

Fast and Efficient Parallel Breadth-First Search with Power-law Graph Transformation

Zite JIANG^{1,2}, Tao LIU², Shuai ZHANG^{1,2}, Mengting YUAN (✉)³ and Haihang YOU (✉)²

1 University of Chinese Academy of Sciences, Beijing 100049, China

2 State Key Laboratory of Computer Architecture, Institute of Computing Technology, Chinese Academy of Sciences, Beijing 100190, China

3 School of Computer Science, Wuhan University, Wuhan 430072, China

Abstract In the big data era, graph computing is widely used to exploit the hidden value in real-world graphs in various scenarios such as social networks, knowledge graphs, web searching, and recommendation systems. However, the random memory accesses result in inefficient use of cache and the irregular degree distribution leads to substantial load imbalance. Breadth-First Search (BFS) is frequently utilized as a kernel for many important and complex graph algorithms. In this paper, we describe a preprocessing approach using Reverse Cuthill-McKee (RCM) algorithm to improve data locality and demonstrate how to achieve an efficient load balancing for BFS. Computations on RCM-reordered graph data are also accelerated with SIMD executions. We evaluate the performance of the graph preprocessing approach on Kronecker graphs of the Green Graph500 benchmark and real-world graphs. Our BFS implementation on RCM-reordered graph data achieves 326.48 MTEPS/W (mega TEPS per watt) on an ARMv8 system, ranking 3rd on the Green Graph500 list in November 2020 (the top two rank uses GPU acceleration).

Keywords Graph Computing, Breadth-First Search, Parallel Algorithms, SIMD

1 Introduction

Data-intensive graph-based computations play an important role in many modern big data graph applications such as

social networks [1], knowledge graphs [2], web searching [3], and recommendation systems [4]. There is an increasing need for accelerating graph processing on modern parallel systems. The solutions to these graph applications typically involve classical graph algorithms such as shortest paths [5], connected components [6], spanning trees [7] and centrality [8]. Breadth-First Search (BFS), an underlying graph algorithm, is frequently utilized as a kernel for these important and more complex graph algorithms [9]. Thus, accelerating BFS on modern parallel systems is crucial to the development of graph applications.

Most real-world graphs are large-scale but unstructured and sparse. One of the most notable characteristics of real-world graphs is the skewed *power law* degree distribution [10]: most vertices have a few neighbors while a few own a large number of neighbors. These characteristics present challenges for efficient parallel graph processing, such as load imbalance, poor locality, and redundant computations.

Apart from modifying the graph programming abstraction or changing the execution models on different architectures, reducing the irregularity of graph data also improves the performance of graph processing [11]. For example, it's well-known that BFS has a bad temporal locality, but it's possible to transform irregular graphs to more regular ones to improve spatial locality and gain more performance.

We employ the Reverse Cuthill-McKee (RCM) algorithm [12–14] to relabel vertices of a graph. The bandwidth of an adjacency matrix representing the graph is reduced, which implies improvements in the spatial locality of the graph data. Also, isolated vertices are excluded to reduce memory con-

sumption. RCM-reordering makes logically related vertices clustered more close physically [15]. However, it introduces more serious load imbalance issues for BFS. We resolve the load imbalance strategies on the analysis of the features of the RCM-reordered graph data. To fully exploit the hardware resources and gain more performance improvements, we employ SIMD optimizations with NEON [16–18] intrinsics on ARM systems.

The key contributions of this paper are:

1) We apply RCM-reordering on graph data to relabel vertex ID. The topology of the graph is unchanged. The isolated vertices are excluded for reducing memory consumption. The regularity of the graph is improved, thus reducing cache misses;

2) We present dynamic load balancing strategies for BFS to solve the load imbalance problem introduced by RCM-reordering;

3) We employ SIMD optimization to gain further performance improvements;

4) We evaluate the performance of BFS on RCM-reordered Kronecker graphs and real-world graphs.

This paper is organized as follows. Section II discusses related work. Section III demonstrates the implementation details and optimization techniques of our optimized hybrid-BFS algorithms. Section IV illustrates the experimental results and Section V presents concluding remarks.

2 Related Work

2.1 Graph Instance and Parallel BFS Algorithms

Power consumption is an important indicator in large-scale graph computing. The Green Graph500 [19] benchmark is proposed to rank supercomputers based on their performance of data-intensive applications. It defines a new rate called traversed edges per second per watt (TEPS/w) as the performance metric. The graph generated by the Green Graph500 benchmark is a scale-free, low-diameter and small-world graph called Kronecker graph, similar to real-world social network graphs like twitter [20] and friendster [21]. BFS is the most important kernel in the Green Graph500 benchmark.

From level-synchronous BFS to hybrid-BFS, parallel algorithms for BFS have been studied for decades on multicore systems [22–26] and distributed memory systems [27–31]. In recent years, most of the novel implementations for parallel BFS follow the structure of hybrid BFS and adapt the algorithm to better fit the underlying parallel architecture.

Given a graph $G = (V, E)$ with n vertices, m edges and a distinguished source vertex s . BFS traverses graph G to discover all reachable vertices from s level by level and produces an implicit *breadth-first spanning tree* with root s by maintaining a *predecessor map*. All vertices at level k are visited before any vertices at level $k + 1$. The BFS *frontier* is defined as the set of vertices in the current level. The complexity of sequential level-synchronous BFS algorithm is $O(m + n)$.

Algorithm 1: Level-synchronous parallel BFS algorithm

Input: $G = (V, E)$: undirected graph.

s : source vertex.

Output: π : predecessor map

```

1  $\pi(\cdot) = -1, \pi(s) = s$ 
2  $visited(\cdot) = false, visited(s) = true$ 
3  $F^{cur} = \{s\}, F^{next} = \emptyset$ 
4 while  $F^{cur} \neq \emptyset$  do
5   for  $v \in F^{cur}$  in parallel do
6     for  $w \in neighbors(v)$  do
7       if  $visited(w) \neq true$  atomic then
8          $visited(w) = true$ 
9          $\pi(w) = v$ 
10         $F^{next} = F^{next} \cup \{w\}$ 
11    $F^{cur} = F^{next}$ 
12    $F^{next} = \emptyset$ 

```

A multithread approach to *level-synchronous* parallel BFS is demonstrated as Alg. 1. We assume that the input graph G is both unweighted and undirected. The output π is a predecessor map where $\pi(v)$ is the predecessor vertex on the shortest path from s to v , or -1 if v is not reachable from s . The *visited* set maintains vertices already visited. F^{cur} is the *current frontier queue* initialized with s and F^{next} is the *next frontier queue* initialized as an empty queue. At each iteration, each thread is assigned with a portion of vertices in F^{cur} , explores all neighbors of each vertex and finds unvisited neighbors and puts it at F^{next} which will be traversed at the next level. Note that examinations on *visited* and updates on both *visited* and F^{next} must be atomic to avoid race conditions and redundant computations.

