

The Reverse Cuthill-McKee Algorithm in Distributed-Memory

Ariful Azad, Mathias Jacquelin, Aydın Buluç, Esmond G. Ng

E-mail: {azad, mjacquelin, abuluc, egng}@lbl.gov

Computational Research Division

Lawrence Berkeley National Laboratory

Abstract—Ordering vertices of a graph is key to minimize fill-in and data structure size in sparse direct solvers, maximize locality in iterative solvers, and improve performance in graph algorithms. Except for naturally parallelizable ordering methods such as nested dissection, many important ordering methods have not been efficiently mapped to distributed-memory architectures. In this paper, we present the first-ever distributed-memory implementation of the reverse Cuthill-McKee (RCM) algorithm for reducing the profile of a sparse matrix. Our parallelization uses a two-dimensional sparse matrix decomposition. We achieve high performance by decomposing the problem into a small number of primitives and utilizing optimized implementations of these primitives. Our implementation attains up to 38x speedup on matrices from various applications on 1024 cores of a Cray XC30 supercomputer and shows strong scaling up to 4096 cores for larger matrices.

I. INTRODUCTION

Reordering a sparse matrix to reduce its bandwidth or profile can speed up many sparse matrix computations [1], [2]. For example, a matrix with a small profile is useful in direct methods for solving sparse linear systems since it allows a simple data structure to be used. It is also useful in iterative methods because the nonzero elements will be clustered close to the diagonal, thereby enhancing data locality. Given a symmetric matrix \mathbf{A} , a bandwidth-reduction ordering aims to find a permutation \mathbf{P} so that the bandwidth of \mathbf{PAP}^T is small. Since obtaining a reordering to minimize bandwidth is an NP-complete problem [3], various heuristics are used in practice such as Cuthill-McKee, Reverse Cuthill-McKee (RCM), and Sloan’s algorithms [4], [5], [6]. This paper solely focuses on the RCM algorithm [5] because, with careful algorithm design, it is amenable to massive distributed-memory parallelism – the primary topic of interest of this paper.

The need for distributed memory RCM algorithms is driven by the move to extreme-scale computing. In a large scale scientific application, the matrix has often been distributed already by the time one arrives at the numerical phases of sparse matrix computations. Hence, it would be a waste of time to gather a distributed matrix onto a single processor to compute an ordering using sequential or multithreaded algorithms. Furthermore, by clustering the nonzeros closer to the diagonal, RCM ordering not only increases cache performance

A two-page abstract of this work has appeared at the SIAM Workshop on Combinatorial Scientific Computing. Short abstracts on SIAM conferences are not archived and do not appear in published proceedings.

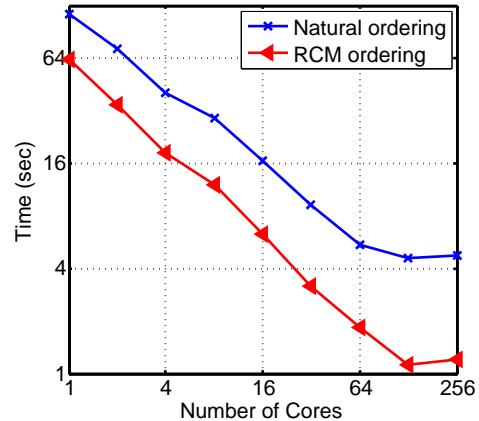


Fig. 1: The time to solve `thermal2` using the conjugate gradient method with block Jacobi preconditioner in PETSc on NERSC/Edison. `thermal2` has 1.2M rows and columns and 4.9M nonzeros. The bandwidths of the original and RCM-permuted matrix are 1,226,000 and 795, respectively.

of iterative solvers, but it can often restrict the communication to resemble more of a nearest-neighbor pattern. Figure 1 illustrates the performance effects of RCM ordering on a preconditioned conjugate gradient solver of the popular PETSc package [7]. Notice that the performance benefit of RCM ordering actually increases as the number of cores increases, possible due to reduced communication costs. Our goal in this paper is therefore to design and develop a scalable distributed-memory parallel implementation of the RCM algorithm [5].

Similar to many sparse matrix computations, RCM ordering has been shown to be a difficult problem to parallelize [8]. The RCM algorithm traverses vertices of a graph level-by-level using a specialized breadth-first search (BFS). Within each level of the traversal tree, vertices are ordered by the labels of their parents and their own degrees. This restriction on how the vertices are numbered in the graph traversal limits available parallelism in the RCM algorithm. Furthermore, RCM is often used for matrices with medium-to-high diameter while most of the existing work on parallel BFS has focused on low-diameter graphs such as synthetic graphs used by the Graph500 benchmark. A higher diameter increases the critical path of level-synchronous BFS algorithms, also limiting available parallelism. Overall, the computational load of the

RCM algorithm is highly dynamic, especially if the graph has high diameter. The problem exacerbates on high concurrency, where load imbalance and communication overhead degrade the performance of the parallel algorithm.

In this paper, we aim to overcome these challenges by using the graph-matrix duality and replacing unstructured graph operations by structured matrix/vector operations. First, we use a specialized sparse matrix-sparse vector multiplication (SpMSPV) whose addition and multiplication operations are overloaded to satisfy the restriction of RCM ordering (see Fig. 2 for an example). The linear algebraic formulation with its operator overloading feature hides the RCM-specific complexity and is able to achieve high performance despite the deterministic nature of the computation. Second, within a BFS level, we replace expensive global sorting of vertices by a number of small sorting tasks each of which is limited only to siblings. These two optimizations lead to a distributed-memory RCM algorithm that scales well on thousands of cores of modern supercomputers.

We make the following contributions:

- We present a scalable distributed-memory algorithm for RCM ordering. Our algorithm relies on a handful of bulk-synchronous parallel primitives that are optimized for both shared-memory and massively parallel distributed memory systems.
- The quality (bandwidth and envelope) of ordering from our distributed-memory implementation is comparable to the state-of-the-art and remains insensitive to the degree of concurrency.
- We provide a hybrid OpenMP-MPI implementation of the RCM ordering that attains up to $38\times$ speedup on matrices from various applications on 1024 cores of a Cray XC30 supercomputer. We provide detailed performance evaluation on up to 4096 cores, which sheds light on the performance bottlenecks and opportunities for future research.

II. PRELIMINARIES

A. Serial Algorithms

Let n be the number of columns in a symmetric matrix \mathbf{A} . Let $f_i(\mathbf{A})$ be the column subscript of the first nonzero element in column i of \mathbf{A} : $f_i(\mathbf{A}) = \min\{j \mid a_{ij} \neq 0\}$. The i -th *bandwidth* $\beta_i(\mathbf{A})$ is defined as $\beta_i(\mathbf{A}) = i - f_i(\mathbf{A})$. The overall *bandwidth* of matrix \mathbf{A} is denoted $\beta(\mathbf{A})$ and is defined as $\beta(\mathbf{A}) = \max\{\beta_i(\mathbf{A}) \mid 1 \leq i \leq n\}$. Using these notations, the envelope $Env(\mathbf{A})$ is defined as:

$$Env(\mathbf{A}) = \{\{i, j\} \mid 0 < j - i \leq \beta_i(\mathbf{A}), 1 \leq i, j \leq n\}.$$

The quantity $|Env(\mathbf{A})|$ is called the profile or envelope size of \mathbf{A} .

