

# Lawrence Berkeley National Laboratory

## LBL Publications

### Title

Butterfly Factorization Via Randomized Matrix-Vector Multiplications

### Permalink

<https://escholarship.org/uc/item/8k17d8xn>

### Journal

SIAM Journal on Scientific Computing, 43(2)

### ISSN

1064-8275

### Authors

Liu, Yang  
Xing, Xin  
Guo, Han  
[et al.](#)

### Publication Date

2021

### DOI

10.1137/20m1315853

Peer reviewed

# BUTTERFLY FACTORIZATION VIA RANDOMIZED MATRIX-VECTOR MULTIPLICATIONS\*

YANG LIU<sup>§†</sup>, XIN XING<sup>‡</sup>, HAN GUO<sup>§</sup>, ERIC MICHIELSSEN<sup>§</sup>, PIETER GHYSELS<sup>†</sup>, AND XIAOYE SHERRY LI<sup>†</sup>

**Abstract.** This paper presents an adaptive randomized algorithm for computing the butterfly factorization of a  $m \times n$  matrix with  $m \approx n$  provided that both the matrix and its transpose can be rapidly applied to arbitrary vectors. The resulting factorization is composed of  $O(\log n)$  sparse factors, each containing  $O(n)$  nonzero entries. The factorization can be attained using  $O(n^{3/2} \log n)$  computation and  $O(n \log n)$  memory resources. The proposed algorithm applies to matrices with strong and weak admissibility conditions arising from surface integral equation solvers with a rigorous error bound, and is implemented in parallel.

**Key word.** Matrix factorization, butterfly algorithm, randomized algorithm, integral operator.

**AMS subject classifications.** 15A23, 65F50, 65R10, 65R20

**1. Introduction.** Butterfly factorizations are an attractive means for compressing highly oscillatory operators arising in many scientific and engineering applications, including the integral equation-based analysis of high-frequency acoustic and electromagnetic scattering problems [20], and the evaluation of Fourier integrals and transforms [1, 27], spherical harmonic transforms [25] and many other types of special function transforms [21]. The butterfly factorization of a  $m \times n$  matrix with  $m \approx n$  exists provided that all judiciously selected submatrices, whose row and column dimensions multiply to  $O(n)$ , are numerically low-rank. Through recursive low-rank factorizations of these submatrices, the operator can be represented as the product of  $O(\log n)$  sparse matrices, each containing  $O(n)$  nonzero entries. The resulting factorization can be rapidly applied to arbitrary vectors using only  $O(n \log n)$  computation and memory resources.

Despite this favorable application cost, the cost of constructing a butterfly representation of a given operator typically scales at least as  $O(n^2)$  [25]. Fortunately, there exist two important categories of operators that allow for fast approximation by a butterfly. (i) Operators that allow each element of their matrix representation to be evaluated in  $O(1)$  operations. This is typically the case when the butterfly factorization applies directly to an oscillatory operator with an explicit formula (e.g., Fourier operators, special transforms, or discretized integral equations) or stored as a full matrix. (ii) Operators with matrix representations that can be applied to arbitrary vectors in quasi-linear, typically  $O(n \log n)$ , complexity. This situation typically arises when re-compressing the composition of highly-oscillatory operators, e.g., composition of Fourier integrals [10, 12], matrix algebras for constructing discretized inverse integral operators [5], or conversion to a butterfly representation from other compression formats (e.g., fast multipole methods (FMM)-like formats).

The butterfly factorization of matrices in category (i) can be constructed using  $O(n \log n)$  computation and memory resources following the low-rank decomposition

---

\*Version of February 11, 2020.

<sup>†</sup>Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA (liuyangzhuan@lbl.gov, pghysels@lbl.gov, xsli@lbl.gov).

<sup>‡</sup>School of Mathematics, Georgia Institute of Technology, GA (xxing33@gatech.edu).

<sup>§</sup>Department of Electrical Engineering and Computer Science, University of Michigan, MI (hanguo@umich.edu, emichiel@umich.edu).

of judiciously selected submatrices using uniform [20], random [19], or Chebyshev [11] proxy points. A wide variety of low-rank decompositions, including the interpolative decomposition (ID) [11], the pseudo skeleton approximation [19], the adaptive cross approximation [24], and singular value decomposition (SVD) [12] can be used for this purpose. This type of butterfly factorization has been extended to multi-scale and multi-dimensional problems [13, 14, 20].

Operators in category (ii), on the other hand, pose bigger challenges to fast butterfly construction algorithms. These algorithms rely on random projection-based algorithms [7, 15] to construct low-rank decompositions of the associated submatrices, typically resulting in higher computation and memory costs than those in category (i). That said,  $O(n \log n)$  algorithms can be developed when the oscillatory operator allows for smooth phase recovery [26] or fast submatrix-vector multiplications (e.g., using FMM-type algorithms). Unfortunately, these requirements are not met when compressing the concatenation of several Fourier operators or when inverting integral equation operators. When the above conditions do not apply, the cost of reconstructing butterfly approximations for operators in category (ii) require  $O(n^{3/2} \log n)$  and  $O(n^{3/2})$  computation and memory resources via nested SVD compression [12].

Recently, a  $O(n^{3/2} \log n)$ -computation but  $O(n \log n)$ -memory algorithm in Category (ii) for general butterfly compressible matrices was introduced in [6]. Unlike the SVD-based algorithm [12] that constructs the factorization from innermost to outermost factors, this algorithm reverses the computation sequence. Consequently, the SVD-based algorithm stores information associated with all random vectors requiring  $O(n^{3/2})$  storage, while the algorithm in [6] only stores information related to subsets of the structured random vectors, thereby requiring only  $O(n \log n)$  storage.

This paper presents a new butterfly reconstruction scheme that provides a three-fold improvement on the algorithm in [6]: (i) The new algorithm is adaptive in nature and permits fast and accurate butterfly reconstructions even when applied to matrices with weak admissibility conditions arising from the discretization of 3D surface integral equation solvers, whereas the previous algorithm results in higher computational complexity. (ii) The algorithm comes with a rigorous error bound obtained using an orthogonal projection argument that grows only weakly with matrix size. (iii) The algorithm can be deployed in parallel. Not surprisingly, the proposed algorithm represents a critical building block for constructing fast iterative and direct solvers for highly-oscillatory problems.

## 2. Preliminary background.

**2.1. Notation.** We use MATLAB notation to denote entries and subblocks of matrices and vectors. For example,  $A(i, j)$  denotes the  $(i, j)$ th entry of matrix  $A$ , and  $A(I, J)$  with index sets  $I$  and  $J$  denotes the subblock of matrix  $A$  with rows and columns with indices in  $I$  and  $J$ , respectively. Moreover,  $\text{diag}(A_1, \dots, A_k)$  denotes a block diagonal matrix with  $k$  diagonal blocks  $A_1, \dots, A_k$ . We assume  $A \in \mathbb{R}^{m \times n}$  but all the algorithms can be trivially extended to complex matrices.

**2.2. Low-rank approximation by projection.** Given a matrix  $A \in \mathbb{R}^{m \times n}$ , we consider a rank- $r$  approximation of  $A$  with accuracy  $O(\epsilon)$  in the *projection form*,

$$A = U(U^T A) + O(\epsilon),$$

where  $U$  is of dimension  $m \times r$  and the symbol “ $T$ ” denotes the transpose of a matrix. The operator  $UU^T$  projects all the columns of  $A$  onto the  $r$ -dimensional subspace  $\text{col}(U)$ . The matrix  $U$  can be computed via singular value or pivoted QR decom-

positions, or randomized methods [15]. We focus on using randomized methods as illustrated in Algorithm 2.1 to construct  $U$ .

---

**Algorithm 2.1** Randomized low-rank approximation method

---

**Input:** Routine to multiply  $A \in \mathbb{R}^{m \times n}$  with arbitrary matrices, rank  $r$ , over-sampling parameter  $p$ , truncation tolerance  $\epsilon$ .

**Output:** Basis matrix  $U$  such that  $A \approx UU^T A$  with rank at most  $(r + p)$ .

- 1: **Step 1:** Form a  $n \times (p + r)$  random matrix  $\Omega$  whose entries are random variables which are independent and identically distributed, following a normal distribution.
  - 2: **Step 2:** Compute  $W = A\Omega$ .
  - 3: **Step 3:** Compute the column-pivoted QR decomposition of  $W$  with truncation tolerance  $\epsilon$  as  $WP = QR$  where  $P$  denotes the permutation,  $Q$  is orthonormal, and  $R$  is upper triangular. Return  $U = Q$ .
- 

This randomized method has been studied extensively in [7]. It can be shown that with high probability the rank- $(r + p)$  approximation constructed by Algorithm 2.1 has approximation error similar to that of the optimal rank- $r$  approximation (see Theorem 10.7 of [7]). Algorithm 2.1 has  $O(mnr)$  computational cost when  $A$  is stored as a full matrix. Moreover, adaptive versions of Algorithm 2.1 that increase the rank estimate  $r$  until the resulting low-rank approximation meets a prescribed approximation accuracy, have been developed [3].

**3. Butterfly factorization.** We consider the butterfly factorization of a matrix  $A = K(T, S)$  defined by a highly-oscillatory operator  $K(\cdot, \cdot)$  and point sets  $S$  and  $T$ . Consider the example of free-space wave interactions between 3D *source* points  $S$  and *target* points  $T$  where  $S$  and  $T$  are non-overlapping (weak admissibility) or well separated (strong admissibility). The entries  $A_{i,j} = K(t_i, s_j) = e^{i2\pi k|t_i - s_j|} / |t_i - s_j|$  for all pairs  $(t_i, s_j) \in T \times S$  represent a discretization of the 3D Helmholtz kernel with wavenumber  $\kappa > 0$ . Let  $|T| = m$  and  $|S| = n$ , and assume that  $m = O(n)$ . Other examples of highly-oscillatory operators includes Green’s function operators for Helmholtz equations with nonconstant coefficients, Fourier transforms, and special function transforms.

**3.1. Hierarchical partitioning.** The point sets  $S$  and  $T$  are recursively partitioned into small subsets by bisection until the finest subsets thus obtained have less than a prescribed number of points  $n_0$ . Many geometry or graph-based partitioning algorithms for this purpose have been studied in [23]. Here, we use bisection and binary trees for simplicity, and denote the two binary trees  $\mathcal{T}_S$  and  $\mathcal{T}_T$ . We further assume both trees  $\mathcal{T}_S$  and  $\mathcal{T}_T$  are complete (all levels are completely filled). These assumptions can be easily relaxed.

We number the levels of  $\mathcal{T}_T$  and  $\mathcal{T}_S$  from the root to the leafs. The root node is at level 0; its children are at level 1, etc. All the leaf nodes are at level  $L$ . At each level  $l$ ,  $\mathcal{T}_T$  and  $\mathcal{T}_S$  both have  $2^l$  nodes.