Real-world graphs such as social networks are usually low-diameter and scale-free. The total number of BFS steps is usually small, e.g. six to eight steps on a Kronecker graph. There are some extremely high degree vertices, causing certain frontiers to grow very fast. When the current frontier reaches its largest size, the majority of the graph has been visited and the following frontiers will shrink. There will be a great number of wasted checks on edges when traversing in

the *top-down* way (like Alg. 1) with a large frontier, which becomes the main bottleneck of level-synchronous parallel BFS.

Algorithm 2: Hybrid parallel BFS algorithm

Input: $G = (V, E)$: undirected graph.
 s : source vertex.
Output: π : predecessor map

```

1  $\pi(\cdot) = -1, \pi(s) = s$ 
2  $visited(\cdot) = false, visited(s) = true$ 
3  $F^{cur} = \{s\}, F^{next} = \emptyset$ 
4  $traversal\_policy = TOP\_DOWN$ 
5 while  $F^{cur} \neq \emptyset$  do
6   if  $traversal\_policy = TOP\_DOWN$  then
7     for  $v \in F^{cur}$  in parallel do
8       for  $w \in neighbors(v)$  do
9         if  $visited(w) \neq true$  atomic then
10            $visited(w) = true$ 
11            $\pi(w) = v$ 
12            $F^{next} = F^{next} \cup \{w\}$ 
13   else
14     for  $v \in V$  in parallel do
15       if  $visited(v) = false$  then
16         for  $w \in neighbors(v)$  do
17           if  $w \in F^{cur}$  then
18              $visited(v) = true$ 
19              $\pi(v) = w$ 
20              $F^{next} = F^{next} \cup \{v\}$ 
21             break
22    $traversal\_policy = update\_traversal\_policy()$ 
23    $F^{cur} = F^{next}$ 
24    $F^{next} = \emptyset$ 

```

Beamer et al. [30] proposed the hybrid BFS algorithm (Alg. 2) which combines the conventional top-down (level-synchronous) approach with a novel bottom-up approach. The hybrid algorithm employs the top-down approach on small frontiers and the bottom-up approach on large frontiers. When the frontier is large, the bottom-up approach will examine every edge attached to an unvisited vertex until finding a visited vertex in the frontier as its parent. Once a vertex has found a parent, it stops checking the rest of its neighbors, and thus the total number of edges examined are reduced (Alg. 2, line 15-21). Beamer et al. also explained a heuristic rule to determine when to switch between the top-down approach and the bottom-up approach by tuning two parameters: α and β .

2.2 Graph Preprocessing and Transformation

Many distributed graph processing systems such as Pregel [32], GraphLab [33], and PowerGraph [10], have been proposed to handle graphs of extremely large scale by exploiting computation resources of clusters. However, there are still many challenges such as load imbalance, memory consumption, and synchronization overhead. Many techniques for graph preprocessing and transformation have been proposed to improve the efficiency of graph processing in recent years.

GraphChi [34], a disk-based system, achieves efficient out-of-memory graph processing by well-designed partitioning and an asynchronous computation model which requires only a very small number of non-sequential accesses to the disk. GridGraph [35] improves the locality and reduces I/O operations with a novel grid representation for graphs and a fast streaming-apply graph processing model. GraphM [36] is an efficient storage system to amortize the storage consumption and the data access overhead between concurrent graph processing jobs. These graph processing systems focus on well-designed graph partitioning, fine-grained memory management, and efficient scheduling policy.

Tigr [11] is a graph transformation framework that can effectively reduce the irregularity of real-world graphs to make these graphs better suited to GPU's SIMD execution. But Tigr changes the topology of a graph, which may affect the convergence of graph processing and alter the final results. Hao et al. [37] proposed a graph ordering algorithm Gorder to keep frequently accessed nodes together locally for minimizing the cache miss ratio. However, Gorder is much slower than RCM [12] and results in comparable performance for the BFS algorithm.

2.3 Reverse Cuthill-McKee Algorithm

Many sparse matrix computations can be accelerated by reordering the matrix to reduce its bandwidth. Similarly, reordering vertices of a graph is crucial to minimize data size, maximize data locality and improve the performance of graph algorithms [15]. Several heuristic reordering algorithms are used in practice since computing the optimal reordering with a minimal bandwidth is NP-complete [38], such as Cuthill-McKee (CM) [39], Reverse Cuthill-McKee (RCM) [12] and Slogan [40].

RCM reordering, a variant of CM reordering, is widely used to reduce the *bandwidth* of a sparse symmetric matrix A . The *bandwidth* [14] of matrix A is defined as below:

$$BW(\mathbf{A}) = \max \left\{ i - j \mid A_{ij} \neq 0, i > j \right\} \quad (1)$$

Obtaining a reordering of rows or columns of \mathbf{A} is equivalent to the process of relabeling vertices of graph $G(\mathbf{A})$ associated with \mathbf{A} . The RCM-reordered matrix usually has a smaller *bandwidth*, where the non-isolated vertices are clustered closer and data locality is improved.

We will illustrate the decrease of cache misses while performing hybrid parallel BFS on RCM-reordered Kronecker graphs in section 4.

2.4 ARM Neon Technology

ARM Neon technology [16], the ARM Advanced SIMD (Single Instruction Multiple Data) architecture extension, provides thirty-two 128-bit vector registers on an ARMv8 system. Each register is capable of containing multiple lanes of data. SIMD instructions are utilized to perform the same operations in parallel on those multiple lanes of data. Many data-intensive applications can benefit from Neon technology, such as multimedia and signal processing, 3D graphics, speech, image processing, where fixed and floating-point performance is critical.

There are several ways to use Neon technology for programmers, including Neon intrinsics, Neon-enabled libraries, auto-vectorization by the compiler and hand-coded Neon assembler. The Neon intrinsics are a set of C and C++ functions supported by the Arm compiler and GCC. These intrinsics give programmers direct access to Neon instructions and offer substantial performance improvement without the need for hand-written assembly code. In this paper, we employ Neon intrinsics to improve the performance of parallel BFS.