Reverse Cuthill-McKee ordering, or RCM, introduced by George [9], is a variant of the Cuthill-McKee ordering [4], which aims at reducing the *bandwidth* of a sparse symmetric matrix \mathbf{A} . This is of particular importance when the matrix is to be stored using a profile based format. Finding a reordering

of the rows/columns of \mathbf{A} corresponds to the process of labeling vertices of the graph $\mathcal{G}(\mathbf{A})$ associated with \mathbf{A} .

RCM repeatedly labels vertices adjacent to the current vertex v_i until all have been labeled, as depicted in Algorithm 1. The algorithm essentially processes vertices levels by levels and labels them in reverse order. The resulting reordered matrix often has a smaller profile [9].

Algorithm 1: Reverse Cuthill-McKee algorithm

```

1  $v_1 \leftarrow r$ 
2 for  $i = 1$  to  $n$  do
3   Find all the unnumbered neighbors of the vertex  $v_i$ 
4   Label the vertices found in increasing order of degree
5 The Reverse Cuthill-McKee ordering is given by
    $w_1, w_2, \dots, w_n$  where  $w_i = v_{n-i+1}$  for  $i = 1, \dots, n$ .
```

For the graph $\mathcal{G}(\mathbf{A}) = (V, E)$ associated with \mathbf{A} , V is the set of vertices and E is the set of edges in $\mathcal{G}(\mathbf{A})$. The number of vertices in $\mathcal{G}(\mathbf{A})$ is denoted $n = |V|$ and m is the number of edges $m = |E|$. The *eccentricity* [10] of a vertex v is defined as $\ell(v) = \max\{d(v, w) \mid w \in V\}$, where $d(v, w)$ is the graph distance between v and w .

The *rooted level structure* [11] of a vertex $v \in V$ is the *partitioning* $\mathcal{L}(v)$ of V satisfying

$$\mathcal{L}(v) = \{L_0(v), L_1(v), \dots, L_{\ell(v)}(v)\},$$

where

$$L_0(v) = \{v\}, \quad L_1(v) = \text{Adj}(L_0(v)),$$

$$L_i(v) = \text{Adj}(L_{i-1}(v)) \setminus L_{i-2}(v), \quad i = 2, 3, \dots, \ell(v).$$

The *length* of $\mathcal{L}(v)$ corresponds to the eccentricity $\ell(v)$ while the *width* $\nu(v)$ of $\mathcal{L}(v)$ is defined by

$$\nu(v) = \max\{|L_i(v)| \mid 0 \leq i \leq \ell(v)\}.$$

Algorithm 2: An algorithm to find a pseudo-peripheral vertex r

```

1  $r \leftarrow \{\text{arbitrary vertex in } V\}$ 
2  $\mathcal{L}(r) \leftarrow \{L_0(r), L_1(r), \dots, L_{\ell(r)}(r)\}$ 
3  $nlvl \leftarrow \ell(r) - 1$ 
4 while  $\ell(r) > nlvl$  do
5    $nlvl \leftarrow \ell(r)$ 
6   (Shrink last level): Choose a vertex  $v$  in  $L_{\ell(r)}(r)$  of
   minimum degree.
7    $\mathcal{L}(v) \leftarrow \{L_0(v), L_1(v), \dots, L_{\ell(v)}(v)\}$ 
8    $r \leftarrow v$ 
```

The first vertex being labeled strongly impacts the bandwidth of the permuted matrix. Experience shows that it is better to start with a node having a large eccentricity. A *peripheral* vertex is a vertex with maximum eccentricity. Finding such vertex is prohibitively expensive, and a common heuristic is to use a *pseudo-peripheral vertex* instead [13], [5]. A *pseudo-peripheral vertex* is a vertex displaying a high

TABLE I: Matrix-algebraic primitives needed for the RCM algorithm.

Function	Arguments	Returns	Serial Complexity	Communication
IND	x : a sparse vector	local indices of nonzero entries of x	$O(nnz(x))$	None
SELECT	x : a sparse vector y : a dense vector $expr$: logical expr. on y assume $size(x) = size(y)$	$z \leftarrow$ an empty sparse vector for $i \in \text{IND}(x)$ if ($expr(y[i])$) then $z[i] \leftarrow x[i]$ return z	$O(nnz(x))$	None
SET	x : a sparse vector y : a dense vector	for $i \in \text{IND}(x)$ $y[i] \leftarrow x[i]$ return y	$O(nnz(x))$	None
SPMSPV	\mathbf{A} : a sparse matrix x : a sparse vector SR: a semiring	return $\mathbf{A} \cdot x$	$\sum_{k \in \text{IND}(x)} nnz(\mathbf{A}(:, k))$	AllGather & AlltoAll on subcommunicator [12]
REDUCE	x : a sparse vector y : a dense vector	$mv =$ maximum value in y for $i \in \text{IND}(x)$ $mv \leftarrow \min\{mv, y[i]\}$	$O(nnz(x))$	AllReduce
SORTPERM	x : a sparse vector y : a dense vector	$T \leftarrow$ an empty array of tuples for $i \in \text{IND}(x)$ $T[i] \leftarrow (x[i], y[i], i)$ sort T in the ascending order and return the permutation	$O(nnz(x) \log nnz(x))$	AllToAll

eccentricity, as close to the graph diameter as possible. Gibbs et al. [14] introduced an algorithm to find such a vertex, which is later refined by George and Liu [13]. The process, given in Algorithm 2, starts with an arbitrary vertex in V and computes its rooted level structure. Then, a vertex in the last level is picked and the corresponding level structure is computed. This process is repeated until the number of levels in the rooted level structure converges. Computing the level structure corresponds to a BFS of the graph $\mathcal{G}(\mathbf{A})$.

B. Previous Work on Parallel RCM

There is a strong connection between BFS, finding a pseudo-peripheral vertex and computing the RCM ordering. However, RCM is often used for matrices with higher diameter than graphs for which parallel BFS is often optimized [12], [15]. In addition, computing the actual RCM ordering is even harder to parallelize because the vertices within each level of the traversal tree have to be ordered by degree.

Computing sparse matrix orderings in parallel have received intermittent attention over the last few decades. While RCM is often used to accelerate iterative solvers, one of its first reported supercomputer-scale implementations was in the context of direct solvers by Ashcraft et al. [16] on a CRAY X-MP. It is hard to compare results from 30 years ago, where both the architectures and the sizes of matrices were significantly different. Karantasis et al. [8] recently studied the shared-memory parallelization of various reordering algorithms including RCM.

III. ALGORITHMS BASED ON MATRIX ALGEBRA

In this section, we present a distributed memory algorithm for computing the RCM ordering that takes advantage of the

equivalence between graph algorithms and sparse matrix algebra to exploit latest hardware platforms. The $nnz()$ function computes the number of nonzeros in its input, e.g., $nnz(\mathbf{x})$ returns the number of nonzeros in \mathbf{x} . We utilize the Matlab colon notation: $\mathbf{A}(:, i)$ denotes the i th column, and $\mathbf{A}(i, :)$ denotes the i th row.