Let  $T_\tau$  be the subset of points in  $T$  corresponding to node  $\tau$  in  $\mathcal{T}_T$ . Obviously,  $T_t = T$  if  $t$  denotes the root node. Furthermore, for any non-leaf node  $\tau \in \mathcal{T}_T$  with children  $\tau_1$  and  $\tau_2$ ,  $T_{\tau_1} \cup T_{\tau_2} = T_\tau$  and  $T_{\tau_1} \cap T_{\tau_2} = \emptyset$ . The same properties hold true for the partitioning of  $S$ , i.e.,  $\{S_\nu\}_{\nu \in \mathcal{T}_S}$ . We assume that both  $\mathcal{T}_T$  and  $\mathcal{T}_S$  have the same depth  $L = O(\log n)$  so that the leaf point subsets have  $O(1)$  points.

**3.2. Complementary low-rank property.**  $A = K(T, S)$  satisfies the *complementary low-rank property* if for any level  $0 \leq l \leq L$ , node  $\tau$  at level  $l$  of  $\mathcal{T}_T$  and a




FIG. 3.1. The partitioning trees  $\mathcal{T}_S$ ,  $\mathcal{T}_T$  and the blocks  $K(T_\tau, S_\nu)$  for a 2-level butterfly factorization. The blocks correspond to level  $l$  of  $\mathcal{T}_T$  and level  $L - l$  of  $\mathcal{T}_S$  with (a)  $l = 0$  (b)  $l = 1$  (c)  $l = 2$ .

node  $\nu$  at level  $(L - l)$  of  $\mathcal{T}_S$ , the subblock  $K(T_\tau, S_\nu)$  is numerically low-rank with rank  $r_{\tau, \nu}$  bounded by a small number  $r$ ;  $r$  is called the (maximum) butterfly rank. See Figure 3.1 for example subblocks of a 2-level butterfly. The complementary low-rank property results from the analytic properties of the Helmholtz kernel [17, 19], and Fourier integrals [1], and has been observed for many other highly-oscillatory operators.

**3.3. Low-rank approximation of the blocks.** A butterfly factorization of  $A$  compresses all blocks  $K(T_\tau, S_\nu)$  with  $l = 0, 1, \dots, L$  for nodes  $\tau$  at level  $l$  of  $\mathcal{T}_T$  and  $\nu$  at level  $(L - l)$  of  $\mathcal{T}_S$  into low-rank form via a nested approach. There exist three different but equivalent approaches: (1) column-wise butterfly factorization, (2) row-wise butterfly factorization, and (3) hybrid butterfly factorization. The proposed algorithm uses the hybrid factorization as it yields the lowest algorithmic complexity and highest parallel efficiency. That said, we will describe all three approaches as the column- and row-wise factorizations also serve as the building blocks in the hybrid factorization.

**3.3.1. Column-wise butterfly factorization.** Each  $K(T_\tau, S_\nu)$  is compressed into rank- $r$  form

$$(3.1) \quad K(T_\tau, S_\nu) \approx U_{\tau, \nu} U_{\tau, \nu}^T K(T_\tau, S_\nu) = U_{\tau, \nu} E_{\tau, \nu},$$

where  $U_{\tau, \nu} \in \mathbb{R}^{|T_\tau| \times r}$  has orthonormal columns and is referred to as *the column basis matrix* associated with  $(\tau, \nu)$ . The compression is performed directly for any  $T_\tau$  at the leaf level  $L$  of  $\mathcal{T}_T$ .

At a non-leaf level  $l < L$ , consider a node  $\tau$  at level  $l$  of  $\mathcal{T}_T$  and a node  $\nu$  at level  $(L - l)$  of  $\mathcal{T}_S$ . Let  $\{\tau_1, \tau_2\}$  be the children of  $\tau$  at level  $(l + 1)$  of  $\mathcal{T}_T$  and  $p_\nu$  be the parent node of  $\nu$  at level  $(L - l - 1)$  of  $\mathcal{T}_S$ . The low-rank approximation (3.1) of  $K(T_\tau, S_\nu)$  is constructed using the low-rank approximations of blocks  $K(T_{\tau_1}, S_{p_\nu})$  and  $K(T_{\tau_2}, S_{p_\nu})$  at level  $(l + 1)$  using the nested compression approach described next.

Since  $T_\tau = T_{\tau_1} \cup T_{\tau_2}$ ,  $K(T_\tau, S_\nu)$  can be split into two blocks, i.e.,

$$(3.2) \quad K(T_\tau, S_\nu) = \begin{bmatrix} K(T_{\tau_1}, S_\nu) \\ K(T_{\tau_2}, S_\nu) \end{bmatrix}.$$

Meanwhile, since  $S_\nu$  is a subset of  $S_{p_\nu}$ , it follows that  $K(T_{\tau_a}, S_\nu)$ , for each child  $\tau_a$  of  $\tau$ , is a subblock of  $K(T_{\tau_a}, S_{p_\nu})$ . Thus, the low-rank approximation  $K(T_{\tau_a}, S_{p_\nu}) \approx$






FIG. 3.2. (a) Column-wise, (b) row-wise, and (c) hybrid factorizations for a 4-level butterfly. The blocks in black multiply to  $U_{t,\nu_1} E_{t,\nu_1}$  in (3.6),  $F_{\tau_1,s} V_{\tau_1,s}^T$ , and  $U_{\tau_1,\nu_1} B_{\tau_1,\nu_1} V_{\tau_1,\nu_1}^T$  in (3.13), respectively.

**3.3.2. Row-wise butterfly factorization.** Each  $K(T_\tau, S_\nu)$  is compressed into rank- $r$  form as

$$(3.9) \quad K(T_\tau, S_\nu) \approx K(T_\tau, S_\nu) V_{\tau,\nu} V_{\tau,\nu}^T = F_{\tau,\nu} V_{\tau,\nu}^T,$$

where  $V_{\tau,\nu} \in \mathbb{R}^{|S_\nu| \times r}$  has orthonormal columns and is referred to as *the row basis matrix* associated with  $(\tau, \nu)$ . Just like for the column-wise factorization, we define the transfer matrix  $W_{\tau,\nu}$  for a non-leaf node  $\nu$  as

$$(3.10) \quad V_{\tau,\nu} = \begin{bmatrix} V_{p_\tau,\nu_1} & \\ & V_{p_\tau,\nu_2} \end{bmatrix} W_{\tau,\nu}.$$

The basis and transfer matrices can be constructed upon applying the column-wise butterfly factorization to  $K(T, S)^T$ , yielding the row-wise butterfly structure

$$(3.11) \quad K(T, S) = F^L (W^L W^{L-1} \dots W^1 V^0)$$

Here the outer factors are  $F^L = \text{diag}(F_{\tau_1,s}, \dots, F_{\tau_{2^L},s})$  with  $s$  denoting the root of  $\mathcal{T}_S$  and  $V^0 = \text{diag}(V_{t,\nu_1}^T, \dots, V_{t,\nu_{2^L}}^T)$ , and the block-diagonal inner factors  $W^l, l = 1, \dots, L$  have blocks  $W_\tau$  for all nodes  $\tau$  at level  $l-1$  of  $\mathcal{T}_T$ . Each  $W_\tau$  consists of  $W_{\tau_1,\nu}$  and  $W_{\tau_2,\nu}$  for all nodes  $\nu$  at level  $L-l+1$  of  $\mathcal{T}_S$ , as

$$(3.12) \quad W_\tau = \begin{bmatrix} W_{\tau_1,\nu_1} & & & & & \\ & W_{\tau_1,\nu_2} & & & & \\ & & \ddots & & & \\ & & & & W_{\tau_1,\nu_{2^{L-l+1}}} & \\ W_{\tau_2,\nu_1} & & & & & \\ & W_{\tau_2,\nu_2} & & & & \\ & & \ddots & & & \\ & & & & W_{\tau_2,\nu_{2^{L-l+1}}} & \end{bmatrix}.$$

Figure 3.2(b) shows an example of a 4-level row-wise factorization.

**3.3.3. Hybrid butterfly factorization.** At any level  $l$  of  $\mathcal{T}_T$ ,  $K(T_\tau, S_\nu)$  with all nodes  $\tau$  at level  $l$  of  $\mathcal{T}_T$  and nodes  $\nu$  at level  $(L-l)$  of  $\mathcal{T}_S$  (referred to as the blocks at level  $l$ ) form a non-overlapping partitioning of  $K(T, S)$ . Fixing  $l$ , we can combine the computed row and column basis matrices  $U_{\tau,\nu}$ ,  $V_{\tau,\nu}$  from both the column-wise and row-wise butterfly factorizations of  $K(T, S)$  above to compress  $K(T_\tau, S_\nu)$  as

$$K(T_\tau, S_\nu) \approx U_{\tau,\nu} U_{\tau,\nu}^T K(T_\tau, S_\nu) V_{\tau,\nu} V_{\tau,\nu}^T = U_{\tau,\nu} B_{\tau,\nu} V_{\tau,\nu}^T.$$

The hybrid butterfly factorization of  $K(T, S)$  is constructed as,

$$\begin{aligned} (3.13) \quad K(T, S) &= \begin{bmatrix} K(T_{\tau_1}, S_{\nu_1}) & K(T_{\tau_1}, S_{\nu_2}) & \cdots & K(T_{\tau_1}, S_{\nu_q}) \\ K(T_{\tau_2}, S_{\nu_1}) & K(T_{\tau_2}, S_{\nu_2}) & \cdots & K(T_{\tau_2}, S_{\nu_q}) \\ \vdots & \vdots & \ddots & \vdots \\ K(T_{\tau_p}, S_{\nu_1}) & K(T_{\tau_p}, S_{\nu_2}) & \cdots & K(T_{\tau_p}, S_{\nu_q}) \end{bmatrix} \\ &\approx \begin{bmatrix} U_{\tau_1, \nu_1} B_{\tau_1, \nu_1} V_{\tau_1, \nu_1}^T & U_{\tau_1, \nu_2} B_{\tau_1, \nu_2} V_{\tau_1, \nu_2}^T & \cdots & U_{\tau_1, \nu_q} B_{\tau_1, \nu_q} V_{\tau_1, \nu_q}^T \\ U_{\tau_2, \nu_1} B_{\tau_2, \nu_1} V_{\tau_2, \nu_1}^T & U_{\tau_2, \nu_2} B_{\tau_2, \nu_2} V_{\tau_2, \nu_2}^T & \cdots & U_{\tau_2, \nu_q} B_{\tau_2, \nu_q} V_{\tau_2, \nu_q}^T \\ \vdots & \vdots & \ddots & \vdots \\ U_{\tau_p, \nu_1} B_{\tau_p, \nu_1} V_{\tau_p, \nu_1}^T & U_{\tau_p, \nu_2} B_{\tau_p, \nu_2} V_{\tau_p, \nu_2}^T & \cdots & U_{\tau_p, \nu_q} B_{\tau_p, \nu_q} V_{\tau_p, \nu_q}^T \end{bmatrix} \\ (3.14) \quad &= (U^L R^{L-1} R^{L-2} \dots R^l) B^l (W^l W^{l-1} \dots W^1 V^0) \end{aligned}$$

where  $\tau_1, \tau_2, \dots, \tau_p$  are the  $p = 2^l$  nodes at level  $l$  of  $\mathcal{T}_T$ , and  $\nu_1, \nu_2, \dots, \nu_q$  are the  $q = 2^{L-l}$  nodes at level  $(L-l)$  of  $\mathcal{T}_S$ .

