

# A DISTRIBUTED-MEMORY HIERARCHICAL SOLVER FOR GENERAL SPARSE LINEAR SYSTEMS

Chao Chen<sup>a,\*</sup>, Hadi Pouransari<sup>b</sup>, Sivasankaran Rajamanickam<sup>c</sup>, Erik G. Boman<sup>c</sup>, Eric Darve<sup>a,b</sup>

<sup>a</sup>*Institute for Computational and Mathematical Engineering, Stanford University, Stanford, CA, USA*

<sup>b</sup>*Department of Mechanical Engineering, Stanford University, Stanford, CA, USA*

<sup>c</sup>*Center for Computing Research, Sandia National Laboratories, NM, USA.*

---

## Abstract

We present a parallel hierarchical solver for general sparse linear systems on distributed-memory machines. For large-scale problems, this fully algebraic algorithm is faster and more memory-efficient than sparse direct solvers because it exploits the low-rank structure of fill-in blocks. Depending on the accuracy of low-rank approximations, the hierarchical solver can be used either as a direct solver or as a preconditioner. The parallel algorithm is based on data decomposition and requires only local communication for updating boundary data on every processor. Moreover, the computation-to-communication ratio of the parallel algorithm is approximately the volume-to-surface-area ratio of the subdomain owned by every processor. We present various numerical results to demonstrate the versatility and scalability of the parallel algorithm.

*Keywords:* Parallel linear solver, Sparse matrix, Hierarchical matrix

---

---

\*Corresponding author

*Email addresses:* cchen10@stanford.edu (Chao Chen), hadip@stanford.edu (Hadi Pouransari), srajama@sandia.gov (Sivasankaran Rajamanickam), egboman@sandia.gov (Erik G. Boman), darve@stanford.edu (Eric Darve)

## 1. Introduction

Solving large sparse linear systems is an important building block – but often a computational bottleneck – in many engineering applications. Large sparse linear systems arise, for example, from local discretization of elliptic partial differential equations. Solving a general linear system that has  $N$  non-zeros with  $O(N)$  computer memory and CPU time is challenging, especially when the underlying physical problem is three-dimensional. Existing methods fall into three categories as follows.

The first category of methods are sparse direct solvers [1] based on Gaussian elimination. These methods organize computation efficiently and leverage fill-reducing ordering schemes. For example, the nested dissection multifrontal algorithm [2, 3] performs elimination according to a specific hierarchical structure of the unknowns. Because of their robustness and efficiency, several state-of-the-art sparse direct solvers have been implemented into software packages, which target sequential [4–6], shared-memory [7, 8] and distributed-memory computers [9, 10]. However, sparse direct solvers generally require  $O(N^2)$  work and  $O(N^{\frac{4}{3}})$  computer memory for a three-dimensional problem of size  $N$ . This quadratic factorization cost and large memory footprint seriously limit the application of sparse direct solvers to truly large-scale problems.

The second category consists of iterative solvers. These solvers require only  $O(N)$  computer memory to store the linear system and are thus more memory-efficient than sparse direct solvers. They can also achieve the optimal time-complexity if the number of iterations is small. For example, the multigrid method [11] is typically the fastest solver for many discretized elliptic PDEs. However, iterative solvers have three disadvantages. First, the convergence of iterative solvers is not guaranteed for general linear systems. For example, the multigrid method may fail to converge for indefinite linear systems and linear systems coming from coupled PDEs. Second, the number of iterations may grow rapidly as the condition number of a linear system increases. Third, the setup phase of some iterative solvers relies on sparse matrix-sparse matrix multiplication, as in the algebraic multigrid method, which is complex to scale [12, 13].

The third category of methods developed for solving sparse linear systems are hierarchical solvers [14–19] and their parallel counterparts [20–23]. The general idea behind these hierarchical solvers is exploiting the low-rank structure of dense matrix blocks that arise during the elimination process

to reduce storage and computational cost. As a result of utilizing the data sparsity, hierarchical solvers, compared with sparse direct solvers, typically have reduced memory footprint and smaller computational complexity. However, some of the existing hierarchical solvers still cannot attain quasilinear complexity for solving three-dimensional problems, and others may be either too complicated to be implemented efficiently or may be restricted to only structured problems.

In this paper, we introduce a new parallel hierarchical solver for solving general sparse linear systems. Our parallel algorithm is based on the sequential algorithm in [19] called LoRaSp, which computes an approximate factorization with sparse block-triangular factors. In particular, our method eliminates the unknowns cluster by cluster and compresses fill-in blocks in a hierarchical fashion. The singular values of these fill-in blocks are often found to decrease geometrically, and general algebraic techniques, such as the SVD, can be used to compute corresponding low-rank approximations. Since the dropping/truncation rule in our method is based on the decay of singular values, it is expected to be more efficient than other level-based or threshold-based rules, which are typically used in the incomplete LU (ILU) factorization [24]. In our method, the bases computed in low-rank approximations serve the same role as restriction and prolongation operators do in the algebraic multigrid method (AMG) [25, 26]. While the construction of restriction and prolongation operators in AMG may require tuning and manual adjustments for a specific linear system, the low-rank bases in our method are computed in a systematic fashion, regardless of the underlying PDE or physical problem.

There are two differences between our method and other hierarchical solvers. First, while most of hierarchical solvers are developed under either the  $\mathcal{H}$ - [27, 28] or the hierarchically semiseparable (HSS) [29, 30] matrix frameworks, our method is built upon the  $\mathcal{H}^2$ -matrix theory [31, 32], which provides a more efficient hierarchical low-rank structure. Under some mild conditions, the computational cost and the memory footprint of our method scale linearly with respect to the problem size, and we observed quasilinear complexity in practice for solving various types of problems. Second, unlike other hierarchical solvers, which are typically combined with the multifrontal algorithm, our method relies on domain partitioning, which naturally leads to a data decomposition scheme in the parallel algorithm. Put another way, these two differences allow our parallel algorithm to have the following three features:

1. Only local communication is required for every processor.
2. The computation-to-communication ratio is approximately the volume-to-surface-area ratio of the subdomain owned by a processor.
3. The bulk of computation is from using sequential dense linear algebra, which has the potential to be significantly accelerated on modern many-core architectures.

To summarize, this paper presents a parallel hierarchical solver for general sparse linear systems, and especially, our work makes the following three major contributions:

1. New derivation of the LoRaSp algorithm, which reveals the structure of calculation and data dependency in the original algorithm;
2. Development of a bulk-synchronous parallel algorithm and a task-based asynchronous parallel algorithm with the optimal scheduling strategy;
3. Development of a coloring scheme to extract maximum concurrency in the execution, and discussion on optimizing load-balancing in the presence of coloring constraints.

The remainder of this paper is organized as follows. Section 2 presents the sequential algorithm, a new derivation of LoRaSp. Section 3 presents the parallel algorithm, focusing on techniques to keep communication local and to maximize concurrency. Section 4 analyzes the computation and communication cost of the parallel algorithm. Section 5 presents numerical results to demonstrate the versatility and parallel scalability of our parallel hierarchical solver.

## 2. Sequential algorithm

This section presents our new derivation of the LoRaSp algorithm. Although the algorithm works for a general sparse linear system  $Ax = b$ , we focus on symmetric positive definite (SPD) systems for ease of presentation. From a high-level perspective, the algorithm computes an approximate factorization of an SPD matrix  $A$  with the following steps.

First, a partitioning of the rows/columns of  $A$  is computed algebraically. Suppose  $\Pi$  is the set of row/column indices (DOFs). A clustering  $\Pi = \cup_i \pi_i$ , where  $\pi_i$  is a cluster of DOFs, can be computed with a graph partitioner, such as METIS/ParMETIS [33], Scotch [34] and Zoltan [35]. Second, some portions of DOFs in every cluster are eliminated as all clusters in  $\Pi = \cup_i \pi_i$

are looped over. Specifically, the fill-in blocks associated with a cluster  $\pi_s$  are compressed, and  $\pi_s$  is split into fine DOFs  $\pi_s^f$  and coarse DOFs  $\pi_s^c$ , where  $\pi_s^f$  involves no fill-in. Then  $\pi_s^f$  is eliminated safely, which does not propagate any existing fill-in (no level-2 fill-in introduced). Third, after all fine DOFs are eliminated, a smaller linear system,  $A_c x_c = b_c$ , corresponding to the coarse DOFs remains, and the previous steps are applied repeatedly until  $A_c$  is small enough to be factorized with the conventional Cholesky factorization.

### 2.1. Low-rank elimination

This subsection illustrates the key step in the algorithm, which exploits the low-rank structure of fill-in blocks to eliminate the fine DOFs in a cluster (if a cluster to be eliminated has no fill-in, e.g., the first eliminated cluster, then all DOFs of this cluster are fine DOFs). Suppose at least one cluster in  $\Pi = \cup_i \pi_i$  has been processed (fine DOFs have been eliminated), and consider the Schur complement,  $\bar{A}$ , which corresponds to the remaining clusters:

$$\bar{A} = \begin{pmatrix} A_{ss} & A_{sn} & A_{sw} \\ A_{sn}^T & A_{nn} & A_{nw} \\ A_{sw}^T & A_{nw}^T & A_{ww} \end{pmatrix}$$

where

- $s$  stands for a cluster  $\pi_s$  to be processed (“self”),
- $n$  stands for neighbor clusters (“neighbor”), and
- $w$  stands for the rest (“well-separated”).

