

# Lawrence Berkeley National Laboratory

## LBL Publications

### Title

A parallel hierarchical blocked adaptive cross approximation algorithm

### Permalink

<https://escholarship.org/uc/item/2df365sj>

### Journal

The International Journal of High Performance Computing Applications, 34(4)

### ISSN

1094-3420

### Authors

Liu, Yang  
Sid-Lakhdar, Wissam  
Rebrova, Elizaveta  
[et al.](#)

### Publication Date

2020-07-01

### DOI

10.1177/1094342020918305

Peer reviewed

# A parallel hierarchical blocked adaptive cross approximation algorithm

Yang Liu<sup>1</sup>, Wissam Sid-Lakhdar<sup>1</sup>, Elizaveta Rebrova<sup>2</sup>, Pieter Ghysels<sup>1</sup> and Xiaoye Sherry Li<sup>1</sup>

## Abstract

This article presents a low-rank decomposition algorithm based on subsampling of matrix entries. The proposed algorithm first computes rank-revealing decompositions of submatrices with a blocked adaptive cross approximation (BACA) algorithm, and then applies a hierarchical merge operation via truncated singular value decompositions (H-BACA). The proposed algorithm significantly improves the convergence of the baseline ACA algorithm and achieves reduced computational complexity compared to the traditional decompositions such as rank-revealing QR. Numerical results demonstrate the efficiency, accuracy, and parallel scalability of the proposed algorithm.

## Keywords

Adaptive cross approximation, singular value decomposition, rank-revealing decomposition, parallelization, multilevel algorithms

## 1. Introduction

Rank-revealing decomposition algorithms are important numerical linear algebra tools for compressing high-dimensional data, accelerating solution of integral and partial differential equations, constructing efficient machine learning algorithms, analyzing numerical algorithms, and so on, as matrices arising from many science and engineering applications oftentimes exhibit numerical rank-deficiency. Despite the favorable  $O(nr)$  memory footprint of such decompositions with  $n$  and  $r$ , respectively, denoting the matrix dimension (assuming a square matrix) and the numerical rank, the computational cost can be expensive. Existing rank-revealing decompositions such as truncated singular value decomposition (SVD), column-pivoted QR (QRCP), CUR decomposition, interpolative decomposition (ID), and rank-revealing LU typically require at least  $O(n^2r)$  operations (Cheng et al., 2005; Gu and Eisenstat, 1996; Mahoney and Drineas, 2009; Voronin and Martinsson, 2017). This complexity can be reduced to  $O(n^2 \log r + nr^2)$  by structured random matrix projection-

cross approximation (ACA) (Bebendorf, 2000; Bebendorf and Grzhibovskis, 2006; Zhao et al., 2005) algorithms can achieve  $O(nr)$  complexity. However, the robustness of these algorithms relies heavily on matrix properties that are not always present in practice.

2. When the matrix can be rapidly applied to arbitrary vectors, algorithms such as randomized SVD, QR, and UTV (T lower or upper triangular) (Feng et al., 2019; Liberty et al., 2007; Martinsson et al., 2019; Xiao et al., 2017) can be utilized to achieve quasi-linear complexity.
3. Finally, given a matrix with missing entries, the low-rank decomposition can be constructed via matrix completion algorithms (Balzano et al., 2010; Candès and Recht, 2009) in quasi-linear time assuming incoherence properties of the matrices (i.e. projection of natural basis vectors onto the space spanned by singular vectors of the matrix should not be very sparse).

based algorithms (Liberty et al., 2007; Voronin and Martinsson, 2017). In addition, faster algorithms are available in the following three scenarios.

1. When each element entry can be computed in  $O(1)$  CPU time with prior knowledge (i.e. smoothness, sparsity, or leverage scores) about the matrix, faster algorithms such as randomized CUR and adaptive

<sup>1</sup> Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA, USA

<sup>2</sup> Department of Mathematics, University of California, Los Angeles, CA, USA

Corresponding author:

Yang Liu, Computational Research Division, Lawrence Berkeley National Laboratory, Berkeley, CA 94720, USA.  
Email: liuyangzhan@lbl.gov

This work concerns the development of a practical algorithm, in application scenario 1, that improves the robustness of ACA algorithms while maintaining reduced complexity for broad classes of matrices.

The partially pivoted ACA algorithm, closely related to LU with rook pivoting (Foster, 1997; Neal and Poole, 1992; Poole and Neal, 1992, 1991), constructs an LU-type decomposition upon accessing one row and column per iteration. For matrices resulting from asymptotically smooth kernels, ACA is a rank-revealing and optimal-complexity algorithm that converges in  $O(r)$  iterations (Bebendorf, 2000). Despite its favorable computational complexity, it is well-known that the ACA algorithm suffers from deteriorated convergence and/or premature termination for non-smooth, sparse, and/or coherent matrices (Heldring et al., 2014). Hybrid methods or improved convergence criteria (e.g. hybrid ACA-CUR, averaging, statistical norm estimation) have been proposed to partially alleviate the problem (Grasedyck and Hackbusch, 2005; Heldring et al., 2015). The main difficulty of leveraging ACA as robust algebraic tools for general low-rank matrices results from ACA's partial pivot-search strategy to attain low complexity. In addition to the above-mentioned remedies, another possibility to improve ACA's robustness is to search for pivots in a wider range of rows/columns without sacrificing too much computational efficiency. Here we consider two different strategies:

1. Instead of searching one row/column per iteration as in ACA, it is possible to search a block of rows/columns to find multiple pivots together.
2. Instead of applying ACA directly on the entire matrix, it is possible to start with compressing submatrices via ACA and then merge the results as one low-rank product.

In extreme cases (e.g. when block size equals matrix dimension or submatrix dimension equals one), these strategies lead to quadratic computational costs. Therefore, it is valuable to address the question: for what matrix kernels and under what block/submatrix sizes will these strategies retain low complexity.

For the first strategy, this work proposes a blocked ACA algorithm (BACA) that extracts a block row/column per iteration to significantly improve convergence of the baseline ACA

algorithms. The blocked version also enjoys higher flop performance as it involves mainly BLAS-3 operations. Compared to the aforementioned remedies, the proposed algorithm provides a unified framework to bal-

H-BACA algorithm is at most  $O(p n_b n r^2)$  assuming the block size in BACA is less than the rank. In other words, the proposed H-BACA algorithm is a general numerical linear algebra tool as an alternative to ACA, SVD, QR, and so on. In addition, the overall algorithm can be parallelized using distributed-memory linear algebra packages such as ScaLAPACK (Blackford et al., 1997) which avoids the difficulty of efficient parallelization of plain ACA algorithms. Numerical results illustrate good accuracy, efficiency, and parallel performance. In addition, the proposed algorithm can be used as a general low-rank compression tool for constructing hierarchical matrices (Rebrova et al., 2018).

## 2. Notation

Throughout this article, we adopt the Matlab notation of matrices and vectors. Submatrices of a matrix  $A$  are denoted  $A(I; J)$ ,  $A(:, J)$ , or  $A(I, :)$  where  $I$  and  $J$  are index sets. Similarly, subvectors of a column vector  $u$  are denoted  $u(I)$ . An index set  $I$  permuted by  $J$  reads  $I/J$ . Transpose, inverse, pseudo-inverse of  $A$  are  $A^t$ ,  $A^{-1}$ ,  $A^\dagger$ .  $\|A\|$  and  $\|u\|$  denote Frobenius norm and 2-norm. Note that  $u$  refers to a  $(n-1)$  column vector. Vertical and horizontal

ance robustness and efficiency. Upon increasing the block size (i.e. the number of rows/columns per iteration), the algorithm gradually changes from ACA to ID. For the second strategy, the proposed algorithm further subdivides the matrix into  $n_b$  submatrices compressed via BACA, followed by a hierarchical merge algorithm leveraging low-rank arithmetic (Grasedyck and Hackbusch, 2003; Hackbusch et al., 2002). The overall cost of this

concatenations of  $A, B$  are  $A; B$  and  $A; B$ . Element-wise multiplication of  $A$  and  $B$  is  $A \cdot B$ . All matrices are real-valued unless otherwise stated. It is assumed for  $A \in \mathbb{R}^{m \times n}$ ,  $m \geq n$ , but the proposed algorithms also apply to complex-valued and tall-skinny/short-fat matrices. We denote truncated SVD as  $[U; S; V; r] \approx SVD(A)$  with  $U \in \mathbb{R}^{m \times r}$ ,  $V^t \in \mathbb{R}^{n \times r}$  column orthogonal,  $S \in \mathbb{R}^{r \times r}$  diagonal, and  $r$  being e-rank defined by  $\min_{k \in \mathbb{N}} \{S_{k+1, k+1} < \epsilon S_{1,1}\}$ . We denote QRCP as  $[Q; \sigma; J] \approx QR(A)$  with  $Q \in \mathbb{R}^{m \times r}$  column orthogonal,  $T \in \mathbb{R}^{r \times n}$  upper triangular,  $J$  being column pivots, and  $\epsilon$  and  $r$  being the prescribed accuracy and rank, respectively. QR without column-pivoting is simply written as  $Q; T \approx QR(A)$ . Cholesky decomposition without pivoting is written as  $Chol(A)$  with  $T$  upper triangular.  $\log n$  means logarithm of  $n$  to the base 2.

