



SAGE: A Storage-Based Approach for Scalable and Efficient Sparse Generalized Matrix-Matrix Multiplication

Myung-Hwan Jang*
Department of Computer Science
Hanyang University
Seoul, Republic of Korea
sugichiin@hanyang.ac.kr

Yunyong Ko*
Department of Computer Science
Hanyang University
Seoul, Republic of Korea
koyunyong@hanyang.ac.kr

Hyuck-Moo Gwon
Department of Computer Science
Hanyang University
Seoul, Republic of Korea
howling6@hanyang.ac.kr

Ikhyeon Jo
Department of Computer Science
Hanyang University
Seoul, Republic of Korea
chilyouth@hanyang.ac.kr

Yongjun Park†
Department of Computer Science
Yonsei University
Seoul, Republic of Korea
yongjunpark@yonsei.ac.kr

Sang-Wook Kim†
Department of Computer Science
Hanyang University
Seoul, Republic of Korea
wook@hanyang.ac.kr

ABSTRACT

Sparse generalized matrix-matrix multiplication (SpGEMM) is a fundamental operation for real-world network analysis. With the increasing size of real-world networks, the single-machine-based SpGEMM approach cannot perform SpGEMM on large-scale networks, exceeding the size of main memory (i.e., *not scalable*). Although the distributed-system-based approach could handle large-scale SpGEMM based on multiple machines, it suffers from severe inter-machine communication overhead to aggregate results of multiple machines (i.e., *not efficient*). To address this dilemma, in this paper, we propose a novel *storage-based* SpGEMM approach (**SAGE**) that stores given networks in storage (e.g., SSD) and loads only the necessary parts of the networks into main memory when they are required for processing via a 3-layer architecture. Furthermore, we point out three challenges that could degrade the overall performance of SAGE and propose three effective strategies to address them: (1) *block-based workload allocation* for balancing workloads across threads, (2) *in-memory partial aggregation* for reducing the amount of unnecessarily generated storage-memory I/Os, and (3) *distribution-aware memory allocation* for preventing unexpected buffer overflows in main memory. Via extensive evaluation, we verify the superiority of SAGE over existing SpGEMM methods in terms of scalability and efficiency.

CCS CONCEPTS

• Information systems → Data management systems.

*Two first authors have contributed equally to this work.

†Corresponding authors.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.
CIKM '23, October 21–25, 2023, Birmingham, United Kingdom
© 2023 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 979-8-4007-0124-5/23/10...\$15.00
<https://doi.org/10.1145/3583780.3615044>

KEYWORDS

sparse matrix multiplication, real-world graphs, network analysis

ACM Reference Format:

Myung-Hwan Jang, Yunyong Ko[1], Hyuck-Moo Gwon, Ikhyeon Jo, Yongjun Park[2], and Sang-Wook Kim. 2023. SAGE: A Storage-Based Approach for Scalable and Efficient Sparse Generalized Matrix-Matrix Multiplication. In *Proceedings of the 32nd ACM International Conference on Information and Knowledge Management (CIKM '23)*, October 21–25, 2023, Birmingham, United Kingdom. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3583780.3615044>

1 INTRODUCTION

Graphs are widely used to model real-world networks, where a node represents an object and an edge does the *pair-wise* relationship between two objects. To analyze real-world networks and discover useful knowledge, many graph algorithms have been studied [5, 14, 16, 20–24, 26, 34, 39, 42, 45], where a graph is generally represented as a *dense* matrix (i.e., 2-D array), and each element $e_{i,j} = 1$ if an edge exists between nodes i and j , and $e_{i,j} = 0$ otherwise. *Generalized matrix-matrix multiplication* (GEMM), one of the key operations in these algorithms, plays a fundamental role to diffuse the node information via network connectivity. Recently, with the growing size of real-world networks, it has become more crucial to efficiently perform GEMM on large-scale graphs [3, 13, 50].

Performing GEMM on real-world graphs represented as dense matrices, however, often requires an unnecessarily large amount of space and computational cost [36, 40]. This is because they tend to follow a *power-law degree distribution* [31] – i.e., a majority of nodes have only a few edges, while a small number of nodes have a large number of edges – so that the number of non-zero elements (existing edges) is much smaller than that of zero elements (non-existing edges) in the matrix for a real-world graph. Figure 1(a) shows the power-law degree distribution of Wikipedia [18] and Figure 1(b) does the matrix density of Wiki-Vote [29, 30], where each point represents a non-zero element, respectively.

To address this issue, *sparse* GEMM (*SpGEMM*) methods have been widely studied [1–3, 9–12, 37, 40, 48], where a graph is represented as a *sparse matrix* containing only existing edges. On one hand, *single-machine-based* methods [1, 36, 46, 47] aim to perform

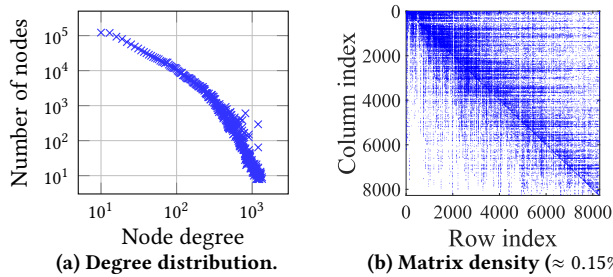


Figure 1: (a) Power-law degree distribution and (b) visualized matrix density of real-world graphs.

SpGEMM efficiently on a single machine by considering the characteristics of real-world graphs, where they assume that the sparse matrix can fit in main memory. Thus, they are unable to process SpGEMM on large-scale graphs exceeding the size of the main memory (i.e., **not scalable**). On the other hand, *distributed-system-based* methods [3, 6, 10, 43] are able to perform SpGEMM on large-scale graphs, exceeding the size of the main memory, by using multiple machines concurrently. These methods, however, require a substantial amount of *inter-machine communication* overhead to aggregate the results from multiple machines, which is often more than 50% of the entire process [3, 25], and a lot of *costs and efforts* to maintain complex distributed systems (i.e., **not efficient**).

To tackle these two limitations together, in this paper, we propose a novel *storage-based* approach for large-scale SpGEMM, named **Storage-based Approach for SpGEMM (SAGE)**. SAGE stores the entire graphs in external storage of a single machine and loads only some parts of the graphs when necessary into the main memory to process them. For efficient data transfers between storage and main memory, we design SAGE with a 3-layer architecture, i.e., (1) storage, (2) in-memory, and (3) operation layers, where each layer closely interacts with other layers. Thus, SAGE (1) is able to process large-scale graphs, exceeding the size of the main memory by using *sufficient capacity of external storage* (i.e., **scalable**) and (2) requires only *intra-machine communication* overhead (storage-memory I/Os), much smaller than the inter-machine communication overhead (i.e., **efficient**).

To improve the performance of SAGE, it is crucial to handle *storage-memory I/Os* efficiently, which are inevitably generated due to the limited size of main memory on a single machine. To this end, first, we point out three important challenges that could lead to serious performance degradation in SAGE: **(C1) Workload balancing**, i.e., *How to distribute workloads of SpGEMM evenly across multiple threads?*, **(C2) Intermediate handling**, i.e., *How to handle a large amount of the storage-memory I/Os generated by the intermediate results of SpGEMM?*, and **(C3) Memory management**, i.e., *How to allocate the main memory space to three buffers to reduce storage-memory I/Os?* (for two input and one output matrices). Then, we propose three effective strategies to address these challenges: (1) *block-based workload allocation* to evenly distribute workloads of SpGEMM into multiple threads for C1, (2) *in-memory partial aggregation* to reduce the amount of unnecessarily generated storage-memory I/Os for C2, and (3) *distribution-aware memory allocation* to adjust the proportions of the three buffers, based on the characteristics of real-world graphs for C3.