Two clusters  $\pi_a$  and  $\pi_b$  are neighbors if  $A(\pi_a, \pi_b) \neq 0$ . In other words, the two clusters are connected in the graph of  $A$ . The definition of *well-separated* is important in  $\mathcal{H}$ - and  $\mathcal{H}^2$ - matrix theories, which implies a partitioning of the clusters in  $n$  and  $w$  above. Although our framework is general enough that it can use different definitions of well-separated criteria, here we focus on a simple definition that two clusters are well-separated if they are not neighbors. This definition implies that well-separated blocks,  $A_{sw}$  and  $A_{ws} = A_{sw}^T$ , are fill-in blocks.

Next, fine DOFs in  $\pi_s$  will be selected and eliminated. It is observed in [19] that the fill-in block,  $A_{sw}$ , is numerically low-rank. With an algebraic method, such as the SVD, we obtain the low-rank approximation

$$A_{sw} = UZ + O(\epsilon) \tag{1}$$

where  $U$  is an  $m \times k$  orthonormal matrix ( $k \ll m$ ). Using a Gram-Schmidt type of orthogonalization procedure, an orthonormal matrix  $V$  of size  $m \times (m - k)$  can be computed such that

$$V^T A_{ss}^{-1} U = 0.$$

We introduce the *sparsification operator*:

$$\mathcal{E}_s = \begin{pmatrix} V^T A_{ss}^{-1} & & & \\ U^T A_{ss}^{-1} & & & \\ & I & & \\ & & I & \end{pmatrix}, \quad (2)$$

which decouples  $\pi_s^f$  from  $\pi_s^c$  and  $w$  clusters as

$$\mathcal{E}_s \bar{A} \mathcal{E}_s^T \approx \begin{pmatrix} V^T A_{ss}^{-1} V & & V^T A_{ss}^{-1} A_{sn} & \\ & U^T A_{ss}^{-1} U & U^T A_{ss}^{-1} A_{sn} & (U^T A_{ss}^{-1} U) Z \\ A_{sn}^T A_{ss}^{-1} V & A_{sn}^T A_{ss}^{-1} U & A_{nn} & A_{nw} \\ Z^T (U^T A_{ss}^{-1} U) & A_{nw}^T & A_{ww} & \end{pmatrix}$$

where the first  $m - k$  rows/columns correspond to the fine DOFs  $\pi_s^f$  and the rest of the DOFs in  $\pi_s$  correspond to the coarse DOFs  $\pi_s^c$ . Note that eliminating  $\pi_s^f$  does not propagate any existing fill-in associated with  $\pi_s$  in the sense of block elimination and low-rank representation (e.g.,  $A_{sn}^T A_{ss}^{-1} V$  is not considered as fill-in). In other words, no level-2 fill-in is introduced from the elimination of  $\pi_s^f$ .

To eliminate  $\pi_s^f$ , we introduce the *elimination operator*:

$$\mathcal{G}_s = \begin{pmatrix} I & & & \\ & I & & \\ -(A_{sn}^T A_{ss}^{-1} V) L^{-T} & & I & \\ & & & I \end{pmatrix} \begin{pmatrix} L^{-1} & & & \\ & I & & \\ & & I & \\ & & & I \end{pmatrix} \quad (3)$$

where  $V^T A_{ss}^{-1} V = LL^T$  is the Cholesky factorization, and the first diagonal block in  $\mathcal{G}_s \mathcal{E}_s \bar{A} \mathcal{E}_s^T \mathcal{G}_s^T$  corresponding to  $\pi_s^f$  is an identity matrix.

We introduce the *permutation operator*:

$$P_s = \begin{pmatrix} I & & & \\ & I & & \\ & & I & \\ & & & I \end{pmatrix} \quad (4)$$

where the four identity matrices (from the first to the last row) in  $P_s$  have the same sizes as the number of DOFs in  $\pi_s^f$ ,  $n$  clusters,  $w$  clusters and  $\pi_s^c$ . Note that  $P_s$  is used to permute rows and columns corresponding to  $\pi_s^c$  to the last in  $\bar{A}$ .

Combining the above three operators, we define the *low-rank elimination operator*:  $\mathcal{W}_s = P_s \mathcal{G}_s \mathcal{E}_s$ , and  $\mathcal{W}_s \bar{A} \mathcal{W}_s^T$  selects and eliminates the fine DOFs in  $\pi_s$  (and permutes  $\pi_s^c$  to be the last indices). Again, applying  $\mathcal{W}_s$  and  $\mathcal{W}_s^T$  on both sides of  $\bar{A}$  does not introduce any level-2 fill-in in  $\bar{A}$ . The pseudo-code of the low-rank elimination step is shown at lines 1520 in Algorithm 1.

## 2.2. Hierarchical solver

The hierarchical solver repeatedly applies the low-rank elimination procedure to all clusters in  $\Pi = \cup_i \pi_i$ , and the result is equivalent to computing an approximate factorization of the original matrix, subject to the error of low-rank approximation. The algorithm the hierarchical solver is shown at lines 114 in Algorithm 1.

In Algorithm 1,  $\pi_0$  is the first cluster to be eliminated. Since no fill-in exists yet,  $\mathcal{E}_0$  is an identity matrix, and applying  $\mathcal{W}_0$  and  $\mathcal{W}_0^T$  is one step of standard block Cholesky factorization. This is also true for any cluster that does not involve any fill-in. Suppose the elimination of  $\pi_0$  introduces some fill-in for  $\pi_1$ . Then, the low-rank elimination procedure is applied to  $\bar{A}$ , the Schur complement in  $\mathcal{W}_0 A \mathcal{W}_0^T = \begin{pmatrix} I & \\ & \bar{A} \end{pmatrix}$ . After applying the low-rank elimination procedure to all clusters in  $\Pi$ , we are left with the coarse DOFs,  $\cup_{i=0}^{m-1} \pi_i^c$ , and we can apply the same process to the coarse DOFs.

Algorithm 1 outputs a factorization of the original matrix  $A$ , which can be used to solve the corresponding linear system, and the solve phase follows the standard forward and backward substitution, as shown in Algorithm 2.

*Relation to LoRaSp.* Algorithm 1 is mathematically equivalent to LoRaSp [19]. However, LoRaSp uses *extended sparsification* (introducing auxiliary variables) and needs to maintain a binary tree structure of all DOFs, all of

---

**Algorithm 1** Hierarchical solver: factorization phase
 

---

```

1: procedure HIERARCHICAL_FACTOR( $A$ )
2:   if the size of  $A$  is small enough then
3:     Factorize  $A$  with the conventional Cholesky factorization
4:   return
5: end if
6:   Partition the graph of  $A$  and obtain vertex clusters  $\Pi = \cup_{i=0}^{m-1} \pi_i$ 
    $\triangleright m$  is chosen to get roughly constant cluster sizes

7:    $A_0 \leftarrow A$ 
8:   for  $i \leftarrow 0$  to  $m - 1$  do
9:      $A_{i+1} \leftarrow \text{LowRank\_Elimination}(A_i, \Pi, \pi_i)$ 
10:  end for  $\triangleright A_m = \mathcal{W}_{m-1} \dots \mathcal{W}_1 \mathcal{W}_0 A \mathcal{W}_0^T \mathcal{W}_1^T \dots \mathcal{W}_{m-1}^T$ 
11:   Extract  $A_c$  from the block diagonal matrix  $A_m \approx \begin{pmatrix} I & \\ & A_c \end{pmatrix}$ 
    $\triangleright A_c$  is the Schur complement for the coarse DOFs
12:    $A_c^{fac} \leftarrow \text{HIERARCHICAL\_FACTOR}(A_c)$ 
    $\triangleright$  Recursive call with a smaller matrix
13:   return  $A^{fac} = \mathcal{W}_0^{-1} \mathcal{W}_1^{-1} \dots \mathcal{W}_{m-1}^{-1} \begin{pmatrix} I & \\ & A_c^{fac} \end{pmatrix} \mathcal{W}_{m-1}^{-T} \dots \mathcal{W}_1^{-T} \mathcal{W}_0^{-T}$ 
    $\triangleright A_c^{fac}$  is not written out explicitly

14: end procedure

15: procedure LOWRANK_ELIMINATION( $A_i, \Pi, \pi_i$ )
16:   Extract  $\bar{A}$  from  $A_i \approx \begin{pmatrix} I & \\ & \bar{A} \end{pmatrix}$ 
17:   Compute the low-rank elimination operator  $\bar{\mathcal{W}}_i = P_i \mathcal{G}_i \mathcal{E}_i$  based on  $\bar{A}$ 
    $\triangleright \mathcal{E}_i, \mathcal{G}_i$  and  $P_i$  are defined in Eq. (2), Eq. (3) and Eq. (4)
18:    $\mathcal{W}_i \leftarrow \begin{pmatrix} I & \\ & \bar{\mathcal{W}}_i \end{pmatrix}$ 
    $\triangleright \mathcal{W}_i$  has the same size as  $A_i$ 
19:   return  $\mathcal{W}_i A_i \mathcal{W}_i^T$ 
20: end procedure
    $\triangleright$  Notation:  $a \leftarrow b$  means assign the value  $b$  to  $a$ , whereas  $a = b$  means they are equivalent