In this level- $l$  hybrid butterfly factorization, these column basis matrices  $U_{\tau,\nu}$  are recursively defined as in (3.5) using column basis and transfer matrices in the lower levels of  $\mathcal{T}_T$ , while the row basis matrices  $V_{\tau,\nu}$  are recursively defined as in (3.10) using row basis and transfer matrices at upper levels of  $\mathcal{T}_T$ . In (3.13), the outer factors  $U^L, V^0$  and inner factors  $R^a, W^a$  are defined in Sections 3.3.1 and 3.3.2, and the inner factor  $B^l$  consists of blocks  $B_{\tau,\nu}$  at level  $l$  in (3.14). Typically, the level  $l$  is set to  $l_m = \lfloor L/2 \rfloor$ . Figure 3.2(c) shows a hybrid factorization with  $L = 4$  and  $l_m = 2$ .

**4. Adaptive butterfly factorization via randomized matrix-vector products.** We propose an algorithm for constructing the butterfly factorization of a matrix  $A = K(T, S)$  using only products of  $A$  and its transpose with random vectors. The proposed Algorithm 4.1 returns a hybrid factorization with prescribed accuracy  $\epsilon$  assuming black-box matrix-vector multiplications. With minor modifications, Algorithm 4.1 also applies to column- and row-wise factorizations albeit with a much higher computational cost. In what follows, we describe the four key components of the algorithm, including

- The computation of  $K(T_\tau, S_\nu) \Omega_\nu$  and  $K(T_\tau, S_\nu)^T \Gamma_\tau$  with random matrices  $\Omega_\nu$  and  $\Gamma_\tau$  using a black-box routine.
- The construction of column basis matrices  $U_{\tau,\nu}$  (or transfer matrices  $R_{\tau,\nu}$  for non-leaf nodes  $\tau$  in  $\mathcal{T}_T$ ) based on matrix-vector multiplications involving  $K(T, S)$ .
- The construction of row basis matrices  $V_{\tau,\nu}$  (or transfer matrices  $W_{\tau,\nu}$  for non-leaf nodes  $\nu$  in  $\mathcal{T}_S$ ) based on the matrix-vector multiplications involving  $K(T, S)^T$ .
- The construction of intermediate matrices  $B_{\tau_a, \nu_b}$  in (3.13).



---

**Algorithm 4.1** Adaptive and randomized hybrid butterfly factorization based on matrix-vector multiplication

---

**Input:** Black-box routine for multiplying  $K(T, S) \in \mathbb{R}^{m \times n}$  and its transpose with arbitrary matrices, over-sampling parameter  $p$ , truncation tolerance  $\epsilon$ , initial rank guess  $r_0$ , binary partitioning trees  $\mathcal{T}_S$  and  $\mathcal{T}_T$  of  $L$  levels.

**Output:**  $K(T, S) \approx (U^L R^{L-1} R^{L-2} \dots R^{l_m}) B^{l_m} (W^{l_m} W^{l_m-1} \dots W^1 V^0)$  with  $l_m = \lfloor L/2 \rfloor$ .

```

1:  $r = r_0$ .
2: while not converged do                                 $\triangleright$  Adaptive computation of  $V_{\tau, \nu}$  at level 0
3:   Form a random matrix  $\Gamma_t \in \mathbb{R}^{|T_t| \times (r+p)}$  for the root node  $t$  of  $\mathcal{T}_T$ .
4:   Compute  $K(T_t, S)^T \Gamma_t$ .
5:   for  $\nu$  at level  $L$  of  $\mathcal{T}_S$  do
6:     Apply Algorithm 2.1 with  $A = K(T_t, S_\nu)^T$  to compute  $V_{t, \nu}$ .
7:   end for
8:   Converge if  $r > \max_{\nu} \{r_{t, \nu}\}$ .                     $\triangleright$  Over all nodes  $\nu$  at leaf level of  $\mathcal{T}_S$ 
9:    $r \leftarrow 2r$ .
10: end while
11: while not converged do                                 $\triangleright$  Adaptive computation of  $U_{\tau, \nu}$  at level  $L$ 
12:   Form a random matrix  $\Omega_s \in \mathbb{R}^{|S_s| \times (r+p)}$  for the root node  $s$  of  $\mathcal{T}_S$ .
13:   Compute  $K(T, S_s) \Omega_s$ .
14:   for  $\tau$  at level  $L$  of  $\mathcal{T}_T$  do
15:     Apply Algorithm 2.1 with  $A = K(T_\tau, S_s)$  to compute  $U_{\tau, s}$ .
16:   end for
17:   Converge if  $r > \max_{\tau} \{r_{\tau, s}\}$ .                     $\triangleright$  Over all nodes  $\tau$  at leaf level of  $\mathcal{T}_T$ 
18:    $r \leftarrow 2r$ .
19: end while
20: for  $l = 1$  to  $l_m$  do                                   $\triangleright$  Computation of  $W_{\tau, \nu}$ 
21:   for  $\tau$  at level  $l$  of  $\mathcal{T}_T$  do
22:     Form a random matrix  $\Gamma_\tau \in \mathbb{R}^{|T_\tau| \times (r+p)}$  with  $r = \max_{\nu} \{r_{p\tau, \nu_1} + r_{p\tau, \nu_2}\}$ .  $\triangleright$ 
     Over all nodes  $\nu$  at level  $L - l$  of  $\mathcal{T}_S$ 
23:     Compute  $K(T_\tau, S)^T \Gamma_\tau$ .
24:     for  $\nu$  at level  $L - l$  of  $\mathcal{T}_S$  do
25:       Compute  $\begin{bmatrix} V_{p\tau, \nu_1}^T & \\ & V_{p\tau, \nu_2}^T \end{bmatrix} K(T_\tau, S_\nu)^T \Gamma_\tau$ , apply Algorithm 2.1 for  $W_{\tau, \nu}$ .
26:     end for
27:   end for
28: end for
29: for  $l = L - 1$  to  $l_m$  do                               $\triangleright$  Computation of  $R_{\tau, \nu}$  and  $B_{\tau, \nu}$ 
30:   for  $\nu$  at level  $L - l$  of  $\mathcal{T}_S$  do
31:     Form a random matrix  $\Omega_\nu \in \mathbb{R}^{|S_\nu| \times (r+p)}$  with  $r = \max_{\tau} \{r_{\tau_1, p\nu} + r_{\tau_2, p\nu}\}$ .  $\triangleright$ 
     Over all nodes  $\tau$  at level  $l$  of  $\mathcal{T}_T$ 
32:     Compute  $K(T, S_\nu) \Omega_\nu$ .
33:     Compute  $V_{\tau, \nu}^T \Omega_\nu$  if  $l = l_m$ .                 $\triangleright$   $V_{\tau, \nu}^T$  is not explicitly computed.
34:     for  $\tau$  at level  $L - l$  of  $\mathcal{T}_T$  do
35:       Compute  $\begin{bmatrix} U_{\tau_1, p\nu}^T & \\ & U_{\tau_2, p\nu}^T \end{bmatrix} K(T_\tau, S_\nu) \Omega_\nu$ , apply Algorithm 2.1 for  $R_{\tau, \nu}$ .
36:       Compute  $B_{\tau, \nu}$  using (4.1) if  $l = l_m$ .
37:     end for
38:   end for
39: end for

```




FIG. 3.3. A 4-level hybrid factorization based on matrix-vector products consists of steps that compute (a)  $V^0$ , (b)  $W^0$ , (c)  $W^1$ , (d)  $U^4$ , (e)  $R^3$ , (f)  $R^2$ , and (g)  $B^2$ . Note that the vectors  $\bar{\Omega}_\nu$  or  $\bar{\Gamma}_\tau$  and the blocks being computed are marked with the same texture. The already-computed blocks needed at each step are plotted in Green.

#### 4.1. Multiplication of $K(T_\tau, S_\nu)$ and $K(T_\tau, S_\nu)^T$ by random matrices.

Here we assume the existence of a black-box program to perform matrix-vector multiplications involving  $K(T, S)$  and  $K(T, S)^T$ . To use [Algorithm 2.1](#) to compute the column/row basis matrices for each block  $K(T_\tau, S_\nu)$ , we multiply  $K(T, S)$  or  $K(T, S)^T$  with *structured random matrices*  $\bar{\Omega}_\nu$  and  $\bar{\Gamma}_\tau$  to obtain the matrices  $K(T_\tau, S_\nu)\bar{\Omega}_\nu$  and  $K(T_\tau, S_\nu)^T\bar{\Gamma}_\tau$  that are used in [Algorithm 2.1](#). As we shall see later,  $\bar{\Omega}_\nu$  and  $\bar{\Gamma}_\tau$  are sub-vectors of  $\bar{\Omega}_\nu$  and  $\bar{\Gamma}_\tau$ . Their entries are random variables which are independent and identically distributed, following a normal distribution.

Fixing a node  $\nu$  at level  $(L-l)$  of  $\mathcal{T}_S$ , we compute  $K(T_\tau, S_\nu)\bar{\Omega}_\nu$  with a  $|S_\nu| \times (r+p)$  matrix  $\bar{\Omega}_\nu$  for all nodes  $\tau$  at level  $l$  of  $\mathcal{T}_T$ , by multiplying  $K(T, S)$  with a sparse  $|S| \times (r+p)$  matrix  $\bar{\Omega}_\nu$  whose only non-zero entries  $\Omega_\nu$  are located on the rows corresponding to  $S_\nu$ . To evaluate all the multiplications  $K(T_\tau, S_\nu)\bar{\Omega}_\nu$  at level  $l$  (of  $\mathcal{T}_T$ ),  $2^{(L-l)}(r+p)$  matrix-vector multiplications by  $K(T, S)$  are needed.