## 3. Algorithm description

### 3.1. Adaptive cross approximation

Before describing the proposed algorithm, we first briefly summarize the baseline ACA algorithm (Zhao et al., 2005). Consider a matrix  $A \in \mathbb{R}^{m \times n}$  of e-rank  $r$ , the ACA algorithm approximates  $A$  by a sequence of rank-1 outer products as

$$A \approx \sum_{k=1}^r u_k v_k^t \quad (1)$$

Define the residual matrix  $E_k$  as

$$E_0 \approx A; \quad E_k \approx E_{k-1} - u_k v_k^t \quad (2)$$

At each iteration  $k$ , the algorithm selects column  $u_k$  (pivot  $j_k$  from remaining columns) and  $v_k^t$  (pivot  $i_k$  row from remaining rows) from the residual matrix  $E_{k-1}$

**Algorithm 1. Adaptive cross approximation algorithm (ACA).**

**input :** Matrix  $A \in \mathbb{R}^{m \times n}$ , relative tolerance  $\epsilon$   
**output:** Low-rank approximation of  $A \approx UV$  with rank  $r$

$U = 0, V = 0, \mu = 0, r_0 = 0, j_1$  is a random column index;

**for**  $k = 1$  **to**  $\min\{m, n\}$  **do**

- 3  $u_k = E_{k-1}(:, j_k) = A(:, j_k) - UV(:, j_k);$
- 4  $i_k = \arg \max_i |u_k(i)|;$
- 5  $u_k \leftarrow u_k / u_k(i_k);$
- 6  $v_k^t = E_{k-1}(i_k, :) = A(i_k, :) - U(i_k, :)V;$
- 7  $j_{k+1} = \arg \max_j |v_k(j)|;$
- 8  $\nu^2 = \|u_k\|_2^2 \|v_k\|_2^2;$
- 9  $\mu^2 \leftarrow \mu^2 + \nu^2 + 2(u_k^t U)(V v_k);$
- 10  $U \leftarrow [U, u_k], V \leftarrow [V; v_k^t], r_k = r_{k-1} + 1;$
- 11 **\_** Terminate if  $\nu < \epsilon \mu.$

algo-

corresponding to an element denoted by  $E_{k-1}(i_k, j_k)$  with sufficiently large magnitude. Note that  $u_k$  and  $v_k$  are  $m \times 1$  and  $n \times 1$  vectors. The partially pivoted ACA algorithm (ACA for short), selecting  $j_k; i_k$  by only looking at previously selected rows and columns, is described as Algorithm 1. Specifically, each iteration  $k$  selects pivot  $i_k$  used in the current iteration and pivot  $j_{k+1}$  for the next iteration (via lines 4 and 7) as

$$i_k = \arg \max_{i \in \{i_1, \dots, i_{k-1}\}} |E_{k-1}(i, j_k)| \quad (3)$$

$$j_{k+1} = \arg \max_{j \in \{j_1, \dots, j_k\}} |E_{k-1}(i_k, j)| \quad (4)$$

and  $j_1$  is a random initial column index. Note that

$i_1, \dots, i_{k-1}, j_1, \dots, j_{k-1}$  are distinct and the algorithm is terminated when  $\nu < \epsilon \mu$ .

$$\mu^2 = \| \sum_{k=1}^r u_k u_k^t U \sum_{k=1}^r V v_k^t V \sum_{k=1}^r u_k u_k^t \|_F \quad (5)$$

and  $\epsilon$  is the prescribed tolerance. Note that each iteration requires only  $O(nr_k)$  flop operations with  $r_k$  denoting currently revealed numerical rank. The overall complexity of partially pivoted ACA scales as  $O(nr^2)$  when the algorithm converges in  $O(r)$  iterations. Despite the favorable complexity, the convergence of ACA for general rank-

deficient matrices is unsatisfactory. For many rank-deficient matrices arising from the numerical solution of PDEs, signal processing and data science, ACA oftentimes either requires  $O(n)$  iterations or exhibits premature termination. First, as ACA does not search the full residual matrices for the largest element, it cannot avoid selection of

### 3.2. Blocked adaptive cross approximation

Instead of selecting only one column and row from the residual matrix in each ACA iteration, we can select a fixed-size block of columns and rows per iteration to improve the convergence and accuracy of ACA. In addition, many BLAS-1 and BLAS-2 operations of ACA become BLAS-3 operations and hence higher flop performance can be achieved.

Specifically, the proposed BACA algorithm factorizes  $A$

$$A = \begin{matrix} & & \times \\ & & nd \\ & \diamond & UV \\ \begin{matrix} \frac{1}{4} \\ k \end{matrix} & & \begin{matrix} \frac{k}{4} \\ 1 \end{matrix} \\ & & V_k \end{matrix}$$

where  $U \in \mathbb{R}^{m \times d_k}$  and  $V_k \in \mathbb{R}^{d_k \times n}$ . In principle, the

algorithm selects a block of  $d$  rows and columns via cross approximations in the residual matrix and then  $d_k \times d$  ones

via rank-revealing algorithms to form a low-rank update at iteration  $k$ . The total number of iterations is approximately  $n/d$  if  $d_k = d$ . Instead of selecting row/column pivots via lines 4 and 7 of Algorithm 1, the proposed algorithm selects row and column index sets  $I_k$  and  $J_k$  by performing QRCP on  $d$  columns (more precisely their transpose) and rows of the residual matrices. This proposed strategy is described in Algorithm 2.

Each BACA iteration is composed of three steps.

- Find block row  $I_k$  and block column  $J_{k+1}$  by QRCP. Starting with a random column index set  $J_1$ , the block row  $I_k$  and the next iteration's block column  $J_{k+1}$  are selected by (lines 4 and 7)

$$I_k = \arg \min_{I \in \mathcal{I}_k} \| [Q_k^t; T_k] \cdot I \cdot [Q_{k-1}^t; T_{k-1}] \|_F \quad (6)$$

$$J_{k+1} = \arg \min_{J \in \mathcal{J}_k} \| [Q_{k-1}; T_{k-1}; I_{k+1}] \cdot J \cdot [Q_k^t; T_k] \|_F \quad (7)$$

deficient matrices is unsatisfactory. For many rank-deficient matrices arising from the numerical solution of PDEs, signal processing and data science, ACA oftentimes either requires  $O(n)$  iterations or exhibits premature termination. First, as ACA does not search the full residual matrices for the largest element, it cannot avoid selection of

smaller pivots for general rank-deficient matrices and may require  $O(n^3)$  iterations. Second, the approximation  $U_k V^T$  in (5)

often causes the premature termination with the selection of smaller pivots. Remedies such as averaged stopping criteria (Zhou et al., 2017), stochastic error estimation (Heldring et al., 2015), ACA (Grasedyck and Hackbusch, 2005), and hybrid ACA (Grasedyck and Hackbusch, 2005) have

Here the algorithm first selects  $d$  skeleton rows from the submatrix  $E_{k-1} ; J_k$  (i.e.  $d$  columns from its transpose) and then selects  $d$  skeleton columns from the submatrix  $E_{k-1} I_k ;$  by leveraging the LAPACK implementation of QRCP as it provides a simple way of greedily selecting well- conditioned columns by examining column norms in the  $R$  factor at each iteration. Note that many other subset selection algorithms exist in both the machine learning and numerical linear algebra communities (e.g. strong rank-revealing QR (Gu and Eisenstat, 1996), spectrum-revealing QR (Feng et al., 2019), and column subset selection problems (Boutsidis et al., 2009)), which ideally pick  $d$  matrix columns with maximum volumes. Note that  $I_k$  excludes rows selected in previous iterations. To efficiently enforce such condition, the

QRCP is performed on the submatrix  $E_{k-1} ; J_k$  of  $E$  excluding previously selected rows rather than directly on  $E_{k-1} ; J_k$ . Similarly,  $J_k$  excludes columns selected in previous iterations. See Figure 1(a) for an illustration of the procedure.  $I_k$  and  $J_{k+1}$  are selected by QRCP on the column

Algorithm 2. Blocked adaptive cross approximation (BACA) algorithm.