3 Implementation Details

3.1 Graph Representation

The CSR (Compressed Sparse Row) representation is space-efficient and extremely fast on a single-node system, which provides constant-time access to neighbors of a vertex [29]. When the graph fits comfortably into memory, it is recommended to use CSR for fast computations. In this paper, we use CSR to store the adjacency matrix representing the graph. There are mainly two arrays in this data structure: *dst* holds neighbors' vertex ID of vertices in the graph; *row_starts* records the offset index of all neighbors of each vertex in the *dst* array. Given a vertex v_0 , the set of neighbors of v_0 in *dst* are:

$$\text{neighbors}(v_0) = \left\{ \text{dst}[i] \mid \text{row_starts}[v_0] \leq i < \text{row_starts}[v_0 + 1] \right\} \quad (2)$$

Like the Kronecker graph in the Green Graph500 benchmark, we assume that the graph is undirected. So each edge is stored twice in both directions. We use 32-bit integers to represent the vertex ID, which is sufficient when the graph scale is at most 30 so that our multicore system with 384 GB memory can handle the Kronecker graph with scale 30.

Algorithm 3: Reverse Cuthill-McKee algorithm

Input: $G = (V, E)$: undirected graph
Output: P : permutation array

- 1 V^+ = all non-isolated vertices in V sorted in increasing order of degree
- 2 *sorted_neighbors*(v) = neighbors of v sorted in increasing order of degree
- 3 $P(0) = V^+(0)$
- 4 *visited*(\cdot) = *false*
- 5 *min_index* = 0
- 6 *slow* = 0, *fast* = 0
- 7 **while** *slow* < $|V^+|$ **do**
- 8 // find next unvisited vertex with the minimal degree
- 9 **for** $i \in [\text{min_index}, |V^+|)$ **do**
- 10 **if** *visited*($V^+(i)$) = *false* **then**
- 11 $P(\text{slow}) = V^+(i)$
- 12 *visited*($V^+(i)$) = *true*
- 13 *min_index* = $i + 1$
- 14 *fast* = *slow* + 1
- 15 **break**
- 16 // explore next connected component and relabel
- 17 // unvisited vertices in increasing order of degree
- 18 **while** *slow* < *fast* **do**
- 19 **for** $v \in \text{sorted_neighbors}(P(\text{slow}))$ **do**
- 20 **if** *visited*(v) \neq *true* **then**
- 21 *visited*(v) = *true*
- 22 $P(\text{fast}) = v$
- 23 *fast*++
- 24 *slow*++
- 25 $P = \text{reverse}(P)$
- 26 $P.\text{append}(V \setminus V^+)$

3.2 RCM Reordering

A big problem with BFS is the bad locality because of the randomness of memory accesses on graph data. By reordering the vertex IDs to better fit in the access pattern of BFS,

we can expect to achieve a higher locality. We apply RCM reordering to reduce the bandwidth of the adjacency matrix of the graph as a preprocessing step before running BFS. The RCM implementation is depicted in Alg. 3.

The first labeled vertex strongly impacts the bandwidth of the permuted adjacency matrix. A frequently employed heuristic is to find a *pseudo-peripheral vertex* with a high *eccentricity* [14]. However, to evaluate on a reproducible permutation array, we simply select the vertex with minimal degree as the first labeled vertex (Alg. 3, line 11).

There are two prerequisite data for RCM:

1. V^+ : All non-isolated vertices sorted in increasing order of degree. V^+ is used to find the vertex with minimal degree in the next connected component.
2. $sorted_neighbors(v)$: Neighbors of v sorted in increasing order of degree.

RCM repeatedly explores all connected components in the graph, like BFS, level by level and relabels neighbors of the current vertex in increasing order of degree. In each connected component, the vertex with a minimal degree is first relabeled as *slow* (Alg. 3 lines 9-15). *slow* and *fast* are both indices of the permutation array P . The RCM kernel repeatedly explores the neighbors of the vertex with an original ID of $P(slow)$, relabels unvisited neighbors and moves *slow* and *fast* forward until the current connected component is fully explored (breaking the while loop at Alg. 3 line 18). Then the consecutive connected components are explored in the same way. After relabeling all vertices in V^+ , P is reversed and all isolated vertices are appended to P . Then P is the permutation array generated by RCM.

We recorded the number of connected components and the range of relabeled vertex ID for each connected component when the RCM-reordering is executed on the Kronecker graph with a scale from 21 to 30. Over 99.9% of the non-isolated vertices are located in the last connected component and almost all other connected components are constituted by only two vertices connected to each other. The end ID of the last connected component is also the number of non-isolated vertices.

The original adjacency matrix, as shown in Fig. 1, looks like the typical sparse symmetric matrix. According to the definition of bandwidth, It is estimated that its bandwidth is very close to $|V|$. After RCM reordering, the structure of the adjacency matrix becomes the shape of “leaf”, which implies that the bandwidth is less than $|V^+|$. It can be considered that bandwidth is strong related to cache misses rate.




Fig. 1 Structure comparison of adjacency matrices before and after RCM-reordering. The left side shows the adjacency matrix for the original Kronecker graph. The right side shows the adjacency matrix for the RCM-reordered graph.

3.3 Load Balancing

Hybrid BFS on Kronecker graphs typically terminates after six to eight steps. These steps are divided into a *growing* phase and a *shrinking* phase according to the evolution of the *frontier* in [25]. Instead, We divide these steps into three phases and analyze the load balancing of each phase:

Phase 1: A top-down phase in the first several steps. The frontier is small and usually contains high-degree vertices, which indicates the next frontier will be much larger. However, the degree distribution of frontier vertices is quite uneven, which will cause severe load imbalance among the threads. Some threads are idle, especially when the number of frontier vertices is less than the number of threads, which makes it impossible to make full use of all CPU cores. A solution to this problem is shown in Alg. 4.

Like Alg. 1, Alg. 4 also uses a static scheduling policy. For each vertex in the current frontier, we divide its neighbors evenly into blocks of the same size and assign them to each thread. The remaining neighbors of the j th vertex are assigned to $(j \bmod t)$ th thread in a Round-Robin way.