Table 1: Comparison of existing SpGEMM methods with SAGE based on the three performance-critical challenges

Method	(C1) WB	(C2) IH	(C3) MM
Intel MKL [46]	✓		
KNL-SpGEMM [36]	✓		
IA-SpGEMM [47]	✓		
SpSUMMA [6]	✓		
Graphulo [17]	✓	✓	
Split-3D [3]	✓		
gRRp [10]	✓	✓	
SAGE (proposed)	✓	✓	✓

We note that although there exists several existing storage-based methods [15, 18, 28, 41, 49], they focus only on the multiplication between sparse matrix (SpM) and dense vector (DV), i.e., $SpMV$ but not $SpGEMM$ that we focus on. More specifically, the existing storage-based SpMV methods handle only “one” large sparse matrix and small dense vectors ($SpM \times DV = DV$), where the dense vectors are much smaller than a sparse matrix ($SpM \gg DV$). On the other hand, SAGE handles “three” large sparse matrices ($SpM \times SpM = SpM$), which implies that it is more challenging to efficiently process storage-memory I/Os in the storage-based SpGEMM than the storage-based SpMV. To the best of our knowledge, this is the first work to successfully perform large-scale SpGEMM by using external storage on a single machine. We believe that SAGE could be a practical alternative to researchers and practitioners who aim to analyze large-scale real-world graphs.

The main contributions of this work are summarized as follows:

- **New Approach:** We propose a new *storage-based* approach for large-scale SpGEMM, SAGE to address the limitations of existing single-machine-based (*scalability*) and distributed-system-based (*efficiency*) approaches simultaneously.
- **Effective Strategies:** We point out three performance-critical challenges and propose effective strategies to tackle them: (1) *block-based workload allocation*, (2) *in-memory partial aggregation*, and (3) *distribution-aware memory allocation*.
- **Extensive Evaluation:** We conduct extensive evaluation, which demonstrates that (1) SAGE is able to perform SpGEMM on large-scale graphs, surpassing the limits of the existing single-machine-based methods, (2) SAGE achieves high SpGEMM performance comparable to or better than the distributed-system-based methods, and (3) each of the proposed strategies of SAGE is effective in improving the performance of SAGE.

2 RELATED WORK

In this section, we review existing SpGEMM methods. Table 1 compares SAGE with existing methods in terms of the three performance-critical challenges: (C1) workload balancing (WB), (C2) intermediate handling (IH), and (C3) memory management (MM).

Single-machine-based approach. A single-machine-based approach [1, 36, 46, 47] aims to improve the performance of SpGEMM, based on the properties of real-world networks on a single machine (e.g., node degree distribution). Intel MKL [46] is an open source library supporting a series of math functions such as SpGEMM and SpMV, which are optimized for Intel CPU architecture. Nagasaka

Table 2: Notations and their descriptions

Notation	Description
M_{in}^1, M_{in}^2	input matrices of SpGEMM
M_{out}	output matrix ($M_{in}^1 \times M_{in}^2$)
$M(i, ;)$	i -th row of matrix M
B_{in}^1, B_{in}^2	buffers for input matrices in the main memory
B_{out}	buffer for output matrix in the main memory
T_{buf}^1, T_{buf}^2	buffer index tables for input matrices M_{in}^1, M_{in}^2
T_{obj}^1, T_{obj}^2	object index tables for input matrices M_{in}^1, M_{in}^2

et al. [36] proposed memory management to keep output results and thread scheduling strategies to resolve the bottlenecks of the entire SpGEMM process. Xie et al. [47] proposed an input-aware auto-tuning framework for SpGEMM (IA-SpGEMM) to reduce the overhead of memory access and sparse accumulation in SpGEMM. These single-machine-based methods, however, assume that the entire input matrices can be loaded in the main memory (i.e., in-memory based SpGEMM) because the input matrices have only existing edges in SpGEMM (i.e., sparse matrix). Therefore, they are not able to process SpGEMM on large-scale graphs exceeding the size of main memory in a single machine.


There are some works using external storage to process large-scale graphs on a single machine [15, 18, 28, 49, 50]. However, they focus on improving *sparse matrix and dense vector multiplication* (SpMV) rather than SpGEMM that our work focuses on. Note that we can obtain the exactly same results as SpGEMM by iteratively performing SpMVs in concept; however, this way is not suitable for processing SpGEMM since the operations for dense vectors can lead to serious performance degradation due to a large amount of memory usage and computation overhead [3, 6, 40].

Distributed-system-based approach. A distributed-system-based approach [3, 6, 10, 17, 43] aims to process SpGEMM on large-scale graphs that cannot be loaded on the main memory of a single machine, via multiple machines in a distributed cluster. This approach splits and stores input matrices into multiple machines, and then processes them in parallel. Buluç and Gilbert [6] proposed SpSUMMA that parallelizes SpGEMM by partitioning input matrices into 2-D grids and distributing them to multiple processors. Hutchison et al. [17] proposed a distributed SpGEMM method, Graphulo that considers the *locality* of data stored in a distributed database to reduce the inter-machine communication overhead. Azad et al. [3] proposed Split-3D that splits input/output matrices into 3-D grids for parallelization of computations and communications. Demirci and Aykanat [10] proposed gRRp, a bipartite graph-based partitioning method for balancing workloads among multiple machines.

Though the distributed-system-based approach is able to perform SpGEMM on large-scale graphs, a substantial amount of *inter-machine communication overhead* is required to aggregate the results from multiple machines, which is non-trivial (more than 50%) [3, 25]. It also requires a lot of *costs and efforts* to maintain the high-performance infrastructure of distributed systems such as fault tolerance and graph partitioning for distributed machines.

3 SAGE: PROPOSED FRAMEWORK

We present a novel storage-based SpGEMM method, named **Storage-based** approach for SpGEMM (**SAGE**). First, we describe the notations and the problem statement (Section 3.1). Then, we present the

**Figure 2: The process of row-wise product.**

architecture and core algorithms of SAGE (Section 3.2), and three novel strategies to address performance-critical challenges (Section 3.3). Finally, we analyze the complexity of SAGE (Section 3.4).

3.1 Notations and Problem Statement

3.1.1 Notations. Table 2 shows the notations used in this paper. We denote two input sparse matrices for SpGEMM as M_{in}^1 and M_{in}^2 , and the output matrix as M_{out} (i.e., $M_{in}^1 \times M_{in}^2$). In a matrix M , $M(i, ;)$ is an i -th row, and $M(i, j)$ is a j -th element in the i -th row, where the length of each row ($|M(i, ;)|$) can vary across rows because M is a sparse matrix. For clarity, we denote buffers and index tables in the main memory as B and T , respectively. B_{in}^1 and B_{in}^2 are the input buffers to hold necessary parts of the input matrices, and B_{out} is the output buffer for the intermediate results.

3.1.2 Problem statement. Given two input sparse matrices of M_{in}^1 and M_{in}^2 , SAGE aims to produce the output sparse matrix M_{out} by performing SpGEMM between the two sparse matrices. When performing SpGEMM, three different types of products can be considered: (1) *inner product* (row-by-column product), (2) *outer product* (column-by-row product), and (3) *row-wise product* (row-by-row product). Although these three products generate exactly the same result, M_{out} , they require different computation overhead and memory space. Specifically, the row-wise product requires not only much smaller computation overhead (e.g., index-matching) than the inner product, but also less memory usage than the outer products [4, 44]. Thus, we adopt the *row-wise product* for performing SpGEMM in SAGE, following existing algorithms [1, 36, 46, 47]. In the row-wise product, the i -th row of the output matrix is computed by Eq. 1,

$$M_{out}(i, ;) = \sum_{j=1}^n M_{in}^1(i, j) \times M_{in}^2(j, ;), \quad (1)$$

where n indicates the number of rows in the input matrices. Figure 2 illustrates the process of the row-wise product computing the i -th row of the output matrix M_{out} in performing SpGEMM. First, every j -th element of the i -th row in the former input matrix, $M_{in}^1(i, j)$, is multiplied by its corresponding j -th row in the latter input matrix, $M_{in}^2(j, ;)$, (highlighted in the same color in Figure 2). Then, the intermediate results are aggregated to compute the final result of the i -th row of the output matrix, $M_{out}(i, ;)$. This process is repeated until all rows of the output matrix are computed.