---

**input** : Matrix  $A \in \mathbb{R}^{m \times n}$ , block size  $d$ , relative tolerance  $\epsilon$

**output**: Low-rank approximation  $UV$  with rank  $r$

- 1  $U = 0, V = 0, r_0 = 0, \mu = 0, J_{-1}$  is a random index set of cardinality  $d$ ;
- 2 **for**  $k = 1$  **to**  $\min(m, n)$  **do**
- 3  $C_k = E_{k-1}(:, J_k) \equiv A(:, J_k) U V(:, J_k)$ ;
- 4  $[Q_k, T_k, I_k] = \text{QR}(C_k, d)$ ,  $I_k$  denotes selected skeleton rows;
- 5  $R_k = E_{k-1}(I_k, :) \equiv A(I_k, :)$ ;
- 6  $[r_k, T_{k+1}, J_{k+1}] = \text{QR}(R_k, d)$ ,  $J_{k+1}$  denotes selected skeleton columns;
- 7  $W_k = E_{k-1}(I_k, J_k) = A(I_k, J_k) U(I_k, :)$ ;
- 8  $[U_k, V_k, d_k, J] = \text{LRID}(C_k, W_k, R_k)$ ;
- 9  $I_k \leftarrow I_k \cup \{1 : d\}, J_k \leftarrow J_k \cup J$ ;
- 10  $r_k = r_{k-1} + d_k$ ;
- 11  $v = \text{LRnorm}(U_k, V_k)$ ;
- 12  $\mu \leftarrow \text{LRnormUp}(U, V, \mu, U_k, V_k, v)$ ;
- 13  $U \leftarrow [U, U_k], V \leftarrow [V; V_k]$ ;
- 14 Terminate if  $v < \mu$ .

15 **Function**  $\text{LRID}(C, W, R)$

**input** :  $C = A(:, J), R = A(I, :)$ ,  $W = A(I, J)$  with  $I, J$  of same cardinality

**output**:  $AUV$  with  $U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{r \times n}$

- 16  $[Q, T, J, r] = \text{QR}(W)$ ;
- 17  $U = C(:, J)$ ;
- 18  $V = T^{-1} Q^T R$ ;
- 19 **return**  $U, V, r, J$

20 **Function**  $\text{LRnorm}(U, V)$

**input** :  $A = UV$

**output**:  $A$

- 21  $T_1 = \text{Chol}(U^T U)$ ;
- 22  $T_2 = \text{Chol}(V V^T)$ ;

23 **return**  $\|T_1^{-1} T_2^{-1}\|_F$ ;

24 **Function**  $\text{LRnormUp}(U, V, v, U_k, V_k, v^-)$

**input** :  $U \in \mathbb{R}^{m \times r}, V \in \mathbb{R}^{r \times n}, U^- \in \mathbb{R}^{m \times r}, V^- \in \mathbb{R}^{r \times n}, v = \|UV\|_F, v^- = \|U^- V^-\|_F$




Figure 1. (a) Selection of  $I_k/J_k$  and form the low-rank update  $U_k V_k$ . (b) Low-rank merge operation.

$E_{k-1} \diamond C_k W_k^y R_k \frac{1}{4} U_k V_k$  (Voronin and Martinsson, 2017) by (9) and (10). Note that the pseudo inverse is computed via rank-revealing QR (also see the LRID algorithm at line 8). The rank-revealing algorithm is needed as the  $d \times d$  block  $W_k$  can be further compressed with rank  $d_k$ . Particularly, for matrices where the ACA algorithm tends to fail, the corresponding  $d \times d$  matrices  $W_k$  in BACA are often rank-deficient. In this case, BACA becomes more robust than ACA as the effective  $d_k$  pivots can still be used to generate  $d$  columns  $J_{k+1}$  for the next iteration (as long as  $d_k > 0$ ). Consequently, the effective rank increase is  $d_k$  and the pivot pair  $I_k, J_k$  is updated in (11) by the column pivots  $J$  of QRCP in (9).

$$[Q; T; J] \frac{1}{4} \text{QR}(W_k); \text{ with } Q \in \mathbb{R}^{d \times d_k} \quad (9)$$

$$\frac{1}{4} C_k(:, J); \frac{1}{4} T_k^{-1} Q^T R_k \quad (10)$$

$$I \cup \{1 : d\}; J \cup J \quad (11)$$



25 **output:**  $s = v^2 + v^{-2} + 2v^1$   $V(i, j)$  with

$$\tilde{U}^t \bar{U}^t = \prod_{i=1}^k \prod_{j=1}^k V^{-t}$$

26 **return**

---

- Compute  $n \times n$   $U$  and update  $n \times n$   $V$

Assuming constant block size  $d$ , the norm of the low-rank update can be computed in  $O(nd^2)$  operations (line

$k$

11) via  $T_U \times \text{Chol}(U_k) \times U_k^T$ ;  $T_V \times \text{Chol}(V_k) \times V_k^T$   $O(12k)$

$k$   $k$   $k$   $k$

and transpose of the row marked in yellow,

respectively. The

column marked in grey is used to select  $I_{k+1}$  in

$n \times n$   $J J^T$   $O(13k)$

the next

$J J$

iteration. For illustration purpose, index sets in Figure 1(a)

Once  $n$  is computed, the norm of  $UV$  can be updated efficiently in  $O(\sum_{k=1}^K r_k d_k)$  operations (line 12) as

$$\sum_{k=1}^K r_k d_k$$

consist of contiguous indices.

- Form the factors of the low-rank product

$U_k V_k$ . Let  $C_k = \frac{1}{4}$

$E_{k-1} \delta$ ;  $J_k \delta$ ,  $R_k \frac{1}{4} E_{k-1} \delta l_k$ ;  $\delta$ , and  $W_k \frac{1}{4} E_{k-1} \delta l_k$ ;

$J_k \delta$ ,

$E_{k-1}$  can be

$$2^2 \quad 2 \text{ p n}^2 \text{ p}$$

approximated by an ID-type decomposition

$i/41 \quad j/41$

\*V ði; jþ

$$V \sim \frac{1}{4} \delta V V^t \text{ } \rho_0 \delta U^t U_k \text{ } \rho$$

where  $r_k$

$$\delta 14 \rho \quad k$$

represents the column dimension of  $U$  at iteration  $k$ . Note that the matrix multiplications in (12) involving  $U$  and  $V$  (and similarly for those involving  $U$  and  $U$ ) can be performed as  $V ; V_k V_k$  to further improve the computational efficiency. Then the algorithm updates  $U, V$  as  $U ; U_k, V ; V_k$  and tests the stopping criterion  $n < \epsilon$ . Note that  $n$ ; with larger  $d$  provides better approximations

to the exact stop criterion compared to those in (5) hence can significantly reduce the chance of premature termination.

We would like to highlight the difference between the proposed BACA algorithm and existing ACA algorithms. First, as BACA selects a block of rows and columns per iteration as opposed to a single row and column in the base-line ACA algorithm, the convergence behavior and flop performance can be significantly improved. In the existing ACA algorithms, convergence can also be improved by leveraging averaged stopping criteria (Zhou et al., 2017) or searching a single pivot in a broader range of rows and columns (e.g. fully pivoted ACA). However, they still find one row or column at a time in each iteration and hence suffer from poor flop performance. Moreover, they cannot utilize strong rank-revealing algorithms to select skeleton rows and columns with better volume (determinant in modulus) qualities. Second, BACA also has important connections to the hybrid ACA algorithm (Grasedyck and Hackbusch, 2005). The hybrid ACA algorithm assumes prior knowledge about the skeleton rows and columns to leverage interpolation algorithms (e.g. ID and CUR) on a skeleton submatrix and use ACA to refine the skeletons. In contrast, BACA uses cross approximations with QRCP to select skeleton rows and columns and uses interpolation algorithms (LRID at line 8) to form the low-rank update in each iteration. In other words, hybrid ACA can be treated as embedding ACA into interpolation algorithms while BACA can be thought of as embedding interpolation algorithms into ACA iterations. In addition, BACA is purely algebraic and requires no prior knowledge of the row/column skeletons or geometrical information about the rows/columns.