```

---



---

**Algorithm 2** Hierarchical solver: solve phase

---

```
1: procedure HIERARCHICAL_SOLVE( $A^{fac}$ ,  $b$ )
2:    $y \leftarrow$  FORWARD_SUBSTITUTION( $A^{fac}$ ,  $b$ )
3:    $x \leftarrow$  BACKWARD_SUBSTITUTION( $A^{fac}$ ,  $y$ )
4:   return  $x$ 
5: end procedure

6: procedure FORWARD_SUBSTITUTION( $A^{fac}$ ,  $b$ )
7:    $y \leftarrow b$ 
8:   for  $i \leftarrow 0$  to  $m - 1$  do
9:      $y \leftarrow \mathcal{W}_i y$   $\triangleright y$  is overwritten
10:  end for  $\triangleright y = (y_c, y_f)$  is of the concatenation of  $y_f$  and  $y_c$ 
11:  Extract  $y_f$  and  $y_c$  from  $y$ 
 $\triangleright y_f$  and  $y_c$  correspond to the fine DOFs and the coarse DOFs
12:   $y_c \leftarrow$  FORWARD_SUBSTITUTION( $A_c^{fac}$ ,  $y_c$ )  $\triangleright y_c$  is overwritten
13:  return  $y = (y_f, y_c)$   $\triangleright$  output the concatenation of  $y_f$  and  $y_c$ 
14: end procedure

15: procedure BACKWARD_SUBSTITUTION( $A^{fac}$ ,  $y$ )
16:    $x \leftarrow y$ 
17:   for  $i \leftarrow m - 1$  to  $0$  do
18:      $x \leftarrow \mathcal{W}_i^T x$   $\triangleright x$  is overwritten
19:   end for  $\triangleright x = (x_f, x_c)$  is of the concatenation of  $x_c$  and  $x_f$ 
20:   Extract  $x_f$  and  $x_c$  from  $x$ 
 $\triangleright x_f$  and  $x_c$  correspond to the fine DOFs and the coarse DOFs
21:    $x_c \leftarrow$  BACKWARD_SUBSTITUTION( $A_c^{fac}$ ,  $x_c$ )  $\triangleright x_c$  is overwritten
22:   return  $x = (x_f, x_c)$   $\triangleright$  Output the concatenation of  $x_f$  and  $x_c$ 
23: end procedure
```

$\triangleright$  Notation:  $a \leftarrow b$  means assign the value  $b$  to  $a$ , whereas  $a = b$  means they are equivalent

---

which are avoided in the new derivation. This new derivation also allows more general partitioning strategies. In the special case when recursive bisection is used to compute the partitioning  $\Pi = \cup \pi_i$ , the new algorithm becomes fairly similar to LoRaSp.

### 2.3. Fill-in property

As we emphasized earlier,

*level-2 fill-in is never created*

in the hierarchical solver. This fill-in property is not only the fundamental reason why the hierarchical solver is efficient, but also the base of the parallel algorithm. To make the statement of fill-in property precise, we introduce a new notation,  $A_B$ , to denote the block sparsity pattern of  $A$  corresponding to a partitioning of its rows/columns. Given a partition  $\Pi = \cup_{i=0}^{m-1} \pi_i$ ,  $A_B$  is an  $m$ -by- $m$  matrix such that

$$A_B(i, j) = \begin{cases} 0 & \text{if the subblock } A(\pi_i, \pi_j) = 0, \\ 1 & \text{else.} \end{cases} \quad (5)$$

With the notation of block sparsity pattern, we restate the fill-in property of the hierarchical solver as below:

*fill-in blocks correspond to non-zeros in  $(A_B)^2$  that are not in  $A_B$ .*

Using ILU terminology, the amount of fill-in never goes beyond ILU(1). Intuitively, in the hierarchical solver, the low-rank elimination operator decouples fine DOFs from fill-in, so eliminating the fine DOFs introduces only level-1 fill-in. See Theorem 5.3 in [19] for a formal proof.

This fill-in property of the hierarchical solver is the base of the parallel algorithm. Given a partition of all DOFs  $\Pi = \cup_i \pi_i$ , two clusters  $\pi_i$  and  $\pi_j$  never “interact” if the distance between them is greater than two. The distance is defined as the length of the shortest path between vertex  $i$  and vertex  $j$  in the graph of  $A_B$ .

## 3. Parallel hierarchical solver

Similar to the parallelization of a multigrid method, our parallel algorithm uses a data decomposition scheme, where every processor owns a subdomain

of the entire grid. Since every processor exchanges only boundary data with its neighbor processors, the communication is always local in the parallel algorithm. Moreover, the computation-to-communication ratio is approximately the volume-to-surface-area ratio of the subdomain. In other words, the amount of communication is one order of magnitude less than the amount of computation when the subdomain is large enough.

### 3.1. Data decomposition

We present our parallel algorithm for the finest grid, which corresponds to the graph  $G$  of  $A_B$  (defined in Eq. (5)). The same algorithm applies for the coarse levels. In the following discussion, we use the term “node” to mean a cluster of DOFs in  $\Pi = \cup_i \pi$ , and the term “edge” to mean a matrix block (the edge between node  $i$  and node  $j$  corresponds to block  $A(\pi_i, \pi_j)$ ).

In the parallel algorithm, the global graph  $G$  is decomposed among all processors  $G = \cup_P G_P$ , where processor  $P$  owns a subgraph  $G_P$ . With this decomposition, nodes owned by one processor can be classified into three categories:

- d1 nodes: boundary nodes,
- d2 nodes: (local) neighbors of d1 nodes that are not on the boundary,
- d3 nodes: the remaining interior nodes.

Fig. 1 shows a simple example to clarify the above definitions. We assume that the matrix  $A$  is distributed by rows among all processors. Every processor  $P$  owns the submatrix corresponding to the local graph  $G_P$  and also stores the edges to neighbor processors. For example, the matrix  $A_0$  owned by processor  $P_0$  in Fig. 1 is

$$A_0 = \begin{pmatrix} A_0^{d3} & A_0^{d3,d2} & & & \\ A_0^{d2,d3} & A_0^{d2} & A_0^{d2,d1} & & \\ & A_0^{d1,d2} & A_0^{d1} & A_0^{d1} & \\ & & & & A_0^{d1} \end{pmatrix}$$

where for the sake of simplicity we assumed that d3 nodes were ordered before d2 nodes and d1 nodes. The coupling with processor  $P_1$  is contained within the  $A_{0,1}^{d1}$  block, since only d1 nodes have edges to processor  $P_1$ .

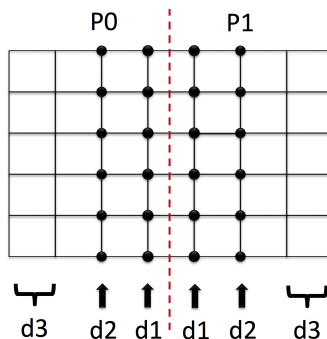


FIG. 1. A two-processors example showing  $d1$  nodes,  $d2$  nodes and  $d3$  nodes on each processor. The structured grid represents the graph of  $A_B$  (defined in Eq. (5)). On each processor, the nodes on the boundary are  $d1$  nodes;  $d2$  nodes are the (local) neighbors of  $d1$  nodes that are not on the boundary; and the remaining interior nodes are  $d3$  nodes.

### 3.2. Bulk-synchronous parallel algorithm

Recall the discussion in Section 2.3 that no level-2 fill-in exists in the algorithm. The implication is that we can process (apply low-rank elimination operators on) two clusters/nodes  $\pi_i$  and  $\pi_j$  in parallel, as long as the distance between them is greater than two. The distance is defined as the length of the shortest path between vertex  $i$  and vertex  $j$  in the graph of  $A_B$ .

This fill-in property motivates us to design the parallel algorithm as follows. First, since the distance between a local  $d3$  node and any remote node is at least three, processing  $d3$  nodes is embarrassingly parallel.

Second, since the distance between a local  $d2$  node and a remote  $d2$  node is at least three,  $d2$  nodes can be processed in parallel. However, communication is required because the shortest distance between a local  $d2$  node and a remote  $d1$  node is two, and the amount of communication is proportional to the number of  $d2$  nodes.

