 7$ | 5.68 e7 | $1.19 e 8$ | 2.50 e 8 | $5.05 e 8$ | $1.01 e 9$ |
|  | Time | $3.60 e-1$ | $5.00 e-1$ | $1.12 e 0$ | $2.33 e 0$ | 4.79 e 0 | $9.85 e 0$ | 1.89 e 1 |
|  | $\sigma$ | $1.28 e 5$ | $2.99 e 5$ | $6.85 e 5$ | 1.41e6 | $2.87 e 6$ | 5.78 e6 | $1.16 e 7$ |
|  | $\gamma_{2}$ | 4.76-15 | $6.55 e-15$ | $1.72 e-13$ | $4.83 e-13$ | $6.49 e-14$ | $2.66 e-13$ | $9.71 e-15$ |
|  | $e_{2}$ | $7.07 e-2$ | $2.37 e-1$ | 7.05 e 0 | $3.17 e 0$ | $7.47 e 1$ | 5.93 e 1 | $7.88 e 2$ |
| $\left\|\begin{array}{c} \text { Matrix } \\ \mathrm{E} \end{array}\right\|$ | $\xi_{\text {constr }}$ | $3.61 e 7$ | $8.60 e 7$ | $1.86 e 8$ | $3.91 e 8$ | $7.99 e 8$ | $1.62 e 9$ | $3.29 e 9$ |
|  | $\zeta_{\text {sol }}$ | 1.05 e 7 | $2.72 e 7$ | $5.95 e 7$ | $1.23 e 8$ | $2.50 e 8$ | 5.05 e 8 | $1.01 e 9$ |
|  | Time | $3.80 e-1$ | $5.40 e-1$ | $1.19 e 0$ | $2.35 e 0$ | $4.87 e 0$ | 1.05 e 1 | 2.03 e 1 |
|  | $\sigma$ | $1.46 e 5$ | $3.32 e 5$ | $6.97 e 5$ | $1.42 e 6$ | 2.88 e6 | 5.78 e6 | $1.16 e 7$ |
|  | $\gamma_{2}$ | 4.40-16 | $9.14 e-16$ | $1.85 e-13$ | $2.67 e-13$ | $5.21 e-13$ | $2.85 e-12$ | $5.01 e-12$ |
|  | $e_{2}$ | 1.56 e 1 | $3.60 e 1$ | $1.44 e 1$ | $4.43 e 1$ | $7.96 e 1$ | $1.73 e 2$ | $2.92 e 2$ |

Example 3. Matrix F: The Toeplitz matrix defined as follows [17]:

$$
t_{k}= \begin{cases}\frac{9}{10}+\frac{1}{10} \operatorname{rand}(1) & \text { if } k=0, \\ -t_{0} & \text { if } k>0, \\ 0 & \text { if }-\frac{n}{2}<k<0, \\ \operatorname{rand}(1) & \text { otherwise }\end{cases}
$$

where $\operatorname{rand}(1)$ generates a random number from uniform distribution in $(0,1)$.
NEW is compared with two methods:
Levinson: The Levinson algorithm in [8].
GEPP: Gaussian elimination with partial pivoting.
For this matrix, both the straightforward GEPP and the Levinson algorithm fail since they produce huge element growth [17]. See Table 5.2. Here, NEW still provides reasonably good accuracies. In comparison, Levinson works only for small $n$, and quickly fails to produce meaningful results. GEPP fails for all the matrix sizes in the table. Thus instead, for GEPP, we show the growth factors which are very large.