It is worth mentioning that the choice of  $d$  affects the trade-off between efficiency and

The BACA algorithm oftentimes exhibits overestimated ranks compared to those revealed by truncated SVD. Therefore, an SVD re-compression step of  $U$  and  $V$  may be needed via first computing a QR of  $U$  and  $V$  as and (14)  $[Q_U ; T_U] \frac{1}{4} QR \delta U \text{ } \rho, [Q_V ; T_V] \frac{1}{4} QR \delta V \text{ } \rho$ , and then a truncated SVD of (Heldring et al., 2015). The result can

robustness of the BACA algorithm. When  $d < r$ , the algorithm requires  $O \delta n r^2 \rho$  operations assuming convergence in  $O \delta r = d \rho$  iterations as each iteration requires  $O n r_k d$  operations. For example, BACA (Algorithm 2) precisely reduces to ACA (Algorithm 1) when  $d = 1$ . In what follows we refer to the base-line ACA algorithm as BACA with  $d = 1$ . On the other hand, BACA converges in a constant number of iterations when  $d = r$ . In the extreme case, BACA reduces to QRCP-based ID when  $d = \min m; n$  (note that the LRID algorithm at line 8 remains the only nontrivial operation). In this case, the algorithm requires  $O n^2 r$  operations but enjoys the provable convergence of QRCP. Detailed complexity analysis of the BACA algorithm will be provided in Section 4.


be viewed as an approximate truncated SVD of  $A$  and we assume this is the output of the BACA algorithm in the rest of this article.

### 3.3. Parallel hierarchical low-rank merge

The distributed-memory implementations of the proposed BACA algorithm and the baseline ACA algorithm can pose performance challenges as straightforward parallelization of all operations in Algorithm 2 and 1 involves many collective communications. To see this, assuming the  $U$  and  $V$  factors in Algorithm 1 follow 1-D block row and column data layouts, then every operation from line 3 to line 9 requires one or more collective communications. Instead, one can assign one process to perform BACA/ACA on submatrices without any communication and then leverage parallel low-rank arithmetic to merge the results into one single low-rank product. To elucidate the proposed algorithm, we first describe the hierarchical low-rank merge algorithm, then outline its parallel implementation.

Given a matrix  $A \in \mathbb{R}^{m \times n}$  with  $m \geq n$ , the algorithm first creates  $L$ -level binary trees for index vectors  $[1; m]$  and  $[1; n]$  with index set  $I_t$  and  $J_n$  for nodes  $t$  and  $n$  at each level, upon recursively dividing each index set into  $I_{t_i} \cup J_{n_j}$  of approximately equal sizes,  $i = 1, 2, j = 1, 2$ . Here,  $t_i$  and  $n_j$  are children of  $t$  and  $n$ , respectively. The leaf and root levels are denoted 0 and  $L$ , respectively. This process generates  $n_b$  leaf-level submatrices of similar sizes. For simplicity, it is assumed  $n_b = 4^L$ . We denote submatrices associated with  $t_i; n_j$  as  $A_{t_i; J_{n_j}}$  and their truncated SVD as  $U_{t_i; n_j}; S_{t_i; n_j}; V_{t_i; n_j}; r_{t_i; n_j}$ . Here  $r_{t_i; n_j}$  is the e-rank of  $A_{t_i; n_j}$ . As submatrices  $A_{t_i; n_j}$  have significantly smaller dimensions than  $A$  (e.g. when  $n_b = O(n^2)$  as an extreme case), both BACA and ACA algorithms become more robust to attain the truncated SVD. Following compression of  $n_b$  submatrices  $A_{t_i; n_j}$  by BACA or ACA at step  $l = 0$ , there are multiple approaches to combine them into one low-rank product including randomized algorithms via applying  $A$  to random matrices, and deterministic algorithms via recursively pair-wise re-compressing the blocks using low-rank arithmetic. Here we choose the deterministic algorithm for simplicity of rank estimation and parallelization. Here, we deploy truncated SVD as the re-compression tool but other tools such as ID, QR, UTV can also be applied. Figure 1(b) illustrates one re-compression operation for transforming

SVDs of  $A_{t_i; n_j}; i = 1, 2; j = 1, 2$  into that of  $A_{t_n}$ . The operation first horizontally compresses SVDs of  $A_{t_i; n_j}; i = 1, 2; j = 1, 2$  at step  $l = 0$  and then vertically compresses the results, that is, SVDs of  $A_{t_i; n}; i = 1, 2$  at step  $l = 1; \dots; L$ . Specifically, the horizontal compression step is composed of one






Figure 2. Parallel hierarchical merge with eight processes. Blocks surrounded by solid lines represent  $A_{tm}$  after compression at each step  $l$ . Blocks surrounded by dashed lines represent ScaLAPACK blocks.

concatenation operation in (15) and one compression operation in (16):

$$U_{t,n}^{-1/4} [U_{t,n_1} S_{t,n_1}; U_{t,n_2} S_{t,n_2}]; V_{t,n}^{-1/4} \text{diag}(\delta V_{t,n_1}, \delta V_{t,n_2}) \quad (15)$$

$$[U_{t,n}; S_{t,n}; V_{t,n}; r_{t,n}] \text{SVD}(U_{t,n}; \epsilon); V_{t,n} \quad V_{t,n}^{-1/2} \quad (16)$$

with  $j = 1; 2$ . Let  $U_{t,n}^{-1/4} V_{t,n}^{-1/4}$  and  $U_{t,n} S_{t,n} V_{t,n}^{-1/4}$  denote the

submatrix before and after the SVD truncation, respectively. Similarly, the vertical compression step can be performed via horizontal merge of  $t$ ;  $j = 1; 2$ . Let  $A^t$

represent the maximum rank  $r_{tn}$  among all blocks at steps

$l = 0; 1; \dots; L$ . Note that the algorithm returns an approximate truncated SVD after  $L$  steps. As an example, the hierarchical merge algorithm with the level count of the hierarchical merge  $L = 2$  and  $n_b = 16$  is illustrated in Figure 2. At step  $l = 0$ , the algorithm compresses all  $n_b$  submatrices with BACA; at step  $l = 0.5; 1.5$ , the algorithm merges every horizontal pair of blocks; similarly at level  $l = 1; 2$ , the algorithm merges every vertical pair of blocks. Note that blocks surrounded by solid lines represent results after compression at each step  $l$ .

The above-described hierarchical algorithm with BACA for leaf-level compressions is dubbed H-BACA (Algorithm 3). In the following, a distributed-memory implementation of the H-BACA algorithm is described. Without

loss of generality, it is assumed that  $m = n2^i$  and  $p = 2^j$ . The proposed parallel implementation first creates two  $\log_2 p$ -level binary trees with  $p$  denoting the total number of MPI processes. One process performs BACA compression of one or two leaf-level submatrices and low-rank

Algorithm 3. Hierarchical low-rank merge algorithm with BACA (H-BACA).

**input** : Matrix  $A \in \mathbb{R}^{m \times n}$ , number of leaf-level subblocks  $n_b$ , block size  $d$  of leaf-level BACA, relative tolerance with rank

Truncated SVD of  $A \approx U \Sigma V^T$

```

1 Create  $L$ -level trees on index vectors  $[1, m]$  and  $[1, n]$  with index set  $I_\tau$  and  $J_\nu$  for nodes  $\tau$  and  $\nu$ 
   at each level,  $L$  are denoted 0 and  $L$ , respectively;
2 for  $l = 0$  to  $L$  do
3   foreach  $A_{\tau\nu} = A(I_\tau, J_\nu)$  at level  $l$  do
4     if leaf-level then
5        $[U_{\tau\nu}, \Sigma_{\tau\nu}, V_{\tau\nu}, r_{\tau\nu}] = \text{BACA}(A_{\tau\nu}, d, \epsilon)$ ;
6     else
7       Let  $\tau_1, \tau_2$  and  $\nu_1, \nu_2$  denote children
8       of  $\tau$  and  $\nu$ ;
       for  $i = 1$  to 2 do