A. Primitives for the RCM algorithm

With an aim to design a scalable parallel algorithm, we redesign the sequential algorithms for RCM ordering (Algorithm 1) and finding a pseudo-peripheral vertex (Algorithm 2) in terms of matrix-algebraic operations. For this purpose, we use the primitives summarized in Table I. A sparse vector is used to represent a subset of vertices. Consider that a subset of n_1 unique vertices V_1 is represented by a sparse vector x . Then, for each vertex v_i in V_1 , there is a nonzero entry in the i -th location of x (i.e., $x[i] \neq 0$), and the number of nonzeros $nnz(x)$ in x is equal to n_1 . The nonzero entries of the sparse vector can store arbitrary values depending on the context of the algorithm. By contrast, a dense vector y stores information about all vertices in the graph, hence the length of y is always n . In sparse matrix computations, we often are interested in where the nonzero elements are in the sparse matrices. Hence, if \mathbf{A} is an $n \times n$ sparse symmetric matrix, then we use a graph $G = (V, E)$ to describe the sparsity structure of \mathbf{A} : $V = \{v_1, v_2, \dots, v_n\}$, where v_i corresponds to row/column i of \mathbf{A} , for $1 \leq i \leq n$, and $\{v_i, v_j\} \in E$ when $\mathbf{A}_{i,j} \neq 0$.

Let x be a sparse vector, y be a dense vector, $expr$ be a logical unary operation and \mathbf{A} be a sparse matrix. $\text{IND}(x)$ returns the indices of nonzero entries in x . $\text{SELECT}(x, y, expr)$ returns every nonzero entry $x[i]$ of x where $expr(y[i])$ is true. $\text{SET}(y, x)$ replaces values of y by corresponding nonzero entries of x (other entries of y remain unchanged). RE-

Algorithm 3: Reverse Cuthill-McKee algorithm in terms of matrix-algebraic operations. **Inputs:** a sparse adjacency matrix \mathbf{A} , a dense vector \mathcal{D} storing the degrees of all vertices, and a pseudo-peripheral vertex r . **Output:** the RCM ordering \mathcal{R} .

```

1  $\mathcal{R} \leftarrow -1$  // Dense vector storing the ordering of all vertices; initialized to -1
2  $\mathcal{L}_{cur} \leftarrow \{r\}$ ,  $\mathcal{L}_{next} \leftarrow \phi$  // Current and next BFS levels (frontier)
3  $\mathcal{R}[r] \leftarrow 0$  // label of  $r$  is set to 0
4  $nv \leftarrow 1$  // Number of vertices labeled so far
5 while  $\mathcal{L}_{cur} \neq \phi$  do
6    $\mathcal{L}_{cur} \leftarrow \text{SET}(\mathcal{L}_{cur}, \mathcal{R})$ 
7    $\mathcal{L}_{next} \leftarrow \text{SPMSPV}(\mathbf{A}, \mathcal{L}_{cur}, \text{SR}=(\text{select2nd}, \text{min}))$  // Visit neighbors of the frontier
8    $\mathcal{L}_{next} \leftarrow \text{SELECT}(\mathcal{L}_{next}, \mathcal{R}, \mathcal{R} = -1)$  // Keep unvisited vertices
9    $\mathcal{R}_{next} \leftarrow \text{SORTPERM}(\mathcal{L}_{next}, \mathcal{D})$  // Lexicographically sorted permutation of vertices in
   the next frontier based on the (parent_order, degree) pair
10   $\mathcal{R}_{next} \leftarrow \mathcal{R}_{next} + nv$  // Global ordering
11   $nv \leftarrow nv + \text{nnz}(\mathcal{R}_{next})$ 
12   $\mathcal{R} \leftarrow \text{SET}(\mathcal{R}, \mathcal{R}_{next})$  // Set orders of the newly visited vertices
13   $\mathcal{L}_{cur} \leftarrow \mathcal{L}_{next}$ 
14 return  $\mathcal{R}$  in the reverse order

```

Algorithm 4: Algorithm to find a pseudo-peripheral vertex r in terms of matrix-algebraic operations. **Inputs:** a sparse adjacency matrix \mathbf{A} and a dense vector \mathcal{D} storing the degrees of all vertices. **Output:** a pseudo-peripheral vertex r .

```

1  $r \leftarrow \{\text{arbitrary vertex in } V\}$ 
2  $\ell \leftarrow 0$ ,  $nlvl \leftarrow \ell - 1$  // Number of levels in current and previous BFS trees
3 while  $\ell > nlvl$  do
4    $\mathcal{L} \leftarrow -1$  // Dense vector storing the BFS level that each vertex belongs to
5    $\mathcal{L}_{cur} \leftarrow \{r\}$ ,  $\mathcal{L}_{next} \leftarrow \phi$  // Current and next BFS levels (frontier)
6    $nlvl \leftarrow \ell$ ,  $\mathcal{L}[r] \leftarrow 0$  // Level of  $r$  is set to 0
7   do
8      $\mathcal{L}_{cur} \leftarrow \text{SET}(\mathcal{L}_{cur}, \mathcal{L})$ 
9      $\mathcal{L}_{next} \leftarrow \text{SPMSPV}(\mathbf{A}, \mathcal{L}_{cur}, \text{SR}=(\text{select2nd}, \text{min}))$  // Visit neighbors
10     $\mathcal{L}_{next} \leftarrow \text{SELECT}(\mathcal{L}_{next}, \mathcal{L}, \mathcal{L} = -1)$  // Keep unvisited vertices
11    if  $\mathcal{L}_{next} \neq \phi$  then
12       $\mathcal{L} \leftarrow \text{SET}(\mathcal{L}, \mathcal{L}_{next})$  // Set levels of newly visited vertices
13       $\mathcal{L}_{cur} \leftarrow \mathcal{L}_{next}$ 
14     $\ell \leftarrow \ell + 1$ 
15  while  $\mathcal{L}_{next} \neq \phi$ 
16   $r \leftarrow \text{REDUCE}(\mathcal{L}_{cur}, \mathcal{D})$  // Find the vertex with the minimum degree

```

$\text{DUCE}(x, y, op)$ considers only the nonzero indices I of the sparse vector x and performs a reduction on the values of the dense vector $y[I]$ using the operation op . In Table I, we provide an example where the reduction operation is to find the minimum value. $\text{SORTPERM}(x, y)$ creates a tuple $(x[i], y[i], i)$ for each nonzero index i in x . The list of tuples are lexicographically sorted and the function returns the sorted permutation (the indices passed to the sorting routine are used to obtain the permutation). Finally, $\text{SPMSPV}(\mathbf{A}, x, SR)$ performs a sparse matrix-sparse vector multiplication over the semiring SR . For the purposes of this work, a semiring is defined over (potentially separate) sets of ‘scalars’, and has its two operations ‘multiplication’ and ‘addition’ redefined. We refer to a semiring by listing its scaling operations, such as the (*multiply*, *add*) semiring. The usual semiring multiply for BFS

is *select2nd*, which returns the second value it is passed. The BFS semiring is defined over two sets: the matrix elements are from the set of binary numbers, whereas the vector elements are from the set of integers.

B. The RCM algorithm using matrix-algebraic primitives

Algorithm 3 describes the RCM algorithm using the operations from Table I. Here, we assume that the graph is connected. The case for more than one connected components can be handled by repeatedly invoking Algorithm 3 for each connected component. Algorithm 3 takes a sparse adjacency matrix \mathbf{A} and a pseudo-peripheral vertex r as inputs, and returns the RCM ordering as a dense vector \mathcal{R} .