Similarly, fixing a node  $\tau$  at level  $l$  of  $\mathcal{T}_T$ , we compute  $K(T_\tau, S_\nu)^T\bar{\Gamma}_\tau$  with a  $|T_\tau| \times (r+p)$  matrix  $\bar{\Gamma}_\tau$  for all nodes  $\nu$  at level  $(L-l)$  of  $\mathcal{T}_S$ , by multiplying  $K(T, S)^T$  with a sparse matrix  $\bar{\Gamma}_\tau \in \mathbb{R}^{|T| \times (r+p)}$  whose only non-zero entries  $\Gamma_\tau$  are located on the rows corresponding to  $T_\tau$ . To evaluate all the multiplications  $K(T_\tau, S_\nu)^T\bar{\Gamma}_\tau$  at level  $l$  (of  $\mathcal{T}_T$ ),  $2^l(r+p)$  matrix-vector multiplications by  $K(T, S)^T$  are needed. In [Algorithm 4.1](#), the products  $K(T_\tau, S_\nu)\bar{\Omega}_\nu$  and  $K(T_\tau, S_\nu)^T\bar{\Gamma}_\tau$  are performed on lines [4](#), [23](#), [13](#), [32](#).

#### 4.2. Computation of $U_{\tau,\nu}$ and $R_{\tau,\nu}$ .

For each leaf node  $\tau$  at level  $L$  of  $\mathcal{T}_T$  and the root node  $s$  of  $\mathcal{T}_S$ , the column basis matrix  $U_{\tau,s}$  can be directly computed using [Algorithm 2.1](#) by evaluating  $K(T_\tau, S_s)\bar{\Omega}_s$ . In [Algorithm 4.1](#) (line 2), the ranks of  $U_{\tau,s}$  are determined by adaptively doubling the size of the random matrices  $\bar{\Omega}_s$  for  $U_{\tau,s}$  until convergence. The iteration is terminated if the rank estimate  $r$  exceeds the

maximum revealed rank  $\max_{\nu} \{r_{t,\nu}\}$  (see line 17). This heuristic stopping criterion is used as a more rigorous criterion requires expensive computation of the approximation error.

For each non-leaf node  $\tau$  at level  $l_m \leq l < L$  of  $\mathcal{T}_T$  and each node  $\nu$  at level  $(L-l)$  of  $\mathcal{T}_S$ , the matrix to be compressed in (3.3) when computing  $R_{\tau,\nu}$  can be expressed as,

$$\begin{bmatrix} E_{\tau_1,\nu} \\ E_{\tau_2,\nu} \end{bmatrix} = \begin{bmatrix} U_{\tau_1,p\nu}^T & \\ & U_{\tau_2,p\nu}^T \end{bmatrix} K(T_\tau, S_\nu) \approx R_{\tau,\nu} E_{\tau,\nu}.$$

Thus,  $R_{\tau,\nu}$  can be computed via Algorithm 2.1 using the products

$$\begin{bmatrix} U_{\tau_1,p\nu}^T & \\ & U_{\tau_2,p\nu}^T \end{bmatrix} K(T_\tau, S_\nu) \Omega_\nu.$$

Note that no rank adaptation is needed for  $R_{\tau,\nu}$  in Algorithm 4.1 as the rank  $r_{\tau,\nu}$  is bounded by  $r_{\tau_1,p\nu} + r_{\tau_2,p\nu}$ . The dimensions of the random matrices  $\Omega_\nu$  therefore are chosen using the rank estimate  $r = \max_{\nu} \{r_{\tau_1,p\nu} + r_{\tau_2,p\nu}\}$  on line 31. This process recursively traverses  $\mathcal{T}_T$  from the leafs to center level  $l_m$  and  $\mathcal{T}_S$  from the root to center level  $l_m$ .

**4.3. Computation of  $V_{\tau,\nu}$  and  $W_{\tau,\nu}$ .** The computation of  $V_{\tau,\nu}$  and  $W_{\tau,\nu}$  resembles the above computation of  $U_{\tau,\nu}$  and  $R_{\tau,\nu}$ , but uses the multiplication results  $K(T_\tau, S_\nu)^T \Gamma_\tau$ .  $V_{\tau,\nu}$  for leaf nodes is computed adaptively on line 11 while  $W_{\tau,\nu}$  for non-leaf nodes is computed using Algorithm 2.1 using the multiplication results

$$\begin{bmatrix} V_{p\tau,\nu_1}^T & \\ & V_{p\tau,\nu_2}^T \end{bmatrix} K(T_\tau, S_\nu)^T \Gamma_\tau.$$

where the dimensions of the random matrices  $\Gamma_\tau$  are chosen using the rank estimate on line 22 without adaptation. This recursive construction starts from the leaf level of  $\mathcal{T}_S$  and ends at center level  $l_m$ .

**4.4. Computation of  $B_{\tau,\nu}$  at level  $l_m$ .** It follows from (3.13) that for each node  $\tau$  at level  $l_m$  of  $\mathcal{T}_T$  and node  $\nu$  at level  $(L-l_m)$  of  $\mathcal{T}_S$ ,  $K(T_\tau, S_\nu)$  is approximated as

$$K(T_\tau, S_\nu) \approx U_{\tau,\nu} B_{\tau,\nu} V_{\tau,\nu}^T.$$

Using the existing multiplication results  $K(T_\tau, S_\nu) \Omega_\nu$ ,  $B_{\tau,\nu}$  can be estimated as

$$\begin{aligned} B_{\tau,\nu} &= \underset{B \in \mathbb{R}^{r_{\tau,\nu} \times r_{\tau,\nu}}}{\operatorname{argmin}} \| (K(T_\tau, S_\nu) \Omega_\nu) - U_{\tau,\nu} B V_{\tau,\nu}^T \Omega_\nu \|_F \\ (4.1) \quad &= U_{\tau,\nu}^T (K(T_\tau, S_\nu) \Omega_\nu) (V_{\tau,\nu}^T \Omega_\nu)^T. \end{aligned}$$

Note that Algorithm 4.1 computes and stores  $K(T_\tau, S_\nu) \bar{\Omega}_\nu$  (or  $K(T_\tau, S_\nu)^T \bar{\Gamma}_\tau$ ) only for one  $\Omega_\nu$  at a time, therefore the algorithm is memory efficient. As an example, Figure 3.3 illustrates the procedure for constructing a 4-level hybrid factorization. Note that the random matrices for the inner factors are structured.

**4.5. Cost Analysis.** Let  $c(n)$  denote the number of operations for the black-box multiplication of  $K(S, T) \Omega$  or  $K(T, S)^T \Omega$  for an arbitrary  $n \times 1$  vector  $\Omega$ . In the best-case scenario,  $c(n) = O(n \log n)$  as  $K(T, S)$  is typically stored in a compressed form using  $O(n \log n)$  storage units; in the worst-case scenario,  $c(n) = O(n^2)$  when the matrix is explicitly stored in full. In what follows, we assume  $c(n) = O(n \log n)$ . Let  $r = \max_{\tau,\nu} \{r_{\tau,\nu}\}$  denote the maximum butterfly rank. Here we analyze the computation and memory costs of Algorithm 4.1 when applied to two classes of butterfly-compressible matrices: (i)  $r$  is constant (up to a logarithmic factor). (ii)  $r = O(n^{1/4})$ .

**4.5.1.  $r$  is constant.** This case typically occurs when the matrix arises from the discretization of 2D surface integral equations exploiting strong or weak admissibility conditions [5, 16, 17], 3D surface integral equation solvers using strong admissibility [6], low-dimensional Fourier operators [12], etc. For example, Figure 4.1(a) shows the center-level partitioning of a 2D curve used in surface integral-based Helmholtz equation solvers in which all blocks have constant rank except for  $O(1)$  ones with rank  $O(\log n)$ . (see [17] for a proof).

As described in subsection 4.1, there are  $2^l(r+p)$  black-box matrix-vector multiplications by  $K(T, S)^T$  at level  $l = 0, \dots, l_m$  and  $2^{(L-l)}(r+p)$  black-box matrix-vector multiplications by  $K(T, S)$  at level  $l = L, \dots, l_m$ . Therefore the black-box multiplications require a total of  $2(r+p)(1+2+\dots+2^{l_m})c(n) = O(rn^{1/2}c(n)) = O(rn^{3/2} \log n)$  operations. It is worth noting that the multiplications on lines 25 and 35 only involve partial factors  $V_{p_r, \nu_a}$  and  $U_{\tau_a, p_\nu}$  and their computational cost is dominated by that of the black-box multiplications. In addition, the algorithm only stores multiplication results for each random matrix of dimensions  $n \times (r+p)$  and the computed butterfly factors. The computation and memory costs of Algorithm 4.1 therefore scale as  $O(n^{3/2} \log n)$  and  $O(n \log n)$ , respectively.

**4.5.2.  $r$  is  $O(n^{1/4})$ .** This case often results from discretizing 3D surface integral equations using weak admissibility. For example, Figure 4.1(b) shows the center-level partitioning of a 3D surface used in surface integral methods for Helmholtz equations. Out of the  $16 \times 16 = 256$  center-level blocks of size  $O(n^{1/2}) \times O(n^{1/2})$ , only 4 blocks have rank  $O(n^{1/4})$  representing interactions between adjacent pairs. As the adjacent pair is typically co-planar, the edge shared by the pair can serve as the proxy surface that represent interactions between the pair. Therefore the rank is bounded by the edge dimension  $O(n^{1/4})$ . We claim, without proof, that except for  $O(n^{1/4})$  blocks with rank at most  $O(n^{1/4})$  at each level, the rest has constant rank.

We only need to focus on the construction of blocks with non-constant ranks as the rest has been analyzed above. We first analyze the cost of black-box multiplications on lines 4 and 23. For all blocks at level  $0 \leq l \leq l_m$ , denote  $N_l = \{\tau | r_{\tau, \nu} \text{ non constant}\}$ . It can be verified that  $|N_l| = O(2^{l/2})$  and each  $\tau \in N_l$  requires  $O(r+p) = O(n^{1/4})$  matrix-vector multiplications. Therefore the black-box multiplications  $K(T_\tau, S)^T \Gamma_\tau$  require a total of  $\sum_{l=0}^{l_m} |N_l| O(n^{1/4}) O(n \log n) = O(n^{1/2} c(n)) = O(n^{3/2} \log n)$  operations. A similar number of operations is required for  $K(T, S_\nu) \Omega_\nu$  on line 13 and 32. In addition, since each computed factor has  $O(n^{1/4})$  blocks of rank  $O(n^{1/4})$ , the storage cost is  $O(n^{1/4}) O(n^{1/2}) L \leq O(n \log n)$ . Overall the computation and memory costs of Algorithm 4.1 are of the same orders as those in the case of constant  $r$ .

**5. Error analysis.** In what follows, we assume that Algorithm 2.1 used in Algorithm 4.1 computes a low-rank approximation satisfying  $\|A - UU^T A\|_F \leq \epsilon \|A\|_F$  given a matrix  $A$  (see [7]).