Last, since  $d1$  nodes are coupled between neighbor processors, we color them in a way such that  $d1$  nodes with the same color can be processed in parallel; the coloring scheme is discussed in Section 3.4. We introduce two notations:  $\mathcal{N}_1(P)$  for the neighbors of processor  $P$  and  $\mathcal{N}_2(P) = \mathcal{N}_1(\mathcal{N}_1(P)) \setminus \{P\}$ , i.e., neighbors of the neighbors of processor  $P$ . Since the distance between a  $d1$  node on processor  $P$  and a remote  $d1$  node owned by a processor in  $\mathcal{N}_2(P)$  can be two, they may need to communicate. The total amount of communication is proportional to the number of  $d1$  nodes.

To be more specific, the communication in the parallel algorithm corresponds to the “right-looking style” communication in a conventional sparse direct solver, and two types of communication are needed:

1. Suppose we eliminate the fine DOFs of a d1 node on processor  $P$ . The resulting Schur complement may contain (1) edges that connect a local node on  $P$  and a remote node on  $P$ 's neighbor, and (2) edges that connect two remote d1 nodes on two different processors in  $\mathcal{N}_2(P)$ . Therefore, every processor communicates only with its neighbors.
2. Consider the application of the sparsification operator (Eq. (2)). A sparsification of a (local) d2 node may lead to communication with a neighbor processor, and a sparsification of a (local) d1 node may involve communication with a processor in  $\mathcal{N}_2(P)$ .

The computation-to-communication ratio is closely related to the geometry of the subdomain on every processor. The communication volume is proportional to the number of d1 nodes (note in particular that local d2 nodes only exchange data with remote d1 nodes, never with remote d2 nodes), while the amount of computation is proportional to the total number of d1 nodes, d2 nodes and d3 nodes. Therefore, the computation-to-communication ratio is approximately the volume-to-surface-area ratio of the subdomain owned by every processor.

The above discussion is summarized in Algorithm 3<sup>1</sup> (some details are skipped, and the focus is on the high level idea of processing d1 nodes, d2 nodes and d3 nodes separately).

### 3.3. Task-based asynchronous parallel algorithm

Algorithm 3 uses a classical bulk-synchronous style of programming. In this subsection, we explore an asynchronous “task-based” approach. Consider the critical path of execution in Algorithm 3: d1 nodes must be processed first, and d2 nodes should be processed next. Therefore, the asynchronous algorithm always attempts to process d1 nodes; if none is ready, d2 nodes are processed; and finally, if all else fail, d3 nodes are processed. Processing d3 nodes does not require communication, which is the fall-back if no d1 node or d2 node can be processed. The pattern of communication is organized as follows.

---

<sup>1</sup>The pseudo-code is in *single program multiple data* (SPMD) style.

---

**Algorithm 3** Bulk-synchronous parallel algorithm

---

1: **procedure** HSOLVER\_BULK\_SYNCHRONOUS(local vertex clusters  $\Pi$ , local subgraph  $G$ , local submatrix  $A$ )  
     $\triangleright$  synchronous communication is used in this algorithm  
     $\triangleright$  this is processor  $P$

2: Partition  $\Pi$  into d1 nodes  $\Pi^{(1)}$ , d2 nodes  $\Pi^{(2)}$ , and d3 nodes  $\Pi^{(3)}$

3: Parallel (distance-2) coloring of d1 nodes  $\triangleright$  discussed in Section 3.4

4: **for** color  $i = 1$  **to** NumColors **do**  $\triangleright$  start processing d1 nodes

5:     **for**  $\pi_j \in \Pi^{(1)}$  with color  $i$  **do**

6:          $A \leftarrow \text{LRE}(A, \Pi, \pi_j)$

7:     **end for**

8:     Communicate with processors in  $\mathcal{N}_2(P)$   $\triangleright$  neighbor of neighbor processors

9: **end for**  $\triangleright$  finish processing d1 nodes

10: **for**  $\pi_i \in \Pi^{(2)}$  **do**  $\triangleright$  start processing d2 nodes

11:      $A \leftarrow \text{LRE}(A, \Pi, \pi_i)$

12: **end for**

13: Communicate with processors in  $\mathcal{N}_1(P)$   $\triangleright$  neighbor processors  
 $\triangleright$  finish processing d2 nodes

14: **for**  $\pi_i \in \Pi^{(3)}$  **do**  $\triangleright$  start processing d3 nodes

15:      $A \leftarrow \text{LRE}(A, \Pi, \pi_i)$

16: **end for**  $\triangleright$  finish processing d3 nodes  
 $\triangleright$  no communication needed for d3 nodes

17: Extract  $\Pi^C$ ,  $G^C$ ,  $A^C$  associated with the coarse level

18: HSOLVER\_BULK\_SYNCHRONOUS( $\Pi^C$ ,  $G^C$ ,  $A^C$ )

19: **end procedure**

---

Attempt to process a d1 node  $\pi^{(1)}$  with color  $i$  on processor  $P$  as follows.

1. Check that the communication for  $\pi^{(1)}$  is complete, which is associated with remote d1 nodes owned by  $\mathcal{N}_2(P)$  that have colors smaller than  $i$ .
2. Apply the sparsification operator (Eq. (2)) on  $\pi^{(1)}$ .
3. **Send** the updated edges to remote d1 nodes owned by  $\mathcal{N}_2(P)$ , and to remote d2 nodes owned by  $\mathcal{N}_1(P)$ .
4. Eliminate the fine DOFs in  $\pi^{(1)}$ .
5. **Send** the updated edges (in the Schur complement) to remote d1 nodes owned by  $\mathcal{N}_1(P)$ .

Otherwise, attempt to process a d2 node  $\pi^{(2)}$  as follows.

1. Check that the communication for  $\pi^{(2)}$  is complete, which is associated with remote d1 nodes owned by  $\mathcal{N}_1(P)$ .
2. Apply the sparsification operator (Eq. (2)) on  $\pi^{(2)}$ .
3. **Send** the updated edges to remote d1 nodes owned by  $\mathcal{N}_1(P)$  (this data is only used at the coarse level).
4. Eliminate the fine DOFs in  $\pi^{(2)}$ .

Otherwise, process a d3 node.

The above discussion is summarized in Algorithm 4<sup>2</sup>, which is based on a critical path analysis of the execution. This algorithm executes the phase with highest priority first (located along the critical path), followed by a phase with medium priority, and low priority. Communication is initiated in an asynchronous manner as soon as the data is available to achieve the maximum concurrency. The large number of executable tasks allows hiding some of the communication latency, thereby minimizing the idle time due to communication.

### 3.4. Coloring of d1 nodes

The d1 nodes are coupled between neighbor processors. To maximize concurrency, we use a graph coloring scheme to assign colors to all d1 nodes such that nodes with the same color can be processed in parallel. For that purpose, a pair of d1 nodes  $\pi_i$  and  $\pi_j$  must have different colors, if the following two conditions hold.

---

<sup>2</sup>The pseudo-code is in SPMD style.

---

**Algorithm 4** Asynchronous parallel algorithm

---

1: **procedure** HSOLVER\_ASYNCHRONOUS(local vertex clusters  $\Pi$ , local subgraph  $G$ , local submatrix  $A$ )

▷ asynchronous communication is used in this algorithm

▷ this is processor  $P$

2: Partition  $\Pi$  into d1 nodes  $\Pi^{(1)}$ , d2 nodes  $\Pi^{(2)}$ , and d3 nodes  $\Pi^{(3)}$

3: Parallel (distance-2) coloring of d1 nodes ▷ discussed in Section 3.4

4: color  $i \leftarrow 1$  ▷ start with the smallest color

5: **while**  $\Pi^{(1)} \cup \Pi^{(2)} \cup \Pi^{(3)}$  is not empty **do**

▷ attempt d1 nodes first

6:   **if**  $\pi_j \in \Pi^{(1)}$  has color  $i$  and communication for  $\pi_j$  is complete **then**

7:      $A \leftarrow \text{LRE}(A, \Pi, \pi_j)$

8:     Send data to  $\mathcal{N}_2(P)$  processors ▷ for remote nodes with colors  $> i$

9:     Pop  $\pi_j$  from  $\Pi^{(1)}$

▷ attempt d2 nodes secondly

10:   **else if** communication for  $\pi_j \in \Pi^{(2)}$  is complete **then**

11:      $A \leftarrow \text{LRE}(A, \Pi, \pi_j)$

12:     Send data to  $\mathcal{N}_1(P)$  processors ▷ for remote d1 nodes

13:     Pop  $\pi_j$  from  $\Pi^{(2)}$

▷ attempt d3 node last

14:   **else if**  $\pi_j \in \Pi^{(3)}$  **then**

15:      $A \leftarrow \text{LRE}(A, \Pi, \pi_j)$

16:     Pop  $\pi_j$  from  $\Pi^{(3)}$

17:   **end if**

18:   If all d1 nodes with color  $i$  have been processed,  $i \leftarrow i + 1$

▷ processing d1 nodes with color  $i$  is on the critical path of execution

19: **end while**

20: Wait for remaining communication

21: Extract  $\Pi^C, G^C, A^C$  associated with the coarse level

22: HSOLVER\_ASYNCHRONOUS( $\Pi^C, G^C, A^C$ )

23: **end procedure**

---



- $\pi_i$  and  $\pi_j$  are owned by different processors.
- the distance between  $\pi_i$  and  $\pi_j$  is less than or equal to two.