The i -th element $\mathcal{R}[i]$ of \mathcal{R} is initialized to -1 and it retains the initial value until the i -th vertex is visited and labeled. Let

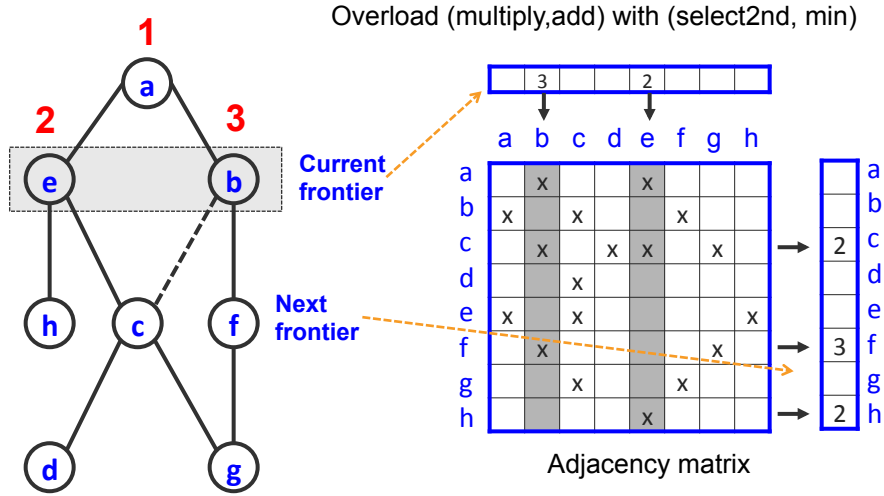


Fig. 2: An example of using specialized sparse matrix-sparse vector multiplication in exploring the next-level vertices in the RCM algorithm. On the left, we show a BFS tree rooted at vertex a . The current frontier is the set of two vertices $\{e, b\}$. Labels of the already explored vertices are shown in red numbers. The adjacency matrix representing the graph is shown on the right. The input sparse vector corresponding to the current frontier is shown at the top of the matrix. The sparse vector has two nonzero entries corresponding to e and b . The values of the nonzero entries are the labels of the vertices. The SpMSpV algorithm over a $(select2nd, min)$ semiring then selects the columns of the matrix corresponding to the nonzero indices of the vector shown in gray columns and retains the minimum product in each row where there exists at least one nonzero entry in the selected columns. The indices of the output vector represent the vertices in the next level while the values in the vector denote the labels of parent vertices. The $(select2nd, min)$ semiring ensures that a child v always attaches itself to a parent with the minimum label among all of v 's visited neighbors. Therefore, vertex e is selected as the parent of vertex c because e is the lowest numbered vertex adjacent to e in the previous level. Previously visited vertex a is not shown in the output vector.

\mathcal{L}_{cur} and \mathcal{L}_{next} be the set of vertices in the current and the next BFS level, respectively. \mathcal{L}_{cur} is called the BFS frontier (the set of current active vertices). The algorithm starts by labeling the pseudo-peripheral vertex r , inserting it into \mathcal{L}_{cur} .

The while loop (lines 5-13 of Algorithm 3) explores the vertices level-by-level until the frontier \mathcal{L}_{cur} becomes empty. Vertices in \mathcal{L}_{cur} have been already labeled in the previous iteration or during initialization. Hence each iteration of the while loop traverses the unvisited neighbors \mathcal{L}_{next} of \mathcal{L}_{cur} and labels vertices in \mathcal{L}_{next} .

At the beginning of the while loop of Algorithm 3, the values of the sparse vector \mathcal{L}_{cur} are set to the labels of vertices. Next, we discover vertices \mathcal{L}_{next} that are adjacent to the current frontier \mathcal{L}_{cur} by multiplying \mathbf{A} by \mathcal{L}_{cur} over the $(select2nd, min)$ semiring (line 7). Here the overloaded multiplication operation $select2nd$ passes the labels of parents to the children and the overloaded addition operation min ensures that each vertex in \mathcal{L}_{next} becomes a child of a vertex in \mathcal{L}_{cur} with the minimum label. This step is explained in Figure 2 with an example. Notice that the specific operator overloading by a $(select2nd, min)$ semiring makes the vertex exploration deterministic, and parallelizing the specialized SpMSpV becomes more challenging than traditional BFS. Thankfully, the linear algebraic formulation with its operator overloading feature hides the RCM-specific complexity and is able to achieve high performance despite the deterministic nature of the computation.

After the vertices in the next level are discovered via SpMSpV, previously labeled vertices are removed from \mathcal{L}_{next} (line 8). The next step is to label the vertices in \mathcal{L}_{next} . For this purpose, we sort vertices in \mathcal{L}_{next} based on the labels of their parents and their own degrees. The sorted permutation of vertices in \mathcal{L}_{next} returned by the SORTPERM function is used to set their labels. In SORTPERM, the indices of vertices are used to break ties when multiple vertices have the same parent and degree. We observe that vertices within each level are already partially sorted by their parents' labels because of the $(select2nd, min)$ semiring used in SPMSpV. Hence, a vertex with a smaller parent's label will always be numbered before another vertex with larger parent's label. This breaks the task of an expensive global sorting to a number of small sorting tasks each of which is limited only to siblings. For example, in Figure 2, vertices c and h will be labeled before f because the parent of c and h has a smaller label than the parent of f . Hence, c and h can be labeled based on their degrees, independent of other vertices.

Finally, we update the labels of the newly visited vertices (line 12) and proceed to the next iteration of the while loop. When the current frontier becomes empty, Algorithm 3 returns with the RCM ordering by reversing the labels of the vertices.

C. Finding the pseudo-peripheral vertex

Algorithm 4 finds a pseudo-peripheral vertex using matrix-algebraic operations from Table I. Similar to Algorithm 2,

Algorithm 4 starts with an arbitrary vertex r . Initially, the number of levels in the current and previous BFSs are set to 0 and -1 , respectively, so that the while loop in line 4 iterates at least twice before termination. Each iteration of the while loop in line 4 runs a full BFS starting with r . As before, \mathcal{L}_{cur} and \mathcal{L}_{next} represent the subsets of vertices in the current and next levels of the BFS. \mathcal{L} stores the BFS level to which each vertex belongs (-1 denotes an unvisited vertex). The do-while loop (lines 8-16) performs the BFS traversal similar to Algorithm 3. The overloaded addition operation of the SpMSpV at line 12 is set to *min* only for convenience. It can be replaced by any equivalent operation because the order of vertices within a level of BFS does not matter in the discovery of a pseudo-peripheral vertex. At the end of the BFS, we find a vertex with the minimum degree from the last non-empty level and make it the root for the next BFS (line 17). The while loop at line 4 terminates when the number of levels in the latest BFS is not greater than the previous BFS. At this point, Algorithm 4 returns the pseudo-peripheral vertex r .

IV. DISTRIBUTED MEMORY ALGORITHM

In order to parallelize the RCM algorithm, we simply need to parallelize the matrix-algebraic functions described in Table I. Since our RCM algorithm relies heavily on SpMSpV and SortPerm, we mostly focus on these two primitives. We note that distributed-memory algorithm for SpMSpV is not a contribution of this paper; hence we cited our prior work where the algorithm was first described. This paper provides an analysis of SpMSpV in the context of RCM. We would like to mention that our SpMSpV is specialized for RCM with a (select2nd, min) semiring and the nonzero elements of the vectors are (order, degree) pairs. A specialized SortPerm algorithm is a new contribution of the paper and described in Section IV-B.