**5.1. Approximation error for the nested basis  $U_{\tau, \nu}$ .** Given any level  $l_m \leq l \leq L$  of  $\mathcal{T}_T$ , matrix  $K(T, S)$  is partitioned into blocks  $\{K(T_\tau, S_\nu)\}$  with  $\tau$  at level  $l$  of  $\mathcal{T}_T$  and  $\nu$  at level  $(L-l)$  of  $\mathcal{T}_S$ , and  $K(T_\tau, S_\nu)$  is approximated by  $U_{\tau, \nu} U_{\tau, \nu}^T K(T_\tau, S_\nu)$  where  $U_{\tau, \nu}$  is the nested basis computed in subsection 4.2. Theorem 5.1 states that this approximation for all blocks at level  $l$  has an error bounded by  $\epsilon \sqrt{L-l+1} \|A\|_F$ .

**THEOREM 5.1.** *Given a level  $l$  in the range  $L \geq l \geq l_m$ , the low-rank approximations of all the blocks  $K(T_\tau, S_\nu)$  at level  $l$  ( $\tau$  at level  $l$  of  $\mathcal{T}_T$  and  $\nu$  at level  $(L-l)$  of  $\mathcal{T}_S$ ) based on the computed column basis matrices  $U_{\tau, \nu}$  have approximation errors*




FIG. 4.1. (a) Center-level geometrical partitioning of a 2D curve for a 4-level butterfly factorization of a weak-admissible matrix. The ranks for all center-level blocks are bounded by  $O(\log n)$ . (b) Center-level geometrical partitioning of a 3D surface for a 8-level butterfly factorization of a weak-admissible matrix. There are  $O(n^{1/4})$  center-level blocks with rank  $O(n^{1/4})$  (with the interaction pairs denoted by the red arrows)

bounded by

$$(5.1) \quad \sum_{\tau, \nu \text{ at level } l} \|K(T_\tau, S_\nu) - U_{\tau, \nu} U_{\tau, \nu}^T K(T_\tau, S_\nu)\|_F^2 \leq (L - l + 1) \epsilon^2 \|K(T, S)\|_F^2,$$

where the summation of  $\tau$  is over the nodes at level  $l$  of  $\mathcal{T}_T$  and the summation of  $\nu$  is over the nodes at level  $(L - l)$  of  $\mathcal{T}_S$ .

*Proof.* We prove the theorem by induction for level  $l$  changing from  $L$  to  $l_m$ . First, for  $l = L$ , the column basis matrix  $U_{\tau, s}$  for block  $K(T_\tau, S_s)$  at level  $L$  of  $\mathcal{T}_T$  is directly computed by [Algorithm 2.1](#). Thus, based on the above relative error threshold assumption, the approximation error is bounded by

$$\|K(T_\tau, S_s) - U_{\tau, s} U_{\tau, s}^T K(T_\tau, S_s)\|_F^2 \leq \epsilon^2 \|K(T_\tau, S_s)\|_F^2.$$

Moreover, summing over all nodes  $\tau$  on both sides of the above inequality proves (5.1) for  $l = L$ .

Assume that (5.1) holds true for level  $(l + 1)$ . For each node  $\tau$  at level  $l$  of  $\mathcal{T}_T$ , let  $\nu_1$  and  $\nu_2$  be siblings at level  $(L - l)$  of  $\mathcal{T}_S$ . Let  $\{\tau_1, \tau_2\}$  be the children of  $\tau$  at level  $(l + 1)$  of  $\mathcal{T}_S$  and  $\nu$  be the parent node of  $\{\nu_1, \nu_2\}$  at level  $(L - l - 1)$  of  $\mathcal{T}_S$ . Note that  $K(T_\tau, S_{\nu_1})$  and  $K(T_\tau, S_{\nu_2})$  are compressed blocks at level  $l$  while  $K(T_{\tau_1}, S_\nu)$  and  $K(T_{\tau_2}, S_\nu)$  are compressed blocks at level  $l + 1$ . Moreover, these two sets of blocks correspond to the same large block in  $K(T, S)$ , i.e.,

$$\begin{bmatrix} K(T_\tau, S_{\nu_1}) & K(T_\tau, S_{\nu_2}) \end{bmatrix} = \begin{bmatrix} K(T_{\tau_1}, S_\nu) \\ K(T_{\tau_2}, S_\nu) \end{bmatrix} = K(T_\tau, S_\nu).$$

Recall that  $U_{\tau, \nu_1}$  is computed in a nested fashion as

$$\begin{aligned} K(T_\tau, S_{\nu_1}) &= \begin{bmatrix} K(T_{\tau_1}, S_{\nu_1}) \\ K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} \approx \begin{bmatrix} U_{\tau_1, \nu} & \\ & U_{\tau_2, \nu} \end{bmatrix} \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} \\ &\approx \begin{bmatrix} U_{\tau_1, \nu} & \\ & U_{\tau_2, \nu} \end{bmatrix} R_{\tau, \nu_1} R_{\tau, \nu_1}^T \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} = U_{\tau, \nu_1} U_{\tau, \nu_1}^T K(T_\tau, S_{\nu_1}), \end{aligned}$$

where  $R_{\tau, \nu_1}$  is computed by applying [Algorithm 2.1](#) to the intermediate matrix above, i.e.,

$$(5.2) \quad \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} \approx R_{\tau, \nu_1} R_{\tau, \nu_1}^T \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix}.$$

The approximation error of block  $K(T_\tau, S_{\nu_1})$  at level  $l$  can be estimated as

$$\begin{aligned}
& \|K(T_\tau, S_{\nu_1}) - U_{\tau, \nu_1} U_{\tau, \nu_1}^T K(T_\tau, S_{\nu_1})\|_F^2 \\
&= \left\| K(T_\tau, S_{\nu_1}) - \begin{bmatrix} U_{\tau_1, \nu} & \\ & U_{\tau_2, \nu} \end{bmatrix} \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} \right\|_F^2 \\
&\quad + \left\| \begin{bmatrix} U_{\tau_1, \nu} & \\ & U_{\tau_2, \nu} \end{bmatrix} \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} - U_{\tau, \nu_1} U_{\tau, \nu_1}^T K(T_\tau, S_{\nu_1}) \right\|_F^2 \\
&= \|K(T_{\tau_1}, S_{\nu_1}) - U_{\tau_1, \nu} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1})\|_F^2 + \|K(T_{\tau_2}, S_{\nu_1}) - U_{\tau_2, \nu} U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1})\|_F^2 \\
&\quad + \left\| \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} - R_{\tau, \nu_1} R_{\tau, \nu_1}^T \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} \right\|_F^2
\end{aligned}$$

where the first equation is due to the orthogonality between the columns from the two bracketed terms, and the second equation follows from the unitary invariance of the Frobenius norm, as  $U_{\tau_1, \nu}$  and  $U_{\tau_2, \nu}$  both have orthonormal columns. The last term above, which is the approximation error of (5.2), is guaranteed to satisfy

$$\epsilon^2 \left\| \begin{bmatrix} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_1}) \\ U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} \right\|_F^2 \leq \epsilon^2 \left\| \begin{bmatrix} K(T_{\tau_1}, S_{\nu_1}) \\ K(T_{\tau_2}, S_{\nu_1}) \end{bmatrix} \right\|_F^2 = \epsilon^2 \|K(T_\tau, S_{\nu_1})\|_F^2.$$

Similarly, we can estimate the approximation error of block  $K(T_\tau, S_{\nu_2})$  at level  $l$  as

$$\begin{aligned}
& \|K(T_\tau, S_{\nu_2}) - U_{\tau, \nu_2} U_{\tau, \nu_2}^T K(T_\tau, S_{\nu_2})\|_F^2 \\
&\leq \|K(T_{\tau_1}, S_{\nu_2}) - U_{\tau_1, \nu} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_{\nu_2})\|_F^2 \\
&\quad + \|K(T_{\tau_2}, S_{\nu_2}) - U_{\tau_2, \nu} U_{\tau_2, \nu}^T K(T_{\tau_2}, S_{\nu_2})\|_F^2 + \epsilon^2 \|K(T_\tau, S_{\nu_2})\|_F^2.
\end{aligned}$$

Assembling the above three inequalities and using  $S_\nu = S_{\nu_1} \cup S_{\nu_2}$ , we obtain

$$\begin{aligned}
& \|K(T_\tau, S_{\nu_1}) - U_{\tau, \nu_1} U_{\tau, \nu_1}^T K(T_\tau, S_{\nu_1})\|_F^2 + \|K(T_\tau, S_{\nu_2}) - U_{\tau, \nu_2} U_{\tau, \nu_2}^T K(T_\tau, S_{\nu_2})\|_F^2 \\
&\leq \|K(T_{\tau_1}, S_\nu) - U_{\tau_1, \nu} U_{\tau_1, \nu}^T K(T_{\tau_1}, S_\nu)\|_F^2 + \|K(T_{\tau_2}, S_\nu) - U_{\tau_2, \nu} U_{\tau_2, \nu}^T K(T_{\tau_2}, S_\nu)\|_F^2 \\
&\quad + \epsilon (\|K(T_\tau, S_{\nu_1})\|_F^2 + \|K(T_\tau, S_{\nu_2})\|_F^2).
\end{aligned}$$

The left hand side of the inequality is the approximation error of the two blocks at level  $l$  while the first two terms on the right hand side are the approximation errors of two blocks at level  $(l+1)$ . Summing over all the nodes  $\tau$  at level  $l$  of  $\mathcal{T}_T$  and all the node pairs  $(\nu_1, \nu_2)$  at level  $(L-l)$  of  $\mathcal{T}_S$  on both sides of this inequality, we obtain,

$$\begin{aligned}
& \sum_{\tau, \nu \text{ at level } l} \|K(T_\tau, S_\nu) - U_{\tau, \nu} U_{\tau, \nu}^T K(T_\tau, S_\nu)\|_F^2 \\
&\leq \sum_{\tau, \nu \text{ at level } (l+1)} \|K(T_\tau, S_\nu) - U_{\tau, \nu} U_{\tau, \nu}^T K(T_\tau, S_\nu)\|_F^2 + \epsilon^2 \|K(T, S)\|_F^2 \\
&\leq (L - (l+1) + 1) \epsilon^2 \|K(T, S)\|_F^2 + \epsilon^2 \|K(T, S)\|_F^2 = (L - l + 1) \epsilon^2 \|K(T, S)\|_F^2. \square
\end{aligned}$$

**5.2. Approximation error for the nested basis  $V_{\tau, \nu}$ .** The overall approximation error for the projection to the row bases is bounded just like that of the column bases shown in [Theorem 5.2](#).