```

$$U_{\tau\nu}^{-1/2} = [U_{\tau_1\nu_1}^{-1/2} \Sigma_{\tau_1\nu_1}^{-1/2} U_{\tau_1\nu_2}^{-1/2} \Sigma_{\tau_1\nu_2}^{-1/2}];$$

$$V_{\tau\nu}^{-1/2} = \text{diag}(V_{\tau_1\nu_1}^{-1/2}, V_{\tau_1\nu_2}^{-1/2}, V_{\tau_2\nu_1}^{-1/2}, V_{\tau_2\nu_2}^{-1/2});$$

$$[U_{\tau\nu}, \Sigma_{\tau\nu}, V_{\tau\nu}, r_{\tau\nu}] \text{SVD}(U_{\tau\nu}^{-1/2} \Sigma_{\tau\nu}^{-1/2} V_{\tau\nu}^{-1/2});$$

$$V_{\tau\nu} \leftarrow V_{\tau\nu} V_{\tau\nu}^{-1/2};$$

$$U_{\tau\nu}^{-1/2} = \text{diag}(U_{\tau_1\nu_1}^{-1/2}, U_{\tau_1\nu_2}^{-1/2}, U_{\tau_2\nu_1}^{-1/2}, U_{\tau_2\nu_2}^{-1/2});$$

$$V_{\tau\nu} = [\Sigma_{\tau_1\nu_1} V_{\tau_1\nu_1}; \Sigma_{\tau_2\nu_1} V_{\tau_2\nu_1}];$$

merge operations from the bottom up until it reaches a submatrix



$$[U_{\tau\nu}, \Sigma_{\tau\nu}, V_{\tau\nu}, r_{\tau\nu}] \leftarrow SVD(V_{\tau\nu}, ); U_{\tau\nu} \leftarrow U_{\tau\nu}^{-1}$$

shared by more than one process. Then, all such blocks are handled by PBLAS and ScaLAPACK with BLACS process grids that aggregate those in corresponding submatrices. Consider the example in Figure 2 with process count  $p = 8$ . The workload of each process is labeled with its process rank and highlighted with one color. The dashed lines represent the ScaLAPACK blocks. First, BACA compressions and merge operations at  $l = 0; 0:5$  are handled locally by one process without any communication. Next, merge operations at  $l = 1; 1:5; 2$  are handled by BLACS grids of  $2 \times 1$ ,  $2 \times 2$ , and  $4 \times 2$ , respectively. For illustration purposes, we select the ScaLAPACK block size in Figure 2 as  $n_0 \times n_0$  where  $n_0$  is the dimension of the finest-level submatrices in the hierarch-

ical merge algorithm and  $n = p^{1/4} n_0$ . In this case, the only

$$U_{\tau\nu};$$

17 return  $U = U_{\tau\nu}, V = V_{\tau\nu}, \Sigma = \Sigma_{\tau\nu}, r = r_{\tau\nu};$

required data redistribution is from step  $l = 1$  to  $l = 5$ . However, the ScaLAPACK block size may be set to much smaller numbers in practice, requiring data redistribution at each row/column re-compression step. Similarly, the requirement of  $m^{1/4} n^{1/4} 2^i$  and  $p^{1/4} 2^j$  is not needed in practice.

#### 4. Cost analysis

In this section, the costs for computation and communication of the proposed BACA and H-BACA algorithms are analyzed.

### 4.1. Computational cost

First, the costs for BACA can be summarized as follows. Assuming BACA converges in  $O(d)$  iterations, each iteration performs entry evaluation from the residual matrices,

ORCP for pivot selection, LRID for forming the and estimation of matrix norms. The entry evaluation computes  $O(dn)$  entries each requiring  $O(d)$  operations; QRCP on block rows requires  $O(d^2n)$  operations; the LRID algorithm requires  $O(d^2n)$  operations; norm estimation requires  $O(dn)$  operations. Summing up these costs, the overall cost for the BACA algorithm is

$$C_{BACA} = \frac{1}{4} O(d^2n) + O(dn) + O(d^2n) + O(dn) = O(d^2n)$$

Here we assume the block size  $d = r$ . Note that when  $d = r$  (e.g.  $d = O(n)$ ), it follows that the worst-case complexity is  $C_{BACA} = O(n^2r)$  by bypassing the pivot selection step that causes the  $nd^2$  term. In practice, one would always avoid the case of  $d = r$ .

Next, the computational costs of the H-BACA algorithm are analyzed. The costs are analyzed for two cases of distributions of the maximum ranks  $s_l$  at each level, that is,  $s_l = r$  (ranks stay constant during the merge) and  $s_l = 2^l r = n_b 2^{l-1} r$  (rank increases by a factor of 2 per level),  $l = 0; 1; \dots; L$ . The constant rank case is often valid for matrices with their numerical ranks independent of matrix dimensions (e.g. random low-rank matrices, matrices representing well-separated interactions from low-frequency and static wave equations and certain quantum chemistry matrices); the increasing-rank case holds true for matrices whose ranks depend polynomially (with order no bigger than 1) on the matrix dimensions (e.g. those arising from

high-frequency wave equations, matrices representing equations, and certain classes of kernel methods on high dimensional data sets). Note that these rank distributions often follow from a proper hierarchical partitioning tree and may not be valid using an arbitrary partitioning tree. From the aforementioned analysis of BACA, the computational

costs for the leaf-level compression  $c_b$  are:

$$n$$

Table 1. Flop counts and communication costs for the leaf-level compression and hierarchical merge operations in Algorithm 3 for two classes of low-rank matrices.<sup>a</sup>

	Constant rank $s_l = r$	Increasing rank $s_l = r = n_b 2^{l-1} \times 2^l$
BACA	$O(nr^2)$	$O(nr^2)$
Merge compute	$O(dnr^2)$	$O(dnr^2)$
Merge communicate	$O(dnr \log^2 p)$	$O(dnr \log^2 p)$

<sup>a</sup> $n$  and  $r$  denote matrix dimension and rank.  $d$  denotes the block size in BACA.  $p$  and  $n_b$  denote number of processes and leaf-level submatrices.  $s_l$  denotes maximum ranks among all level- $l$  submatrices.

summarized in Table 1. Note that the costs of the BACA algorithm can also be extracted from Table 1 upon setting  $n_b = 1$ . Not surprisingly, the hierarchical merge algorithm

induces a computational overhead of at most  $p$  when ranks stay constant; the leaf-level compression can have a  $1/p$  reduction factor for the increasing rank case and  $n_b$  overhead for the constant rank case.

For completeness, the comparison between the proposed BACA, H-BACA algorithms (assuming  $d = r_0$ ) and existing ACA algorithms are given in Table 2. In contrast to existing ACA algorithms that select one pivot at a time, BACA and H-BACA select  $d$  and  $n_b d$  pivots simultaneously. As such, H-BACA is the most robust algorithm among all listed here. Not surprisingly, H-BACA can

induce a computational overhead of  $p/n_b$ .

### 4.2. Communication cost

As the leaf-level BACA compression requires no communication, only the communication costs for the hierarchical merge operations are analyzed here. Since the merge operations

may introduce an  $O(d^2 n_b)$  computational overhead, create more the process count  $p/n_b$ . Consider the parallelization of one level  $l$  merge operation, that is, transformation of the SVDs of  $A_{t_j n_j}; j = 1; 2; \dots; L$  into that of  $A_{tn}$  shown in Figure 1(b). Let  $p_j = 4^l$  denote the number of processes sharing one level  $l$  block  $A_{tn}$ ,  $l = 0; \dots; L$ . The horizontal compression step in (15, 16) requires redistribution from the process grids of size  $p$

$$c_b \approx O(p_0^2 n_b^3); \text{ if } d \leq s_0$$

$$\delta 18p$$

which represent the complexity with ACA

when  $n_b \approx 1$ . Let  $n = n_{t;n}$  denote the size of submatrices  $A$  at level  $l$ . The computational costs  $c_m$  of hierarchical merge operations can be estimated as

$$\sum_{l=1}^L c_m \approx O(4^{L-l} n_l s^2 p^l)$$

$$\delta 19p$$

Accounting for the two cases of rank distributions, the computational costs for the leaf-level BACA and hierarchical merge operations of the H-BACA algorithm are

sharing  $A_{t;n_j}; i \approx 1; 2; j \approx 1; 2$

to the process grids of size  $2p_{l-1}$  sharing  $A_{t;n}; i \approx 1; 2$ .

After redistribution, each process grid involves

PDGESVD function in ScaLAPACK (see (16)) to compute the new rank after the combination in (15), and a PDGEMM function in PBLAS to multiply factors  $V_{t;n}$  with  $V_{-t;n}$  (see 16). Similarly, the vertical compression step

requires redistribution from the process grids of size

$2p_{l-1}$  sharing  $A_{t;n}; i \approx 1; 2$  to the process grids of size  $p_l$

sharing  $A_{t;n}$ , and calling PDGESVD and PDGEMM in the new grids. Let the pair

$[\# \text{messages}, \text{volume}]$  denote the communication cost including the number of

messages and the number of words transferred along the critical path. Then the

communication costs for each

Table 2. Comparisons between proposed BACA, H-BACA algorithms and existing ACA algorithms.<sup>a</sup>

Algorithm	ACA/ACA <sup>b</sup>	Hybrid-ACA	BACA	H-BACA
Cost (constant rank)	$O(nr^2p)$	$O(nr^2p)$	$O(nr^2p)$	$n_b d$
$nr^2$	$O(nr^2p)$	$O(nr^2p)$	$O(nr^2p)$	$O(nr^2)$
Cost (increasing rank)	$O(m^2)$	$O(m^2)$	No	Yes
Preselection of submatrices	No	No	No	No

<sup>a</sup> Note that the algorithms show increasing robustness from left to right.