TABLE 5.2
Example 3: Numerical results for Matrix $F$, where $m=80, \tilde{r}=60$ in NEW, $\rho_{\text {grow }}$ is the pivot growth factor in GEPP, and NaN and Inf are outputs in MATLAB to mean certain failures.

| $n$ |  | $2^{5} \times 10$ | $2^{6} \times 10$ | $2^{7} \times 10$ | $2^{8} \times 10$ | $2^{9} \times 10$ | $2^{10} \times 10$ | $2^{11} \times 10$ |
| :--- | :--- | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
| NEW | $\xi_{\text {constr }}$ | $1.03 e 8$ | $2.38 e 8$ | $5.32 e 8$ | $1.05 e 9$ | $2.18 e 9$ | $4.45 e 9$ | $8.84 e 9$ |
|  | $\xi_{\text {sol }}$ | $1.91 e 7$ | $5.76 e 7$ | $1.38 e 8$ | $2.95 e 8$ | $6.08 e 8$ | $1.23 e 9$ | $2.48 e 9$ |
|  | $\gamma_{2}$ | $3.55 e-17$ | $9.85 e-15$ | $1.66 e-12$ | $5.99 e-12$ | $4.02 e-11$ | $4.30 e-10$ | $5.66 e-9$ |
|  | $e_{2}$ | $2.80 e-15$ | $3.89 e-12$ | $3.18 e-10$ | $1.43 e-8$ | $2.98 e-7$ | $1.05 e-5$ | $3.35 e-4$ |
| Levinson | $\gamma_{2}$ | $7.02 e-2$ | NaN | NaN | NaN |  |  |  |
|  | $e_{2}$ | $8.28 e 31$ | NaN | NaN | NaN |  |  |  |
| GEPP | $\rho_{\text {grow }}$ | $1.34 e 48$ | $1.78 e 96$ | $4.16 e 192$ | Inf |  |  |  |

6. Conclusions and future work. In this paper, we study superfast structured Toeplitz solutions based on an SPRR factorization and randomized sampling. The Cauchy-like matrix $\mathcal{C}$ can be approximated quickly by an HSS matrix in $O(n)$ flops for block compression and $O\left(n \log ^{2} n\right)$ flops for FFTs. One solver we propose is significantly faster than an existing one based on SSS methods and some other methods for large $n$. Strong rank-revealing factorizations are used for stability. We also briefly show an HSS construction procedure with precomputations, where the compression costs only $O\left(\log ^{3} n\right)$ after FFTs. This construction procedure is especially useful if $n$ is very large and if there are multiple Toeplitz matrices of the same size.

The discussions here provide the possibility of developing new fast versions. For example, it is possible to compute the products such as $\mathbf{C} X$ and $\Phi_{i}$ in (3.9) just once in the precomputation. After this, the HSS construction for $\mathcal{C}$ involves just $O(n)$ cost block compression. This will appear in our future work. The current version without precomputation gives superior results.

This work also shows a practical way of using randomized sampling in the development of new efficient structured methods.

Acknowledgments. The authors are very grateful to the two anonymous referees for their valuable suggestions. The authors would also like to thank J. P. Boyd for some test examples and M. Stewart for his code in [40]. Some other test codes used are from G. Rodriguez's webpage (http://bugs.unica.it/~gppe/soft/).