For best performance such a coloring should both minimize the number of colors and be load-balanced (the number of nodes of a given color should be roughly constant across all processors). However, even with only the first objective, the coloring problem is NP-hard. Fortunately, fast linear-time greedy heuristics work well in practice for graph coloring. One option is to bring the boundary graph to a single processor and compute the coloring sequentially, but this is not scalable in terms of memory. We therefore use the parallel distance-2 coloring [36] routine in Zoltan [35]. One example of the coloring result is shown in Fig. 2.

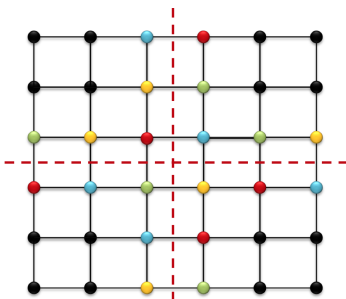


FIG. 2. A four-processors example of the coloring, where the grid is split onto four processors. A pair of nodes that are at a distance less than two from each other must have different colors, unless they are owned by the same processor. With this coloring result, all four processors are able to process a subset of the  $d_1$  nodes concurrently. As seen in this example, however, load imbalance is hard to reduce. Each processor has one node for three given colors and two nodes for the last fourth color, leading to a factor of 2/1 load imbalance during the loop over colors.

#### 4. Complexity analysis

In this section, we first review the computational cost and the memory consumption of LoRaSp, and clarify the assumptions for its linear complexity. Based on these results, we then analyze the parallel algorithm in terms of the computation complexity, communication complexity and memory consumption. The computational cost and the memory consumption of LoRaSp have been analyzed in [19]. Below, we rephrase Theorem 5.4 in [19], which summarizes the results.

*Linear complexity conditions.* The computational cost of the factorization and the memory consumption in LoRaSp scales as  $O(Nr^2)$  and  $O(Nr)$ , respectively, with respect to the problem size  $N$  if the following two conditions hold. First, for every cluster of DOFs, the number of neighbors is bounded by a constant. Second, the largest cluster size at the first level (level 0),  $r$ , is bounded by a constant; and  $r_i$ , the largest cluster size at level  $i$ , satisfies the relationship that  $r_i < \alpha^i r$  ( $i > 0$ ), where  $0 < \alpha < 2^{1/3}$ .

For the rest of this section, we assume the linear complexity conditions always hold. We also assume the linear system is evenly distributed among all processors in the parallel algorithm. Since the computation and the storage are evenly split across all processors, the computational cost and memory consumption are  $O(Nr^2/p)$  and  $O(Nr/p)$  on every processor, where  $p$  is the number of processors.

In the parallel algorithm, communication happens at two places. First, all processors exchange boundary data until the problem size becomes small enough to be factorized efficiently with the sequential Cholesky factorization. Second, the small matrix is gathered to a single processor with a reduction tree and the sequential Cholesky factorization is performed.

For every processor, the amount of communication is dominated by that from the first phase, which is proportional to the number of boundary clusters. Since the matrix associated with a cluster has at most  $r^2$  entries, the total amount of communication is

$$O\left(r^2 \cdot \left(\frac{N}{rp}\right)^{2/3}\right) = O\left(\left(\frac{Nr^2}{p}\right)^{2/3}\right)$$

where  $\frac{N}{rp}$  is the number of clusters on every processor, and  $\left(\frac{N}{rp}\right)^{2/3}$  is the surface area (assuming an underlying 3D subdomain).

The number of messages sent by every processor in the first communication phase is roughly the number of colors at all levels of grids. If the number of colors at every level is bounded by a constant, then the number of messages is proportional to the number of levels, i.e.,  $O(\log(N/(rp)))$ . In the second communication phase, the number of messages needed is  $O(\log p)$  for the conventional Cholesky factorization. Therefore, the total number of messages sent by every processor is

$$O\left(\log\left(\frac{N}{rp}\right)\right) + O(\log p).$$

To measure parallel scalability, we use standard definitions of parallel speedup, strong scaling efficiency and weak scaling efficiency as follows. Let  $T(N, p)$  be the wall-clock time to factor a matrix of size  $N$  on  $p$  processors. The parallel speedup is defined as

$$S(N, p) = \frac{T(N, p_0)}{T(N, p)} \quad (6)$$

where  $p_0$  is the smallest number of processors on which the baseline is obtained. The efficiency of strong scaling (the total problem size is fixed as the number of processors increases) is

$$E_s = \frac{S(N, p) \times p_0}{p}. \quad (7)$$

The efficiency of weak scaling (the total problem size increases proportionally to the number of processors, or the problem size per processor is fixed) is

$$E_w = \frac{T(N p_0, p_0)}{T(N p, p)}. \quad (8)$$

Following the above definitions, we can derive corresponding formulas for our parallel algorithm. As an example, the weak scaling efficiency of our parallel algorithm is

$$E_w = \frac{N}{N + \log N + \log p}$$

where  $N$  is the problem size per processor,  $p_0 = 1$  is assumed, and the cluster size  $r$ , a constant, is not shown. Therefore, a constant efficiency  $E_w$  with a fixed problem size per processor should not be expected, even if the parallel algorithm is implemented perfectly with perfect underlying hardware (no network saturation, fixed diameter, etc.). In our case, maintaining a fixed  $E_w$  implies that the total problem size should increase at least as fast as  $O(p \log(p))^3$ .

---

<sup>3</sup> $O(p \log(p))$  is the iso-efficiency function of our parallel algorithm [37]. This efficiency result is similar to that of computing parallel reduction.

## 5. Numerical results

This section presents various results to show the versatility and parallel scalability of the asynchronous parallel algorithm in Algorithm 4. Our implementation is based on the algorithm in Section 2 and the LoRaSp algorithm. To run on distributed-memory machines, our code is written with the *message passing interface* (MPI). Sequential results are obtained by running our code with one MPI rank. Unless otherwise stated, all tests were run at the NERSC Edison supercomputer<sup>4</sup>. Each node of Edison has two 12-core Intel “Ivy Bridge” processors and nodes are connected by a Cray Aries high-speed interconnect with Dragonfly topology.

### 5.1. Linear scalability results

In this subsection, we present results for solving sequences of problems to demonstrate the scalability of the hierarchical solver. In practice, the linear complexity conditions in section 4 may not hold, but we shall see that the total cost of factorization and solve scales almost linearly as the problem size increases.

The results presented in this subsection are the number of iterations and the total time (factorization + solve) for solving three PDEs (discretized with the seven-point stencil on the unit cube). The first is Poisson’s equation:

$$-\Delta u(x) = f(x).$$

The second is variable-coefficient Poisson’s equation:

$$-\nabla \cdot (a(x)\nabla u(x)) = f(x),$$

where  $a(x)$  is a quantized high-contrast random field generated in the following way: (1) initialize  $a(x)$  randomly with a uniform distribution; (2) convolve the initial  $a(x)$  with Gaussian distribution of deviation  $4h$ , where  $h$  is the stencil spacing; and (3) set  $a(x)$  to  $10^2$  if it is larger than 0.5, otherwise  $10^{-2}$ . We chose a quantized high-contrast random field because these are problems known to be difficult to solve using iterative methods. Although the performance worsens with our hierarchical solver, the algorithm still remains very efficient. For Poisson’s equation and variable-coefficient Poisson’s

---

<sup>4</sup><http://www.nersc.gov/users/computational-systems/edison/configuration/>

equation, we use our solver as a preconditioner for CG with a tolerance of  $10^{-12}$ .

The third is the Helmholtz equation:

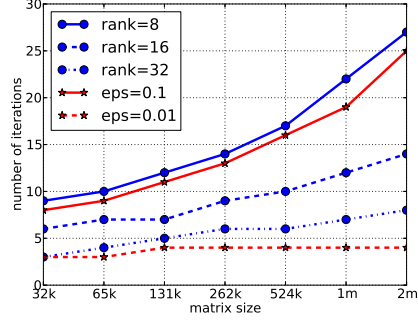
$$(-\Delta - k^2)u(x) = f(x),$$

where  $k$  is the wave number. We fix the number of DOFs per wave length and increased  $k$  proportionally to the number of DOFs in each dimension. As the frequency increases, hierarchical solvers will eventually break down because the ranks required to reach a good accuracy become too large for the method to be efficient. We fixed the resolution to 32 DOFs per wavelength, and the frequencies increased from  $f = 1\text{Hz}$  ( $32^3$  grid) to  $f = 4\text{Hz}$  ( $128^3$  grid). For the Helmholtz equation, we use our solver as a preconditioner for GMRES with a lower tolerance of  $10^{-3}$  because the Helmholtz equation is often solved with relatively low accuracy in applications such as seismic imaging.