3.2 Architecture and Algorithm

SAGE stores the entire graphs in storage (e.g., SSD), loads only the necessary parts of the graphs into the main memory, and processes them in parallel by using multiple threads. Thus, *storage-memory I/Os* are inevitably generated during the process of storage-based SpGEMM, which could become a *bottleneck* of the overall SpGEMM process. To handle the inevitable overhead efficiently, in this section, we propose a novel 3-layer architecture for efficient data transfers




Figure 3: Overview of SAGE with the 3-layer architecture (storage, in-memory, and operation layers).

between storage and main memory (Section 3.2.1), and describe how SAGE processes SpGEMM based on the architecture (Section 3.2.2).

3.2.1 Architecture. As illustrated in Figure 3, SAGE consists of three layers: (1) storage, (2) in-memory, and (3) operation layers, where each layer closely interacts with other layers to efficiently process storage-memory I/Os.

Storage layer. This layer is in charge of managing the data stored in external storage: two input matrices (M_{in}^1 and M_{in}^2), an output matrix (M_{out}), and intermediate results ($M_{out}(i, :)$). This storage layer handles read/write requests from the in-memory layer. We define a fixed-size unit standard I/O as a *block* and store the input matrices in a set of blocks. Each block usually contains multiple objects, each of which stores each row of the input matrices (i.e., a node and its related edges). In a real-world graph following the power-law degree distribution, there are a few nodes with a very high degree, which are too large to be stored in a block. For the exceptional nodes, we divide and store them in multiple blocks.

In-memory layer. This layer is in charge of (1) loading blocks required for the next operations into main memory, (2) aggregating the intermediate results from the operation layer to generate the final result, and (3) writing the final result to the storage layer. For this process, we define three buffers and four index tables:

- Input buffers (B_{in}^1 and B_{in}^2): storing rows of the input matrices (i.e., blocks) required for the row-wise product.
- Output buffer (B_{out}): storing the output result of a row-wise product computed from the operation layer.
- Buffer index tables (T_{buf}^1 and T_{buf}^2): indicating blocks loaded in the two input buffers of B_{in}^1 and B_{in}^2 .
- Object index tables (T_{obj}^1 and T_{obj}^2): indicating the mappings of blocks and objects in external storage.

The buffer index tables (T_{buf}^*) are referred to check whether the required rows are already loaded on the main memory. If the rows are not on the main memory, this layer requests ‘the block having the required rows’ to the storage layer. In the object index tables (T_{obj}^*), the i -th column represents the indices of objects stored in the i -th block. For memory efficiency, we sort objects by their indices and store the only two indices of the first and last objects. The object index tables (T_{obj}^*) are pinned in the main memory to quickly load the necessary blocks. Note that the pinned object index tables do not affect the entire performance of SAGE since it requires only a small amount of memory less than 0.01% of the original graph.

Operation layer. This layer is in charge of performing actual multiplications using the data in B_{in}^1 and B_{in}^2 . This layer is performed as follows: (1) reading two rows to be multiplied in B_{in}^1 and B_{in}^2 by referring to T_{buf}^1 and T_{buf}^2 , (2) performing a row-wise product between them, and (3) writing the result (the intermediate result of $M_{out}(i, :)$) to the output buffer. This process is repeated until the input matrices are processed. The intermediate results are aggregated in the in-memory layer to generate the final result $M_{out}(i, :)$, and then stored in the external storage.

3.2.2 Algorithm and performance consideration. Next, let us describe how SAGE performs SpGEMM based on the 3-layer architecture. Algorithm 1 shows the entire process of SAGE (logical view). Given the two input matrices, which are split and stored in multiple blocks, and their object index tables, (**Data loading**) SAGE first loads the two rows to be multiplied into the input buffers (LoadData(\cdot) in lines 10-18). Specifically, (1) the operation layer requests the two rows to the in-memory layer, (2) the in-memory layer accesses the input buffer tables to check whether the blocks having the required rows are already in the input buffers, and (3) the in-memory layer sends the blocks having the required rows to the

Algorithm 1: SAGE: STORAGE-BASED SPGEMM

Input: Input matrices M_{in}^1, M_{in}^2 , Object index tables T_{obj}^1, T_{obj}^2
Output: Output matrix \mathcal{H}^*

```

1 Function SAGE( $M_{in}^1, M_{in}^2, T_{obj}^1, T_{obj}^2$ ):
2    $T_{buf}^1, T_{buf}^2, M_{out}, t\_row \leftarrow \emptyset$ 
3   for  $M_{in}^1(i, :) \in M_{in}^1$  do // processed in multiple threads
4     LoadData( $i, T_{obj}^1, T_{buf}^1$ )
5     for  $M_{in}^2(j, :) \in M_{in}^2$  do // i-th row computation
6       LoadData( $j, T_{obj}^2, T_{buf}^2$ )
7        $B_{out} \leftarrow B_{out} \cup \{t\_row\}, t\_row \leftarrow M_{in}^1(i, j) \times M_{in}^2(j, :)$ 
8        $M_{out}(i, :) \leftarrow \sum_{tmp \in B_{out}} tmp$  // result aggregation
9   return  $M_{out}$ 
10 Function LoadData( $i, T_{obj}, T_{buf}$ ):
11 for  $j = 0, 1, \dots, |T_{obj}| - 1$  do
12    $(n_1, n_2) \leftarrow T_{obj}[j]$  // first/last nodes in a block
13   if  $n_1 \leq i \leq n_2$  then
14      $b\_index \leftarrow j$ 
15     break
16 if  $b\_index \notin T_{buf}$  then // storage-memory I/O
17    $T_{buf} \leftarrow T_{buf} \cup \{b\_index\}$ 
18 return
    
```

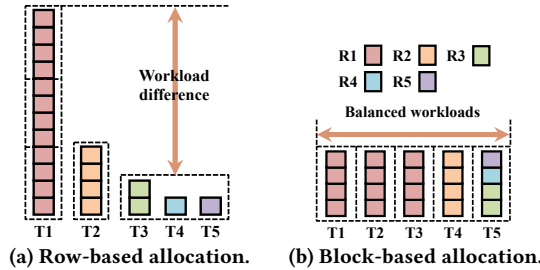


Figure 4: Comparison of (a) the row-based workload allocation and (b) the block-based workload allocation.

operation layer if they are already in the input buffers, or otherwise, requests the blocks to the storage layer. (**Multiplication**) Then, for the two rows loaded in the main memory, SAGE computes the i -th row of the output matrix, $M_{out}(i, ;)$, by Eq. 1 (lines 3-8), where the intermediate results (t_row) are stored in B_{out} and aggregated to generate the final result $M_{out}(i, ;)$ in the in-memory layer. Finally, the in-memory layer sends (i.e., writes) the final results of the output matrix to the storage layer.

SAGE employs a *multi-thread based row-wise product* to accelerate SpGEMM. Basically, SAGE assigns each row of the former input matrix to each thread, rather than the latter input matrix. Thus, the outer **for** loop in Algorithm 1 (lines 3-8) is performed by multiple threads in parallel. In this way, SAGE can benefit from the parallelism of SpGEMM since the computation of each thread is *independent* to each other. Furthermore, to maximize the parallelism, we adopt a *block-based workload allocation* strategy to distribute workloads into multiple threads evenly (elaborated in Section 3.3.1).

3.3 Challenges and Strategies

We then point out three critical challenges leading to serious performance degradation in SAGE and propose effective strategies to address them: (1) block-based workload allocation, (2) in-memory partial aggregation, and (3) distribution-aware memory allocation.