## REFERENCES

[1] G. S. Ammar and W. B. Gragg, Superfast solution of real positive definite Toeplitz systems, SIAM J. Matrix Anal. Appl., 9 (1988), pp. 61-76.
[2] A. Aricò, and G. Rodriguez, A fast solver for linear systems with displacement structure, Numer. Algorithms, 55 (2010), pp. 529-556.
[3] T. Bella, V. Olshevsky, and M. Stewart, A nested product decomposition of a quasiseparable matrix, submitted.
[4] A. W. Bojańczyk, R. P. Brent, F. R. de Hoog, and D. R. Sweet, On the stability of the Bareiss and related Toeplitz factorization algorithms, SIAM J. Matrix Anal. Appl., 16 (1995), pp. 40-57.
[5] J. P. Boyd and K. W. Gildersleeve, Numerical experiments on the condition number of the interpolation matrices for radial basis functions, Appl. Numer. Math., 61 (2011), pp. 443459.
[6] J. R. Bunch, Stability of methods for solving Toeplitz systems of equations, SIAM J. Sci. Statist. Comput., 6 (1985), pp. 349-364.
[7] R. H. Chan, J. G. Nagy, and R. J. Plemmons, Displacement preconditioners for Toeplitz least squares iterations, Electron. Trans. Numer. Anal., 2 (1994), pp. 1-13.
[8] T. F. Chan and P. C. Hansen, A look-ahead Levinson algorithm for general Toeplitz systems, IEEE Trans. Signal Process., 40 (1992), pp. 1079-1090.
[9] R. H. Chan and M. K. NG, Conjugate gradient methods for Toeplitz systems, SIAM Rev., 38 (1996) pp. 427-482.
[10] S. Chandrasekaran, M. Gu, and T. Pals, A fast ULV decomposition solver for hierarchically semiseparable representations, SIAM J. Matrix Anal. Appl., 28 (2006), pp. 603-622.
[11] S. Chandrasekaran, P. Dewilde, M. Gu, T. Pals, X. Sun, A.-J. van der Veen, and D. White, Some fast algorithms for sequentially semiseparable representations, SIAM J. Matrix Anal. Appl., 27 (2005), pp. 341-364.
[12] S. Chandrasekaran, M. Gu, X. Sun, J. Xia, and J. Zhu, A superfast algorithm for Toeplitz systems of linear equations, SIAM J. Matrix Anal. Appl., 29 (2007), pp. 1247-1266.
[13] S. Chandrasekaran and A. H. Sayed, Stabilizing the fast generalized Schur algorithm, SIAM J. Matrix Anal. Appl., 17 (1996), pp. 950-983.
[14] G. Cybenko, The numerical stability of the Levinson-Durbin algorithm for Toeplitz systems of equations, SIAM J. Sci. Statist. Comput., 1 (1980), pp. 303-319.
[15] I. Gohberg, T. Kailath, and V. Olshevsky, Fast Gaussian elimination with partial pivoting for matrices with displacement structure, Math. Comp., 64 (1995), pp. 1557-1576.
[16] G. H. Golub and C. V. Loan, Matrix Computations, 3rd ed., The Johns Hopkins University Press, Baltimore, MD, 1996.
[17] M. Gu, Stable and efficient algorithms for structured systems of linear equations, SIAM J. Matrix Anal. Appl., 19 (1998), pp. 279-306.
[18] M. Gu and S. C. Eisenstat, Efficient algorithms for computing a strong-rank revealing $Q R$ factorization, SIAM J. Sci. Comput., 17 (1996), pp. 848-869.
[19] M. Gu and J. XiA, A numerically stable and superfast algorithm for solving Toeplitz systems of linear equations, Presentation in the SIAM Conference on Applied Linear Algebra, Monterey, CA, 2009.
[20] N. Halko, P. G. Martinsson, and J. A. Tropp, Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, SIAM Rev., 53 (2011), pp. 217-288.
[21] G. Heinig, Inversion of generalized Cauchy matrices and other classes of structured matrices, in Linear Algebra for Signal Processing, IMA Vol. Math. Appl. 69, Springer, New York, 1995, pp. 95-114.
[22] G. Heinig, Solving Toeplitz systems after extension and transformation, Calcolo, 33 (1996), pp. 115-129.
[23] G. Heinig and F. Hellinger, Displacement structure of pseudoinverses, Linear Algebra Appl., 197/198 (1994), pp. 623-649.
[24] G. Heinig and K. Rost, Algebraic methods for Toeplitz-like matrices and operators, Oper. Theory Adv. Appl. 13, Birkhäuser Verlag, Basel, 1984, pp. 109-127.