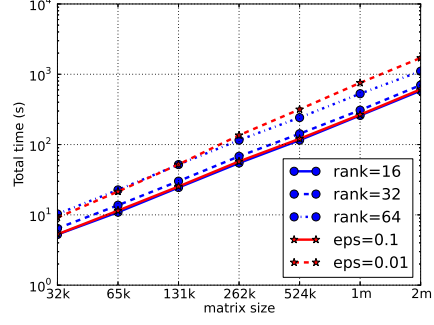
We investigated two options to select the rank: (a) dynamic rank based on the singular values of the fill-in blocks, which is computed for a given user-prescribed error tolerance  $\epsilon$  (error in the 2-norm), and (b) a fixed value of rank  $\mathcal{K}$ . As Fig. 3 shows, the former typically gives better quality preconditioners but may be more expensive, while the latter puts a strict upper limit on the memory usage and the factorization cost. In Fig. 3, the total time, which is the sum of the factorization time and the solve time, scales almost linearly with respect to the problem size for all three PDEs.

### 5.2. Parallel scalability and scalability bottleneck

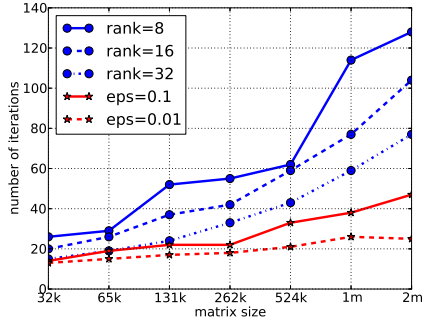
In this subsection, we show the parallel scalability of the hierarchical solver. We present both the sequential factorization time and the corresponding parallel results on up to 256 processors (16 processors per node). The definition weak scaling efficiency is given in Section 4. Fig. 4 shows the parallel factorization time and the weak scaling efficiency for solving the three PDEs. Note that we stopped the scaling experiments when the number of DOFs per processor is fewer than 8 thousand. At that scale, the communication costs dominate. The hierarchical solver achieves an average speedup of 45 in factorizing three two-million sized test problems using 256 processors.



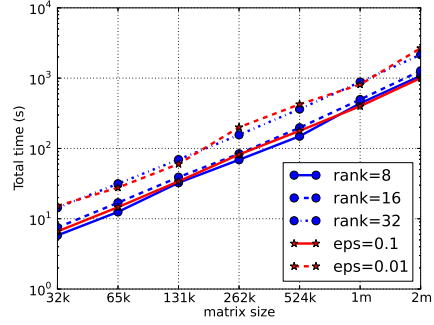
(a) Iteration number (Poisson)



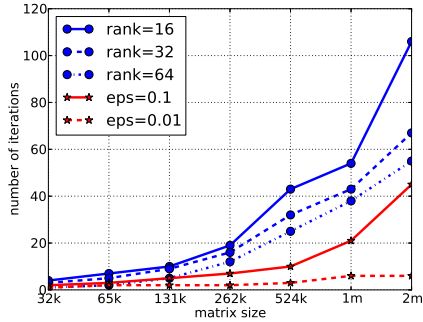
(b) Total time (Poisson)



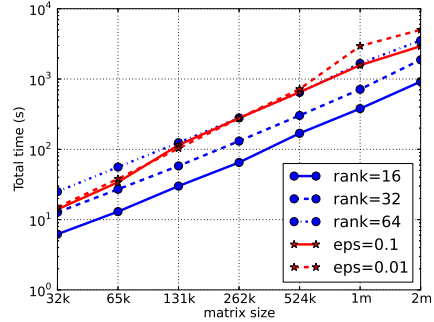
(c) Iteration number (VC-Poisson)



(d) Total time (VC-Poisson)

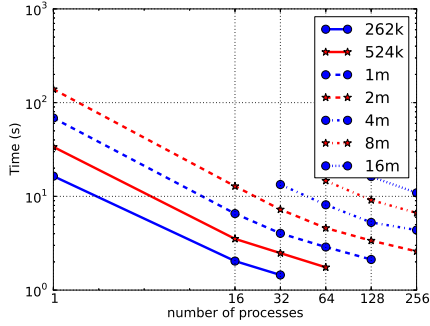


(e) Iteration number (Helmholtz)

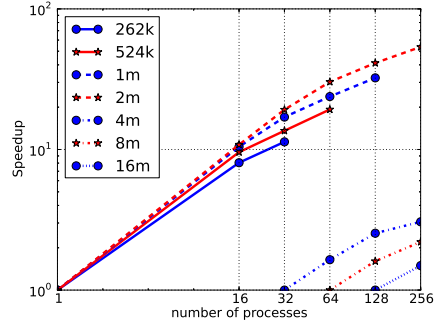


(f) Total time (Helmholtz)

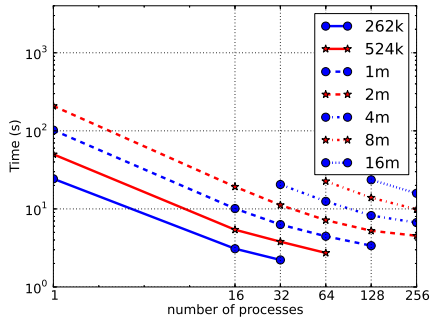
FIG. 3. Number of iterations and the total time (factorization + solve) for solving three PDEs. Different low-rank truncation criteria have been used: fixing the rank  $K$  and fixing the truncation error  $\epsilon$ . The convergence tolerance is  $10^{-12}$  for the first two problems and is  $10^{-3}$  for the third problem.



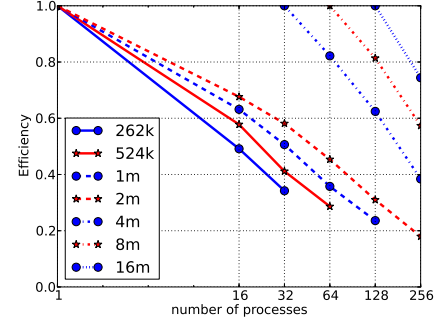
(a) Factorization time (Poisson)



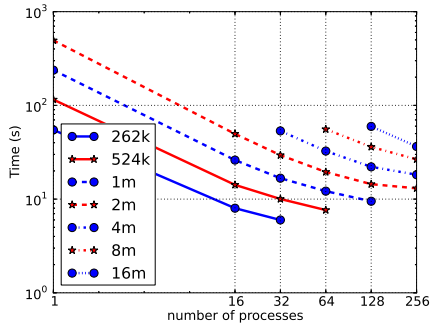
(b) Parallel speedup (Poisson)



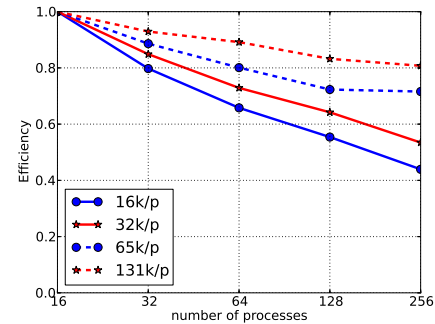
(c) Factorization time (VC-Poisson)



(d) Strong scaling (VC-Poisson)



(e) Factorization time (Helmholtz)



(f) Weak scaling (Helmholtz)

FIG. 4. *Parallel scalability results: factorization time and parallel speedups for solving Poisson's equation ( $\mathcal{K} = 8$ ), factorization time and strong scaling efficiency for solving the variable-coefficient Poisson's equation ( $\mathcal{K} = 16$ ), and factorization time and weak scaling efficiency for solving the Helmholtz equation ( $\mathcal{K} = 32$ ).*

To understand the performance bottleneck, we show breakup of the factorization time in Fig. 5. Since the parallel scalability results are almost the same for solving the three PDEs, we present only the results for solving the Helmholtz equation ( $\mathcal{K} = 32$ ). The strong scaling test corresponds to a fixed problem size of four-million DOFs on all processors, and the weak scaling test corresponds to a fixed problem size of 131-thousand DOFs per processor. As Fig. 5 shows, in the strong scaling test, the computation time for d1 nodes, d2 nodes and d3 nodes decreases by half when the number of processors doubles, whereas d1 communication time and ‘Other’ time remains almost constant. In the weak scaling test, the computation time for d1 nodes, d2 nodes and d3 nodes remains almost constant when the number of processors increases, whereas d1 communication time and ‘Other’ time increases.