3.3.1 Workload balancing. To efficiently perform SpGEMM on large-scale graphs, existing works generally adopt multi-thread based computation [1, 3, 6, 36, 46, 47], where each thread is in charge of computing different rows (nodes) of the output matrix in parallel (i.e., row-based workload allocation). Thus, the workload of each thread is decided, depending on the number of elements of each row. In real-world graphs, however, nodes (rows) have a quite different number of edges (elements) as real-world graphs tend to follow the power-law degree distribution. It implies that the workloads of threads could be significantly different from each other. Such workload differences among threads can adversely affect the overall performance of SpGEMM by decreasing the benefit of parallel processing. Figure 4(a) shows an example of row-base workload allocation for five rows with different lengths (R1-R5) stored in five blocks (dotted boxes), where a large amount of workloads for (R1), stored in three blocks, are assigned on a single thread (T1) (i.e., over-loaded), while only a small amount of workloads for (R3, R4, and R5), stored in a single block, are split and assigned on three threads (T3-T5) (i.e., under-loaded). Thus, it is important to *distribute workloads evenly into threads* to maximize the benefit of parallelizing SpGEMM (**Challenge 1**).




Figure 5: The amount of storage-memory I/Os could be significantly reduced by the in-memory partial aggregation.

To tackle the challenge of workload balancing, we propose a simple yet effective strategy that distributes workloads into multiple threads in a block-based manner (**block-based workload allocation**). Specifically, SAGE equipped with the block-based workload allocation assigns “the same number of blocks” to each thread, rather than the same number of rows with different numbers of edges, thereby distributing the workloads into threads evenly. Note that this strategy is applied to the outer **for** loop (line 3) in Algorithm 1, where the rows of the former input matrix stored in each block are assigned together to the same thread. Figure 4(b) shows an example of our block-based workload allocation, where the large row (R1), stored in three blocks, is split and assigned on three threads (T1-T3), and the small rows (R3-R5), grouped in a single block, are assigned together on a single thread (T5). As a result, all five threads have the same amount of workloads (one block per thread).

Theoretically, our block-based strategy improves the workload difference among threads, compared to the row-based one, $O(|r|) \rightarrow O(b)$, where $|r|$ is the # of elements in a row r and b (fixed) is # of elements in a block (in general, $|r|_{max} \gg b$).

3.3.2 Intermediate handling. As described in Section 3.2.2, the i -th row of the output matrix ($M_{out}(i, ;)$) is computed by aggregating the intermediate results of multiple row-wise products – the inner **for** loop (lines 5-8) in Algorithm 1. For the aggregation, the intermediate results should be loaded on the output buffer. The output buffer, however, may not be sufficient to store all intermediates due to the limited size of the main memory, and it could be overflowed by a large amount of intermediates. In that case, the intermediates should be flushed into the storage multiple times to process the remaining row-wise products. They are then loaded back into the main memory to be aggregated in the end once the remaining row-wise products are completed. In other words, a large amount of storage-memory I/Os could be generated by the intermediates, which could degrade the performance of SpGEMM significantly. Therefore, it is critical to *handle a large amount of storage-memory I/Os generated by the intermediates results* (**Challenge 2**).

To address this challenge, we propose to column-wisely aggregate the intermediate results of each row directly in the output buffer (**in-memory partial aggregation**). Figure 5 illustrates the effect of our in-memory partial aggregation strategy, where the intermediates with the same column indices are aggregated in the main memory before they are flushed to storage. Accordingly, SAGE is able to reduce a large amount of storage-memory I/Os unnecessarily generated by the intermediates. Furthermore, we implement

hash-based index matching in SAGE to quickly find and aggregate the intermediates with the same column index by following [36]. However, we have observed that SpGEMM on extremely large-scale graphs may still cause even the partially aggregated intermediates to be overflowed. In this case, SAGE stores them in storage and aggregate them after finishing all of the row-wise products.

3.3.3 Memory management. SAGE manages three buffers in the main memory (i.e., two input and an output buffers). Due to the limited size of the main memory, the amount of storage-memory I/Os can vary depending on the proportions of the three buffers. For example, (**Case 1**) if the input buffers are set too small, the data loaded into the input buffers are processed in only a few iterations and replaced frequently, incurring lots of storage-memory I/Os. On the other hand, (**Case 2**) if the output buffer is set too small, it is easily overflowed by a small number of intermediate results, which also causes storage-memory I/Os frequently. Therefore, it is crucial to *adjust the proportions of the three buffers* to reduce the amount of storage-memory I/Os (**Challenge 3**).

Toward this challenge, we propose a practical strategy that adjusts the proportions of the three buffers, based on the characteristics of real-world graphs (**distribution-aware allocation**). First, SAGE sets the size of the former input buffer B_{in}^1 as $(b \times t)$, where b is the size of a block and t is the number of threads. Note that the size of $(b \times t)$ is sufficient for the former input buffer to load the blocks that all threads can process simultaneously since SAGE allocates workloads to threads in a *block-based manner*.

Second, SAGE determines the size of the output buffer B_{out} . Unfortunately, it is infeasible to predict the exact size of the output row in advance as it depends on data characteristics. Some previous works [1, 36, 46, 47] have tried to pre-compute the theoretical maximum size of the intermediate results for each row, and set the output buffer size as the maximum size in order to prevent the output buffer from being overflowed. However, it is inappropriate for SpGEMM on real-world graphs following the power-law degree distribution since the space required for the intermediate results of a majority of rows is often much smaller than the theoretical maximum size (i.e., inefficient memory usage). Thus, we argue that the output buffer with a size much smaller than the maximum output size is practically sufficient to process the intermediate results in most cases. Based on this intuition, we can compute the maximum size of the intermediate results, max_row , and sets the size of the output buffer as $(\alpha \times max_row \times t)$, considering the degree distribution of real-world graphs (i.e., *distribution-aware allocation*), where α is the hyperparameter to adjust the output buffer size.

Lastly, SAGE allocates the remaining portion of the main memory to the latter input buffer B_{in}^2 . Formally, given the capacity of the main memory C , a block size b , the number of threads t , and the hyperparameter α , the main memory is allocated as:

$$SAGE_mem(C, b, t, \alpha) = \begin{cases} b \times t, & B_{in}^1 \\ C - (|B_{in}^1| + |B_{out}|), & B_{in}^2 \\ \alpha \times max_row \times t, & B_{out}. \end{cases} \quad (2)$$

3.4 Complexity Analysis

The computational cost of SAGE derives mainly from (1) loading input matrices, (2) row-wise products, (3) aggregating the intermediate results, and (4) storing the final results in the output matrix. For brevity, we consider SpGEMM between the two sparse matrices representing the same real-world graph $G = (N, E)$, where N and E is a set of nodes and edges, respectively. The input loading requires the time complexity of $O(|N|^2)$ for two input matrices. For $|N|$ many rows, the computational cost of each row-wise products is $O(|N| \times \tilde{r})$, where \tilde{r} is the average number of elements in each row, where $\tilde{r} \ll |N|$ since a real-world graph follows the power-law degree distribution. Thus, the overall complexity of row-wise products is $O(|N|^2 \times \tilde{r})$. The aggregation of the intermediate results requires $O(|N|^2 \times \tilde{r})$, i.e., (1) hash-based index matching and (2) element-wise aggregation. Lastly, storing the final results in the output matrix requires the complexity of $O(|N|^2 \times \tilde{r})$. As a result, the overall time complexity of SAGE is $O(|N|^2 \times \tilde{r})$, which is comparable to the time complexity of common SpGEMM methods. We will empirically evaluate the scalability of SAGE with the increasing size of graphs in Section 4.2.3.

4 EXPERIMENTS

In this section, we comprehensively evaluate SAGE by answering the following evaluation questions (EQs):

- **EQ1:** Does SAGE improve the performance of SpGEMM, compared to the existing single-machine-based methods?
- **EQ2:** Does SAGE improve the performance of SpGEMM, compared to the existing distributed-system-based methods?
- **EQ3:** How does the SpGEMM performance of SAGE scale up with the increasing size of graphs?
- **EQ4:** How effective are the proposed strategies of SAGE in improving the performance of SpGEMM?