**THEOREM 5.2.** *Given a level  $l$  in the range  $L \geq l \geq l_m$ , the low-rank approximations of all the blocks  $K(T_\tau, S_\nu)$  at level  $l$  ( $\tau$  at level  $l$  of  $\mathcal{T}_T$  and  $\nu$  at level  $(L-l)$  of  $\mathcal{T}_S$ ) based on the computed row basis matrices  $V_{\tau,\nu}$  have approximation errors bounded as*

$$(5.3) \quad \sum_{\tau, \nu \text{ at level } l} \|K(T_\tau, S_\nu) - K(T_\tau, S_\nu)V_{\tau,\nu}V_{\tau,\nu}^T\|_F^2 \leq (l+1)\epsilon^2 \|K(T, S)\|_F^2,$$

where the summation of  $\tau$  is over the nodes at level  $l$  of  $\mathcal{T}_T$  and the summation of  $\nu$  is over the nodes at level  $(L-l)$  of  $\mathcal{T}_S$ .

**5.3. Approximation error for the overall factorization.** Combining the above two error analyses for the column- and row-wise butterfly factorizations, the overall approximation error of a hybrid butterfly factorization at any level  $l$  can be bounded as shown in [Theorem 5.3](#)

**THEOREM 5.3.** *Given center level  $l_m = \lfloor L/2 \rfloor$ , the low-rank approximations of all the blocks  $K(T_\tau, S_\nu)$  at level  $l_m$  ( $\tau$  at level  $l_m$  of  $\mathcal{T}_T$  and  $\nu$  at level  $(L-l_m)$  of  $\mathcal{T}_S$ ) based on the computed basis matrices  $U_{\tau,\nu}$  and  $V_{\tau,\nu}$  have an overall approximation error bounded by*

$$(5.4) \quad \sum_{\tau, \nu \text{ at level } l} \|K(T_\tau, S_\nu) - U_{\tau,\nu}U_{\tau,\nu}^TK(T_\tau, S_\nu)V_{\tau,\nu}V_{\tau,\nu}^T\|_F^2 \leq (L+2)\epsilon^2 \|K(T, S)\|_F^2,$$

where the summation of  $\tau$  is over the nodes at level  $l_m$  of  $\mathcal{T}_T$  and the summation of  $\nu$  is over the nodes at level  $(L-l_m)$  of  $\mathcal{T}_S$ .

*Proof.* Every block  $K(T_\tau, S_\nu)$  has its approximation error bounded as

$$\begin{aligned} & \|K(T_\tau, S_\nu) - U_{\tau,\nu}U_{\tau,\nu}^TK(T_\tau, S_\nu)V_{\tau,\nu}V_{\tau,\nu}^T\|_F^2 \\ & \leq \|K(T_\tau, S_\nu) - U_{\tau,\nu}U_{\tau,\nu}^TK(T_\tau, S_\nu)\|_F^2 \\ & \quad + \|U_{\tau,\nu}U_{\tau,\nu}^TK(T_\tau, S_\nu) - U_{\tau,\nu}U_{\tau,\nu}^TK(T_\tau, S_\nu)V_{\tau,\nu}V_{\tau,\nu}^T\|_F^2 \\ & \leq \|K(T_\tau, S_\nu) - U_{\tau,\nu}U_{\tau,\nu}^TK(T_\tau, S_\nu)\|_F^2 + \|K(T_\tau, S_\nu) - K(T_\tau, S_\nu)V_{\tau,\nu}V_{\tau,\nu}^T\|_F^2. \end{aligned}$$

Using [\(5.1\)](#) and [\(5.3\)](#) in the above equation, inequality [\(5.4\)](#) is proven.  $\square$

**6. Parallelization.** This section outlines a distributed-memory implementation of [Algorithm 4.1](#). We first consider the task of parallelizing a butterfly-vector multiplication assuming that the butterfly factors are stored in a distributed fashion. Note that this is the dominant computational task in the proposed algorithm as the black-box multiplications on lines [4](#), [23](#), [13](#), [32](#) often involve existing butterfly representations, and the explicit multiplications on lines [33](#), [25](#), [35](#) involve partial butterfly factors. A good parallelization strategy for these two types of multiplications therefore is paramount to the parallel implementation of the proposed randomized butterfly reconstruction scheme.

Without loss of generality, we assume that  $r_{\tau,\nu} = r$  for some constant  $r$ , the number of butterfly levels  $L$  is even,  $n = r2^L$ , and the number of processes  $1 \leq p \leq 2^L$  is a power of two. To estimate the algorithm's communication cost, we only analyze the number of messages and communication volume as we assume the time to communicate a message of size  $m$  between two processes is  $\alpha + \beta m$  where  $\alpha$  and  $\beta$  represent message latency and inverse bandwidth, respectively [\[9\]](#).



**6.1. Column/Row-wise factorization.** We first describe the parallelization scheme for the column-wise butterfly factorization studied in [22]. The parallel data layout can be described as follows: (i) starting from level  $L$  of  $\mathcal{T}_T$ , one process stores  $2^L/p$  consecutive blocks  $U_{\tau,s}$  following a 1D-row layout. (ii) At levels  $l = L-1, \dots, 1$ , let  $\{\tau_1, \tau_2\}$  be the children of  $\tau$  at level  $l$  of  $\mathcal{T}_T$  and  $\{\nu_1, \nu_2\}$  be the children of  $\nu$  at level  $L-l-1$  of  $\mathcal{T}_S$ . Consider the *combined* transfer matrix:

$$(6.1) \quad R_{\tau,\nu} = \begin{bmatrix} R_{\tau_1,\nu} & \\ & R_{\tau_2,\nu} \end{bmatrix} [R_{\tau,\nu_1} \ R_{\tau,\nu_2}]$$

where  $R_{\tau_a,\nu}$  can be replaced by  $U_{\tau_a,\nu}$  if  $l = L-1$ . The parallelization scheme stores  $R_{\tau_a,\nu}$  (or  $U_{\tau_a,\nu}$  if  $l = L-1$ ) and  $R_{\tau,\nu_a}$  on the same process. (iii) At level  $l = 0$ , the blocks  $E_{t,\nu}$  and  $R_{t,\nu}$  are stored on the same process. Figure 6.1(a) illustrates the data layout for a 4-level column-wise butterfly factorization using 4 processes. Note that at level  $l = 0$ , the layout of  $E_{t,\nu}$  is similar to the 1D-row layout for  $U_{\tau,s}$ , but in an index-reversed order.

When multiplying a (partial) butterfly stored as described above by a vector stored using the 1D-row layout, an all-to-all communication is required to convert the vector from the 1D-row layout to the index-reversed layout of  $E_{t,\nu}$ . To this end, each process communicates  $\min\{2^L/p, p-1\}$  messages of total size  $r2^L/p$ . For the example in Figure 6.1(a), process 0 (light green) needs to scatter the locally stored part of  $\Omega$  of size  $2^L/p \times r$  to all other processes before the multiplication operation with  $E^0$  can take place. After that, there are  $\log p$  levels requiring pair-wise exchanges of the intermediate multiplication results  $R^l \dots R^0 E^0 \Omega$  (e.g., after multiplication with  $R_0$  and  $R_1$  in Figure 6.1(a)). For each exchange operation, reductions involving messages of size  $r2^L/p$  between two processes are performed (note: broadcast is conducted for multiplying the transpose  $K(T, S)^T$ ). For example, considering the multiplication with  $R_{\tau,\nu_1}$  (stored on process 0 in light green) and  $R_{\tau,\nu_2}$  (stored on process 2 in dark green) in the first diagonal block of  $R^0$ , the local multiplication results of size  $2r \times r$  require a reduction between processes 0 and 2 before the local multiplication with the first diagonal block of  $R^1$  is performed. The communication costs are listed in Table 6.1.

The parallelization of the row-wise butterfly factorization can be described similarly: At level  $l = 0$  of  $\mathcal{T}_0$ ,  $V_{t,\nu}$  is stored using the 1D-row layout; at levels  $l = 1, \dots, L$ ,  $W_{\tau,\nu_a}$  (or  $V_{\tau,\nu_a}$  if  $l = 1$ ) and  $W_{\tau_a,\nu}$  are stored on the same process. The communication costs for the row-wise parallelization are the same as those for the column-wise parallelization.

**6.2. Hybrid factorization.** Since Algorithm 4.1 computes a hybrid factorization, it is more convenient and efficient to combine the column-wise and row-wise parallel data layouts. Specifically, the factors  $U^L R^{L-1} R^{L-2} \dots R^{l_m}$  are stored using the column-wise layout whereas those of  $W^{l_m} W^{l_m-1} \dots W^1 V^0$  adhere to the row-wise layout. The block  $B_{\tau,\nu}$  is handled by the same process as  $W_{\tau,\nu}$  in  $W^{l_m}$ . When multiplying the (partial) butterfly with a vector  $\Omega$ , all-to-all communication is needed for the intermediate result  $B^{l_m} W^{l_m} W^{l_m-1} \dots W^1 V^0 \Omega$ . In contrast to the column/row-wise layout, the number of levels requiring pair-wise exchange is only  $\max\{0, \log p^2/2^L\}$ . Note that no exchange is needed if  $p \leq 2^{L/2}$ . As an example, Figure 6.1(b) shows the multiplication of a 4-level hybrid butterfly using 4 processes. Here,  $p = 2^{L/2}$  and the only communication is the all-to-all operation that occurs after multiplying with  $B^2$  to switch from the row- to the column-wise layout. Table 6.1 lists the communication costs for the hybrid data layout. Clearly, the multiplication



|            | exchange   |                                   | all-to-all       |                              |
|------------|--|-----------------------------------|------------------|------------------------------|
|            | volume   | message count                     | volume           | message count                |
| column/row | $\frac{r2^L \log p}{p}$                          | $\log p$                          | $\frac{r2^L}{p}$ | $\min\{\frac{2^L}{p}, p-1\}$ |
| hybrid     | $\frac{r2^L}{p} \max\{0, \log \frac{p^2}{2^L}\}$ | $\max\{0, \log \frac{p^2}{2^L}\}$ | $\frac{r2^L}{p}$ | $\min\{\frac{2^L}{p}, p-1\}$ |

TABLE 6.1

Communication volume and message counts for one matrix-vector multiplication.

using the hybrid factorization requires less communication than multiplication with the column/row-wise factorization.

Now we can summarize the proposed parallelization strategy for Algorithm 4.1 as follows:

- The black-box multiplications on lines 4, 23, 13, 32 follow the hybrid layout if they involve existing parallel butterfly representations.
- The computed butterfly factors follow the hybrid layout for each process that applies Algorithm 2.1 on lines 25 and 35 locally for the blocks it is in charge of.
- The explicit multiplications with computed butterfly blocks on lines 33, 25, 35 follow the hybrid layout but may not involve all the processes.




FIG. 6.1. Parallel data layout for distributing a 4-level (a) column-wise and (b) hybrid factorization with 4 processes. Each color represents one process. The arrows denote places where exchange and all-to-all communications are needed for multiplying the butterfly with a 1D row distributed vector. Note that no exchange is needed for the hybrid butterfly.