For a vertex v_0 , the neighbors assigned to thread i and the remaining neighbors are both obtained by calculating offset in dst with thread ID in constant time as below:

$$\begin{aligned} neighbors_i(v_0, t) &= \\ & \left\{ dst[j] \mid startpos_i \leq j < startpos_i + workload \right\} \\ remaining_neighbors(v_0, t) &= \\ & \left\{ dst[j] \mid startpos_t \leq j < row_starts[v_0 + 1] \right\} \end{aligned} \quad (3)$$

where

$$\begin{aligned} workload &= \frac{row_starts[v_0 + 1] - row_starts[v_0]}{t} \\ startpos_k &= row_starts[v_0] + k \times workload \end{aligned}$$

Algorithm 4: Load-balanced top-down BFS

Input: $G = (V, E)$: undirected graph.
 F^{cur} : current frontier.
 t : number of threads.
 $visited$: the set of visited vertices.
 π : predecessor map.
Output: F^{next} : next frontier.

```

1  $F^{next} = \emptyset$ 
2 for all threads  $T_i$  in parallel do
3   for  $j \in [0, F^{cur}]$  do
4      $v = F^{cur}[j]$ 
5      $W = neighbors_i(v, t)$ 
6     if  $i = j \bmod t$  then
7        $W = W \cup remaining\_neighbors(v, t)$ 
8     for  $w \in W$  do
9       if  $visited(w) \neq true$  atomic then
10         $visited(w) = true$ 
11         $\pi(w) = v$ 
12         $F^{next} = F^{next} \cup \{w\}$ 

```




Fig. 2 Comparison of distribution of edges among threads in a top-down step (scale 26, edgfactor 16, level 3). The edges are nearly evenly distributed to each thread for Alg. 4. However, for Alg. 1, most edges are assigned to the first several threads.

Fig. 2 shows the load balancing of *top-down* kernel in Alg. 1 and Alg. 4. This experiment selected a Kronecker graph with scale 26 and edgfactor 16. We collect data on the distribution of the number of edges at level 3, a *top-down* step, where 8364088 edges are traversed.

Phase 2: A *bottom-up* phase in the middle several steps. The frontier is extremely larger than the ones in the other two phases. The benefits of RCM-reordering are best reflected in this phase, where most of the vertices are visited.

In the *bottom-up* step, for each unvisited vertex, the *bottom-up* kernel explores each of its neighbors until one neighbor in the frontier is found. The data structure of the frontier and next frontier is transformed into a bitmap so that we can check if a vertex is in the frontier and insert the vertex into the next frontier in constant time. The parameters $visited$, F^{cur} and F^{next} are all bitmaps of length $|V|$. Each of these bitmaps is cache line aligned and divided into blocks of cache line size (64 bytes in this paper) without the risks of race conditions and *false sharing*. And there is no need to use locks or atomic operations to keep thread safety.

Assuming there are t threads, a simple static scheduling policy divides V evenly into t partitions and assigns one partition to each thread. This scheduling policy is efficient on random degree distribution before RCM-reordering but will cause load imbalance after RCM-reordering. We scrutinize the RCM-reordered graph data and it features an ascending degree distribution.

We use a descending partitioning policy to reduce differences in workload between partitions since the distribution of degrees of the RCM-reordered graph is ascending. One partition contains many blocks of cache line size. The sequence of the number of blocks in each partition is a descending arithmetic sequence. To achieve better load balancing, the number of generated partitions is much larger than the number of threads. At runtime, each thread continuously steals the next unprocessed partition and executes the *bottom-up* kernel on that partition until all partitions have been processed.

We achieve good load balancing in the *bottom-up* phase with the combination of static partitioning and dynamic work-stealing as demonstrated in Alg. 5. A parameter λ is introduced to determine the number of partitions:

$$|S| = \lambda \times t \quad (4)$$

$S[s]$ is the next partition to be processed. Each thread continuously steals a partition through a *fetch_add* operation on s and executes the *bottom-up* kernel until the value of s is equal to or greater than $|S|$.

Algorithm 5: Load-balanced bottom-up BFS

Input: $G = (V, E)$: undirected graph.
 F^{cur} : current frontier (bitmap).
 t : number of threads.
 λ : partition factor.
 $visited$: the set of visited vertices (bitmap).
 π : predecessor map.

Output: F^{next} : next frontier (bitmap).

```

1  $F^{next} = \emptyset$ 
2  $S = get\_partitions(G, \lambda, t)$ 
3  $s = 0$ 
4 for all threads  $T_i$  in parallel do
5    $pos = s.fetch\_add(1)$ 
6   if  $pos \geq |S|$  then
7     break
8   for  $v \in S[pos]$  do
9     if  $visited(v) = false$  then
10      for  $w \in neighbors(v)$  do
11        if  $w \in F^{cur}$  then
12           $visited(v) = true$ 
13           $\pi(v) = w$ 
14           $F^{next} = F^{next} \cup \{v\}$ 
15        break

```

Phase 3: A *top-down* phase in the last several steps. The frontier is small but relatively larger than in phase 1. Most of the vertices in the frontier are low-degree. It is efficient enough to employ the *top-down* kernel in Alg. 1 in this phase.

3.4 Workload Reduction

The major bottleneck of the hybrid BFS algorithm is the *bottom-up* step. Reducing the workload is another way to accelerate the *bottom-up* step apart from load balancing.

We employ two methods to reduce the workload in a *bottom-up* step in Alg. 6.

1) *degree-aware BFS:* The *degree-aware* BFS [26] is proposed to reduce the number of traversed edges in a *bottom-up* step. It suggests that most traversed edges are concentrated in the first *bottom-up* step and the number of traversed edges is affected by the ordering of the degree of each vertex's neighbors. And the descending ordering strategy is a better choice. As shown in Alg. 6, the neighbors of each vertex v is sorted in descending degree and separated into the highest-degree neighbor $neighbors^+(v)$ and the resting neighbors $neighbors^-(v)$. The original loop of the *bottom-up* kernel (Alg. 5, lines 8-15) is also separated into two loops (Alg. 6, lines 11-17 and lines 19-26).

Algorithm 6: Load-balanced bottom-up BFS with workload reduction

Input: $G = (V, E)$: undirected graph.
 F^{cur} : current frontier (bitmap).
 t : number of threads.
 λ : partition factor.
 $visited$: the set of visited vertices (bitmap).
 π : predecessor map.

Output: F^{next} : next frontier (bitmap).