4.1 Experimental Setup

Datasets. We use widely used 26 real-world datasets [27, 32] with various sizes, ranging from 10MB at the minimum to 60GB at the maximum, and synthetic datasets. For the synthetic datasets, we use the **Recursive MATrix generator (R-MAT)** [7] to generate synthetic graphs following a power-law degree distribution. R-MAT takes several parameters: (1) a benchmark type (Graph500, SSCA, or ER), (2) the scale of the graph n , (3) the edge factor e , and (4) the skewness parameters a , b , c , and d , determining the skewness of the generated graph. Based on the input parameters, R-MAT generates a matrix $M \in \mathbb{R}^{2^n \times 2^n}$ with $e \times 2^n$ elements. We use the same parameters as used in [3, 6, 10].

Evaluation protocol. To comprehensively evaluate SAGE, we consider two different types of SpGEMM: (1) SpGEMM between “the same two matrices” ($M \times M$) and (2) SpGEMM between “two different matrices” ($M_1 \times M_2$ or $M \times M^T$). The first type (1) is used in many graph algorithms such as finding all-pairs shortest paths [8], self-similarity joins [16], and summarization of sparse datasets [38], and the second type (2) is used in collaborative filtering [33], triangle counting [5], and similarity joins of two different sparse graphs [38]. For the evaluation metric, we use *the execution time of SpGEMM* of each method because the goal of this work is to improve the performance of graph algorithms by accelerating SpGEMM.

Table 3: Comparison of the single-machine-based approach with SAGE on 18 real-world datasets with various sizes.

Dataset	Poisson3D	Enron	Epinions	Sphere	Filter3D	598a	Torso2	Cop20k	Cage12
Intel MKL	0.1131	0.5792	0.6343	0.2521	0.3921	0.2660	0.0940	0.4120	0.3669
SAGE	0.3206	0.7497	0.7001	0.5335	0.7436	0.5626	0.3976	0.7433	0.6331
Difference (sec.)	0.2075↑	0.1705↑	0.0658↑	0.2814↑	0.3515↑	0.2966↑	0.3036↑	0.3313↑	0.2662↑
Dataset	Slashdot	Gowalla	Pocec	Youtube	Livejournal	Wikipedia	UK-2005	SK-2005	Yahoo
Intel MKL	1.9890	7.8299	12.8732	N/A, OOM	N/A, OOM	N/A, OOM	N/A, OOM	N/A, OOM	N/A, OOM
SAGE	1.7558	6.772	10.6391	38.734	94.8283	2,276	616	3,709	12,994
Difference (sec.)	0.2332↓	1.0579↓	2.2341↓	-	-	-	-	-	-

Competing methods. We compare SAGE with the following SpGEMM methods in our experiments.

- **Intel MKL** [46]: the Intel MKL, a state-of-the-art single-machine-based method, is the Intel CPU optimized open-source library supporting a series of math functions such as SpGEMM and SpMV, which was also used in [1, 36, 40, 44, 47].
- **SpSUMMA** [6]: SpSUMMA is a distributed-system-based method that parallelizes SpGEMM by partitioning input matrices into 2-D grids and processing them using multiple machines.
- **Graphulo** [17]: Graphulo stores input matrices in a distributed system exploiting the data locality to reduce the inter-machine communication overhead.
- **gRRp** [10]: gRRp adopts a bipartite graph-based partitioning method for balancing workloads among multiple machines in a distributed system.

Implementation details. We run all experiments on a single machine equipped with Intel i7-7700K CPU, 64GB main memory, and 4TB SSD as the storage. We set the number of threads t as 4 to fully utilize all four physical cores in the CPU. As explained in Section 3.3.3, SAGE allocates the main memory into three buffers (i.e., two input and one output matrices) by the *distributed-aware allocation* (Eq. 2), where we set the hyperparameter α as 12.5%. We will empirically evaluate the impact of the hyperparameter α on the performance of SAGE in Section 4.2.4.

4.2 Experimental Results

4.2.1 Single-machine-based approach (EQ1). We first compare SAGE with the single-machine-based approach, Intel MKL [46], in terms of the SpGEMM performance (i.e., execution time).

Setup. To comprehensively evaluate SAGE, we use 18 real-world datasets with various sizes, 12 relatively small datasets that can be loaded on the main memory and 6 large datasets that exceed the limit of the main memory capacity. Then, we (1) perform SpGEMM using each method on the datasets for 5 times and (2) measure the average execution time of SpGEMM of each method. Note that we limit the main memory size to 16GB to rigorously evaluate the ability of SAGE to handle the storage-memory I/Os.

Results and analysis. Table 3 shows the results. First, SAGE provides the comparable performance to or even “better” than the Intel MKL, the in-memory based method. Specifically, SAGE (1) finishes SpGEMM slightly slower (only less than 0.4 sec.) than the Intel MKL in 9 datasets which are small enough to load the entire data on the main memory and (2) even outperforms the Intel MKL in

Table 4: Statistics of synthetic datasets

Benchmark	Scale n	# of nodes	# of edges
Graph500 (Graphulo and gRRp)	15	32.7K	524K
	16	65.6K	1.04M
	17	131K	2.09M
	18	262K	4.19M
SSCA (SpSUMMA)	21	2.09M	16.7M
	22	4.19M	33.5M
	23	8.38M	67.1M
	24	16.7M	134M

Slashdot, Gowalla, and Pocec datasets. We highlight that these improvements over the Intel MKL are significant, given that the Intel MKL is an *in-memory based SpGEMM* method loading and processing the entire data on the main memory (i.e., no storage-memory I/Os) and specially designed for the Intel CPU architecture, while SAGE has to handle the inherent additional overhead for data transfers between the main memory and external storage as described in Section 3.2. Second, SAGE successfully performs SpGEMM on very large real-world datasets (i.e., from Youtube to Yahoo datasets in Table 3) that the Intel MKL could not handle due to the limited capacity of the main memory.

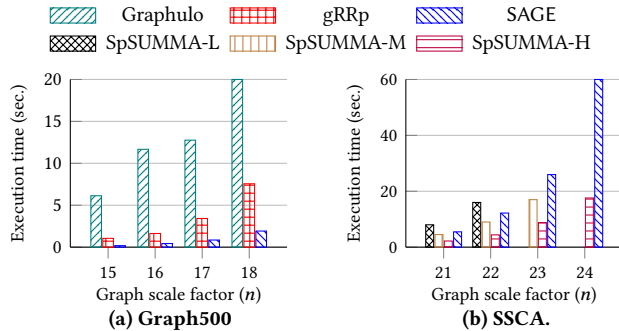
As a result, these results demonstrate that SAGE is able not only (1) to perform SpGEMM on very large graphs by utilizing the sufficient capacity of external storage (i.e., *scalable*), but also (2) to efficiently handle storage-memory I/Os by addressing the three challenges, as described in Section 3.3.

4.2.2 Distributed-system-based approach (EQ2). In this experiment, we compare SAGE with three state-of-the-art distributed-system-based SpGEMM methods: Graphulo [17], gRRp [10], and Sparse SUMMA (SpSUMMA) [6]. We use their experimental results reported in [6, 10] by following [18, 19, 28, 35, 41] since we use the exactly same datasets as in [6, 10] and it is not feasible to construct the identical distributed systems that the existing works used.

Setup. We use both real-world and synthetic datasets, which were used in [6, 10]. Specifically, for comparison with Graphulo and gRRp, we use 13 real-world and 4 synthetic datasets, generated by Graph500 benchmark with the scale factor varying from 15 to 18, where we set the edge factor $e = 16$ and the skewness parameters $a = 0.57$, $b = 0.19$, $c = 0.19$, and $d = 0.05$ by following [10]. For comparison with SpSUMMA, we use 4 synthetic datasets, generated by SSCA benchmark with the scale factor varying n from 21 to 24, where we set the edge factor $e = 8$ and the skewness parameters $a = 0.6$ and $b = c = d = 0.4/3$ by following [6]. Table 4 shows the statistics of the synthetic datasets used in this experiment.