**7. Numerical results.** This section provides several examples to demonstrate the accuracy and efficiency of the proposed randomized algorithm. The accuracy of the proposed algorithms is characterized by

$$(7.1) \quad \text{error} = \frac{\|A\Omega - (U^L R^{L-1} R^{L-2} \dots R^l) B^l (W^l W^{l-1} \dots W^1 V^0) \Omega\|_F}{\|A\Omega\|_F}$$

with a random testing matrix  $\Omega$  of 16 columns. All experiments are performed on the Cori Haswell machine at NERSC, which is a Cray XC40 system and consists of 2388 dual-socket nodes with Intel Xeon E5-2698v3 processors running 16 cores per socket. The nodes are configured with 128 GB of DDR4 memory clocked at 2133 MHz. Unless stated otherwise, all experiments use one Cori node.

**7.1. Exact butterfly factorization.** We first apply the algorithm to “recover” a matrix with known butterfly factorization. In what follows, we use the symbol “ $\bar{a}$ ” to differentiate the blocks and factors of the known butterfly from the computed ones. Let  $A$  be a  $n \times n$  matrix with a given  $L$ -level butterfly factorization  $A = (\bar{U}^L \bar{R}^{L-1} \bar{R}^{L-2} \dots \bar{R}^1) \bar{B}^l (\bar{W}^l \bar{W}^{l-1} \dots \bar{W}^1 \bar{V}^0)$  that has  $r_{\tau,\nu} = r$  for all blocks at all levels for some constant  $r$ . The blocks  $\bar{W}_{\tau,\nu}$ ,  $\bar{R}_{\tau,\nu}$ , and  $\bar{B}_{\tau,\nu}$  have dimensions  $r \times 2r$ ,  $2r \times r$  and  $r \times r$ , respectively. We choose the matrix dimension  $n = 2^{L+3}$  such that  $\bar{U}_{\tau,s}$  and  $\bar{V}_{t,\nu}$  have dimensions  $8 \times r$ . The blocks  $\bar{U}_{\tau,s}$ ,  $\bar{V}_{t,\nu}$ ,  $\bar{W}_{\tau,\nu}$ ,  $\bar{R}_{\tau,\nu}$  are constructed as random unitary matrices, and the entries of  $\bar{B}_{\tau,\nu}$  are random variables which are independent and identically distributed, following a normal distribution. We use this explicit representation to perform “black-box” matrix-vector multiplications and apply [Algorithm 4.1](#) to retrieve the butterfly factorization.

The memory costs for varying  $n$  with fixed  $r = 8$  when using the proposed algorithm and the randomized SVD-based reference algorithm in [\[12\]](#) are plotted in [Figure 7.1\(a\)](#). We set  $p = 2$  and  $r_0 = 4$  in [Algorithm 4.1](#). The proposed algorithm requires only  $O(n \log n)$  memory as opposed to the reference algorithm, which requires  $O(n^{1.5})$  memory.

Next, the performance of the proposed parallelization scheme is demonstrated by applying [Algorithm 4.1](#) to the butterfly representation with  $n = 2.56 \times 10^6$  and  $r = 10$ . [Figure 7.1\(b\)](#) plots the runtime of [Algorithm 4.1](#) for process counts 16 to 2048. The runtime for a single matrix-vector multiplication when the input butterfly representation follows hybrid and column-wise data layouts is also shown. As predicted by [Table 6.1](#), the hybrid data layout yields a substantially lower communication cost.




FIG. 7.1. (a) Memory for applying [Algorithm 4.1](#) and the reference algorithms to an exact butterfly representation. (b) Runtime for applying [Algorithm 4.1](#) and matrix-vector multiplications to the exact butterfly representation for varying process counts.

**7.2. 2D Helmholtz kernel.** Next, consider the following wave scattering example. Let  $C^1$  and  $C^2$  denote two disjoint curves. Suppose  $C^1$  and  $C^2$  are partitioned into  $n$  and  $m$  constant-sized segments  $C_i^1$  and  $C_j^2$ ,  $i = 1, \dots, n$ ,  $j = 1, \dots, m$ . Let  $k_0$  denote the wavenumber, then  $m, n = O(k_0)$ . Consider the following  $m \times n$  matrix  $A$

$$(7.2) \quad A = Z^{21} (Z^{11})^{-1}$$

$$(7.3) \quad Z_{i,j}^{11} = \int_{C_j^1} H_0^{(2)}(k_0 |\rho_i^1 - \rho|) d\rho, \quad i, j = 1, \dots, n$$

$$(7.4) \quad Z_{i,j}^{21} = \int_{C_j^2} H_0^{(2)}(k_0 |\rho_i^2 - \rho|) d\rho, \quad i = 1, \dots, m, \quad j = 1, \dots, n$$

| $n$     | $L$ | $\epsilon$ | $r$ | error    | Time (sec) | Memory (MB) |
|---------|-----|------------|-----|----------|------------|-------------|
| 20000   | 9   | 1E-03      | 10  | 3.63E-04 | 3.61E+01   | 1.68E+01    |
| 80000   | 11  | 1E-03      | 10  | 4.22E-04 | 3.27E+02   | 7.71E+01    |
| 320000  | 13  | 1E-03      | 10  | 4.98E-04 | 3.27E+03   | 3.47E+02    |
| 1280000 | 15  | 1E-03      | 10  | 7.38E-04 | 3.02E+04   | 1.58E+03    |
| 20000   | 9   | 1E-04      | 12  | 2.45E-05 | 7.80E+01   | 3.13E+01    |
| 80000   | 11  | 1E-04      | 12  | 2.39E-05 | 7.21E+02   | 1.32E+02    |
| 320000  | 13  | 1E-04      | 12  | 2.37E-05 | 6.46E+03   | 5.99E+02    |
| 1280000 | 15  | 1E-04      | 12  | 3.26E-05 | 5.87E+04   | 2.64E+03    |
| 20000   | 9   | 1E-05      | 14  | 8.09E-06 | 8.10E+01   | 3.71E+01    |
| 80000   | 11  | 1E-05      | 14  | 8.39E-06 | 7.82E+02   | 1.71E+02    |
| 320000  | 13  | 1E-05      | 14  | 9.11E-06 | 7.17E+03   | 7.92E+02    |
| 1280000 | 15  | 1E-05      | 14  | 9.57E-06 | 6.17E+04   | 3.56E+03    |

TABLE 7.1

Time, memory and measured error for computing a hybrid factorization of (7.2) using the proposed algorithm with varying matrix size  $n$  and tolerance  $\epsilon$ .

Here,  $H_0^2(\cdot)$  is the zeroth-order Hankel function of the second kind,  $\rho$  is a position vector on  $C^1$ ,  $\rho_i^1$  and  $\rho_i^2$  denote the center of segment  $C_i^1$  and  $C_i^2$ . In principle, the  $n \times n$  matrix  $(Z^{11})^{-1}$  relates the equivalent source on  $C^1$  to the incident fields on  $C^1$ , and the  $m \times n$  matrix  $Z^{21}$  computes fields observed on  $C^2$  scattered by the source on  $C^1$ . The matrix  $A$  in (7.2) resembles the *scattering matrix* as it relates the incident fields on  $C^1$  to its scattered fields on  $C^2$  [8, 16]. In what follows, we seek a butterfly-compressed representation of the matrix  $A$ .

In this example, suppose  $C^1$  and  $C^2$  are two parallel lines with length  $L$  and their respective distance  $L$ . It is well known that  $Z^{11}$  and its inverse have compressed representations using  $\mathcal{H}$ -matrix or other low-rank factorization-based hierarchical techniques, requiring at most  $O(n \log n)$  computation and memory resources [2, 4, 18]. In addition,  $Z^{21}$  can be compressed by the butterfly factorization requiring  $O(n \log n)$  computation and memory resources [16, 17]. Therefore,  $A$  and its transpose can be applied to any vector in  $O(n \log n)$  operations irrespective of wavenumber  $k_0$ .

To compute a butterfly-compressed representation of  $A$  using the proposed algorithm, the lengths of line segments  $C_i^1$  and  $C_i^2$  are set to approximately  $0.05\lambda$  with  $\lambda = 2\pi/k_0$  denoting the wavelength. The sizes of the leaf-level point sets  $T_\tau$  and  $S_\nu$  are set to approximately 39. The matrices  $(Z^{11})^{-1}$  and  $Z^{21}$  are compressed respectively with the  $\mathcal{H}$ -matrix and butterfly factorization with a high accuracy (tolerance  $10^{-8}$ ), respectively. Note that these compressed representations are used instead of  $A$  in (7.1) and in black-box multiplications.

The computational results with butterfly level  $L = 9, 11, 13, 15$  for different tolerances  $\epsilon = 10^{-3}, 10^{-4}, 10^{-5}$  are listed in Table 7.1. We set  $p = 2$  and  $r_0 = 4$  in Algorithm 4.1. The measured error scales proportional to  $L$  as predicted in Theorem 5.3. For  $\epsilon = 10^{-3}$ , the computationally most expensive case (when  $n = 1.28 \times 10^6$ ) requires about 9 h CPU time and 1.6 GB memory. Moreover, the memory cost and the factorization time obey the predicted  $O(n \log n)$  and  $O(n^{1.5} \log n)$  scaling estimates.

The computation time and memory cost with  $l = 8, \dots, 15$  and  $\epsilon = 10^{-3}$  are plotted in Figure 7.2 using the proposed algorithm and the reference algorithm in [10]. The proposed algorithm is slightly slower than the reference algorithm as it requires slightly more testing vectors to reduce the memory cost. That said, the proposed algorithm




FIG. 7.2. (a) Computation time and (b) memory for applying [Algorithm 4.1](#) and the reference algorithms to [\(7.2\)](#).

requires much less memory than the reference algorithm (see [Figure 7.2\(b\)](#)).

**7.3. 3D Helmholtz kernel.** Finally, we consider the wave interactions between two semi-sphere surfaces  $C^1$  and  $C^2$  of unit radius adjacent to each other. Each semi-sphere is discretized via the Nystrm method into  $n$  sample points. We seek a butterfly factorization of the following  $n \times n$  matrix  $A$ :

$$(7.5) \quad A_{i,j} = \frac{e^{i2\pi\kappa|\rho_i - \rho_j|}}{|\rho_i - \rho_j|}$$

with  $\rho_i$  and  $\rho_j$  denoting sample point  $i$  and  $j$  on  $C^1$  and  $C^2$ , respectively. The wavenumber  $\kappa$  is set such that  $n = 50\kappa^2/\pi$  represents approximately 10 sample points per wavelength. We first compute  $A \approx (\bar{U}^L \bar{R}^{L-1} \bar{R}^{L-2} \dots \bar{R}^1) \bar{B}^l (\bar{W}^l \bar{W}^{l-1} \dots \bar{W}^1 \bar{V}^0)$  with  $\epsilon = 10^{-6}$  and use the result for [\(7.1\)](#) and the black-box multiplications.