(BLACS) grid redistribution, PDGEMM and PDGESVD during the hierarchical merge are  $O(1)$ ;  $O(n_s p)$ ,  $O(n_s p)$ , and  $O(n_s \log p)$ ,  $O(n_s \log p)$ , respectively (Blackford et al., 1997). Recall that  $n_i/2^l n = p_i$  and  $s_i$  denote the size and rank of submatrices and note that  $n_i, s_i$ . Therefore the communication cost of the hierarchical merge (and H-BACA) can be estimated as

$$\begin{aligned}
 & \sum_{l=0}^{L-1} 4O(n_s \log p) p_i^{1/4} \\
 & \sum_{l=0}^{L-1} 2^{l-1} n_i^{1/4} s_i^{1/4} p_i^{1/4} \\
 & \sum_{l=0}^{L-1} 2^{l-1} n_i^{1/4} s_i^{1/4} p_i^{1/4} \\
 & \sum_{l=0}^{L-1} 2^{l-1} n_i^{1/4} s_i^{1/4} p_i^{1/4}
 \end{aligned}$$

Consider the two cases of rank distributions, that is,  $s_i = 2^{l-1} r$  and  $s_i = 2^{l-1} r$ , the overall communication costs of H-BACA are  $O(n \log^2 p)$ ;  $O(n \log^2 p)$  and  $v_m/4 O(n \log p)$ ;  $O(n \log p)$ , respectively (see Table 2).

### 5. Numerical results

This section presents several numerical results to demonstrate the accuracy and efficiency of the proposed H-BACA algorithm. The matrices in all numerical examples are generated from the following kernels:

1. Gaussian kernel:  $A_{ij} = \exp(-\frac{\|x_i - x_j\|^2}{2h^2})$ ,  $i, j = 1, \dots, 2n$ . Here  $h$  is the Gaussian width, and  $x_i \in \mathbb{R}^{8 \times 1}$  and  $\mathbb{R}^{784 \times 1}$  are feature vectors in one subset of the SUSY and MNIST Data Sets from the UCI Machine Learning

4. Frontal3D kernel:  $A$  is a dense frontal matrix that arises from the multifrontal sparse elimination for the finite-difference frequency-domain solution of the homogeneous-coefficient Helmholtz equation inside a unit cube.
5. Polynomial kernel:  $A_{ij} = \sum_{k=0}^p x_i^k x_j^k$ . Here  $x_i, x_j \in \mathbb{R}^{50 \times 1}$  are points from a randomly generated data set, and  $h$  is a regularization parameter.
6. Product-of-random kernel:  $A = UV$  with  $U \in \mathbb{R}^{n \times r}$  and  $V \in \mathbb{R}^{r \times n}$  being random matrices with i.i.d. entries

Note that the EFIE2D, EFIE3D, and Frontal3D kernels result in complex-valued matrices. Throughout this section, we refer to ACA as a special case of BACA when  $d = 1$ . In all examples except for the product-of-random kernel, the algorithm is applied to the off-diagonal submatrix  $A_{12} \in \mathbb{R}^{n_1 \times n_2}$  assuming rows/columns of  $A$  have been properly permuted (e.g. by a KD-tree partitioning scheme). Note that the permutation may yield a hierarchical matrix representation of  $A$ , but in this article we only use the permutation to define the partition trees for H-BACA compression of one off-diagonal subblock of  $A$  with H-BACA. 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 at 2133 MHz.

### 5.1. Convergence

Repository (Dheeru and Karra Taniskidou, 2017), respectively. Note that the Gaussian kernel permits low-rank compression as shown in (Bach, 2013; Musco and Musco, 2017; Wang et al., 2018).

2. EFIE2D kernel:  $A_{ij} = \frac{1}{4\pi\epsilon_0} H_0^{(2)}(k|x_i - x_j|)$  resulting from the Nyström discretization of the electric field integral equation (EFIE) for electromagnetic scattering from 2-D curves. Here  $H_0^{(2)}$  is the second kind Hankel function of order 0,  $k$  is the free-space wave-number,  $x_i, x_j \in \mathbb{R}^{2 \times 1}$  are discretization points (15 points per wavelength) of two 2-D parallel strips of length 1 and distance 1.
3. EFIE3D kernel:  $A$  is obtained by the Galerkin method for EFIE to analyze electromagnetic scattering from 3-D surfaces.

First, the convergence of the proposed BACA algorithm is investigated using several matrices: Gaussian-SUSY matrices with  $n = 5000$ ,  $h = 1:0; 0:2$ , an EFIE3D matrix for a unit sphere with  $n = 21788$  and approximately 20 points per wavelength, and a Frontal3D matrix with  $n = 1250$  and 10 points per wavelength. The corresponding e-ranks are  $r = 4683; 1723; 1488; 718$  for  $h = 10^{-6}$ . The residual histories versus revealed ranks  $r_k$ , at each iteration  $k$  of BACA with  $1 \leq d \leq 256$  are plotted in Figure 3. The residual error is defined as  $\|U_j V_j - F\|_F = \|UV - F\|_F$  from (13). As a reference, the singular value spectra  $S(k); k=1; \dots; n$  computed from  $U; S; V$ ;  $r$  SVD  $A$ ;  $e$  are also plotted.

For the Gaussian-SUSY matrices, the baseline ACA algorithm ( $d = 1$ ) behaves poorly with smaller  $h$  due to the exponential decay of the Gaussian kernel. As a result, the matrix becomes increasingly sparse and coherent for small  $h$  particularly for high dimensional data sets. In fact, ACA




Figure 3. Convergence history of BACA for the (a) Gaussian-SUSY kernel with  $h \frac{1}{4} 1:0$ ,  $n \frac{1}{4} 5000$ ,  $e \frac{1}{4} 10^{-6}$ ,  $r \frac{1}{4} 4683$ , (b) Gaussian-SUSY kernel with  $h \frac{1}{4} 0:2$ ,  $n \frac{1}{4} 5000$ ,  $e \frac{1}{4} 10^{-6}$ ,  $r \frac{1}{4} 1723$ , (c) EFIE3D kernel for a unit sphere with  $n \frac{1}{4} 21; 788$ ,  $e \frac{1}{4} 10^{-6}$ ,  $r \frac{1}{4} 1488$ , and (d) Frontal3D kernel with  $n \frac{1}{4} 1250$ ,  $e \frac{1}{4} 10^{-6}$ ,  $r \frac{1}{4} 718$ .


constantly selects smaller pivots and the residual exhibits wild oscillations particularly for smaller  $h$  (e.g. when  $h \frac{1}{4} 0:2$  in Figure 3(b)). Similarly, the analytical and numerical Green's functions respectively for the EFIE3D (Figure 3(c)) and Frontal3D (Figure 3(d)) matrices are not asymptotically smooth for ACA to converge rapidly. For all examples in Figure 3, significant portions of the residual curves lie below the singular value spectra which causes premature iteration termination for certain given residual errors. In stark contrast, the proposed BACA algorithm ( $d \frac{1}{4} 32; 64; 100; 128; 256$ ) shows increasingly smooth residual histories residing above the singular value spectra as the block size  $d$  increases. Although BACA may over-estimate the matrix ranks particularly for larger  $d$ , the SVD re-compression step mentioned in Section 3.2 can effectively reduce the ranks.

$1:0; 0:2$ , one EFIE3D matrix for a unit sphere with  $n \frac{1}{4} 1707$  and approximately 20 points per wavelength, and a Frontal3D matrix with  $n \frac{1}{4} 1250$  and 10 points per wavelength. The relative

## 5.2. Accuracy

Next, the accuracy of the H-BACA algorithm is demonstrated using the following matrices: two Gaussian-SUSY matrices with  $n \frac{1}{4} 5000$ ,  $h \frac{1}{4}$

Frobenius-norm error  $\|A - UV\|_F = \sqrt{\sum_j \lambda_j^2}$  is computed for changing number of leaf-level submatrices  $n_b$  and block size  $d$ . When  $h = 1:0$  for the Gaussian-SUSY matrix (Figure 4(a)), the H-BACA algorithms achieve desired accuracies ( $e^{-2}, 10^{-6}, 10^{-10}$ ) using the baseline ACA ( $d = 1$ ), and BACA ( $d = 32$ ) when  $n_b = 1$  and the hierarchical merge operation only causes slight error increases as  $n_b$  increases. However when  $h = 0:2$  for the Gaussian-SUSY matrix (Figure 4(b)), all data points for H-BACA with  $d = 1$  fail due to the wildly oscillating residual histories. In contrast, H-BACA with  $d = 32$  achieves significantly better accuracies for most data points particularly as  $n_b$  increases. For the EFIE3D (Figure 4(c)) and Frontal3D (Figure 4(d)) matrices, H-BACA with  $d = 32$  achieves comparable accuracies as H-BACA with  $d = 1$  for most data points. Note that  $d = 32$  is significantly better than  $d = 1$  when the prescribed residual error is large ( $e^{-2}$ ). This agrees with the residual histories in Figure 3(c) and (d) as they lie below the singular value spectra when iteration count  $k$  is small.