- 1 $neighbors^+(v)$ = the highest-degree neighbor of v
- 2 $neighbors^-(v)$ = neighbors of v sorted in decreasing order of degree (excluding $neighbors^+(v)$)
- 3 $F^{next} = \emptyset$
- 4 $S = get_partitions(G, \lambda, t)$
- 5 $s = 0$
- 6 **for all threads** T_i **in parallel do**
- 7 $pos = s.fetch_add(1)$
- 8 **if** $pos \geq |S|$ **then**
- 9 **break**
- 10 $shrink(S[pos])$
- 11 **for** $v \in S[pos]$ **do**
- 12 **if** $visited(v) = false$ **then**
- 13 **for** $w \in neighbors^+(v)$ **do**
- 14 **if** $w \in F^{cur}$ **then**
- 15 $visited(v) = true$
- 16 $\pi(v) = w$
- 17 $F^{next} = F^{next} \cup \{v\}$
- 18 $shrink(S[pos])$
- 19 **for** $v \in S[pos]$ **do**
- 20 **if** $visited(v) = false$ **then**
- 21 **for** $w \in neighbors^-(v)$ **do**
- 22 **if** $w \in F^{cur}$ **then**
- 23 $visited(v) = true$
- 24 $\pi(v) = w$
- 25 $F^{next} = F^{next} \cup \{v\}$
- 26 **break**

2) *shrinking partitions*: Once a thread steals a partition, it shrinks the partition before entering the *bottom-up* kernel (Alg. 6 line 10,18). Each partition $[m, n]$ is shrunk to $[m', n']$ when all vertices in $[m, m')$ and $(n', n]$ are visited. The shrink operation does not introduce additional workload in the current step and helps reducing workload in the consecutive *bottom-up* steps. Compared with the original graph representation, a partition is more likely to shrink after RCM-reordering because the visited vertices are spatially more closer to each other with the bandwidth reduction.

Fig. 3 shows the size of all partitions at each *bottom-up* step with a scale from 21 to 30. With 40 threads and λ set to 20, the RCM-reordering enabled hybrid BFS demonstrates significant and stable reduction on partitions, especially from the 4th step to the 5th step. However, the RCM-reordering disabled version hardly reduces partitions with a scale greater than 23. In addition, the size of all partitions at the first bottom-up step of the RCM-reordering enabled version is much smaller than that of the RCM-reordering disabled version because all isolated vertices are already excluded.

The *bottom-up* workload is reduced with the above two strategies.

3.5 SIMD Optimizations

There are almost no floating-point arithmetic operations in the BFS algorithm. Most of the operations are simple integer operations and bit manipulations. There are a large number of memory accesses, where read operations are much more than write operations. To make better use of the memory bandwidth, we locate two types of memory access intensive operations in our program and accelerate them with NEON intrinsics.

1) *memory copy*: Memory copy happens when (1) thread-local next frontier is copied to global frontier and (2) the *bottom-up* next frontier is transformed into the *top-down* frontier. We implemented a *memcpy_128()* function which encapsulates two 128-bit load and store NEON intrinsics (*vst1q_u64()* and *vld1q_u64()*).

2) *logical or*: In the *bottom-up* kernel, the newly visited vertex v is updated into two bitmaps: *visited* (Alg. 5 line 13) and F^{next} (Alg. 5 line 15). We update the *visited* outside the for loop at Alg. 5 line 9 through 128-bit logical or operations with the NEON intrinsic *vorrq_u64()*. But this method is only suitable for more regular graphs such that a preprocessing on graph data like RCM-reordering is required.

The NEON architecture provides full unaligned support for NEON data access. However, for the Cortex-A8 proces-

or, specifying 128-bit or greater alignment saves one cycle per NEON instruction. We make most of the data 64 bytes aligned to take advantage of both cache efficiency and NEON acceleration.

3.6 Parameter Tuning

The parameters α and β for hybrid BFS were tuned with the methodology as described in [25]. We empirically set the ranges for α and β to $[1, 128]$ and $[1, 32]$, then collect performance data of BFS. The optimal value of (α, β) is determined to be $(64, 8)$ by grid searching.

Another parameter λ is tuned with 40 threads on Kronecker graphs with a scale from 21 to 30 and edgfactor 16. In general, the optimum value of λ increases as the scale is larger. For graphs with scale less than 26, when λ is greater than ten, the mean GTEPS decreases as λ increases, because bigger λ means more small partitions, which leads to more frequent work-stealing among threads. For graphs with a scale greater than 25, with λ bigger than 20, the mean GTEPS fluctuates in a small range around 90 percent of the best performance for this scale. Thus, the value of λ is recommended to be less than or close to 10 with a scale smaller than 26 and more than 20 with a scale bigger than 25.

4 PERFORMANCE EVALUATION

4.1 Evaluation Platform

We evaluate the performance of our algorithm on an ARMv8 system. The experimental environments are listed in Table 2.

4.2 Graph Instances

The performance of our implementation is measured on Kronecker graphs generated by the official Green Graph500 benchmark. Initiator parameters for the Kronecker graph generator (A, B, C, D) are set to $(0.57, 0.19, 0.19, 0.05)$. The graph size is determined by two parameters: *scale* and *edgfactor*. The total number of vertices is $N = 2^{scale}$ and the number of edges is $M = N \times edgfactor$. There are some self-loops, duplicated edges and isolated vertices in the graph.

4.3 Performance Overview

Ablation experiments are tested for each optimization strategy as shown in Fig. 4. In fact, every optimization can improve the performance independently and works well in kronecker graphs, twitter [21] and friendster [20].

