

Fast Dynamic Updates and Dynamic SpGEMM on MPI-Distributed Graphs

Alexander van der Grinten*, Geert Custers†, Duy Le Thanh*, and Henning Meyerhenke*

*Humboldt-Universität zu Berlin Berlin, Germany

†Delft University of Technology, Delft, Netherlands

Abstract—Sparse matrix multiplication (SpGEMM) is a fundamental kernel used in many diverse application areas, both numerical and discrete. For example, many algebraic graph algorithms rely on SpGEMM in the tropical semiring to compute shortest paths in graphs. Recently, SpGEMM has received growing attention regarding implementations for specific (parallel) architectures. Yet, this concerns only the static problem, where both input matrices do not change. In many applications, however, matrices (or their corresponding graphs) change over time. Although recomputing from scratch is very expensive, we are not aware of any dynamic SpGEMM algorithms in the literature.

In this paper, we thus propose a batch-dynamic algorithm for MPI-based parallel computing. Building on top of a distributed graph/matrix data structure that allows for fast updates, our dynamic SpGEMM reduces the communication volume significantly. It does so by exploiting that updates change far fewer matrix entries than there are non-zeros in the input operands. Our experiments with popular benchmark graphs show that our approach pays off. For batches of insertions or removals of matrix entries, our dynamic SpGEMM is substantially faster than the static algorithms in the state-of-the-art competitors CombBLAS, CTF and PETSc.

I. INTRODUCTION

Sparse general matrix-matrix multiplication, usually denoted by SpGEMM, is a fundamental computational kernel for many applications in various areas – for instance scientific computing (such as algebraic multigrid [1]), machine learning with sparse DNNs [2], and data analytics (such as clustering [3]). Besides its use in numerical contexts, numerous graph algorithms (such as triangle counting [4]) make use of SpGEMM due to the correspondence between graphs and matrices [5], [6]. Two other popular applications are shortest paths with multiple sources and graph contraction [7]. Thus, our results apply to both the matrix and the graph perspective as well as for numerical and discrete problems.

In fact, we focus on algorithms and data structures for *dynamic* SpGEMM in this paper. Note that dynamic problems are very common in data analytics applications – for example, consider the continuously changing inputs in areas such as recommender systems, online social networks, or time-dependent mobility networks. Moreover, the use of dynamic SpGEMM is also conceivable for deep learning on changing/growing data sets. Frequently recomputing analytics results (or SpGEMM as one step in the workflow) after changes in the input data is typically very time-consuming or even infeasible for large inputs. It is thus desirable to use a dynamic algorithm that can update the previous result in a more cost-efficient manner. Also,

numerous large-scale applications require (or at least benefit from) the speed and memory of HPC platforms. Yet, despite growing research interest in optimizing static SpGEMM for various architectures and as part of MPI-based tools (see Section II for a brief overview), no dynamic approach for SpGEMM has been published so far – neither sequential nor parallel/distributed.

Outline and Contributions: We thus first propose an MPI-parallel data structure for fast dynamic updates to sparse matrices and graphs (Section IV). This data structure stores the adjacency matrix in a 2D distribution over the processes. Based on this data structure, we propose a parallel dynamic SpGEMM algorithm in Section V. More precisely, we provide two variants, one for algebraic updates (update addition equals the one in the SpGEMM semiring) and one for general updates. The main idea behind the algorithm is to limit the communication between processes at the potential cost of a mild increase in other operations. To achieve this goal, we change the way block submatrices are communicated (and aggregated) between processes while using our dynamic matrices. Our experimental results (see Section VII) show that our approach pays off. A faster redistribution yields much better insertion, update, and deletion times for the data structure compared to the state-of-the-art tools CombBLAS, CTF and PetSc. Regarding SpGEMM, the dynamic algorithm for algebraic updates (which performs better than the general one) is between 3.41 and 6.18 times faster on average than a static SpGEMM computation with CombBLAS (the best competitor) – depending on the batch size with up to some *millions* of non-zero updates. Finally note that related work is presented in Section II, preliminaries and notation in Section III, while concluding remarks are made in Section VIII.

II. RELATED WORK

Sparse graph and matrix computations in general have received considerable research attention in the last decades. One more recent appealing idea is to express graph computations by linear algebraic operations [5], which has led to the GraphBLAS initiative and standard [6], [8]. Compared to sparse matrix-vector multiplication (SpMV), the more challenging SpGEMM has been covered to a lesser extent in the literature so far – as pointed out by Winter et al. [9]. Recently, however, the optimization of static SpGEMM algorithms for specific parallel architectures is on the rise, e.g., for multithreaded CPUs [10], [11] GPUs [9], CPU/GPU combinations [12],

[13], and other accelerators [10]. A possible reason for this spike could be the use of SpGEMM in deep learning with sparse DNNs, as described in Ref. [2]. Rather than providing a comprehensive overview over static SpGEMM, our description focuses on efforts for MPI – as only these are directly comparable to our approach. For more details, the interested reader is referred to a recent systematic literature review on SpGEMM [7].

MPI-based frameworks for sparse graph computations and/or sparse linear algebra include CombBLAS 2.0 [14], Trilinos [15], CTF [16], LAMA [17], and PETSc [18]. In our experiments, we compare against recent versions of CombBLAS, CTF and PETSc.