[25] I. C. F. Ipsen, C. T. Kelley, and S. R. Pope, Rank-deficient nonlinear least squares problems and subset selection, SIAM J. Numer. Anal., 49 (2011), pp. 1244-1266.
[26] T. Kailath, S. Y. Kung, and M. Morf, Displacement ranks of matrices and linear equations, J. Math. Anal. Appl., 68 (1979), pp. 395-407.
[27] T. Kailath and A. H. Sayed, Fast algorithms for generalized displacement structures, in Recent Advances in Mathematical Theory of Systems, Control, Networks, and Signal Pro-
cessing, Vol. II, H. Kimura and S. Kodama, eds., Mita Press, Tokyo, 1992, pp. 27-32.
[28] T. Kailath and A. H. Sayed, eds., Fast Reliable Algorithms for Matrices with Structure, SIAM, Philadelphia, 1999.
[29] E. Liberty, F. Woolfe, P.-G. Martinsson, V. Rokhlin, and M. Tygert, Randomized algorithms for the low-rank approximation of matrices, Proc. Natl. Acad. Sci. USA, 104 (2007), pp. 20167-20172.
[30] L. Lin, J. Lu, and L. Ying, Fast construction of hierarchical matrix representation from matrix-vector multiplication, J. Comput. Phys., 230 (2011), pp. 4071-4087.
[31] P. G. Martinsson, A fast randomized algorithm for computing a hierarchically semiseparable representation of a matrix, SIAM J. Matrix Anal. Appl., 32 (2011), pp. 1251-1274.
[32] P. G. Martinsson, V. Rokhlin, and M. Tygert, A Randomized Algorithm for the Approximation of Matrices, Technical Report 1361, Department of Computer Science, Yale University, New Haven, CT, 2006.
[33] P. G. Martinsson, V. Rokhlin, and M. Tygert, A fast algorithm for the inversion of general Toeplitz matrices, Comput. Math. Appl., 50 (2005), pp. 741-752.
[34] M. K. Ng and R. H. Chan, Scientific applications of iterative Toeplitz solvers, Calcolo, 33 (1996), pp. 249-267.
[35] J. Olkin, Linear and Nonlinear Deconvolution Problems, Ph.D. Thesis, Rice University, Houston, TX, 1986.
[36] V. OlShevski, ed., Fast Algorithms for Structured Matrices: Theory and Applications, Contemp. Math. 323, Amer. Math. Soc., Providence, RI, 2003.
[37] V. Y. Pan and G. Qian, Randomized preprocessing of homogeneous linear systems of equations, Linear Algebra Appl., 432 (2010), pp. 3272-3318.
[38] V. Y. Pan, On computations with dense structured matrices, Math. Comp., 55 (1990), pp. 179190.
[39] G. Strang, A proposal for Toeplitz matrix calculations, Stud. Appl. Math., 74 (1986), pp. 171176.
[40] M. Stewart, A superfast Toeplitz solver with improved numerical stability, SIAM J. Matrix Anal. Appl., 25 (2003), pp. 669-693.
[41] D. R. SWEET, The use of pivoting to improve the numerical performance of algorithms for Toeplitz matrices, SIAM J. Matrix Anal. Appl., 14 (1993), pp. 468-493.
[42] W. F. Trench, Numerical solution of the eigenvalue problem for Hermitian Toeplitz matrices, SIAM J. Matrix Anal. Appl., 10 (1989), pp. 135-146.
[43] E. E. Tyrtyshnikov, Incomplete cross approximation in the mosaic-skeleton method, Computing, 64 (2000), pp. 367-380.
[44] M. Van Barel, G. Heinig, and P. Kravanja, A stabilized superfast solver for nonsymmetric Toeplitz systems, SIAM J. Matrix Anal. Appl., 23 (2001), pp. 494-510.
[45] J. XiA, On the complexity of some hierarchical structured matrix algorithms, SIAM J. Matrix Anal. Appl., 33 (2012), pp. 388-410.
[46] J. Xia, S. Chandrasekaran, M. Gu, and X. S. Li, Fast algorithms for hierarchically semiseparable matrices, Numer. Linear Algebra Appl., 17 (2010), pp. 953-976.


[^0]:    *Received by the editors April 25, 2011; accepted for publication (in revised form) by M. Van Barel April 17, 2012; published electronically August 2, 2012.
    http://www.siam.org/journals/simax/33-3/83198.html
    ${ }^{\dagger}$ Department of Mathematics, Purdue University, West Lafayette, IN 47907 (xiaj@math.purdue. edu, yxi@math.purdue.edu). The first author's research was supported in part by NSF grants DMS1115572 and CHE-0957024.
    ${ }^{\ddagger}$ Department of Mathematics, University of California, Berkeley, CA 94720 (mgu@math.berkeley. edu).