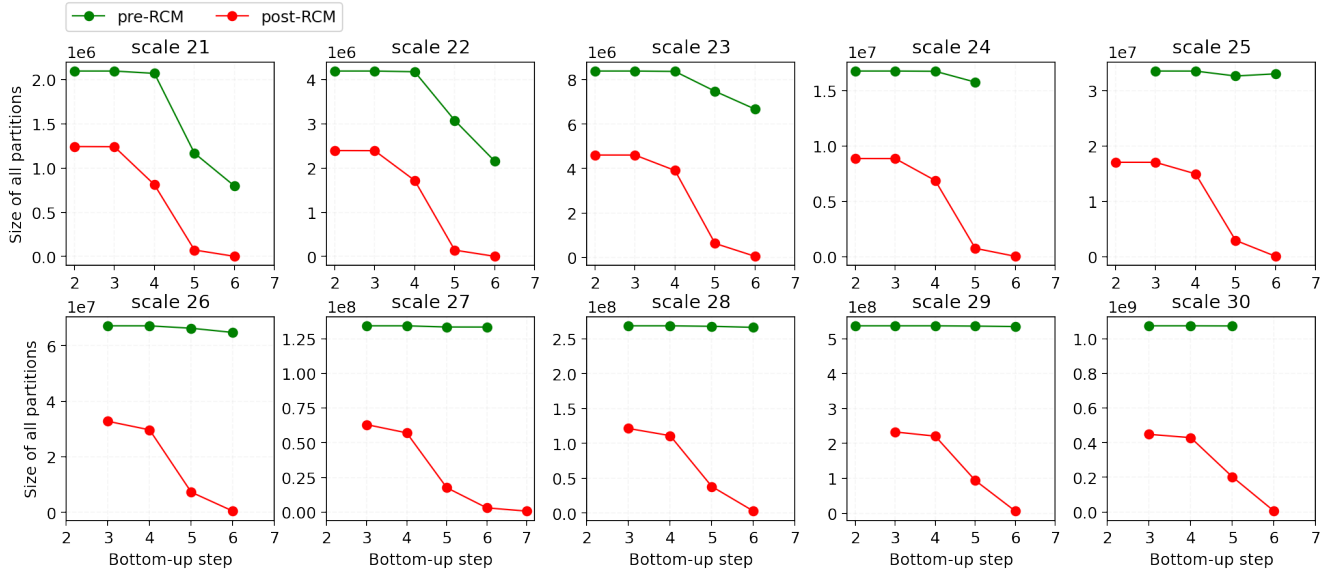


Fig. 3 Size of all partitions at each bottom-up step before and after RCM-reordering. Each size is the average of values collected from the 64 rounds BFS (40 threads, edgefactor 16, λ 20). The initial size of all partitions for the RCM-reordered graph is smaller than that of the original Kronecker graph because all isolated vertices are excluded. The former shows a significant reduction for all scales while the latter rarely reduces when the scale is greater than 24.




Fig. 4 Ablation experiments on different optimizations. 1. *baseline*: hybrid bfs; 2. *rmv-iso*: baseline + remove isolated vertices; 3. *rcm*: rmv-iso + rcm; 4. *rcm-lb*: rcm+load balance; 5. *rcm-simd*: rcm-lb + simd.

Table 1 cache miss comparison before and after RCM-reordering

scale	cache references			cache misses			cache miss rate (%)	
	pre-RCM (10^9)	post-RCM (10^9)	reduction (%)	pre-RCM (10^9)	post-RCM (10^9)	reduction (%)	pre-RCM	post-RCM
21	4.5783	1.6449	64.07	0.2111	0.0474	77.55	4.59	2.88
22	9.6919	3.2398	66.57	0.4654	0.1127	75.78	4.80	3.47
23	18.9666	6.2313	67.15	0.8536	0.2795	67.26	4.49	4.48
24	38.3984	11.8044	69.26	1.6832	0.5653	66.42	4.38	4.79
25	80.5952	23.4695	70.88	3.6657	1.4753	59.75	4.55	6.29
26	161.2918	47.5949	70.49	6.8352	3.7127	45.68	4.23	7.80
27	331.3712	89.7782	72.91	15.6341	6.6334	57.57	4.71	7.39
28	648.0242	165.1345	74.52	28.5306	11.5033	59.68	4.40	6.96
29	1344.9668	348.2868	74.10	55.2038	32.1911	41.69	4.10	9.23
30	2600.5620	676.3802	73.99	111.8545	68.3663	38.88	4.30	10.09

Table 2 Evaluation Platform

Processor	Qualcomm Centriq 2434, 2.3 GHz
Socket	1
Cores	40
Threads	40
Cache(L1/L2/L3)	3.75 MB / 10 MB / 50 MB
SIMD	128 bits
Mem.	384 GB (DDR4 2666 MHz)
TDP	110 W
OS, Compiler	Centos 7, gcc-7.3.1

4.4 Cache Miss Reduction

Table 3 shows the cache misses reduction after applying RCM-reordering on original Kronecker graphs. Our hybrid BFS algorithm runs on Kronecker graphs with a scale from 21 to 30 and edgfactor 16 using a single thread. We collect statistics about cache references and cache misses during running BFS on 64 randomly selected source vertices (after generating Kronecker graph and RCM-reordering) and calculate the corresponding cache miss rate.

Both cache references and cache misses are decreased a lot with RCM-reordering enabled. The cache references decrease at an ascending speed and the cache misses decrease at a descending speed. However, when RCM-reordering is enabled, the cache miss rate becomes high as scale increases and exceeds that when RCM-reordering is disabled with a scale greater than 24. The reason is that the percentage of non-isolated vertices decreases as the scale increases and cache references are greatly reduced. Nevertheless, the performance is improved since the absolute amount of cache references and cache misses instead of the cache miss rate account for the running time.

4.5 Strong Scaling

Strong scaling for original hybrid BFS and RCM-reordering enabled hybrid BFS on a Kronecker graph with scale 30 and edgfactor 16 are illustrated in Table 4 and Fig. 6. Compared with original hybrid BFS, RCM-reordering enabled hybrid BFS demonstrates a speedup over 4 times and better scalability.

Table 3 Strong Scaling comparison (scale 30, edgfactor 16)

Threads	Hybrid BFS		RCM + Hybrid BFS	
	GTEPS	Speedup	GTEPS	Speedup
1	0.235	1.0	0.965	1.0
2	0.414	1.8	1.963	2.0
4	0.846	3.6	4.011	4.2
8	1.626	6.9	7.821	8.1
16	3.256	13.9	15.284	15.8
32	6.068	25.8	29.115	30.2
40	7.283	31.0	39.834	41.3

The strong scaling of the RCM-reordering enabled hybrid BFS with 40 threads is approximately 41 times faster than sequential computation, achieving a superlinear scaling. The reason is that the number of edges precisely traversed is not as large as the total number of edges generated by the Kronecker generator with plenty of duplicated edges.

4.6 Partial RCM

Admittedly, an RCM-based algorithm will increase the overall algorithm time, and the overhead of RCM is much greater than BFS time in Table 4.

Partial RCM is proposed to reduce RCM time cost while maintaining good memory access locality. A p -partial RCM is to only reorder the top- p percents vertices rather than to reorder all vertices. Intuitively, as the ratio increases, efficiency