As pointed out by Buluç et al. [19], the most common data structures for sparse matrices and graphs in the context of algebraic operations are variants of compressed sparse row (CSR) or column (CSC). This also holds for the parallel case, where the distribution of the matrix over the processes is an important aspect. The two most common distributions are 1D (vertex-based in graphs terms) and 2D (edge-based) [20], but also 2.5D and 3D are in use [21]. In particular, for highly irregular inputs, a 2D distribution fares better than 1D due to better load balancing [22]. A 3D implementation with multiple levels of parallelism [21] was shown to outperform previous SpGEMM codes for MPI parallelism. Parallel distributions can lead to hypersparse (sub-)matrices on some processes, for which doubly compressed versions of CSR/CSC are known [20]. In the same work, Buluç and Gilbert [20] also propose sparse SUMMA as their “algorithm of choice” for static parallel SpGEMM due to its flexibility in terms of matrix size. SUMMA is also used as part of the 3D algorithm by Azad et al. [21] (see above). We argue in Section V why sparse SUMMA is not well suited for a dynamic approach.

The potential of (batch-)dynamic algorithms – algebraic or not – for other graph/matrix problems has been shown in numerous recent works, among many others for core maintenance with matchings [23], all-pair shortest paths and betweenness centrality [24], as well as various other centrality measures [25], [26].

III. PRELIMINARIES AND NOTATION

Our work is built on top of MPI. We denote the number of MPI processes by p . Within each MPI process, we also make use of shared memory (i.e., OpenMP) parallelism.

We consider matrices over arbitrary semirings; for an introduction of matrix computations over semirings (in the context of graphs), we refer the reader to Ref. [6]. Common examples include the $(+, \cdot)$ semiring, the (\wedge, \vee) semiring over Boolean values, or the $(\min, +)$ semiring (which is often used in shortest path algorithms). We usually use the symbol 0 to refer to the neutral element of the semiring.

Matrices are denoted by \mathbf{A} , \mathbf{B} , \mathbf{C} and similar. The (u, v) -th entry of \mathbf{A} is denoted by $a_{u,v}$. Our framework will distribute blocks of matrices to each MPI process in a 2D process grid of size $\sqrt{p} \times \sqrt{p}$. In this context, $\mathbf{A}_{i,j}$ refers to the block of \mathbf{A} that is located on MPI process (i, j) , where $i, j \in \{1, \dots, \sqrt{p}\}$. We assume that matrices (e.g., input matrices for SpGEMM)

are *sparse*, i.e., that $nnz(\mathbf{A}) \ll n \cdot m$ for an $n \times m$ matrix \mathbf{A} . Some matrices that arise in our algorithms will be *hypersparse*, i.e., $nnz(\mathbf{A}) \ll n$ and $nnz(\mathbf{A}) \ll m$. Like most other sparse matrix frameworks, we differentiate between *structural* (non-)zeros and *numerical* (non-)zeros. In particular, an entry $a_{u,v}$ of a matrix \mathbf{A} is considered to be structurally non-zero if $a_{u,v}$ is present in our sparse matrix data structure, even if $a_{u,v} = 0$. In particular, a structural non-zero can still have the numerical value of zero. The term *non-zero* will always be used for structural non-zeros throughout this paper. Structural zeros (i.e., entries of the matrix that are not present in our data structures), are always implicitly equal to the (additive) neutral element of the semiring (e.g., ∞ for $(\min, +)$).

SpGEMM (sparse generalized matrix multiplication) is the problem of computing $\mathbf{C} = \alpha\mathbf{A}\mathbf{B} + \beta\mathbf{X}$, where \mathbf{A} and \mathbf{B} may optionally be transposed. Since the main challenge of any SpGEMM algorithm is computing the matrix multiplication $\mathbf{A}\mathbf{B}$, most of this paper will be focused on that computation, although we also give details about handling transposition in Section V-C.

IV. DATA STRUCTURES FOR DYNAMIC DISTRIBUTED GRAPHS

Our framework stores each matrix in a fully distributed way. We employ a 2D distribution of the matrix, i.e., each MPI process stores a block of the matrix. Like other distributed sparse matrix/graph frameworks (such as CombBLAS [14]), we support square process grids of size $\sqrt{p} \times \sqrt{p}$ for this purpose.

To store these blocks locally within each process, we distinguish between *dynamic* and *static* matrices. Dynamic matrices support efficient in-place operations (such as insertions, deletions, matrix addition or other element-wise transformations). To store dynamic matrices, we use the DHB data structure [27] which is based on adjacency arrays, together with a per-row hash table that maps column indices to locations in the adjacency array. This data structure allows us to quickly discover whether a given index pair (i, j) has a non-zero value; it also enables us to efficiently update a local entry of the matrix in $\mathcal{O}(1)$ expected time.

We store static matrices in a CSR data structure. For various hypersparse static matrices that arise in our algorithms, we use a *doubly compressed CSR* (= DCSR) matrix instead, similarly to the DCSC matrix that is used by CombBLAS [28]. DCSR (or DCSC) further decrease the memory overhead of a sparse matrix by storing only row pointer for non-zero rows (or columns, respectively). Doubly compressed storage layouts come at the cost of not being able to directly index (i.e., in $\mathcal{O}(1)$ time) into a specific row (or column) anymore. However, they can substantially decrease communication volume when hypersparse matrices need to be communicated. Since none of our algorithms needs to index into a static CSR or DCSR (i.e., no search for an index is ever necessary), we do not sort these storage layouts in any way (and we also do not maintain hash tables for these layouts).