A. Data distribution and storage

We use the Combinatorial BLAS (CombBLAS) framework [17], which distributes its sparse matrices on a 2D $p_r \times p_c$ processor grid. Processor $P(i, j)$ stores the submatrix \mathbf{A}_{ij} of dimensions $(m/p_r) \times (n/p_c)$ in its local memory. Vectors are also distributed on the same 2D processor grid. CombBLAS supports different formats to store its local submatrices. In this work, we use the CSC format as we found it to be the fastest for the SpMSpV operation with very sparse vectors, which is often the case for matrices where RCM is commonly used. CombBLAS uses a vector of {index, value} pairs for storing sparse vectors. To balance load across processors, we applied a random symmetric permutation \mathbf{P} to the input matrix \mathbf{A} before running the RCM algorithm. Most methods in numerical analysis can operate on a symmetrically permuted matrix instead of the original matrix. Therefore, we argue that the whole distributed-memory solver can work on the permuted matrix \mathbf{PAP}^T , and we do not include the permutation time in the reported results.

B. Analysis of the distributed algorithm

We measure communication by the number of *words* moved (W) and the number of *messages* sent (S). The cost of communicating a length m message is $\alpha + \beta m$ where α is the latency and β is the inverse bandwidth, both defined relative to the cost of a single arithmetic operation. Hence, an algorithm that performs F arithmetic operations, sends S messages, and moves W words takes $T = F + \alpha S + \beta W$ time.

Since the amount of work is variable across BFS iterations, we analyze the aggregate cost over the whole BFS. For ease of analysis, matrix and vector nonzeros are assumed to be i.i.d. distributed. We also assume a square processor grid $p_c = p_r = \sqrt{p}$. Number of complete breadth-first searches is denoted by |iters|. The value of |iters| is exactly one in Algorithm 3, but it can be more than one in Algorithm 4.

We leverage the 2D SpMSpV algorithms implemented in CombBLAS. The complexity of the parallel SpMSpV algorithm has been analyzed in this context recently [18], hence we just state the result here:

$$T_{\text{SpMSpV}} = O\left(\frac{m}{p} + \beta\left(\frac{m}{p} + \frac{n}{\sqrt{p}}\right) + |\text{iters}| \alpha \sqrt{p}\right).$$

As described in Section III-B, our RCM algorithm needs a specialized SORTPERM function that sorts and labels siblings based on their degrees. Observe that the number of children in the next frontier \mathcal{L}_{next} is less than or equal to the number of parents (vertices with at least one child) in the current frontier \mathcal{L}_{cur} . Hence, we designate a bucket for each parent (i.e., $nnz(\mathcal{L}_{cur})$ buckets in total) and place children to their parent's bucket. These buckets are distributed among processors. Each processor sort its assigned buckets locally, which eliminate the need for a global distributed sorting. More specifically, the i -th processor is responsible for sorting vertices whose parent's labels fall in the range

$$\left[\frac{nnz(\mathcal{L}_{cur})}{p}i, \frac{nnz(\mathcal{L}_{cur})}{p}(i+1)\right).$$

Hence, processors exchange tuples using AllToAll and the responsible processors sort tuples locally. After sorting, the vertex indices associated with the tuples work as the inverse permutation. Processors conduct another round of AllToAll (only the indices) to obtain the sorted permutation.

The total number of nonzeros sorted is exactly the sum of the frontier sizes over all iterations, which is $O(n)$. Hence the per-process cost of SORTPERM is upper bounded by

$$T_{\text{SORTPERM}} = O\left(\frac{n \log n}{p} + \beta \frac{n}{p} + |\text{iters}| \alpha p\right),$$

using personalized all-to-all [19]. We found our specialized bucket sort to be faster than state-of-the-art general sorting libraries, such as HykSort [20].

V. RESULTS

A. Experimental Platform

We evaluated the performance of parallel RCM algorithm on Edison, a Cray XC30 supercomputer at NERSC. In Edison,

Name	Dimensions	BW (pre-RCM)
Description	Nonzeros	BW (post-RCM)
Spy Plot	Pseudo-diameter	
nd24k 3D mesh problem	72K×72K 29M	68,114 10,294 14
Ldoor structural prob.	952K×952K 42.49M	686,979 9,259 178
Serena gas reservoir simulation	1.39M×1.39M 64.1M	81,578 81,218 58
audikw_1 structural prob	943K×943K 78M	925,946 35,170 82
dielFilterV3real higher-order finite element	1.1M×1.1M 89.3M	1,036,475 23,813 84
Flan_1565 3D model of a steel flange	1.6M×1.6M 114M	20,702 20,600 199
Li7Nmax6 nuclear configuration interaction calculations	664K×664K 212M	663,498 490,000 7
Nm7 nuclear configuration interaction calculations	4M×4M 437M	4,073,382 3,692,599 5
nlpkkt240 Sym. indefinite KKT matrix	78M×78M 760M	14,169,841 361,755 243

Fig. 3: Structural information on the sparse matrices used in our experiments. All matrices, except two, are from the University of Florida sparse matrix collection [21]. Li7Nmax6 and Nm7 [22] are from nuclear configuration interaction calculations.

nodes are interconnected with the Cray Aries network using a Dragonfly topology. Each compute node is equipped with 64 GB RAM and two 12-core 2.4 GHz Intel Ivy Bridge processors, each with 30 MB L3 cache. We used Cray’s MPI implementation, which is based on MPICH2. We used OpenMP for intra-node multithreading and compiled the code with gcc 5.2.0 with `-O2 -fopenmp` flags. In our experiments, we only used square process grids because rectangular grids are not supported in CombBLAS [17]. The skewness of the processor grid in SpMSPV has been subject to previous work in the context of breadth-first search [23] where the authors found that square grids are close to optimal in practice (within 10% of the best configuration in most cases).

When p cores were allocated for an experiment, we created a $\sqrt{p/t} \times \sqrt{p/t}$ process grid, where t was the number of threads per process. In our hybrid OpenMP-MPI implementation,

TABLE II: The bandwidth and runtime of the shared-memory RCM implementation in SpMP. SpMP did not finish in 30 minutes for Nm7. On Nm7 and nlpkkt240 matrices, our distributed-memory implementation ran out of memory on a single node of Edison.

Graphs	BW	SpMP Runtime (sec)			Distributed RCM Runtime (sec)		
		1t	6t	24t	1t	6t	24t
nd24k	10,608	0.26	0.06	0.03	1.45	0.38	0.12
ldoor	9,099	3.25	0.52	0.28	4.63	1.52	0.74
Serena	85,229	1.64	0.49	0.66	7.75	2.26	1.08
audikw_1	34,202	1.31	0.34	0.16	7.31	1.99	0.81
dielFilterV3real	25,436	1.99	0.73	0.46	8.63	2.37	0.95
Flan_1565	20,849	1.86	0.44	0.17	12.11	3.88	1.35
Li7Nmax6	443,991	4.62	1.48	0.87	20.28	4.91	2.85
Nm7	-	-	-	-	-	-	-
nlpkkt240	346,556	57.21	25.17	9.92	-	-	-

all MPI processes performed local computation followed by synchronized communication rounds. Local computation in every matrix-algebraic kernel was fully multithreaded using OpenMP. Only one thread in every process made MPI calls in the communication rounds.