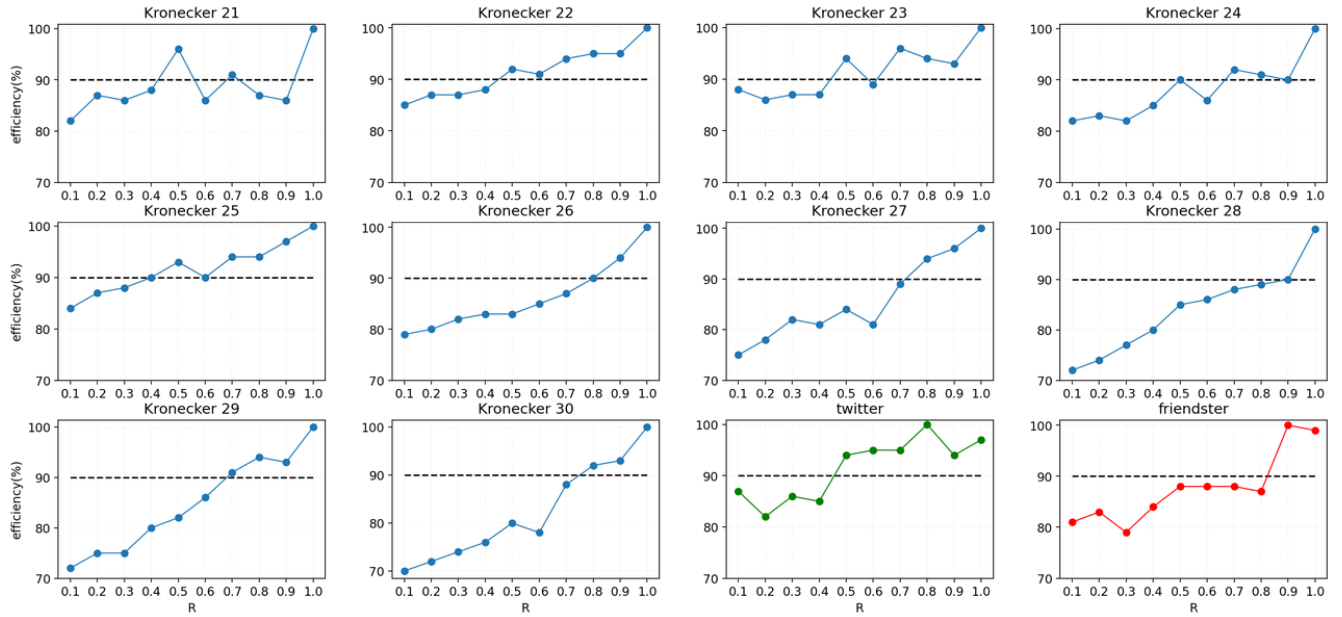


Fig. 5 Partial RCM on Kronecker, twitter and friendster graphs.

Table 4 RCM time cost in BFS

graph	RCM time(s)	RCM BFS time(s)	BFS time(s)
<i>kronecker</i>	836.707	0.425441	1.08366
<i>twitter</i>	35.2149	0.0885213	0.134154
<i>friendster</i>	50.2221	0.226269	0.342534

will increase. And the reduction in time is linearly related to the reduction in proportion.

However, considering the influence of vertices proportion from 0.1 to 1 on the efficiency of RCM BFS, the results show that partial RCM have better performance with proper ratio value in twitter and friendster graphs in Fig. 5. The efficiency of each ratio is obtained by averaging multiple experiments, and the 100% efficiency corresponds to the fastest p -partial RCM in a test case.

The experiments show that p -partial RCM still maintains good performance in BFS search.

5 CONCLUSION

In this paper, inspired by research on the Reverse Cuthill-McKee algorithm for reducing the bandwidth of a sparse matrix, we combine hybrid BFS with RCM-reordering to achieve better efficiency for BFS. The RCM-reordering demonstrates great advantages such as the reduction in mem-

ory accesses, improvements on data locality and significant workload reduction between consecutive *bottom-up* steps. In addition, we employ SIMD optimizations to fully exploit the hardware resources. Our optimized BFS implementation achieves 3 times speedup over the approach without RCM-reordering and shows a near-linear strong scaling on an ARMv8 system with NEON support. The optimizations also works well in real-world graphs. The resulting performance of 39.83 GTEPS is able to rank 83th on the Graph500 list in November 2020. Also, with an average power of 122 watts, our implementation achieves a performance of 326.48 MTEPS/W and ranks 3rd on the Green Graph500 list in November 2020. However, the top 2 systems utilize gpus to acquire higher power efficiency.


With the performance improvements for the BFS algorithm, the results of this work demonstrate potential performance improvement for more complex traversing-based graph algorithms, such as shortest paths, connected components, and spanning trees. Therefore, this work is also valuable to accelerate many real-world graph applications and achieve higher throughput for graph processing.

References

1. John Scott. Social network analysis. *Sociology*, 22(1):109–127, 1988.
2. Yankai Lin, Zhiyuan Liu, Maosong Sun, Yang Liu, and Xuan Zhu. Learning entity and relation embeddings for knowledge graph comple-