### 5.3. Efficiency

This subsection provides six examples to verify the computational complexity estimates in Table 1. H-BACA




Figure 4. Measured error of H-BACA with  $e \frac{1}{4} 10^{-2}$ ;  $10^{-6}$ ;  $10^{-10}$  for the (a) Gaussian-SUSY kernel with  $h \frac{1}{4} 1:0$ ,  $n \frac{1}{4} 5000$ , (b) Gaussian-SUSY kernel with  $h \frac{1}{4} 0:2$ ,  $n \frac{1}{4} 5000$ , (c) EFIE3D kernel for a unit sphere with  $n \frac{1}{4} 1707$ , and (d) Frontal3D kernel with  $n \frac{1}{4} 1250$ .

with leaf-level ACA ( $d = 1$ ) and BACA ( $d = 8$ ;  $16$ ;  $32$ ;  $64$ ;  $128$ ) is tested for the following matrices: one Gaussian-SUSY matrix with  $n \frac{1}{4} 50; 000$ ,  $h \frac{1}{4} 1:0$ ,  $e \frac{1}{4} 10^{-2}$ , one Gaussian-MNIST matrix with  $n \frac{1}{4} 5000$ ,  $h \frac{1}{4} 3:0$ ,  $e \frac{1}{4} 10^{-2}$ , one EFIE3D matrix for a unit sphere with  $n \frac{1}{4} 26268$ ,  $e \frac{1}{4} 10^{-6}$  and 20 points per wavelength, one Frontal3D matrix with  $n \frac{1}{4} 1250$ ,  $e \frac{1}{4} 10^{-6}$  and 10 points per wavelength, one polynomial matrix with  $n \frac{1}{4} 10; 000$ ,  $h \frac{1}{4} 0:2$ ,  $e \frac{1}{4} 10^{-4}$ , and one product-of-random matrix with  $n \frac{1}{4} 2500$ ,  $e \frac{1}{4} 10^{-4}$ . The corresponding e-ranks are 298, 137, 1488, 788, 450, and 1000, respectively. It can be validated that the hierarchical merge operation attains increasing ranks for the Gaussian, EFIE3D and Frontal3D matrices, and relatively constant ranks for the polynomial, and product-of-random matrices. All examples use one process except that the Gaussian-SUSY example uses 16 processes. The CPU times are measured and plotted in Figure 5.

Table 1 predicts that H-BACA exhibits increasing (with a factor of  $n_b$ ) and constant

time when  $s_i$  stays constant and increases, respectively. Note that the rank assumption


$s_i \ll r$  leading to the  $O(n_b^p)$  computational overhead may

not be fully observed for practical values of  $n_b$  and  $n$ . Given one matrix,  $s_i$  may stay approximately constant for a



limited number of subdivision levels  $l$ . For example,  $s_l$  stay constant for bottom levels of EFIE3D and Frontal3D matrices, and top levels of Polynomial and product-of-random matrices. This agrees with the observed scalings (w.r.t  $n_b$ ) in Figure 5(c) to (f). As a reference, the  $O(n_b^3)$  curves are plotted and only small ranges of  $n_b$  exhibit the  $O(n_b^3)$  overhead. For the Gaussian matrices, we even observe nonincreasing CPU time w.r.t.  $n_b$  when  $n_b$  is not too big (see Figure 5(a) and (b)).

The effects of varying block size  $d$  also deserve further discussions. First, larger block size  $d$  can significantly improve the robustness of H-BACA for the Gaussian matrices. For example, H-BACA does not achieve desired accuracies due to premature termination for all data points on the  $d = 1$  curve in Figure 5(a) and  $d = 1$  and  $d = 8$  curves in Figure 5(b). In contrast, H-BACA with larger  $d$  attains desired accuracies. Second, larger block size  $d$  results in reduced CPU time for the Polynomial and Frontal3D matrices due to better BLAS performance (see Figure 5(d) and (e)). For the other tested matrices, no significant performance differences have been observed by changing block size  $d$ . However, for matrices with ranks  $s_0 \leq d$ , larger  $d$  and  $n_b$  can introduce significant overheads.



1/4

1/4  
1/4




Figure 5. Computation time of H-BACA with varying  $n_b$  and  $d$  for the (a) Gaussian-SUSY kernel with  $h \approx 1:0$ ,  $n \approx 50; 000$ ,  $e \approx 10^{-2}$ ,  $r \approx 298$ , (b) Gaussian-MNIST kernel with  $h \approx 3:0$ ,  $n \approx 5000$ ,  $e \approx 10^{-2}$ ,  $r \approx 137$ , (c) EFIE3D kernel for a unit sphere with  $n \approx 26; 268$ ,  $e \approx 10^{-4}$ ,  $r \approx 1488$ , (d) Frontal3D kernel with  $n \approx 1250$ ,  $e \approx 10^{-6}$ ,  $r \approx 1788$ , (e) polynomial kernel with  $h \approx 0:2$ ,  $n \approx 10; 000$ ,  $e \approx 10^{-4}$ ,  $r \approx 450$ , and (f) product-of-random kernel with  $n \approx 2500$ ,  $r \approx 1000$ . Note that the data points where the algorithm fails are shown as triangular markers without lines.

#### 5.4. Parallel performance

Finally, the parallel performance of the H-BACA algorithm is demonstrated via strong scaling studies with the EFIE2D, EFIE3D, product-of-random and Gaussian matrices with process counts  $p \in \{8, \dots, 1024\}$ . For the EFIE2D matrices,  $n \approx 160; 000$  and the wavenumbers are chosen such that the e-ranks with  $e \approx 10^{-4}$  are 937 and 107, respectively. For the EFIE3D matrices for a unit square,  $n \approx 21; 788$  and the wavenumbers are chosen such that the e-ranks with  $e \approx 10^{-6}$  are 1007 and 598, respectively. For the product-of-random matrices,  $n \approx 10; 000$  and the inner dimension of the product is set to  $r$

2000 and 800, respectively. For the Gaussian matrices with a randomly generated data set of

dimension 50 and  $n \approx 10,000$ , we choose  $h \approx 1:0$  and  $h \approx 1:6$  such that the e-ranks with  $\epsilon = 10^{-3}$  are 2106 and 191, respectively. In all examples, the block size and number of leaf-level subblocks in H-BACA are chosen as  $d \approx 8$  and  $n_b \approx p$ .

As the reference, we compare to a straightforward parallel implementation of the baseline ACA algorithm which essentially parallelize every operation in ACA with collective MPI communications.

For all examples, the parallel ACA algorithm stops scaling when  $p$  is sufficiently large (see Figure 6). In contrast, the proposed parallel H-BACA algorithm scales up to  $p \approx 1024$ . In most examples, H-BACA achieves better

$\approx 1/4$   
 $p \approx 1/4$  and  $p \approx 1/4$ . The ScaLAPACK block size is set to




Figure 6. Computation time of H-BACA with varying process counts for the (a) EFIE2D kernel with  $n \approx 160,000$ ,  $\epsilon \approx 10^{-4}$ ,  $r \approx 107$  and  $937$ , (b) EFIE3D kernel for a unit square with  $n \approx 21,788$ ,  $\epsilon \approx 10^{-6}$ ,  $r \approx 598$  and  $1007$ , (c) product-of-random kernel with  $n \approx 10,000$ ,  $r \approx 800$  and  $2000$ , and (d) Gaussian kernel for a randomly generated data set with  $h \approx 1:0, 1:6$ ,  $\epsilon \approx 10^{-3}$ ,  $r \approx 2106$  and  $191$ . Note that for the Gaussian matrix with  $r \approx 191$ , ACA fails to provide accurate results and is not plotted.

parallel efficiencies with larger ranks due to better process utilization during the hierarchical merge operation. We also note that ACA outperforms H-BACA for the Product-of-random matrices with small process count  $p$  (and  $n_b$ ). This is partially attributed to the  $O(p n_b)$  overhead observed in Figure 5(f).