The computational results with butterfly level  $L = 8, 10, 12, 14$  for different tolerances  $\epsilon = 10^{-2}, 10^{-3}, 10^{-4}$  are listed in [Table 8.1](#). We set  $p = 4$  and  $r_0 = 64$  in [Algorithm 4.1](#). These experiments use 2 Cori nodes with a total of 64 MPI processes. The proposed algorithm achieves the desired accuracies predicted by [Theorem 5.3](#). The computation time, memory and observed rank with  $\epsilon = 10^{-2}$  are plotted in [Figure 8.1](#) using the proposed algorithm. Despite the  $O(n^{0.25})$  rank scaling (see [Figure 4.1\(b\)](#)) as an illustration), [Algorithm 4.1](#) still attains  $O(n^{1.5} \log n)$  computation and  $O(n \log n)$  memory complexities as estimated in [subsection 4.5.2](#). In addition, the estimated memory usage with the reference algorithm is also plotted in [Figure 8.1\(b\)](#), the proposed algorithm requires much less memory in comparison.

**8. Conclusion and discussion.** This paper presented a fast and memory-efficient randomized algorithm for computing the butterfly factorization of a matrix assuming the availability of a black-box algorithm for applying the matrix and its transpose to a vector. The proposed algorithm applies the matrix and its transpose to structured random vectors to reconstruct the orthonormal row and column bases of judiciously-selected low-rank blocks of the (assumed) butterfly-compressible matrix. The algorithm only requires  $O(n^{1.5} \log n)$  computation and  $O(n \log n)$  storage resources for matrices arising from the integral equation based discretization of both 2D and 3D Helmholtz problems using either weak or strong admissibility separation criteria. The accuracy of the proposed algorithm only weakly depends on the number of butterfly levels. The computation time of the algorithm can be reduced leveraging distributed-memory parallelism. We expect that the proposed algorithm will play an important role in constructing both dense and sparse, fast and

| $n$     | $L$ | $\epsilon$ | $r$ | error    | Time (sec) | Memory (MB) |
|---------|-----|------------|-----|----------|------------|-------------|
| 23806   | 8   | 1E-02      | 85  | 1.10E-02 | 8.03E-01   | 3.78E+01    |
| 94024   | 10  | 1E-02      | 122 | 1.23E-02 | 5.69E+00   | 1.56E+02    |
| 379065  | 12  | 1E-02      | 173 | 1.53E-02 | 4.78E+01   | 6.60E+02    |
| 1510734 | 14  | 1E-02      | 250 | 1.62E-02 | 3.67E+02   | 2.82E+03    |
| 23806   | 8   | 1E-03      | 136 | 1.07E-03 | 1.59E+00   | 7.82E+01    |
| 94024   | 10  | 1E-03      | 190 | 1.25E-03 | 1.04E+01   | 3.23E+02    |
| 379065  | 12  | 1E-03      | 277 | 1.46E-03 | 8.79E+01   | 1.37E+03    |
| 1510734 | 14  | 1E-03      | 408 | 1.59E-03 | 6.83E+02   | 5.90E+03    |
| 23806   | 8   | 1E-04      | 175 | 7.93E-05 | 2.53E+00   | 1.32E+02    |
| 94024   | 10  | 1E-04      | 250 | 8.85E-05 | 1.71E+01   | 5.56E+02    |
| 379065  | 12  | 1E-04      | 356 | 1.01E-04 | 1.56E+02   | 2.39E+03    |
| 1510734 | 14  | 1E-04      | 522 | 1.12E-04 | 1.47E+03   | 1.04E+04    |

TABLE 8.1

Time, memory and measured error for computing a hybrid factorization of (7.5) using the proposed algorithm with varying matrix size  $n$  and tolerance  $\epsilon$ .

parallel, hierarchical matrix-based direct solvers for high-frequency wave equations. The code described here is part of the Fortran/C++ solver package ButterflyPACK (<https://github.com/liuyangzhuan/ButterflyPACK>), freely available online. The integration of butterfly factorizations into the sparse solver package STRUMPACK is currently in progress.

**Acknowledgments.** This research was supported in part by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the U.S. Department of Energy Office of Science and the National Nuclear Security Administration, and in part by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Scientific Discovery through Advanced Computing (SciDAC) program through the FASTMath Institute under Contract No. DE-AC02-05CH11231 at Lawrence Berkeley National Laboratory.

This research used resources of the National Energy Research Scientific Computing Center (NERSC), a U.S. Department of Energy Office of Science User Facility operated under Contract No. DE-AC02-05CH11231.




FIG. 8.1. (a) Computation time (left y-axis), observed rank (right y-axis) and (b) memory for applying Algorithm 4.1 and the reference algorithms to (7.5).

## REFERENCES

- [1] E. CANDS, L. DEMANET, AND L. YING, *A Fast Butterfly Algorithm for the Computation of Fourier Integral Operators*, *Multiscale Modeling & Simulation*, 7 (2009), pp. 1727–1750.
- [2] S. CHANDRASEKARAN, P. DEWILDE, M. GU, W. LYONS, AND T. PALS, *A fast solver for hss representations via sparse matrices*, *SIAM Journal on Matrix Analysis and Applications*, 29 (2007), pp. 67–81.
- [3] C. GORMAN, G. CHVEZ, P. GHYSELS, T. MARY, F.-H. ROUET, AND X. S. LI, *Robust and accurate stopping criteria for adaptive randomized sampling in matrix-free hierarchically semiseparable construction*, *SIAM Journal on Scientific Computing*, 41 (2019), pp. S61–S85.
- [4] L. GRASEDYCK AND W. HACKBUSCH, *Construction and arithmetics of H-matrices*, *Computing*, 70 (2003), pp. 295–334.
- [5] H. GUO, J. HU, AND E. MICHIELSSEN, *On mlmda/butterfly compressibility of inverse integral operators*, *IEEE Antennas and Wireless Propagation Letters*, 12 (2013), pp. 31–34.
- [6] H. GUO, Y. LIU, J. HU, AND E. MICHIELSSEN, *A Butterfly-Based Direct Integral-Equation Solver Using Hierarchical LU Factorization for Analyzing Scattering From Electrically Large Conducting Objects*, *IEEE Transactions on Antennas and Propagation*, 65 (2017), pp. 4742–4750.
- [7] N. HALKO, P. MARTINSSON, AND J. TROPP, *Finding structure with randomness: probabilistic algorithms for constructing approximate matrix decompositions*, *SIAM Review*, 53 (2011), pp. 217–288.
- [8] S. HAO, P. MARTINSSON, AND P. YOUNG, *An efficient and highly accurate solver for multi-body acoustic scattering problems involving rotationally symmetric scatterers*, *Computers & Mathematics with Applications*, 69 (2015), pp. 304 – 318.
- [9] R. W. HOCKNEY, *The communication challenge for mpp: Intel paragon and meiko cs-2*, *Parallel Comput.*, 20 (1994), p. 389398.
- [10] L. HRMANDER, *Fourier integral operators. i*, *Acta Math.*, 127 (1971), pp. 79–183.
- [11] Y. LI AND H. YANG, *Interpolative Butterfly Factorization*, *SIAM Journal on Scientific Computing*, 39 (2017), pp. A503–A531.
- [12] Y. LI, H. YANG, E. R. MARTIN, K. L. HO, AND L. YING, *Butterfly factorization*, *Multiscale Modeling & Simulation*, 13 (2015), pp. 714–732.
- [13] Y. LI, H. YANG, AND L. YING, *A Multiscale Butterfly Algorithm for Multidimensional Fourier Integral Operators*, *Multiscale Modeling & Simulation*, 13 (2015), pp. 614–631.
- [14] Y. LI, H. YANG, AND L. YING, *Multidimensional butterfly factorization*, *Applied and Computational Harmonic Analysis*, (2017).
- [15] E. LIBERTY, F. WOOLFE, P.-G. MARTINSSON, V. ROKHLIN, AND M. TYGERT, *Randomized algorithms for the low-rank approximation of matrices*, *Proceedings of the National Academy of Sciences*, 104 (2007), pp. 20167–20172.
- [16] Y. LIU, H. GUO, AND E. MICHIELSSEN, *An HSS matrix-inspired butterfly-based direct solver for analyzing scattering from two-dimensional objects*, *IEEE Antennas and Wireless Propagation Letters*, 16 (2017), pp. 1179–1183.
- [17] Y. LIU AND H. YANG, *A hierarchical butterfly lu preconditioner for two-dimensional electromagnetic scattering problems involving open surfaces*, *Journal of Computational Physics*, 401 (2020), p. 109014.
- [18] P. G. MARTINSSON AND V. ROKHLIN, *A fast direct solver for boundary integral equations in two dimensions*, *Journal of Computational Physics*, 205 (2005), pp. 1–23.
- [19] E. MICHIELSSEN AND A. BOAG, *Multilevel evaluation of electromagnetic fields for the rapid solution of scattering problems*, *Microwave and Optical Technology Letters*, 7 (1994), pp. 790–795.
- [20] E. MICHIELSSEN AND A. BOAG, *A multilevel matrix decomposition algorithm for analyzing scattering from large structures*, *IEEE Transactions on Antennas and Propagation*, 44 (1996), pp. 1086–1093.
- [21] M. O’NEIL, F. WOOLFE, AND V. ROKHLIN, *An algorithm for the rapid evaluation of special function transforms*, *Applied and Computational Harmonic Analysis*, 28 (2010), pp. 203 – 226. Special Issue on Continuous Wavelet Transform in Memory of Jean Morlet, Part I.
- [22] J. POULSON, L. DEMANET, N. MAXWELL, AND L. YING, *A Parallel Butterfly Algorithm*, *SIAM Journal on Scientific Computing*, 36 (2014), pp. C49–C65.
- [23] E. REBROVA, G. CHÁVEZ, Y. LIU, P. GHYSELS, AND X. S. LI, *A study of clustering techniques and hierarchical matrix formats for kernel ridge regression*, in *IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, 2018.
- [24] J. M. TAMAYO, A. HELDRING, AND J. M. RIUS, *Multilevel adaptive cross approximation (mlaca)*, *IEEE Transactions on Antennas and Propagation*, 59 (2011), pp. 4600–4608.
- [25] M. TYGERT, *Fast algorithms for spherical harmonic expansions, iii*, *Journal of Computational*

- Physics, 229 (2010), pp. 6181 – 6192.
- [26] H. YANG, *A unified framework for oscillatory integral transforms: When to use nufft or butterfly factorization?*, Journal of Computational Physics, 388 (2019), pp. 103 – 122.
- [27] L. YING, *Sparse Fourier Transform via Butterfly Algorithm*, SIAM Journal on Scientific Computing, 31 (2009), pp. 1678–1694.