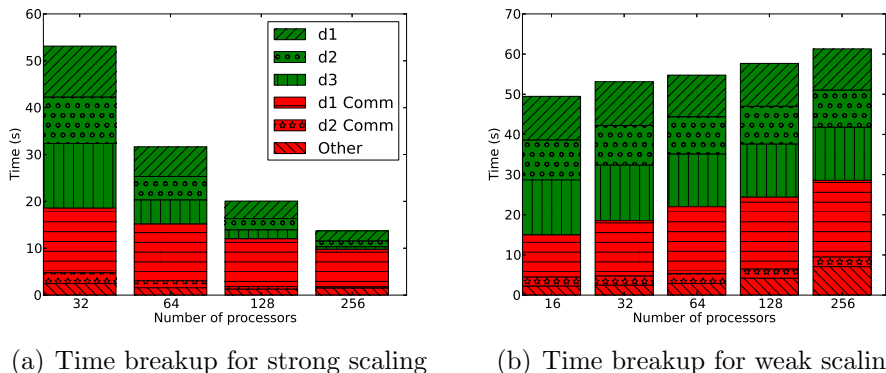


FIG. 5. Breakup of parallel factorization time for solving the Helmholtz equation ( $\mathcal{K} = 32$ ). “d1”, “d2” and “d3” stand for corresponding computational cost. “d1 Comm” and “d2 Comm” stand for corresponding communication cost. “Other” is mainly the cost of computing d1 coloring.

Computing the coloring of d1 nodes affects the parallel scalability in two aspects. First, the time spent in computing the coloring does not scale, as shown in Fig. 5. Second, more importantly, the quality of the coloring result is poor, which results in more d1 communication than necessary. Recall the hierarchical solver needs a coloring such that every pair of d1 nodes have different colors if they satisfy two conditions: (1) they are within distance two from each other and (2) they belong to two different processors. Unfortunately, we are not aware of any existing algorithm that solves this coloring problem. In our implementation, the standard distance-2 coloring algorithm



is used, which ignores the second condition and solves a much more difficult problem. As a result, the output is an unnecessarily stronger coloring result with much more colors than needed. Although the current implementation for computing the coloring is not efficient, the cost can be amortized for solving a sequence of linear systems with the same sparsity pattern, since the coloring result depends only on the graph (symbolic pattern) of the original matrix.

### 5.3. Comparison with sparse direct solver

In this subsection, we present results of comparing our parallel solver to SuperLU-Dist [10, 38], a state-of-the-art parallel sparse direct solver, on 16 processors. Timing results and the corresponding memory footprints are shown in Fig. 6. For Poisson’s equation and the variable-coefficient Poisson’s equation, our solver was used as a preconditioner ( $\mathcal{K} = 8$  and 16, respectively) for CG with a tolerance of  $10^{-12}$ . For the Helmholtz equation, our solver was used as a preconditioner ( $\mathcal{K} = 32$ ) for GMRES with a tolerance of  $10^{-3}$ .

In Fig. 6, we can see that for a small problem size, e.g., 262k ( $128^3$ ), our solver was not competitive with SuperLU-Dist. But for large problem sizes, our solver was much faster and used less memory than SuperLU-Dist.

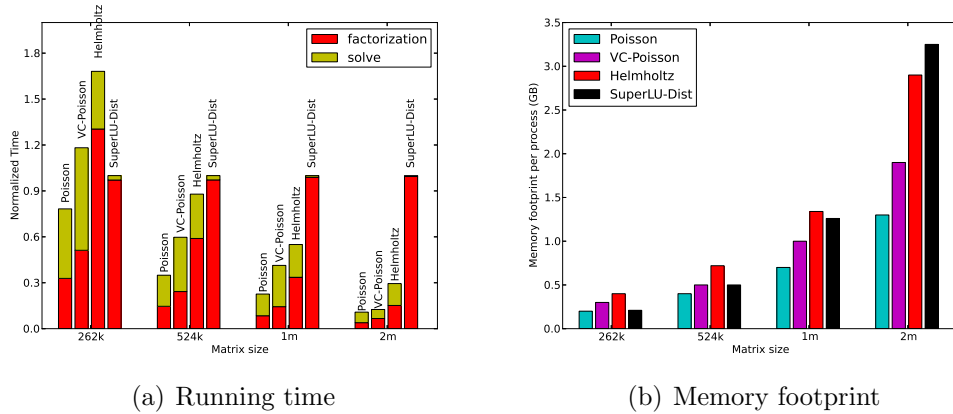


FIG. 6. Comparison with SuperLU-Dist on 16 processors. For each problem size, timing results are normalized by that of SuperLU-Dist. The timing result and the memory footprint of SuperLU-Dist was almost the same for three PDEs.

Note that the memory footprint of our solver is relatively high. The reason is that in our implementation, (sparse) nonzero blocks in the original sparse linear system are stored using a dense-matrix data structure (for ease of coding), and this choice leads to the large memory consumption. Optimizing this memory usage requires some software-engineering work and is on our to-do list.

#### 5.4. General linear systems from various applications

In this subsection, we present results<sup>5</sup> for general linear systems from a number of applications, such as electro-physiological model of a torso, numerical weather prediction and atmospheric modeling, geo-mechanical model of earth crust with underground deformation, and gas reservoir simulation for CO<sub>2</sub> sequestration. These problems are from the SuiteSparse Matrix Collection (formerly known as the University of Florida Sparse Matrix Collection)<sup>6</sup> and the matrices include SPD matrices from unstructured meshes, non-symmetric matrices and indefinite matrices. Important properties of these matrices are shown in Table 1.

TABLE. 1. *General matrices from various applications*

Matrices	size	# nonzero	symbolic pattern symmetry	numeric value symmetry	positive definite
torso3	0.26M	4.4M	no	no	no
atmosmodd	1.3M	8.8M	yes	no	no
Geo_1438	1.4M	60.2M	yes	yes	yes
Serena	1.4M	64.1M	yes	yes	yes

We compared the timing results of SuperLU-Dist and our parallel solver for five low-rank truncation epsilons ( $\epsilon = 0.8, 0.4, 0.2, 0.1$  and  $0.05$ ). Our solver is used as a preconditioner for CG/GMRES, depending on if the matrix is SPD or not, and the convergence tolerance is  $10^{-12}$ .

The comparisons are shown in Fig. 7. In Fig. 7, as the low-rank truncation error  $\epsilon$  decreases (more accurate factorizations), the factorization time

---

<sup>5</sup>tests in this subsection were run on the Vesper machine (vesper@sandia.gov) at the Sandia National Laboratories. Vesper uses AMD many-core processors and has 128GB of main memory.

<sup>6</sup><https://sparse.tamu.edu/>

increases whereas the solve time typically decreases because of the reduced CG/GMRES iteration numbers. As a result, the total solve time achieves its optimal with  $\epsilon$  being around 0.2 or 0.4 for these matrices. The speedups of our parallel solver compared with SuperLU-Dist are 1.5x, 4.8x, 2.8x and 4.3x.

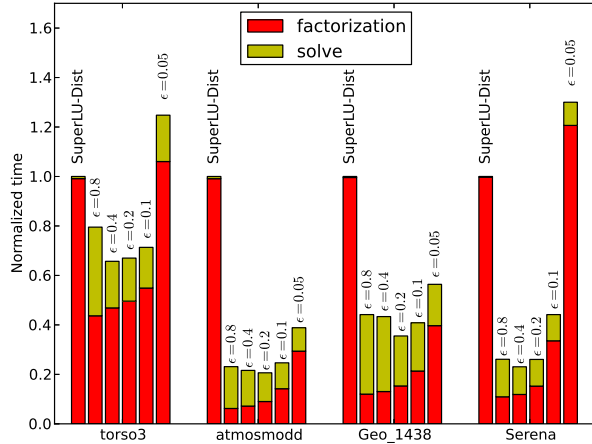


FIG. 7. Timing results for different matrices from the University of Florida Sparse Matrix Collection. For each matrix, the results are normalized with respect to the total solve time (factorization + solve) of SuperLU-Dist.

## 6. Conclusions and future work

We have presented the first parallel algorithm and implementation of the LoRaSp hierarchical solver. Although we derive the algorithm only for SPD matrices in Section 2, the algorithm extends to general matrices. As demonstrated by the numerical experiments, the parallel solver works for non-symmetric matrices, indefinite matrices and matrices from unstructured grids.

The linear complexity conditions are presented in Section 4. In practice, these conditions may not be satisfied perfectly, but the total solve time (factorization + solve) is observed to scale almost linearly for numerical experiments in Section 5.1. To solve large linear systems that cannot be stored on a single compute node, distributed memory parallel computing is necessary. We have shown that our solver achieves good speedups on up to 256 processors.

Several directions for future research are as follows.

- Our graph coloring of the boundary is conservative and other strategies may reduce the number of colors and/or execution time. For example, one could color the processor graph instead of the boundary nodes, which would give more coarse-grained communication but likely more idle time (poor load balance).
- We plan an MPI+X implementation that exploits thread parallelism for dense linear algebra. We believe such an approach may improve the parallel scalability on a large number of compute nodes.
- A variant of LoRaSp has been shown to work successfully for dense matrices [39]. Our parallel algorithm could be extended to solve dense linear systems.