B. Matrix Suite

The sparse matrix test suite used in our experiments are shown in Figure 3. These matrices came from a set of real applications, where either a sparse system $Ax = b$ is solved or an eigenvalue problem $Ax = \lambda x$ is solved. The matrices were chosen to represent a variety of different structures and nonzero densities. Since RCM is only well defined on symmetric matrices, all matrices are symmetric. In the last column of Figure 3, the original (pre-RCM) as well as final (post-RCM) bandwidth of the matrix are shown. In the majority of the cases, RCM effectively reduces the bandwidth. Serena and Flan_1565 seem to be the only two matrices where RCM was ineffective in that regard. For Nm7 and nlpkkt240 matrices, our distributed-memory implementation ran out of memory when executed on a single node of Edison because of MPI buffers and other overheads. However, the amount of memory available on that node is sufficient to hold the matrix itself.

C. Shared-memory performance

Our implementation is fully multithreaded to take advantage of the shared-memory parallelism available within a node of modern supercomputers. Here, we compare the quality and runtime of our algorithm with the RCM implementation in SpMP (Sparse Matrix Pre-processing) by Park et al. [24], which is based on optimization from [25] and on the algorithm presented in [8]. The results from SpMP is shown in Table II. For four out of eight matrices where SpMP was able to compute RCM, the RCM ordering from our distributed-memory algorithm (shown in Figure 3) yields smaller bandwidths than SpMP. SpMP is faster than our implementation in shared-memory due to our distributed-memory parallelization overheads. However, SpMP sometimes loses efficiency across NUMA domains. For example, SpMP slows down for Serena on 24 cores compared to 6 cores.

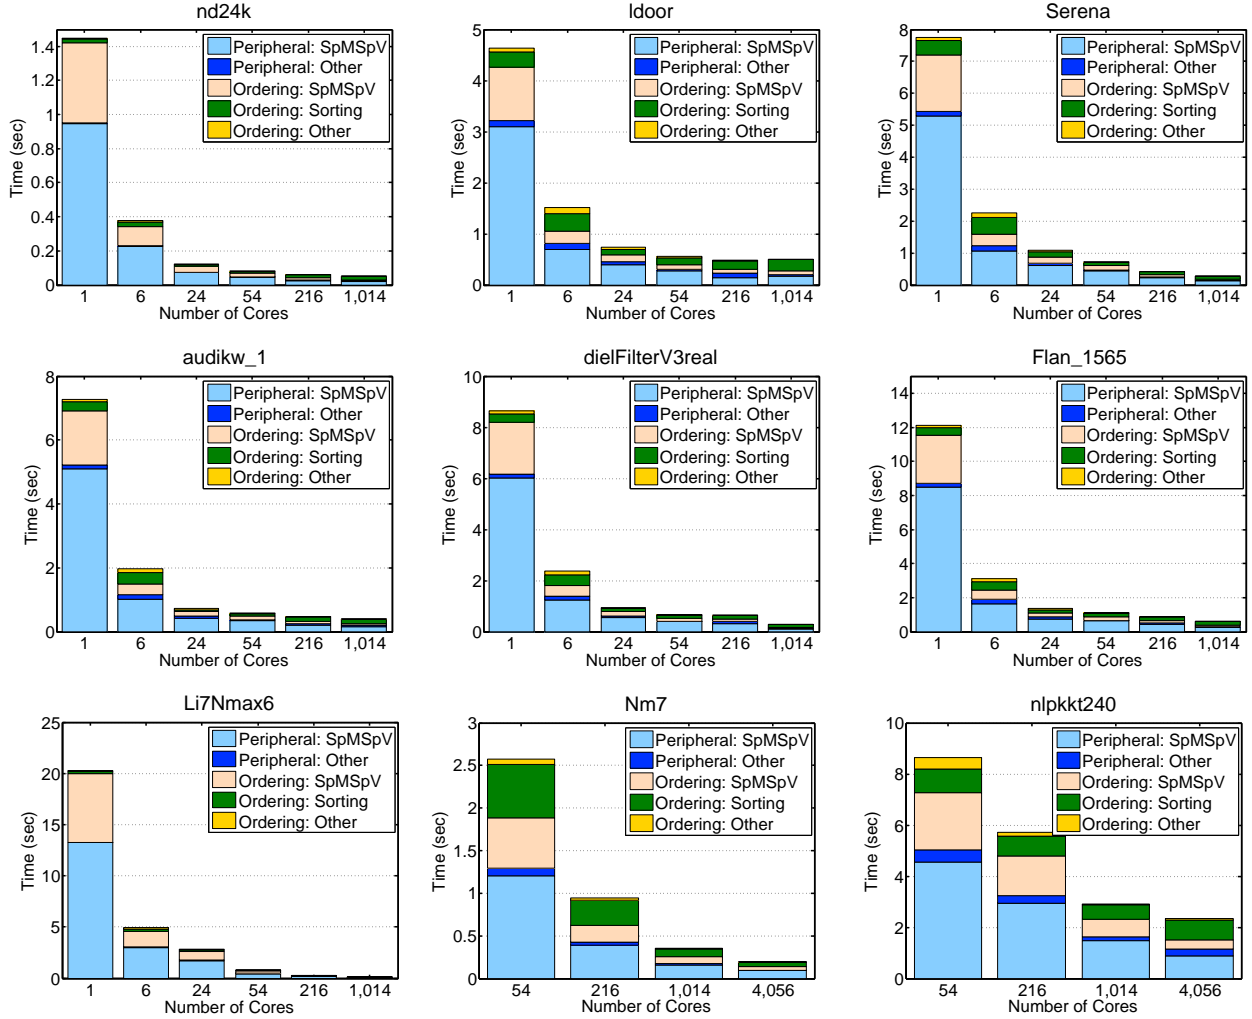


Fig. 4: Runtime breakdown of distributed-memory RCM on different graphs on Edison. Six threads per MPI process are used when the number of cores is greater than or equal to six. Nm7 and nlpkkt240 ran out of memory on a single node of Edison.

One very important thing to note is that in order to compute an ordering using a serial or multithreaded implementation of the RCM algorithm such as SpMP, the matrix structure has to be gathered on a single node. Indeed, in many real-life applications, the matrix is already distributed and this mandatory communication step has a non-negligible cost. For example, it takes over 9 seconds to gather the nlpkkt240 matrix from being distributed over 1024 cores into a single node/core. This time is approximately $3\times$ longer than computing RCM using our algorithm on the same number of cores. One of the key benefit of our approach is that it does not require this step. Adding those times to the time required to compute the ordering itself using SpMP makes our approach highly competitive and often faster.

D. Distributed-memory performance

We ran the distributed-memory RCM algorithm on up to 4096 cores of Edison. All performance results shown in this section used six threads per MPI process as it performed the best. Figure 4 shows the strong scaling of the distributed-

memory RCM algorithm for nine graphs from Table 3. To better understand the performance, we break down the runtime into five parts at each concurrency where the height of a bar denotes the total runtime of identifying a pseudo-peripheral vertex and computing the RCM ordering.