A. Dynamic Updates

As a *dynamic update*, we consider the modification of a non-zero entry of an existing matrix. An update can consist of the insertion of a new non-zero, the deletion of an existing non-zero, or the change of the numerical value of a matrix entry. In many cases, such updates can be represented by adding an *update matrix* to the original matrix; however, this is not true in general semirings such as $(\min, +)$, where the min operator can only decrease (but not increase) the value of non-zeros of the matrix. We assume that MPI processes can generate updates independently and without knowledge of the distribution of data across the MPI process grid.

While our algorithms and data structures for dynamic updates are different compared to the state of the art in MPI-distributed sparse matrix processing, not many changes at the interface are necessary to support dynamic updates. To insert new entries into a matrix \mathbf{A} , we first build an update matrix \mathbf{A}^* that contains exactly the new values of all non-zeros that should be updated in \mathbf{A} . Afterwards, we either add \mathbf{A}^* to \mathbf{A} (if the update can be represented by addition), or we call a $\text{MERGE}(\mathbf{A}, \mathbf{A}^*)$ procedure that replaces all values (i, j) of \mathbf{A} by their corresponding entry in \mathbf{A}^* if the (i, j) -th entry of \mathbf{A}^* is non-zero. To delete entries, we support a $\text{MASK}(\mathbf{A}, \mathbf{A}^*)$ procedure that removes all entries (i, j) from \mathbf{A} for which \mathbf{A}^* is non-zero. Similar operations (optimized for the static case) are already present in most state-of-the-art sparse matrix frameworks.

The construction of \mathbf{A}^* involves communication to redistribute data across MPI processes. After \mathbf{A}^* is constructed, however, all dynamic update operations (i.e., matrix addition, MERGE , MASK) operate only on local blocks of the matrices; no communication is involved. All update operations can be implemented efficiently if \mathbf{A} is stored as a dynamic matrix and \mathbf{A}^* is stored in DCSR layout. To ensure that these matrices are indeed stored in these layouts, our framework requires the user to mark dynamic matrices and update matrices appropriately.

B. Distribution of Update Matrices across MPI Processes

Since MPI processes can generate updates independently and without knowledge of the data distribution, we need to be able to quickly redistribute a set of updates to the MPI process that stores the matrix block that is affected by the update. Our redistribution uses straightforward techniques; however, we found our algorithm to outperform existing MPI-parallel implementations in practice (in particular, this holds for implementations that rely on comparison-based sorting and a single ALLTOALL , see Section VII). To redistribute updates, we represent them as tuples (i, j, x) , where (i, j) is the index pair that should be updated to a new value x . Our redistribution routine takes an array of such tuples as input on each MPI process. We first exchange tuples across the rows of the process grid. In a second step, we exchange these tuples (which are now on the correct row of the process grid) across the columns of the grid. An ALLTOALL communication call is used to transfer the data. Before each ALLTOALL , we use counting sort to group the tuples by their destination rank. While ALLTOALL is expensive in general,

our two-step approach ensures that this call only affects \sqrt{p} processes. Likewise, our counting sort has to consider only \sqrt{p} buckets. We use OpenMP parallelism to efficiently insert updates into local dynamic matrices (i.e., into adjacency arrays and associated hash tables). To enable this procedure, we use counting sort to group updates (i, j) according to $(i \bmod T)$, where T is the number of (shared-memory) threads. Afterwards, we handle updates with different $(i \bmod T)$ in parallel.

Analysis: Let nnz be the number of input tuples that need to be redistributed. We assume that all processes initially have nnz/p of these tuples. Applying a random permutation to the input tuples (before invoking redistribution) allows us to assume that load is evenly distributed even after redistribution [29]. Due to the ALLTOALL call, our redistribution has a communication latency of \sqrt{p} and bandwidth requirements of $\mathcal{O}(\text{nnz}/p)$, since each process sends at most as much data as it locally has. We can assume that inserting into a local dynamic matrix (i.e., into an adjacency list + hash table) can be done in expected amortized $\mathcal{O}(1)$ time. Hence, the (expected amortized) local computation time is $\mathcal{O}(\text{nnz}/p + \sqrt{p})$, i.e., it is dominated by our counting sort that groups non-zeros by their destination process.

V. DYNAMIC DISTRIBUTED SPGEMM

We now consider the problem of dynamically updating the result of an SpGEMM operation. In particular, we want to compute $\mathbf{C} = \mathbf{A}\mathbf{B}$, where \mathbf{A} is a $n \times k$ matrix and \mathbf{B} is a $k \times m$ matrix. Since our work concerns the dynamic variant of this matrix multiplication, we assume that \mathbf{A} and \mathbf{B} have been updated by inserting and/or removing non-zero entries. These updates result in new matrices \mathbf{A}' and \mathbf{B}' , respectively. The *dynamic matrix multiplication* problem now consists of computing $\mathbf{C}' = \mathbf{A}'\mathbf{B}'$, given the result \mathbf{C} of the multiplication before updates were applied. We consider two cases of the dynamic matrix multiplication problem:

- **Algebraic updates.** We say that \mathbf{A}' is obtained by *algebraic* updates if $\mathbf{A}' = \mathbf{A} + \mathbf{A}^*$ is the sum of \mathbf{A} and an update matrix \mathbf{A}^* , where the sum uses the same addition operator as the semiring that is used in the SpGEMM. If \mathbf{A}' and \mathbf{B}' are obtained by algebraic updates, we can exploit the distributive property of matrix multiplication to compute \mathbf{C}' as $\mathbf{C}' = (\mathbf{A} + \mathbf{A}^*)(\mathbf{B} + \mathbf{B}^*) = \mathbf{A}\mathbf{B} + \mathbf{A}\mathbf{B}^* + \mathbf{A}^*\mathbf{B} + \mathbf{A}^*\mathbf{B}^* = \mathbf{C} + \mathbf{A}^*\mathbf{B}' + \mathbf{A}\mathbf{B}^*$. Throughout the remainder of this paper, we refer to the last two terms of the latter expression as \mathbf{C}^* , i.e.,

$$\mathbf{C}^* := \mathbf{A}^*\mathbf{B}' + \mathbf{A}\mathbf{B}^*. \quad (1)$$

The main goal of our algorithm for algebraic updates is to compute \mathbf{C}^* without performing extensive communication (i.e., without communicating \mathbf{B}' or \mathbf{A}).¹

We note that if the semiring that is used for the matrix multiplication is a ring, this covers all possible updates

¹In the formulation of this paper, we always compute \mathbf{C}^* as $\mathbf{A}^*\mathbf{B}' + \mathbf{A}\mathbf{B}^*$. Our algorithms can also be rewritten to perform the computation as $\mathbf{C}^* = \mathbf{A}^*\mathbf{B} + \mathbf{A}'\mathbf{B}^*$. Both expressions are equally suitable for our algorithms; they also yield identical algorithmic complexities.

Algorithm 1 MPI-parallel dynamic SpGEMM, algebraic updates

```

    ▷ Executes in parallel on all MPI processes  $(i, j)$ 
send  $\mathbf{A}^*_{i,j}$  to process  $(j, i)$ , receive  $\mathbf{A}^*_{j,i}$ 
send  $\mathbf{B}^*_{i,j}$  to process  $(j, i)$ , receive  $\mathbf{B}^*_{j,i}$ 
for  $k = 1, \dots, \sqrt{p}$  do
    broadcast  $\mathbf{A}^*_{k,i}$  over  $i$ -th process row
    broadcast  $\mathbf{B}^*_{j,k}$  over  $j$ -th process column
     $\mathbf{X}^i_{k,j} \leftarrow \mathbf{A}^*_{k,i} \mathbf{B}'_{i,j}$  ▷ local
     $\mathbf{Y}^j_{i,k} \leftarrow \mathbf{A}_{i,j} \mathbf{B}^*_{j,k}$  ▷ local
    aggregate  $\mathbf{X}^i_{k,j}$  to  $\mathbf{X}_{k,j}$  on process  $(k, j)$ 
    aggregate  $\mathbf{Y}^j_{i,k}$  to  $\mathbf{X}_{i,k}$  on process  $(i, k)$ 
end for
 $\mathbf{C}'_{i,j} \leftarrow \mathbf{C}_{i,j} + \mathbf{X}_{i,j} + \mathbf{Y}_{i,j}$  ▷ local

```

since \mathbf{A}^* can simply be computed as $\mathbf{A}^* = \mathbf{A}' - \mathbf{A}$ in rings (but not in general semirings).

- **General updates.** If the algebraic update condition does not hold, we say that \mathbf{A}' and \mathbf{B}' are obtained by *general* updates. General updates can update matrices in ways that are “incompatible” with the semiring used for the matrix multiplication. For example, they can set values to 0 in the (\vee, \wedge) semiring or increase values of non-zeros of the matrix in the $(\min, +)$ semiring.

Even in the general case, we can define update matrices \mathbf{A}^* and \mathbf{B}^* that contain only the non-zeros of \mathbf{A}' and \mathbf{B}' that have changed compared to \mathbf{A} and \mathbf{B} . For our algorithms, only the structure (and not the numerical values) of \mathbf{A}^* and \mathbf{B}^* will be relevant. Given these matrices, we can compute \mathbf{C}^* as in Eq. (1). While $\mathbf{C}' \neq \mathbf{C} + \mathbf{C}^*$ in the general case, \mathbf{C}' can still only differ from \mathbf{C} at entries that are non-zero in \mathbf{C}^* .

Note that deletions of non-zeros in the input matrices can be handled in our algorithm for general updates by removing the corresponding entries from \mathbf{A}' or \mathbf{B}' while adding a structural non-zero to \mathbf{A}^* or \mathbf{B}^* to indicate that the corresponding entries in the input matrices have changed.

In both the algebraic and the general case, we assume that update matrices (i.e., \mathbf{A}^* and \mathbf{B}^*) are hypersparse (while \mathbf{A}' and \mathbf{B}' are in general sparse, but not hypersparse). Our algorithm works for any number of non-zeros in \mathbf{A}^* and \mathbf{B}^* ; however, our analysis of compute- and communication time relies on hypersparsity.