## Acknowledgments

Chen, Pouransari, and Darve were supported in part by the U.S. Department of Energy's National Nuclear Security Administration under Award Number de-na0002373-1; Boman, Rajamanickam, Chen and Darve were supported in part by Sandia's Laboratory Directed Research and Development (LDRD) program. Sandia National Laboratories is a multimission laboratory managed and operated by National Technology and Engineering Solutions of Sandia, LLC., a wholly owned subsidiary of Honeywell International, Inc., for the U.S. Department of Energy's National Nuclear Security Administration under contract DE-NA-0003525. This research used resources of the National Energy Research Scientific Computing Center, a DOE Office of Science User Facility supported by the Office of Science of the U.S. Department of Energy under Contract No. DE-AC02-05CH11231.

## References

- [1] T. A. Davis, S. Rajamanickam, W. M. Sid-Lakhdar, A survey of direct methods for sparse linear systems, *Acta Numerica* 25 (2016) 383–566.
- [2] A. George, Nested dissection of a regular finite element mesh, *SIAM Journal on Numerical Analysis* 10 (1973) 345–363.
- [3] I. S. Duff, A. M. Erisman, J. K. Reid, *Direct Methods for Sparse Matrices*, Clarendon press Oxford, 1986.

- [4] Y. Chen, T. A. Davis, W. W. Hager, S. Rajamanickam, Algorithm 887: CHOLMOD, supernodal sparse cholesky factorization and update/downdate, *ACM Transactions on Mathematical Software (TOMS)* 35 (2008) 22.
- [5] T. A. Davis, Algorithm 832: UMFPACK V4. 3—an unsymmetric-pattern multifrontal method, *ACM Transactions on Mathematical Software (TOMS)* 30 (2004) 196–199.
- [6] J. W. Demmel, S. C. Eisenstat, J. R. Gilbert, X. S. Li, J. W. H. Liu, A supernodal approach to sparse partial pivoting, *SIAM J. Matrix Analysis and Applications* 20 (1999) 720–755.
- [7] A. Kuzmin, M. Luisier, O. Schenk, Fast methods for computing selected elements of the Green’s function in massively parallel nanoelectronic device simulations, in: F. Wolf, B. Mohr, D. Mey (Eds.), *Euro-Par 2013 Parallel Processing*, volume 8097 of *Lecture Notes in Computer Science*, Springer Berlin Heidelberg, 2013, pp. 533–544.
- [8] J. W. Demmel, J. R. Gilbert, X. S. Li, An asynchronous parallel supernodal algorithm for sparse Gaussian elimination, *SIAM J. Matrix Anal. Appl.* 20 (1999) 915–952.
- [9] P. R. Amestoy, I. S. Duff, J.-Y. L’Excellent, J. Koster, MUMPS: A general purpose distributed memory sparse solver, in: *Applied Parallel Computing. New Paradigms for HPC in Industry and Academia*, Springer, 2001, pp. 121–130.
- [10] X. S. Li, J. W. Demmel, SuperLU\_DIST: A scalable distributed-memory sparse direct solver for unsymmetric linear systems, *ACM Transactions on Mathematical Software (TOMS)* 29 (2003) 110–140.
- [11] J. Xu, Iterative methods by space decomposition and subspace correction, *SIAM review* 34 (1992) 581–613.
- [12] P. Lin, M. T. Bettencourt, S. P. Domino, T. C. Fisher, M. Hoemmen, J. J. Hu, E. T. Phipps, A. Prokopenko, S. Rajamanickam, C. Siefert, S. R. Kennon, Towards extreme-scale simulations with second-generation Trilinos, *Parallel Processing Letters* (2014).
- [13] P. Lin, M. Bettencourt, S. Domino, T. Fisher, M. Hoemmen, J. Hu, E. Phipps, A. Prokopenko, S. Rajamanickam, C. Siefert, et al., Towards extreme-scale simulations with next-generation Trilinos: a low Mach fluid application

- case study, in: Parallel & Distributed Processing Symposium Workshops (IPDPSW), 2014 IEEE International, IEEE, pp. 1485–1494.
- [14] L. Grasedyck, R. Kriemann, S. Le Borne, Domain decomposition based-LU preconditioning, *Numerische Mathematik* 112 (2009) 565–600.
  - [15] J. Xia, S. Chandrasekaran, M. Gu, X. S. Li, Superfast multifrontal method for large structured linear systems of equations, *SIAM Journal on Matrix Analysis and Applications* 31 (2009) 1382–1411.
  - [16] K. L. Ho, L. Ying, Hierarchical interpolative factorization for elliptic operators: Differential equations, *Communications on Pure and Applied Mathematics* (2015).
  - [17] P. Amestoy, C. Ashcraft, O. Boiteau, A. Buttari, J.-Y. L’Excellent, C. Weisbecker, Improving multifrontal methods by means of block low-rank representations, *SIAM Journal on Scientific Computing* 37 (2015) A1451–A1474.
  - [18] A. Aminfar, S. Ambikasaran, E. Darve, A fast block low-rank dense solver with applications to finite-element matrices, *Journal of Computational Physics* 304 (2016) 170–188.
  - [19] H. Pouransari, P. Coulier, E. Darve, Fast hierarchical solvers for sparse matrices using extended sparsification and low-rank approximation, *SIAM Journal on Scientific Computing* 39 (2017) A797–A830.
  - [20] R. Kriemann, Parallel-matrix arithmetics on shared memory systems, *Computing* 74 (2005) 273–297.
  - [21] P. Ghysels, X. S. Li, F.-H. Rouet, S. Williams, A. Napov, An efficient multi-core implementation of a novel HSS-structured multifrontal solver using randomized sampling, *arXiv preprint arXiv:1502.07405* (2015).
  - [22] S. Wang, X. S. Li, F.-H. Rouet, J. Xia, M. V. De Hoop, A parallel geometric multifrontal solver using hierarchically semiseparable structure, *ACM Transactions on Mathematical Software (TOMS)* 42 (2016) 21:1–21:21.
  - [23] Y. Li, L. Ying, Distributed-memory hierarchical interpolative factorization, *arXiv preprint arXiv:1607.00346* (2016).
  - [24] Y. Saad, ILUT: a dual threshold incomplete LU factorization, *Numerical linear algebra with applications* 1 (1994) 387–402.

- [25] A. Brandt, Algebraic multigrid theory: The symmetric case, *Applied mathematics and computation* 19 (1986) 23–56.
- [26] K. Stüben, A review of algebraic multigrid, *Journal of Computational and Applied Mathematics* 128 (2001) 281–309.
- [27] W. Hackbusch, A sparse matrix arithmetic based on  $\mathcal{H}$ -matrices. Part I: Introduction to  $\mathcal{H}$ -matrices, *Computing* 62 (1999) 89–108.
- [28] W. Hackbusch, B. N. Khoromskij, A sparse  $\mathcal{H}$ -matrix arithmetic., *Computing* 64 (2000) 21–47.
- [29] J. Xia, S. Chandrasekaran, M. Gu, X. S. Li, Fast algorithms for hierarchically semiseparable matrices, *Numerical Linear Algebra with Applications* 17 (2010) 953–976.
- [30] S. Chandrasekaran, M. Gu, T. Pals, A fast ULV decomposition solver for hierarchically semiseparable representations, *SIAM Journal on Matrix Analysis and Applications* 28 (2006) 603–622.
- [31] W. Hackbusch, S. Börm, Data-sparse approximation by adaptive  $\mathcal{H}^2$ -matrices, *Computing* 69 (2002) 1–35.
- [32] W. Hackbusch,  $\mathcal{H}^2$ -matrices, in: *Hierarchical Matrices: Algorithms and Analysis*, Springer, 2015, pp. 203–240.
- [33] G. Karypis, V. Kumar, A fast and high quality multilevel scheme for partitioning irregular graphs, *SIAM Journal on scientific Computing* 20 (1998) 359–392.
- [34] C. Chevalier, F. Pellegrini, Pt-scotch: A tool for efficient parallel graph ordering, *Parallel computing* 34 (2008) 318–331.
- [35] E. G. Boman, U. V. Catalyurek, C. Chevalier, K. D. Devine, The Zoltan and Isorropia parallel toolkits for combinatorial scientific computing: partitioning, ordering, and coloring, *Scientific Programming* 20 (2012) 129–150.
- [36] D. Bozdağ, U. Çatalyürek, A. H. Gebremedhin, F. Manne, E. G. Boman, F. Özgüner, Distributed-memory parallel algorithms for distance-2 coloring and related problems in derivative computation, *SIAM J. Sci. Comput.* 32 (2010) 2418–2446.

- [37] A. Y. Grama, A. Gupta, V. Kumar, Isoefficiency: Measuring the scalability of parallel algorithms and architectures, *IEEE Parallel Distrib. Technol.* 1 (1993) 12–21.
- [38] X. Li, J. Demmel, J. Gilbert, L. Grigori, M. Shao, I. Yamazaki, SuperLU Users' Guide, Technical Report LBNL-44289, Lawrence Berkeley National Laboratory, 1999. <http://crd.lbl.gov/~xiaoye/SuperLU> Last update: August 2011.
- [39] P. Coulier, H. Pouransari, E. Darve, The inverse fast multipole method: Using a fast approximate direct solver as a preconditioner for dense linear systems, *SIAM Journal on Scientific Computing* 39 (2017) A761–A796.