Table 5: Comparison of the distributed-system-based methods with SAGE on 13 real-world datasets with various sizes.


Dataset	Sphere	Filter3D	598a	Torso2	Cage12	144	Wave	Majorbasis	Scircuit	Mac	Offshore	Mario	Tmt_sym
Graphulo	7.37	12.91	9.57	4.81	13.28	11.97	11.85	8.07	6.36	4.63	18.86	9.07	13.55
gRRp	<u>1.50</u>	<u>3.71</u>	<u>1.65</u>	<u>0.86</u>	<u>2.63</u>	<u>2.35</u>	<u>1.87</u>	<u>1.52</u>	<u>1.02</u>	<u>1.31</u>	<u>3.89</u>	<u>1.70</u>	<u>3.34</u>
SAGE	0.36	0.77	0.34	0.11	0.54	0.48	0.41	0.25	0.18	0.25	0.93	0.30	0.62
Difference (sec.)	1.14↓	2.94↓	1.31↓	0.75↓	2.09↓	1.87↓	1.46↓	1.27↓	0.84↓	1.06↓	2.96↓	1.40↓	2.72↓

**Figure 6: Comparison of the distributed-system-based methods with SAGE on synthetic datasets.**

Then, we (1) perform SpGEMM using SAGE on the datasets for 5 times and (2) measure the average execution time of SpGEMM of each method. We note that Graphulo and gRRp have been conducted on the distributed system that consists of 10 machines equipped with 2 Intel Xeon E5-2690 v4 processors connected via the DGS-3120-24TC Ethernet [10] and SpSUMMA has been conducted on the Franklin XT4 system, where each machine is equipped with a quad-core AMD Opteron processor and connected with each other via a 6.4 GB/s network [6].

Results and analysis. Table 5 shows the results of Graphulo, gRRp, and SAGE on 13 real-world datasets. The results reveal that SAGE consistently outperforms the state-of-the-art distributed-system-based methods in all real-world datasets. Specifically, SAGE finishes the SpGEMM by up to 7.7× faster and gRRp, the best performer among competitors. Whereas, Graphulo provides the poor SpGEMM performance on real-world graphs (e.g., by up to 43× slower than SAGE in Torso2 dataset). This is because Graphulo, which adopts the outer product, suffers from the significant overhead of indexing and aggregating a large amount of the intermediate results of outer-product based SpGEMM as described in Section 3.1. Although gRRp provides the performance of SpGEMM much better than Graphulo, SAGE still significantly outperforms gRRp. We argue that these performance improvement of SAGE over Graphulo and gRRp is significant, considering they perform SpGEMM on the high-performance distributed system that consists of 10 machines [10]. Thus, these results imply that SAGE is able to efficiently handle large-scale SpGEMM in a single machine, requiring only intra-machine communication overhead (i.e., storage-memory I/Os), much smaller than the inter-machine communication overhead.

Figure 6 shows the results on synthetic datasets, generated by two benchmarks (Graph500 and SSCA), where the x -axis represents the graph scale factor n of each synthetic dataset and the y -axis represents the execution time (sec.). Similar to the results on real-world datasets, Figure 6(a) shows that SAGE significantly outperforms the

**Figure 7: The SpGEMM execution time of SAGE with the increasing sizes of synthetic graphs.**

two distributed-system-based methods, Graphulo and gRRp, across all graph scales. Moreover, as shown in Figure 6(b)¹, SAGE successfully performs SpGEMM on very large-scale graphs (the scale factor $n \geq 23$) that SpSUMMA-L and SpSUMMA-M fail to handle in a single machine. As a result, These results demonstrate that SAGE is able to process large-scale real-world graphs successfully with only limited computing resources of a single machine, rather than those of costly distributed-systems (i.e., cost-effective).

4.2.3 Scalability (EQ3). In this experiment, we evaluate the scalability of SAGE with the increasing sizes of graphs.

Setup. We use synthetic datasets, generated by two different benchmarks (Graph500 and ER) [3] with various graph scales n ranging from 18 to 28, where the scale factor n decides the number of nodes in a generated graph (i.e., $|N| = 2^n$). For each generated graph, we also consider two different edge factors to control the number of edges (i.e., one for a high-degree graph and the other for a low-degree graph). Specifically, we set the edge factor for the high-degree and low-degree graphs as $1.5 \times e$ and $0.5 \times e$, respectively². Then, we perform SpGEMM (type (2)) on the generated graphs for 5 times and measure the average execution time (sec.).

Results and analysis. Figure 7 shows the results, where the x -axis represents the scale factor n , the y -axis represents the execution time (sec. in log-scale), and the red dotted vertical line represents the limits of existing single-machine-based methods. We observe that the execution time of SAGE tends to increase by 4.5 times as the scale factor increases by 2 (i.e., a graph gets 4× larger). This result demonstrates that the SpGEMM performance of SAGE successfully scales up to very large-scale graphs, exceeding the capacity of the main memory (red dotted vertical line), and SAGE provides (almost) linear scalability with the increasing sizes of graphs.

¹Buluç and Gilbert [6] used three different number of cores in its system with 3 (SpSUMMA-L), 4 (SpSUMMA-M), and 9 (SpSUMMA-H) machines, respectively.

²We could not show the results of high-degree graphs with graph scale 28 since R-MAT generator could not generate them due to the integer overflow.

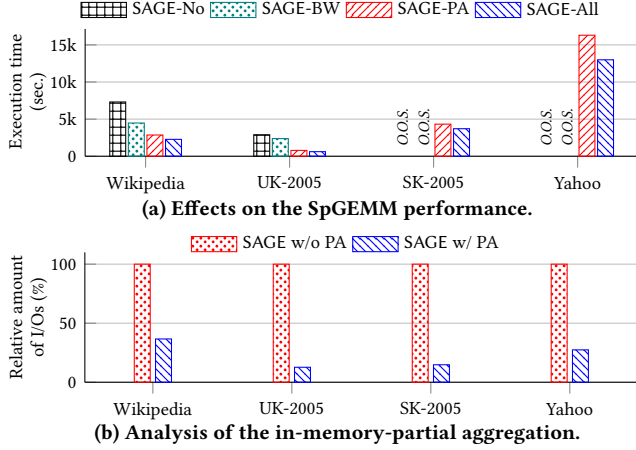


Figure 8: Effects of our proposed strategies on SAGE.

4.2.4 **Ablation study (EQ4).** Lastly, in this section, we evaluate the effectiveness of the proposed strategies of SAGE individually.

Setup (1). We first evaluate the first two strategies: (1) block-based workload allocation and (2) in-memory partial aggregation. To rigorously evaluate the two strategies, we (1) use four largest real-world datasets (i.e., Wikipedia, UK-2005, SK-2005, and Yahoo datasets [18]), which could not be loaded into main memory, and (2) limit the size of the main memory to 2GB to incur a large amount of storage-memory I/Os. We compare the following four versions of SAGE:

- SAGE-No: a baseline without any optimizations
- SAGE-BW: the version with the block-based workload allocation
- SAGE-PA: the version with the in-memory partial aggregation
- SAGE-All: the final version with both optimizations

We perform SpGEMM (type (1)) using each version and measure the execution time (sec.). For this experiment, we set the hyperparameter $\alpha = 12.5\%$ for the main memory buffer allocation.

Results and analysis (1). As shown in Figure 8(a), each of the proposed strategy is effective in improving the performance of SpGEMM and SAGE-All consistently provides the best performance in all datasets (up to 4.7 \times over SAGE-no). Also, for the two largest datasets (i.e., SK-2005 and Yahoo), SAGE-no and SAGE-BW fail to perform SpGEMM because the size of intermediate results even exceeds the capacity of the 4TB external storage, where ‘O.O.S.’ indicates ‘out of storage’. Interestingly, SAGE-PA consistently outperforms SAGE-BW in all cases, which implies that, in case of large-scale SpGEMM, handling storage-memory I/Os could be more critical than balancing workloads, in particular when the main memory size of a single machine is limited.