A. Algorithm for Algebraic Updates

As detailed above, we can compute $\mathbf{C}' = \mathbf{C} + \mathbf{C}^*$ in the algebraic case, with \mathbf{C}^* as in Eq. (1). The standard algorithm to evaluate the matrix multiplications $\mathbf{X} := \mathbf{A}^* \mathbf{B}'$ and $\mathbf{Y} := \mathbf{A} \mathbf{B}^*$ in MPI-distributed frameworks is the SUMMA algorithm (= *scalable universal matrix multiplication*). SUMMA performs \sqrt{p} rounds. In each round, it broadcasts blocks of the left-hand side of the matrix multiplication over the rows of the process grid, and blocks of the right-hand side over the columns of the grid. This procedure ensures that process (i, j) will receive exactly the blocks (i, k) of the left-hand

side and (k, j) of the right-hand side that are required to form block (i, j) of the result i.e., exactly the block of the output matrix that is supposed to reside on process (i, j) . Due to the way data is broadcasted in SUMMA, the aggregation of partial results into block (i, j) of the result is entirely local. SUMMA has an efficient communication pattern if both left-hand and right-hand sides have similar numbers of non-zeros. However, in our case (e.g., when computing $\mathbf{A}^* \mathbf{B}'$), one of the matrices (namely, \mathbf{A}^*) is expected to have far fewer non-zeros than the other matrix, since \mathbf{A}^* represents updates that usually affect only a small subset of \mathbf{A} .

Algorithm 1 depicts the pseudocode of our algebraic algorithm for dynamic SpGEMM. Instead of relying on the usual SUMMA algorithm, our algorithm can be seen as a combination of two passes of input-stationary SUMMA [30], without materializing intermediate results. In particular, our algorithm avoids broadcasting blocks of \mathbf{A} and \mathbf{B}' . This comes at the cost of an additional non-local aggregation step. However, since we expect \mathbf{A}^* , \mathbf{B}^* , and \mathbf{C}^* to be sparser than \mathbf{A} and \mathbf{B}' , it reduces the overall communication volume. The pseudocode of our algorithm is given in Section VI. Like SUMMA, we operate in \sqrt{p} rounds. In each round, we broadcast blocks of \mathbf{A}^* across rows of the process grid and blocks of \mathbf{B}^* across columns (and we do not have to broadcast \mathbf{A} nor \mathbf{B}' at all). In particular, in the k -th round, we broadcast block $\mathbf{A}^*_{k,i}$ across the i -th row of the process grid (and block $\mathbf{B}^*_{j,k}$ across the j -th row), for all i (and j). These broadcasts are visualized in Figure 1a, for matrix \mathbf{A}^* . To make this possible, we first need to perform one round of point-to-point communication to move $\mathbf{A}^*_{k,i}$ and $\mathbf{B}^*_{j,k}$ to the right process row and column, respectively. Since process (k, i) communicates with its transposed rank (i, k) in this round (and likewise for process (j, k)), each process only needs to communicate with a single peer. After the broadcasts are done, we compute the multiplications $\mathbf{X}^i_{k,j} := \mathbf{A}^*_{k,i} \mathbf{B}'_{i,j}$ and $\mathbf{Y}^j_{i,k} := \mathbf{A}_{i,j} \mathbf{B}^*_{j,k}$ locally on MPI process (i, j) . This is done on all processes (i, j) in parallel. Finally, we compute $\mathbf{X}_{k,j} := \sum_{i=1}^{\sqrt{p}} \mathbf{X}^i_{k,j}$ and $\mathbf{Y}_{i,k} := \sum_{j=1}^{\sqrt{p}} \mathbf{Y}^j_{i,k}$ by aggregating all $\mathbf{X}^i_{k,j}$ and $\mathbf{Y}^j_{i,k}$ on process (k, j) and (i, k) , respectively (see Figure 1c).

Analysis: Like SUMMA, our algorithm requires a communication latency of $\mathcal{O}(\sqrt{p} \log p)$ since we perform \sqrt{p} rounds of collective communication over \sqrt{p} processes (assuming that broadcast and aggregation steps are implemented with a latency of $\mathcal{O}(\log p)$). We require a communication bandwidth of $\mathcal{O}(\max(\text{nnz}(\mathbf{A}^*) + \text{nnz}(\mathbf{B}^*), \text{nnz}(\mathbf{C}^*)) / \sqrt{p})$, whereas SUMMA requires $\mathcal{O}((\text{nnz}(\mathbf{A}) + \text{nnz}(\mathbf{B}')) / \sqrt{p})$. In particular, our algorithm requires less bandwidth if \mathbf{A}^* , \mathbf{B}^* and \mathbf{C}^* are considerably sparser than \mathbf{A} and \mathbf{B}' . Let flops denote the number of scalar multiplications required to form \mathbf{C}^* . SUMMA requires $\mathcal{O}(\text{flops}/p)$ time for local computation, while our algorithm requires $\mathcal{O}((\text{flops} \log p)/p)$ (assuming that a $(\log p)$ -round parallel reduction is used for aggregation). Our algorithm consumes at most $\mathcal{O}(\max(\text{nnz}(\mathbf{A}'), \text{nnz}(\mathbf{B}'), \text{nnz}(\mathbf{C}'))/p)$ memory.