Overall, the parallel H-BACA algorithm can achieve reasonably good parallel performances for rank-deficient matrices with modest to large numerical ranks. Not surprisingly, the parallel runtime is dominated by that of ScaLAPACK computation and possible redistributions between each re-compression step as analyzed in Section 4. Also note that the leaf-level BACA compression is embarrassingly parallel for all test cases.


This article presents a parallel and algebraic ACA-type matrix decomposition algorithm given that any matrix entry can be evaluated in  $O(1)$  time. Two proposed strategies, BACA and H-BACA, are leveraged to improve the robustness and parallel efficiency of the (baseline) ACA algorithm for general rank-deficient matrices.

## 6. Conclusion

First, the BACA algorithm searches for blocks of row/ column pivots via QRCP on the column/row submatrices at each iteration. The blocking nature of BACA provides a closer estimation of the true residual error and reduces the chance of selecting smaller pivots when compared to ACA. Therefore, BACA exhibits a much smoother and more reliable convergence history. Moreover, blocked operations also benefit from higher flop performance compared to non-blocked ones. For a rank-deficient matrix with dimension  $n$  and e-rank  $r$ , the computational cost of BACA is  $O(nr^2)$  assuming the block size constant and iteration count  $O(r)$ .

Second, the H-BACA algorithm divides the matrix into  $n_b$  similar-sized submatrices each compressed with BACA and then hierarchically merges the results using low-rank arithmetic. Depending on the rank behaviors of submatrices during the merge, the H-BACA may have a computational overhead of  $O(p)n$  yielding the overall computational cost at most  $O(nr^2p/n_b)$ . The H-BACA algorithm can be parallelized with distributed-memory machines by assigning each process to one submatrix and leveraging PBLAS and ScaLAPACK for the hierarchical merge operation. Such parallelization strategy yields a much more favorable communication cost when compared to the straightforward parallelization of ACA/BACA with

$\delta$   
 $\delta p$



collective MPI routines. Not surprisingly, good parallel performance can be achieved for matrices with modest to large numerical ranks which increases process utilization for each merge operation.

In contrast to the baseline ACA algorithm, the proposed algorithms exhibit improved robustness and favorable parallel performance with low computational overheads for broad ranges of matrices arising from many science and engineering applications.

#### Authors' note

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


#### Declaration of conflicting interests

The author(s) declared no potential conflicts of interest with respect to the research, authorship, and/or publication of this article.

#### Funding

The author(s) disclosed receipt of the following financial support for the research, authorship, and/or publication of this article: This research was supported in part by the Exascale Computing Project (17-SC-20-SC), a collaborative effort of the US Department of Energy Office of Science and the National Nuclear Security Administration, and in part by the US 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.

#### ORCID iD

Yang Liu  <https://orcid.org/0000-0003-3750-1178>

#### References

- Bach F (2013) Sharp analysis of low-rank kernel matrix approximations. In: *Proceedings of the 26th annual conference on learning theory, Proceedings of machine learning research*, vol. 30. Princeton, USA: PMLR, 12-14 June 2013, pp. 185-209.
- Balzano L, Nowak R and Recht B (2010) Online identification and tracking of subspaces from highly incomplete information. In: *2010 48th annual Allerton conference on communication,*

*control, and computing (Allerton)*, pp. 704-711. DOI: 10.1109/ALLERTON.2010.5706976.

Bebendorf M (2000) Approximation of boundary element matrices. *Numerische Mathematik* 86(4): 565-589.

Bebendorf M and Grzhibovskis R (2006) Accelerating Galerkin BEM for linear elasticity using adaptive cross approximation. *Mathematical Methods in the Applied Sciences* 29(14): 1721-1747.

- Blackford LS, Choi J, Cleary A, et al. (1997) *ScaLAPACK Users' Guide*. Philadelphia: Society for Industrial and Applied Mathematics.
- Boutsidis C, Mahoney MW and Drineas P (2009) An improved approximation algorithm for the column subset selection problem. In: *Proceedings of the twentieth annual ACM-SIAM symposium on discrete algorithms, SODA '09*, Philadelphia: Society for Industrial and Applied Mathematics, New York, New York, USA, 4-6 January 2009, pp. 968-977.
- Candès EJ and Recht B (2009) Exact matrix completion via convex optimization. *Foundations of Computational Mathematics* 9(6): 717.
- Cheng H, Gimbutas Z, Martinsson P, et al. (2005) On the compression of low rank matrices. *SIAM Journal on Scientific Computing* 26(4): 1389-1404.
- Dheeru D and Karra Taniskidou E (2017) UCI machine learning repository. Available at: <http://archive.ics.uci.edu/ml> (accessed 2 December 2014).
- Feng Y, Xiao J and Gu M (2019) Flip-flop spectrum-revealing QR factorization and its applications to singular value decomposition. *Electronic Transactions on Numerical Analysis* 51: 469-494.
- Foster LV (1997) The growth factor and efficiency of Gaussian elimination with rook pivoting. *Journal of Computational and Applied Mathematics* 86(1): 177-194.
- Grasedyck L and Hackbusch W (2003) Construction and Arithmetics of H-matrices. *Computing* 70(4): 295-334.
- Grasedyck L and Hackbusch W (2005) Hybrid cross approximation of integral operators. *Numerische Mathematik* 101(2): 221-249.
- Gu M and Eisenstat S (1996) Efficient algorithms for computing a strong rank-revealing QR factorization. *SIAM Journal on Scientific Computing* 17(4): 848-869.
- Hackbusch W, Grasedyck L and Börm S (2002) An introduction to hierarchical matrices. *Mathematica Bohemica* 127(2): 229-241.
- Heldring A, Ubeda E and Rius JM (2014) On the convergence of the ACA algorithm for radiation and scattering problems. *IEEE Transactions on Antennas and Propagation* 62(7): 3806-3809.
- Heldring A, Ubeda E and Rius JM (2015) Stochastic estimation of the Frobenius norm in the ACA convergence criterion. *IEEE Transactions on Antennas and Propagation* 63(3): 1155-1158.
- Liberty E, Woolfe F, Martinsson PG, et al. (2007) Randomized algorithms for the low-rank approximation of matrices. *Proceedings of the National Academy of Sciences* 104(51): 20167-20172.
- Mahoney MW and Drineas P (2009) CUR matrix decompositions for improved data analysis. *Proceedings of the National Academy of Sciences* 106(3): 697-702.
- Martinsson PG, Quintana-Ortí and Heavner N (2019) randUTV: a blocked randomized algorithm for computing a rank-revealing UTV factorization. *ACM Transactions on Mathematical Software* 45(1): 1-26.
- Musco C and Musco C (2017) Recursive sampling for the Nystro"m method. In: *Advances in Neural Information*

- Processing Systems* 30. New York: Curran Associates, Inc., pp. 3833–3845.
- Neal L and Poole G (1992) A geometric analysis of Gaussian elimination. II. *Linear Algebra and its Applications* 173: 239–264.
- Poole G and Neal L (1991) A geometric analysis of Gaussian elimination. I. *Linear Algebra and its Applications* 149: 249–272.
- Poole G and Neal L (1992) Gaussian elimination: When is scaling beneficial? *Linear Algebra and its Applications* 162–164: 309–324.
- Rebrova E, Chavez G, Liu Y, et al. (2018) A study of clustering techniques and hierarchical matrix formats for kernel ridge regression. *2018 IEEE international parallel and distributed processing symposium workshops (IPDPSW)*, Vancouver, Canada, 21–25 May 2018, pp. 883–892.
- Voronin S and Martinsson PG (2017) Efficient algorithms for CUR and interpolative matrix decompositions. *Advances in Computational Mathematics* 43(3): 495–516.
- Wang R, Li Y and Darve E (2018) On the numerical rank of radial basis function kernels in high dimensions. *SIAM Journal on Matrix Analysis and Applications* 39(4): 1810–1835.
- Xiao J, Gu M and Langou J (2017) Fast parallel randomized QR with column pivoting algorithms for reliable low-rank matrix approximations. In: *2017 IEEE 24th international conference on high performance computing (HiPC)*, pp. 233–242. DOI: 10.1109/HiPC.2017.00035.
- Zhao K, Vouvakis MN and Lee JF (2005) The adaptive cross approximation algorithm for accelerated method of moments computations of EMC problems. *IEEE Transactions on Electromagnetic Compatibility* 47(4): 763–773.
- Zhou H, Zhu G and Kong W, et al. (2017) An upgraded ACA algorithm in complex field and its statistical analysis. *IEEE Transactions on Antennas and Propagation* 65(5): 2734–2739.