Our distributed algorithm scales up to 1024 cores for all of the nine graphs as shown in Figure 4. The RCM algorithm attains the best speedup of $38\times$ for Li7Nmax6 and $27\times$ on nd24k on 1024 cores. By contrast, it achieves $10\times$ and $13\times$ speedups for ldoor and Flan_1565. The sharp drop in parallel efficiency on these graphs is due to their relatively high diameters. The level-synchronous nature of our BFS incurs high latency costs and decreases the amount of work per processor on high-diameter graphs. Figure 4 shows that SpMSPV is usually the most expensive operation on lower concurrency. However, SORTPERM starts to dominate on high concurrency because it performs an AllToAll among all processes, which has higher latency. The size of the matrix also contributes to the scalability of the distributed-memory RCM algorithm. For example, the largest two matrices in our

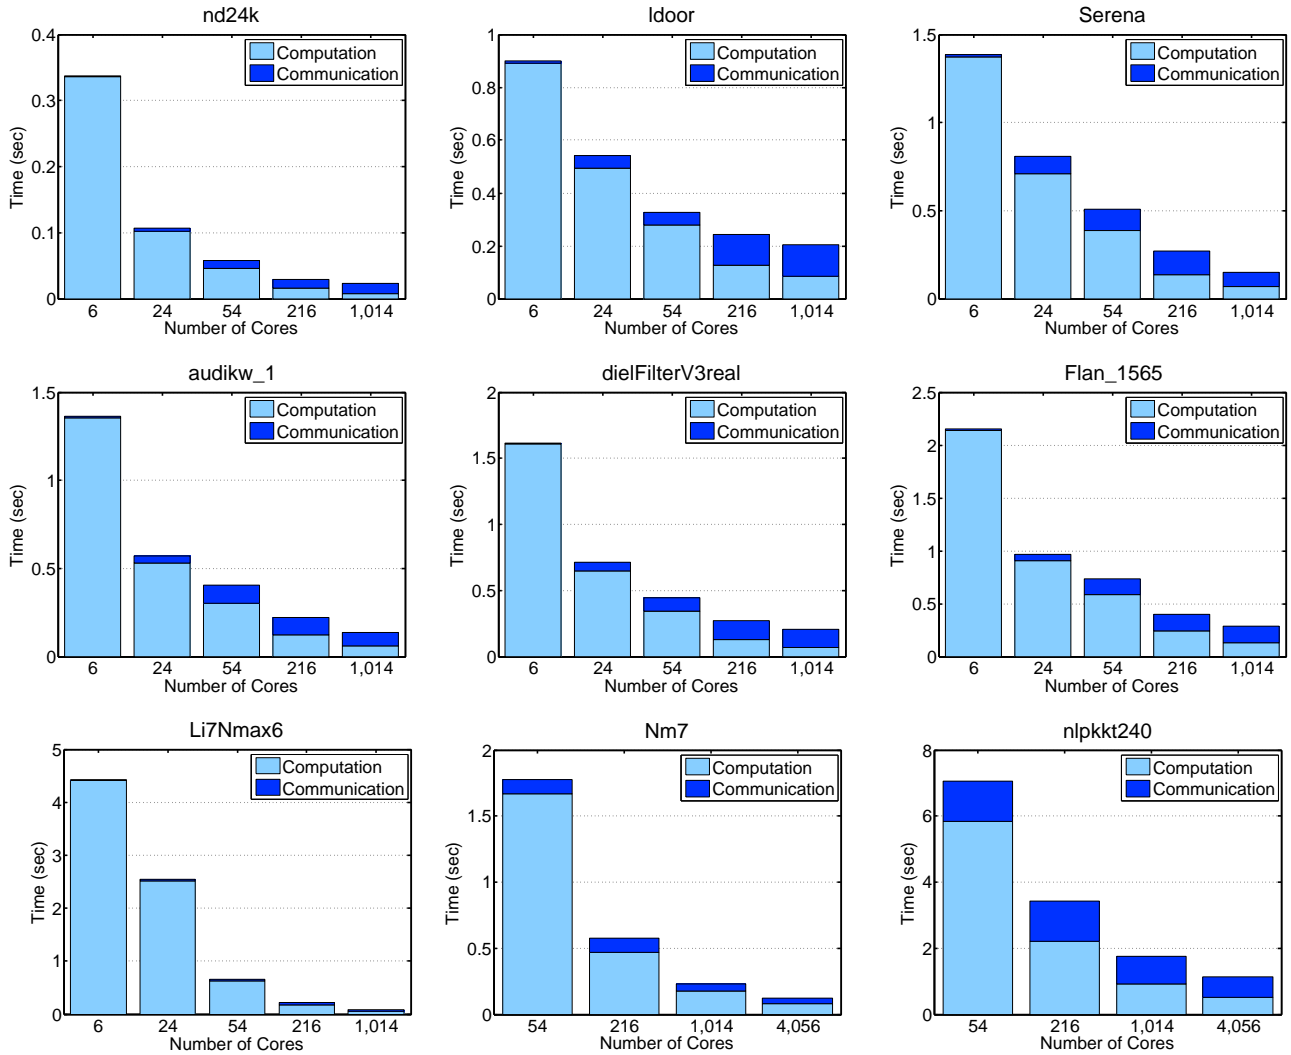


Fig. 5: Computation and communication time in all SpMSpV calls on Edison. Six threads per MPI process are used.

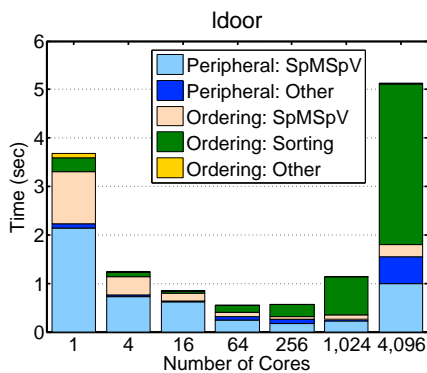


Fig. 6: Breakdown of runtime when computing RCM ordering for `ldoor` using one thread per MPI process on Edison.

test set (`Nm7` and `nlpkkt240`) continue to scale on more than 4K cores whereas smaller problems do not scale beyond 1K cores. This demonstrates that our algorithm can scale on even higher core count if larger matrices are given as inputs.

Since the distributed-memory SpMSpV is the most expen-

sive step of our RCM algorithm, we investigate its performance more closely. Figure 5 shows the breakdown of computation and communication time of the distributed SpMSpV primitive. We observe that on lower concurrency, SpMSpV is dominated by its computation, as expected. The communication time of SpMSpV starts to dominate the computation time on higher concurrency, and the crossover point where the communication becomes more expensive than computation depends on the properties of the matrices. Graphs with higher diameters have higher communication overhead than graphs with low diameters. For example, `ldoor` has one of the highest pseudo-diameters among all problems we considered. Hence, the RCM algorithm on `ldoor` becomes communication bound on 1K cores. By contrast, low-diameter graphs with sizes similar to `ldoor` remain compute bound on 1K cores and scale beyond 1K cores on Edison.

We have also experimented with a flat MPI approach. This non-threaded RCM implementation had higher communication overhead and ran slower than the multithreaded implementation, especially on higher concurrencies. For example, the

flat MPI implementation took $5\times$ longer to compute RCM on the 1door matrix using 4096 cores of Edison as shown in Figure 6.