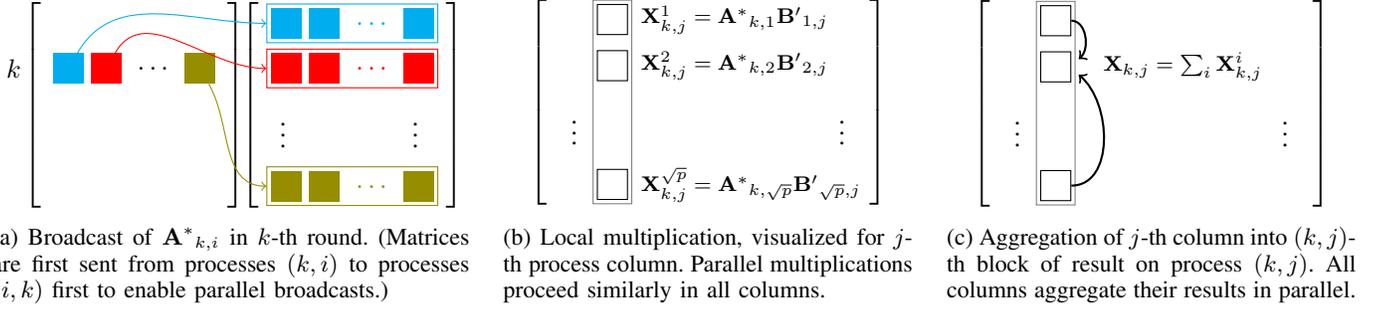


Figure 1: Communication pattern of MPI-parallel dynamic SpGEMM algorithm for algebraic updates. Between broadcast and aggregation, all MPI processes (i, j) compute $\mathbf{A}^*_{k,i} \mathbf{B}'_{i,j}$ locally. Figures depict computation of $\mathbf{A}^* \mathbf{B}'$; computation of $\mathbf{A} \mathbf{B}^*$ proceeds similarly.

Algorithm 2 MPI-parallel dynamic SpGEMM, general updates

```

 $\mathbf{C}^*, \mathbf{F}^* \leftarrow \text{COMPUTEPATTERN}(\mathbf{A}, \mathbf{A}^*, \mathbf{B}, \mathbf{B}^*)$ 
 $\mathbf{E}_{i,j} \leftarrow \mathbf{F}_{i,j} \oplus \mathbf{F}^*_{i,j}$  masked at  $\mathbf{C}^*_{i,j}$  ▷ local
aggregate( $\oplus$ )  $\mathbf{E}_{i,j}$  to  $\mathbf{R}_i$  over  $j$ -th process row
 $\mathbf{A}^{\mathbf{R}}_{i,j} \leftarrow \mathbf{A}'$  filtered by  $\mathbf{R}_i$  ▷ local
send  $\mathbf{A}^{\mathbf{R}}_{i,j}$  to process  $(j, i)$ , receive  $\mathbf{A}^{\mathbf{R}}_{j,i}$ 
for  $k = 1, \dots, \sqrt{p}$  do
  broadcast  $\mathbf{A}^{\mathbf{R}}_{k,i}$  over  $i$ -th process row
  broadcast  $\mathbf{C}^*_{k,j}$  over  $j$ -th process column
   $\mathbf{Z}^i_{k,j}, \mathbf{H}^i_{k,j} \leftarrow \mathbf{A}^{\mathbf{R}}_{k,i} \mathbf{B}'_{i,j}$  masked at  $\mathbf{C}^*_{k,j}$ 
  ▷ local, returns the new Bloom filter in  $\mathbf{H}^i_{k,j}$ 
  aggregate  $\mathbf{Z}^i_{k,j}$  to  $\mathbf{Z}_{k,j}$  over  $j$ -th process column
  aggregate( $\oplus$ )  $\mathbf{H}^i_{k,j}$  to  $\mathbf{H}_{k,j}$  over  $j$ -th process column
end for
merge  $\mathbf{Z}_{i,j}$  into  $\mathbf{C}_{i,j}$ , masked at  $\mathbf{C}^*_{i,j}$ , to obtain  $\mathbf{C}'_{i,j}$  ▷ local
merge  $\mathbf{H}_{i,j}$  into  $\mathbf{F}_{i,j}$ , masked at  $\mathbf{C}^*_{i,j}$ , to obtain  $\mathbf{F}'_{i,j}$  ▷ local

```

B. Algorithm for General Updates

For general updates, it is not enough to communicate the update matrices among MPI processes. Since general updates are not restricted in any way, computing the result of the dynamic matrix multiplication can require arbitrary entries of the new input matrices \mathbf{A}' and \mathbf{B}' (including entries that were *not* changed in comparison to \mathbf{A} and \mathbf{B}). Fortunately, we can still reduce the communication volume and computational cost compared to a static recomputation since not all entries of \mathbf{C}' need to be recomputed. Our approach is to use a masked SpGEMM for this purpose that we specialize for the problem of dynamically updating \mathbf{C}' . For any masked SpGEMM, it is straightforward to identify the rows of \mathbf{A}' and the columns of \mathbf{B}' that can contribute to a given set of entries in \mathbf{C}' (by just considering the appropriate rows/columns that have non-zero entries in the mask). However, our algorithm can further restrict the data that we need to exchange by considering only some *columns* of \mathbf{A}' and some *rows* of \mathbf{B}' . For this purpose, we use a Bloom filter that remembers which of the terms

$a_{i,k} b_{k,j}$ contribute to any $c_{i,j}$.² More specifically, our Bloom filter is a matrix \mathbf{F} that stores an ℓ -bit bitfield in each entry $f_{i,j}$, where ℓ is a constant (in practice, we use $\ell = 64$). While computing $\mathbf{C} = \mathbf{A} \mathbf{B}$, we set the $(k \bmod \ell)$ -th bit of $f_{i,j}$ to 1 if the term $a_{i,k} b_{k,j}$ contributes to the value of $c_{i,j}$ (and $f_{i,j} = 0$ if there is no such k). Given this bitmask, we can later recover a superset of the columns of \mathbf{A} (or rows of \mathbf{B}) that contributed to $c_{i,j}$.