- tion. In *Twenty-ninth AAAI conference on artificial intelligence*, 2015.
3. Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. The pagerank citation ranking: Bringing order to the web. Technical report, Stanford InfoLab, 1999.
 4. Michael J Pazzani and Daniel Billsus. Content-based recommendation systems. In *The adaptive web*, pages 325–341. Springer, 2007.
 5. Giorgio Gallo and Stefano Pallottino. Shortest path algorithms. *Annals of operations research*, 13(1):1–79, 1988.
 6. Daniel S. Hirschberg, Ashok K. Chandra, and Dilip V. Sarwate. Computing connected components on parallel computers. *Communications of the ACM*, 22(8):461–464, 1979.
 7. Ronald L Graham and Pavol Hell. On the history of the minimum spanning tree problem. *Annals of the History of Computing*, 7(1):43–57, 1985.
 8. Linton C Freeman. Centrality in social networks conceptual clarification. *Social networks*, 1(3):215–239, 1978.
 9. Koji Ueno, Toyotaro Suzumura, Naoya Maruyama, Katsuki Fujisawa, and Satoshi Matsuoka. Extreme scale breadth-first search on supercomputers. In *2016 IEEE International Conference on Big Data (Big Data)*, pages 1040–1047. IEEE, 2016.
 10. Joseph E Gonzalez, Yucheng Low, Haijie Gu, Danny Bickson, and Carlos Guestrin. Powergraph: Distributed graph-parallel computation on natural graphs. In *Presented as part of the 10th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 12)*, pages 17–30, 2012.
 11. Amir Hossein Nodehi Sabet, Junqiao Qiu, and Zhijia Zhao. Tigr: Transforming irregular graphs for gpu-friendly graph processing. *ACM SIGPLAN Notices*, 53(2):622–636, 2018.
 12. Wai-Hung Liu and Andrew H Sherman. Comparative analysis of the cuthill–mckee and the reverse cuthill–mckee ordering algorithms for sparse matrices. *SIAM Journal on Numerical Analysis*, 13(2):198–213, 1976.
 13. Wing-Man Chan and Alan George. A linear time implementation of the reverse cuthill–mckee algorithm. *BIT Numerical Mathematics*, 20(1):8–14, 1980.
 14. Ariful Azad, Mathias Jacquelin, Aydin Buluç, and Esmond G Ng. The reverse cuthill–mckee algorithm in distributed-memory. In *2017 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 22–31. IEEE, 2017.
 15. Konstantinos I Karantasis, Andrew Lenharth, Donald Nguyen, Mara J Garzaran, and Keshav Pingali. Parallelization of reordering algorithms for bandwidth and wavefront reduction. In *SC’14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 921–932. IEEE, 2014.
 16. Neon technology. <https://developer.arm.com/architectures/instructionsets/simd-isas/neon>, 2020.
 17. Hwajeong Seo, Zhe Liu, Johann Großschädl, and Howon Kim. Efficient arithmetic on arm-neon and its application for high-speed rsa implementation. *Security and Communication Networks*, 9(18):5401–5411, 2016.
 18. Danilo Câmara, Conrado PL Gouvêa, Julio López, and Ricardo Dahab. Fast software polynomial multiplication on arm processors using the neon engine. In *International Conference on Availability, Reliability, and Security*, pages 137–154. Springer, 2013.
 19. Graph500. <https://graph500.org>, 2020.
 20. Twitter follower network. <https://snap.stanford.edu/data/twitter-2010.html>, 2020.
 21. J Yang and J Leskovec. Defining and evaluating network communities based on ground-truth. In *2012 IEEE 12th International Conference on Data Mining*, pages 745–754, 2012.
 22. David A Bader and Kamesh Madduri. Designing multithreaded algorithms for breadth-first search and st-connectivity on the cray mta-2. In *2006 International Conference on Parallel Processing (ICPP’06)*, pages 523–530. IEEE, 2006.
 23. Virat Agarwal, Fabrizio Petrini, Davide Pasetto, and David A Bader. Scalable graph exploration on multicore processors. In *SC’10: Proceedings of the 2010 ACM/IEEE International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–11. IEEE, 2010.
 24. Rudolf Berrendorf and Mathias Makulla. Level-synchronous parallel breadth-first search algorithms for multicore and multiprocessor systems. 2014.
 25. Yuichiro Yasui, Katsuki Fujisawa, and Kazushige Goto. Numa-optimized parallel breadth-first search on multicore single-node system. In *2013 IEEE International Conference on Big Data*, pages 394–402. IEEE, 2013.
 26. Yuichiro Yasui, Katsuki Fujisawa, and Yukinori Sato. Fast and energy-efficient breadth-first search on a single numa system. In *International Supercomputing Conference*, pages 365–381. Springer, 2014.
 27. Aydin Buluç and Kamesh Madduri. Parallel breadth-first search on distributed memory systems. In *Proceedings of 2011 International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–12, 2011.
 28. Andy Yoo, Edmond Chow, Keith Henderson, William McLendon, Bruce Hendrickson, and Umit Catalyurek. A scalable distributed parallel breadth-first search algorithm on bluegene/l. In *SC’05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing*, pages 25–25. IEEE, 2005.
 29. Aydin Buluç, Scott Beamer, Kamesh Madduri, Krste Asanovic, and David Patterson. Distributed-memory breadth-first search on massive graphs. *arXiv preprint arXiv:1705.04590*, 2017.
 30. Scott Beamer, Krste Asanovic, and David Patterson. Direction-optimizing breadth-first search. In *SC’12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis*, pages 1–10. IEEE, 2012.
 31. Scott Beamer, Aydin Buluç, Krste Asanovic, and David Patterson. Distributed memory breadth-first search revisited: Enabling bottom-up search. In *2013 IEEE International Symposium on Parallel & Distributed Processing, Workshops and Phd Forum*, pages 1618–1627. IEEE, 2013.
 32. Grzegorz Malewicz, Matthew H Austern, Aart JC Bik, James C Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 135–

- 146, 2010.
33. Yucheng Low, Joseph Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, and Joseph M Hellerstein. Distributed graphlab: A framework for machine learning in the cloud. *arXiv preprint arXiv:1204.6078*, 2012.
 34. Aapo Kyrola, Guy Blelloch, and Carlos Guestrin. Graphchi: Large-scale graph computation on just a {PC}. In *Presented as part of the 10th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 12)*, pages 31–46, 2012.
 35. Xiaowei Zhu, Wentao Han, and Wenguang Chen. Gridgraph: Large-scale graph processing on a single machine using 2-level hierarchical partitioning. In *2015 {USENIX} Annual Technical Conference ({USENIX}{ATC} 15)*, pages 375–386, 2015.
 36. Jin Zhao, Yu Zhang, Xiaofei Liao, Ligang He, Bingsheng He, Hai Jin, Haikun Liu, and Yicheng Chen. Graphm: an efficient storage system for high throughput of concurrent graph processing. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–14, 2019.
 37. Hao Wei, Jeffrey Xu Yu, Can Lu, and Xuemin Lin. Speedup graph processing by graph ordering. In *Proceedings of the 2016 International Conference on Management of Data*, pages 1813–1828, 2016.
 38. Ch H Papadimitriou. The np-completeness of the bandwidth minimization problem. *Computing*, 16(3):263–270, 1976.
 39. Elizabeth Cuthill and James McKee. Reducing the bandwidth of sparse symmetric matrices. In *Proceedings of the 1969 24th national conference*, pages 157–172, 1969.
 40. SW Sloan. An algorithm for profile and wavefront reduction of sparse matrices. *International Journal for Numerical Methods in Engineering*, 23(2):239–251, 1986.




Zite Jiang received the BSc degree in computer science from Hangzhou Dianzi University, China, in 2019. He is currently working toward the MS Degree in computer science at Institute of Computing Technology, Chinese Academy of Sciences, China. His research interest is graph computing.



Tao Liu received the BSc degree in computer science and technology from the University of Science and Technology Beijing in 2015, and the MS Degree in computer science at Insti-


tute of Computer Technology, Chinese Academy of Sciences, China in 2018. His re-

search interest is distributed



graph computing. Shuai Zhang received the BSc degree in computer science and technology from Nanjing University of Posts and Telecommunications, China in 2018. He is currently working toward the Phd Degree in computer science at Insti-

tute of Computer Technology, Chinese Academy of Sciences, China. His research interest is distributed graph computing and protein structure prediction.



Mengting Yuan received the MS and Ph.D. degree from Wuhan University, Wuhan, China, in 2006. He is currently an associate professor at School of Computer Science, Wuhan University. His current research interests are programming languages, compiler technology, and software engineering.



Haihang You is a professor at Institute of Computing Technology, Chinese Academy of Sciences, China. Prior joining ICT, Dr. You was research scientist at National Institute of Computational Sciences at Oak Ridge National Laboratory and Innovative Computing Laboratory at University of Tennessee, America. Dr. You's research interest is in the field of high performance computing and high throughput computing, specifically parallel algorithm, numerical algorithm, performance optimization and autotuning.