For more analysis of the in-memory partial aggregation, we measure the amount of storage-memory I/Os generated by the intermediate results during SpGEMM of SAGE with/without the in-memory partial aggregation (SAGE w/ PA and SAGE w/o PA in Figure 8(b)). Figure 8(b) shows the relative amount of storage-memory I/Os, where our in-memory partial aggregation significantly reduces the amount of storage-memory I/Os (up to 87% in UK-2005). This result demonstrates that in performing storage-based SpGEMM on large-scale graphs, a substantial amount of storage-memory I/Os could be generated; thus, it is crucial to efficiently handle these I/Os in order to achieve good performance, as we claimed in Section 3.3.




Figure 9: Effects of distribution-aware memory allocation.

Setup (2). We also evaluate our distribution-aware memory allocation. As described in in Section 3.3.3, we set (1) the former input buffer $B_{in}^1 = b \times t$, where b is the size of a block (1MB) and t is the number of threads; and (2) the output buffer $B_{out} = \alpha \times max_row \times t$, where max_row here indicates the maximum size of the final results for each row. Then, we perform SpGEMM with varying (1) the numbers of threads (2, 4, and 8), (2) the size of B_{in}^1 (1MB to 16MB), and (3) the hyperparameter α (6.25% to 50%).

Results and analysis (2). Figure 9(a) shows that the performance of SAGE gets improved as the number of threads and the size of B_{in}^1 increase, where the performance improvement converges at a specific point where the size of B_{in}^1 is equal to $b \times t$, which implies that the buffer size is sufficient to load the blocks for that all threads can process simultaneously, as explained in Section 3.3.3. Figure 9(b) shows the performance of SAGE with respect to the size of main memory and the hyperparameter α . The performance of SAGE tends to be improved as α and the size of the main memory increase, but the performance improvement gets smaller as α increases, which implies that about 12.5% ~ 25% of the maximum size is enough for B_{out} to process the intermediate results of a majority of rows. As a result, we believe that our distribution-aware memory allocation is effective and practical in large-scale SpGEMM.

5 CONCLUSIONS

In this paper, we point out the limitations of previous approaches for SpGEMM: (i) the single-machine-based approach cannot handle large-scale graphs surpassing the capacity of the main memory (i.e., **not scalable**), and (ii) the distributed-system-based approach requires a substantial amount of inter-machine communication overhead (i.e., **not efficient**). To address both challenges of scalability and efficiency, we propose a novel storage-based approach to SpGEMM (**SAGE**) with the 3-layer architecture to efficiently handle storage-memory I/Os. We further identify three important challenges that could cause serious performance degradation in storage-based SpGEMM and propose three effective strategies to address them: (1) block-based workload allocation, (2) in-memory partial aggregation, and (3) distribution-aware memory allocation. Through comprehensive evaluation, we demonstrate the superiority of SAGE in terms of (1) *scalability*, (2) *efficiency*, and (3) *effectiveness*.

6 ACKNOWLEDGMENTS

This is a joint work between Samsung Electronics Co., Ltd and Hanyang University. This work was supported by Institute of Information & communications Technology Planning & Evaluation (IITP) grant funded by the Korea government (MSIT) (No.2022-0-00352 and No.RS-2022-00155586).