Our algorithm is given in Algorithm 2. We first compute \mathbf{C}^* as given in Eq. (1). This computation is done using the algorithm for algebraic updates. We modify that algorithm to also compute a Bloom filter \mathbf{F}^* such that bit $(k \bmod \ell)$ is set in $f^*_{i,j}$ whenever the terms $a^*_{i,k} b'_{k,j}$ or $a'_{i,k} b^*_{k,j}$ contribute to $c^*_{i,j}$. This computation is denoted by COMPUTEPATTERN in our pseudocode. We remark that we do not require the values of \mathbf{C}^* for our algorithm; computing the sparsity structure of \mathbf{C}^* is enough.

Given \mathbf{F}^* , we compute the matrix $\mathbf{F} \oplus \mathbf{F}^*$, where \oplus denotes bitwise or. We form the matrix \mathbf{E} by keeping only the entries of $\mathbf{F} \oplus \mathbf{F}^*$ that are non-zero in \mathbf{C}^* . This matrix acts as a Bloom filter that can be used to select a superset of the columns of \mathbf{A}' and rows of \mathbf{B}' that are needed to compute \mathbf{C}' ; considering \mathbf{F}^* is needed to account for columns/rows that are required due to new non-zeros in \mathbf{A}' and \mathbf{B}' (compared to \mathbf{A} and \mathbf{B}). We reduce \mathbf{E} over its rows (via bitwise or). The result is a vector \mathbf{R} such that bit $(k \bmod \ell)$ is set in r_i if there is any column j of \mathbf{C}' such that $a'_{i,k} b'_{k,j}$ is required to compute $c'_{i,j}$. We now extract the rows i of \mathbf{A}' such that r_i is non-zero and only extract the columns k such that bit $(k \bmod \ell)$ is set in r_i . This yields the matrix $\mathbf{A}^{\mathbf{R}}$. While it would also be possible to filter (and broadcast, in the next step of the algorithm) \mathbf{B}' instead of \mathbf{A}' , we chose \mathbf{A}' because our matrices are locally stored row-wise. Hence, we can efficiently extract specific rows of \mathbf{A} and then discard a subset of the columns from these rows.

The remainder of the algorithm proceeds similarly to the algorithm for the algebraic case. We broadcast $\mathbf{A}^{\mathbf{R}}$ over the rows of the process grid (similarly to the algebraic case, see Figure 1a). To be able to make use of \mathbf{C}^* as an output mask

²A Bloom filter was previously used by Azad et al. [31] for (non-dynamic) masked SpGEMM. However, while their approach uses a Bloom filter to exclude rows of the left-hand side that need to be communicated, we use the exact mask (and not the Bloom filter) to avoid communicating rows of the left-hand side, and use the Bloom filter to additionally exclude columns.

during the local multiplication, we broadcast this matrix over the columns of the process grid and perform a local masked matrix multiplication. This local multiplication also produces an updated Bloom filter (called $\mathbf{H}_{k,j}^i$ in the pseudocode). Finally, we aggregate both the updated entries of the matrix (called $\mathbf{Z}_{k,j}^i$ in Algorithm 2) and updated entries of the Bloom filter onto MPI process (i, j) .

Analysis: Following a similar analysis as in the algebraic case, we find that our communication latency is $\mathcal{O}(\sqrt{p} \log p)$ (which is identical to the SUMMA algorithm). The communication bandwidth is dominated by the broadcasts, i.e., $\mathcal{O}((nnz(\mathbf{A}^R) + nnz(\mathbf{C}^*)) / \sqrt{p})$. Local computation costs are $\mathcal{O}(\text{flops} \log p) / p$, where flops is the number of scalar multiplications required to form $\mathbf{A}^R \mathbf{B}'$. While $nnz(\mathbf{A}^R) = nnz(\mathbf{A}')$ in the worst case, we expect that our Bloom filter allows us to discard many non-zeros of \mathbf{A}' without considering them in the computation.

C. Handling Transposition

Given an efficient local algorithm for SpGEMM that supports transposition, our algorithms can naturally be extended to the case where \mathbf{A} and/or \mathbf{B} are transposed. In the algorithm for algebraic updates, we can simply replace $\mathbf{A}^*_{k,i}$ and $\mathbf{B}^*_{j,k}$ by $\mathbf{A}^*_{i,k}$ and/or $\mathbf{B}^*_{k,j}$ if \mathbf{A} and/or \mathbf{B} are transposed. Furthermore, if \mathbf{A} and/or \mathbf{B} are transposed, we need to broadcast \mathbf{B}^* over rows and/or \mathbf{A}^* over columns of the process grid. In some cases, this allows us to get rid of the initial send/receive call since the blocks that are broadcasted are already on the right process row and/or column. Finally, the local matrix multiplication algorithm has to take transposition into account.

A similar strategy can be applied in the case of general updates. We note that in the transposed case, the Bloom filter can be used to discard rows of \mathbf{A}' and columns of \mathbf{B}' , respectively. Hence, it is still possible to filter the matrices efficiently, even if transposition is applied.

VI. IMPLEMENTATION DETAILS

A. Algebraic SpGEMM

Our implementation stores \mathbf{C}' as a dynamic matrix. \mathbf{A}^* and \mathbf{B}^* are stored in DCSR format as we expect them to be hypersparse. We also use the DCSR format when broadcasting blocks of matrices.