VI. CONCLUSION AND FUTURE WORK

We have introduced the first distributed-memory implementation of the RCM algorithm. In particular, we have described how the RCM algorithm can be reformulated using the sparse matrix / graph duality, so that the parallel implementation can be accomplished using a small set of parallel primitives that have been highly optimized. Our experiments have shown that the distributed-memory implementation of the RCM algorithm presented in this work scales well up to 1024 processors for smaller matrices, and even to 4096 cores for the largest problems we have evaluated.

We have shown that this new approach is overall faster than the traditional approach that gathers the matrix structure on a single node, followed by the application of a serial or a multithreaded implementation of the RCM algorithm, and then redistributing the permuted matrix. More importantly, our approach removes the memory bottleneck that may be caused by having to store the entire graph with a single node.

A scalable implementation of the RCM algorithm is of prime importance for many iterative solvers that benefit from reordering the matrix. We have assessed this on a sample problem using the conjugate gradient method with block Jacobi preconditioner available in PETSc.

Immediate future work involves finding alternatives to sorting (i.e. global sorting at the end, or not sorting at all and sacrifice some quality). Longer term work would investigate alternative BFS formulations that are not level-synchronous, as the existing approach has trouble scaling to high-diameter graphs.

ACKNOWLEDGMENTS

We thank Hari Sundar for sharing the source code of HykSort. This work is supported by the Applied Mathematics and the SciDAC Programs of the DOE Office of Advanced Scientific Computing Research under contract number DE-AC02-05CH11231. We used resources of the NERSC supported by the Office of Science of the DOE under Contract No. DE-AC02-05CH11231.

REFERENCES

- [1] I. S. Duff, J. K. Reid, and J. A. Scott, "The use of profile reduction algorithms with a frontal code," *International Journal for Numerical Methods in Engineering*, vol. 28, no. 11, pp. 2555–2568, 1989.
- [2] I. S. Duff and G. A. Meurant, "The effect of ordering on preconditioned conjugate gradients," *BIT Numerical Mathematics*, vol. 29, no. 4, pp. 635–657, 1989.
- [3] C. H. Papadimitriou, "The NP-completeness of the bandwidth minimization problem," *Computing*, vol. 16, no. 3, pp. 263–270, 1976.
- [4] E. Cuthill and J. McKee, "Reducing the bandwidth of sparse symmetric matrices," in *Proc. of 24th national conference*. ACM, 1969, pp. 157–172.
- [5] A. George and J. W.-H. Liu, *Computer Solution of Large Sparse Positive Definite Systems*. Englewood Cliffs, New Jersey: Prentice-Hall Inc., 1981.
- [6] S. Sloan, "An algorithm for profile and wavefront reduction of sparse matrices," *International Journal for Numerical Methods in Engineering*, vol. 23, no. 2, pp. 239–251, 1986.
- [7] S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, K. Rupp, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang, "PETSc users manual," Argonne National Laboratory, Tech. Rep. ANL-95/11 - Revision 3.7, 2016. [Online]. Available: <http://www.mcs.anl.gov/petsc>
- [8] K. I. Karantasis, A. Lenharth, D. Nguyen, M. J. Garzaran, and K. Pingali, "Parallelization of reordering algorithms for bandwidth and wavefront reduction," in *International Conference for High Performance Computing, Networking, Storage and Analysis (SC'14)*. IEEE, 2014, pp. 921–932.
- [9] J. A. George, "Computer implementation of the finite element method," Ph.D. dissertation, Stanford, CA, USA, 1971.
- [10] B. Claude, *The Theory of Graphs and its Applications*. John Wiley & Sons, 1962.
- [11] I. Arany, L. Szoda, and W. Smyth, "An improved method for reducing the bandwidth of sparse symmetric matrices," in *IFIP Congress (2)*, 1971, pp. 1246–1250.
- [12] A. Buluç and K. Madduri, "Parallel breadth-first search on distributed memory systems," in *International Conference for High Performance Computing, Networking, Storage and Analysis (SC'11)*. ACM, 2011, pp. 65:1–65:12.
- [13] A. George and J. W. H. Liu, "An implementation of a pseudoperipheral node finder," *ACM Transactions on Mathematical Software (TOMS)*, vol. 5, no. 3, pp. 284–295, Sep. 1979.
- [14] N. E. Gibbs, W. G. Poole, Jr, and P. K. Stockmeyer, "An algorithm for reducing the bandwidth and profile of a sparse matrix," *SIAM Journal on Numerical Analysis (SINUM)*, vol. 13, no. 2, pp. 236–250, 1976.
- [15] F. Checconi, F. Petrini, J. Willcock, A. Lumsdaine, A. R. Choudhury, and Y. Sabharwal, "Breaking the speed and scalability barriers for graph exploration on distributed-memory machines," in *International Conference for High Performance Computing, Networking, Storage and Analysis (SC'12)*. IEEE, 2012, pp. 1–12.
- [16] C. C. Ashcraft, R. G. Grimes, J. G. Lewis, B. W. Peyton, H. D. Simon, and P. E. Bjørstad, "Progress in sparse matrix methods for large linear systems on vector supercomputers," *International Journal of High-Performance Computing Applications (IJHPCA)*, vol. 1, no. 4, pp. 10–30, 1987.
- [17] A. Buluç and J. R. Gilbert, "The Combinatorial BLAS: Design, implementation, and applications," *International Journal of High-Performance Computing Applications (IJHPCA)*, vol. 25, no. 4, 2011.
- [18] A. Azad and A. Buluç, "Distributed-memory algorithms for maximum cardinality matching in bipartite graphs," in *IEEE International Parallel & Distributed Processing Symposium (IPDPS)*, 2016.
- [19] J. Bruck, C.-T. Ho, S. Kipnis, E. Upfal, and D. Weathersby, "Efficient algorithms for all-to-all communications in multiport message-passing systems," *IEEE Transactions on Parallel and Distributed Systems (TPDS)*, vol. 8, no. 11, pp. 1143–1156, 1997.
- [20] H. Sundar, D. Malhotra, and G. Biros, "HykSort: a new variant of hypercube quicksort on distributed memory architectures," in *International Conference on Supercomputing (ICS)*. ACM, 2013, pp. 293–302.
- [21] T. A. Davis and Y. Hu, "The university of florida sparse matrix collection," *ACM Transactions on Mathematical Software (TOMS)*, vol. 38, no. 1, p. 1, 2011.
- [22] H. M. Aktulga, A. Buluç, S. Williams, and C. Yang, "Optimizing sparse matrix-multiple vectors multiplication for nuclear configuration interaction calculations," in *IEEE International Parallel & Distributed Processing Symposium (IPDPS)*. IEEE Computer Society, 2014.
- [23] A. Buluç, S. Beamer, K. Madduri, K. Asanović, and D. Patterson, "Distributed-memory breadth-first search on massive graphs," in *Parallel Graph Algorithms*, D. Bader, Ed. CRC Press, Taylor-Francis, 2017 (in press).
- [24] J. Park et al., "SpMP: Sparse Matrix Pre-processing." [Online]. Available: <https://github.com/jspark1105/SpMP>
- [25] J. Chhugani, N. Satish, C. Kim, J. Sewall, and P. Dubey, "Fast and efficient graph traversal algorithm for CPUs: Maximizing single-node efficiency," in *IEEE International Parallel & Distributed Processing Symposium (IPDPS)*. IEEE, 2012, pp. 378–389.