REFERENCES

- [1] Kadir Akbudak and Cevdet Aykanat. 2017. Exploiting locality in sparse matrix-matrix multiplication on many-core architectures. *IEEE Transactions on Parallel and Distributed Systems* 28, 8 (2017), 2258–2271.
- [2] Michael J Anderson, Narayanan Sundaram, Nadathur Satish, Md Mostofa Ali Patwary, Theodore L Willke, and Pradeep Dubey. 2016. Graphpad: Optimized graph primitives for parallel and distributed platforms. In *2016 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 313–322.
- [3] Ariful Azad, Grey Ballard, Aydin Buluc, James Demmel, Laura Grigori, Oded Schwartz, Sivan Toledo, and Samuel Williams. 2016. Exploiting multiple levels of parallelism in sparse matrix-matrix multiplication. *SIAM Journal on Scientific Computing* 38, 6 (2016), C624–C651.
- [4] Daehyeon Baek, Soojin Hwang, Taekyung Heo, Daehoon Kim, and Jaehyuk Huh. 2021. InnerSP: A Memory Efficient Sparse Matrix Multiplication Accelerator with Locality-Aware Inner Product Processing. In *2021 30th International Conference on Parallel Architectures and Compilation Techniques (PACT)*. IEEE, 116–128.
- [5] Luca Becchetti, Paolo Boldi, Carlos Castillo, and Aristides Gionis. 2008. Efficient semi-streaming algorithms for local triangle counting in massive graphs. In *Proceedings of the ACM International Conference on Knowledge Discovery and Data Mining*. ACM, 16–24.
- [6] Aydin Buluc and John R Gilbert. 2012. Parallel sparse matrix-matrix multiplication and indexing: Implementation and experiments. *SIAM Journal on Scientific Computing* 34, 4 (2012), C170–C191.
- [7] Deepayan Chakrabarti, Yiping Zhan, and Christos Faloutsos. 2004. R-MAT: A recursive model for graph mining. In *Proceedings of the 2004 SIAM International Conference on Data Mining*. SIAM, 442–446.
- [8] Paolo D'Alberto and Alexandru Nicolau. 2007. R-Kleene: A high-performance divide-and-conquer algorithm for the all-pair shortest path for densely connected networks. *Algorithmica* 47, 2 (2007), 203–213.
- [9] Timothy A Davis. 2019. Algorithm 1000: SuiteSparse: GraphBLAS: Graph algorithms in the language of sparse linear algebra. *ACM Transactions on Mathematical Software (TOMS)* 45, 4 (2019), 1–25.
- [10] Gunduz Vehbi Demirci and Cevdet Aykanat. 2020. Scaling sparse matrix-matrix multiplication in the accumulo database. *Distributed and Parallel Databases* 38, 1 (2020), 31–62.
- [11] Mehmet Deveci, Christian Trott, and Sivasankaran Rajamanickam. 2017. Performance-portable sparse matrix-matrix multiplication for many-core architectures. In *2017 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*. IEEE, 693–702.
- [12] Felix Gremse, Kerstin Kupper, and Uwe Naumann. 2018. Memory-efficient sparse matrix-matrix multiplication by row merging on many-core architectures. *SIAM Journal on Scientific Computing* 40, 4 (2018), C429–C449.
- [13] Rong Gu, Yun Tang, Chen Tian, Hucheng Zhou, Guanru Li, Xudong Zheng, and Yihua Huang. 2017. Improving execution concurrency of large-scale matrix multiplication on distributed data-parallel platforms. *IEEE Transactions on Parallel and Distributed Systems* 28, 9 (2017), 2539–2552.
- [14] Masoud Reyhani Hamedani and Sang-Wook Kim. 2017. JacSim: An accurate and efficient link-based similarity measure in graphs. *Information Sciences* 414 (2017), 203–224.
- [15] Wook-Shin Han, Sangyeon Lee, Kyungyeol Park, Jeong-Hoon Lee, Min-Soo Kim, Jinha Kim, and Hwanjo Yu. 2013. TurboGraph: A fast parallel graph engine handling billion-scale graphs in a single PC. In *Proceedings of the ACM international conference on knowledge discovery and data mining (KDD)*. 77–85.
- [16] Guoming He, Haijun Feng, Cuiping Li, and Hong Chen. 2010. Parallel SimRank computation on large graphs with iterative aggregation. In *Proceedings of the ACM international conference on knowledge discovery and data mining (SIGKDD)*. 543–552.
- [17] Dylan Hutchison, Jeremy Kepner, Vijay Gadepally, and Adam Fuchs. 2015. Graphulo implementation of server-side sparse matrix multiply in the Accumulo database. In *2015 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, 1–7.
- [18] Yong-Yeon Jo, Myung-Hwan Jang, Sang-Wook Kim, and Sunju Park. 2019. RealGraph: A graph engine leveraging the power-law distribution of real-world graphs. In *Proceedings of the World Wide Web Conference*. ACM, 807–817.
- [19] Sang-Woo Jun, Andy Wright, Sizhuo Zhang, Shuotao Xu, et al. 2018. GrafBoost: Using accelerated flash storage for external graph analytics. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*. IEEE, 411–424.
- [20] Yoonsuk Kang, Jun Seok Lee, Won-Yong Shin, and Sang-Wook Kim. 2020. Cr-graph: Community reinforcement for accurate community detection. In *Proceedings of the 29th ACM International Conference on Information & Knowledge Management*. 2077–2080.
- [21] George Karypis and Vipin Kumar. 1998. Multilevel k-way partitioning scheme for irregular graphs. *J. Parallel and Distrib. Comput.* 48, 1 (1998), 96–129.
- [22] Jeremy Kepner, Peter Aaltonen, David Bader, Aydin Buluc, Franz Franchetti, John Gilbert, Dylan Hutchison, Manoj Kumar, Andrew Lumsdaine, Henning Meyerhenke, et al. 2016. Mathematical foundations of the GraphBLAS. In *2016 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, 1–9.
- [23] Jeremy Kepner and John Gilbert. 2011. *Graph algorithms in the language of linear algebra*. SIAM.
- [24] Jon M Kleinberg. 1999. Authoritative sources in a hyperlinked environment. *J. ACM* 46, 5 (1999), 604–632.
- [25] Yunyong Ko, Kibong Choi, Jiwon Seo, and Sang-Wook Kim. 2021. An In-Depth Analysis of Distributed Training of Deep Neural Networks. In *Proceedings of the IEEE International Parallel and Distributed Processing Symposium (IPDPS)*. IEEE, 994–1003.
- [26] Yunyong Ko, Jae-Seo Yu, Hong-Kyun Bae, Yongjun Park, Dongwon Lee, and Sang-Wook Kim. 2021. MASCO: A Quantization Framework for Efficient Matrix Factorization in Recommender Systems. In *2021 IEEE International Conference on Data Mining (ICDM)*. IEEE, 290–299.
- [27] Scott P Kolodziej, Mohsen Aznaveh, Matthew Bullock, Jarrett David, Timothy A Davis, Matthew Henderson, Yifan Hu, and Read Sandstrom. 2019. The suitesparse matrix collection website interface. *Journal of Open Source Software* 4, 35 (2019), 1244.
- [28] Aapo Kyrola, Guy E Blelloch, and Carlos Guestrin. 2012. GraphChi: Large-scale graph computation on just a pc. In *Proceedings of the USENIX symposium on operating systems design and implementation (OSDI)*. 31–46.
- [29] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. 2010. Predicting positive and negative links in online social networks. In *Proceedings of the 19th international conference on World wide web*. 641–650.
- [30] Jure Leskovec, Daniel Huttenlocher, and Jon Kleinberg. 2010. Signed networks in social media. In *Proceedings of the SIGCHI conference on human factors in computing systems*. 1361–1370.
- [31] Jure Leskovec, Jon Kleinberg, and Christos Faloutsos. 2007. Graph evolution: Densification and shrinking diameters. *ACM Transactions on Knowledge Discovery from Data* 1, 1 (2007), 1–41.
- [32] Jure Leskovec and Andrej Krevl. 2014. SNAP Datasets: Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data>.
- [33] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon.com recommendations: Item-to-item collaborative filtering. *IEEE Internet computing* 7, 1 (2003), 76–80.
- [34] Hang Liu, H Howie Huang, and Yang Hu. 2016. ibfs: Concurrent breadth-first search on gpus. In *Proceedings of the 2016 International Conference on Management of Data*. 403–416.
- [35] Steffen Maass, Changwoo Min, Sanidhya Kashyap, Woonhak Kang, Mohan Kumar, and Taesoo Kim. 2017. Mosaic: Processing a trillion-edge graph on a single machine. In *Proceedings of the Twelfth European Conference on Computer Systems*. 527–543.
- [36] Yusuke Nagasaka, Satoshi Matsuoka, Ariful Azad, and Aydin Buluc. 2018. High-performance sparse matrix-matrix products on intel knl and multicore architectures. In *Proceedings of the 47th International Conference on Parallel Processing Companion*. 1–10.
- [37] Yusuke Nagasaka, Akira Nukada, and Satoshi Matsuoka. 2017. High-performance and memory-saving sparse general matrix-matrix multiplication for nvidia pascal gpu. In *2017 46th International Conference on Parallel Processing (ICPP)*. IEEE, 101–110.
- [38] Carlos Ordonez, Yiqun Zhang, and Wellington Cabrera. 2016. The Gamma matrix to summarize dense and sparse data sets for big data analytics. *IEEE Transactions on Knowledge and Data Engineering* 28, 7 (2016), 1905–1918.
- [39] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. 1999. *The PageRank citation ranking: Bringing order to the web*. Technical Report. Stanford InfoLab.
- [40] Subhankar Pal, Jonathan Beaumont, Dong-Hyeon Park, Apurva Amarnath, Siying Feng, Chaitali Chakrabarti, Hun-Seok Kim, David Blaauw, Trevor Mudge, and Ronald Dreslinski. 2018. Outerspace: An outer product based sparse matrix multiplication accelerator. In *2018 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 724–736.
- [41] Amitabha Roy, Ivo Mihailovic, and Willy Zwaenepoel. 2013. X-Stream: Edge-centric graph processing using streaming partitions. In *Proceedings of the ACM symposium on operating systems principles (SOSP)*. 472–488.
- [42] Robert Sedgewick and Kevin Wayne. 2011. *Algorithms*. Addison-Wesley professional.
- [43] Oguz Selvitopi, Gunduz Vehbi Demirci, Ata Turk, and Cevdet Aykanat. 2019. Locality-aware and load-balanced static task scheduling for MapReduce. *Future Generation Computer Systems* 90 (2019), 49–61.
- [44] Nitish Srivastava, Hanchen Jin, Jie Liu, David Albonese, and Zhiru Zhang. 2020. Matraprot: A sparse-sparse matrix multiplication accelerator based on row-wise product. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 766–780.
- [45] Kenji Suzuki, Isao Horiba, and Noboru Sugie. 2003. Linear-time connected-component labeling based on sequential local operations. *Computer Vision and Image Understanding* 89, 1 (2003), 1–23.
- [46] Endong Wang, Qing Zhang, Bo Shen, Guangyong Zhang, Xiaowei Lu, Qing Wu, and Yajuan Wang. 2014. Intel math kernel library. In *High-Performance Computing on the Intel® Xeon Phi™*. Springer, 167–188.

- [47] Zhen Xie, Guangming Tan, Weifeng Liu, and Ninghui Sun. 2019. IA-SpGEMM: An input-aware auto-tuning framework for parallel sparse matrix-matrix multiplication. In *Proceedings of the ACM International Conference on Supercomputing*. 94–105.
- [48] Zhekai Zhang, Hanrui Wang, Song Han, and William J Dally. 2020. Sparch: Efficient architecture for sparse matrix multiplication. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*. IEEE, 261–274.
- [49] Da Zheng, Disa Mhembere, Randal Burns, Joshua Vogelstein, Carey E Priebe, and Alexander S Szalay. 2015. FlashGraph: Processing billion-node graphs on an array of commodity SSDs. In *Proceedings of the USENIX conference on file and storage technologies (FAST)*. 45–58.
- [50] Da Zheng, Disa Mhembere, Vince Lyzinski, Joshua T Vogelstein, Carey E Priebe, and Randal Burns. 2016. Semi-external memory sparse matrix multiplication for billion-node graphs. *IEEE Transactions on Parallel and Distributed Systems* 28, 5 (2016), 1470–1483.