Our local multiplication uses Gustavson’s row-wise sparse matrix multiplication algorithm [32]. We use shared-memory parallelism to parallelize the computation of different rows of the result. Each shared-memory thread uses a sparse accumulator based on a dynamic array combined with a hash table for this purpose. We concatenate all output rows into a DCSR to form $\mathbf{X}_{k,j}^i$ and $\mathbf{Y}_{i,k}^j$.

Since $\mathbf{X}_{k,j}^i$ and $\mathbf{Y}_{i,k}^j$ are expected to have different sparsity patterns, we cannot use a straightforward MPI REDUCE call to aggregate them. Instead, we use an approach based on a custom reduce-scatter implementation for sparse matrices. Since the output of that aggregation (i.e., \mathbf{C}') is a dynamic matrix that supports efficient local updates, we do not need any auxiliary data structure (such as a SPA [33] or similar) during aggregation.

Table I: List of real-world instances.

Instance	Source	Type	n	nnz
LiveJournal	SNAP	Social	4 M	86 M
orkut	SNAP	Social	3 M	234 M
tech-p2p	Network Repository	Peer-to-Peer	5 M	295 M
indochina	Network Repository	Web	7 M	304 M
sinaweibo	Network Repository	Social	58 M	522 M
uk2002	Network Repository	Web	18 M	529 M
wikipedia	Network Repository	Web	27 M	1 088 M
PayDomain	Network Repository	Web	42 M	1 165 M
uk2005	Network Repository	Web	39 M	1 581 M
webbase	Network Repository	Web	118 M	1 736 M
twitter	Network Repository	Social	41 M	2 405 M
friendster	SNAP	Social	124 M	3 612 M

B. General SpGEMM

As in our algorithm for algebraic updates, we store the result matrices \mathbf{C}' and \mathbf{F}' as dynamic matrices and all intermediate matrices as DCSR.

To efficiently compute the local masked matrix multiplication, we locally build a hash table that stores the indices (u, v) of all non-zeros in $\mathbf{C}^*_{k,j}$ on process (i, j) . While this duplicates the work of building the same hash table on multiple MPI processes, we found it to be faster than broadcasting the hash table itself (instead of $\mathbf{C}^*_{k,j}$) in preliminary experiments; this is caused by the fact that the hash table is considerably larger than $nnz(\mathbf{C}^*_{k,j})$ due to its empty slots. We then perform a variant of Gustavson’s row-wise sparse matrix multiplication that uses the hash table to check whether (u, v) is non-zero in \mathbf{C}^* before adding an index pair (u, v) to the sparse aggregator.

VII. EXPERIMENTS

In this section, we present experiments to evaluate the performance of our algorithms in practice. We have implemented our algorithms in C++. The code of our algorithms will be published as open source software once this paper is accepted. We use CombBLAS 2.0, CTF 1.35 and PETSc 3.17.1 as state-of-the-art competitors. We note that CombBLAS seems to outperform both CTF and PETSc on our benchmarks; this result is in line with results by the authors of CombBLAS [14].

A. Experimental Setup

The experiments are performed on a 16-node compute cluster. Each compute node of the cluster features two Intel Xeon 6126 CPUs with 12 cores per CPU, and 192 GB RAM. The cluster is connected using 100 Gbit Intel Omni-Path Architecture interconnects. For CombBLAS, CTF and our algorithms, we run 4 MPI processes per node (i.e., two MPI processes per CPU socket) as these frameworks require a square processor grid. In this configuration, we run 6 OpenMP threads per MPI process. For PETSc, we found that a single MPI process per node (and 24 threads per MPI process) yields the best performance.³ Except in the parallel scalability experiments, we use all 16 compute nodes in each experiment.

³In experiments that have a fixed input size per MPI process (e.g., weak scaling experiments below), we adjust the input size of PETSc to ensure that all competitors operate on the same number of non-zeros.

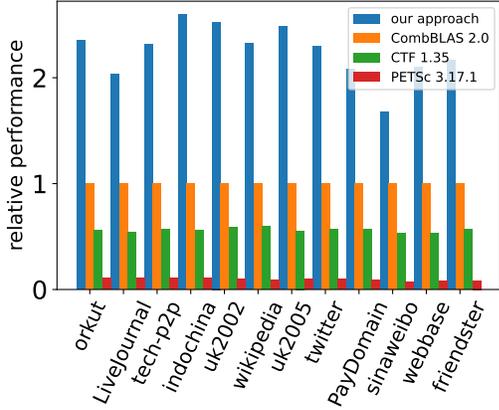


Figure 2: Construction

Figure 3: Matrix construction performance on real-world graphs. Performance is measured relative to CombBLAS.

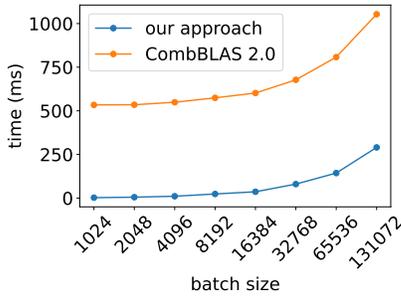


Figure 4: Mean insertion performance on real-world graphs. CTF and PETSc are considerably slower than our algorithm (at least $55.15\times$ slower, and $460.83\times$ slower, respectively) and not plotted.

Our experimental data consists of several large graphs – ranging from 86 million to almost 4 billion edges – and is shown in Table I. All instances were downloaded from the SNAP datasets [34] and Network Repository [35]. We always read graphs as undirected when creating adjacency matrices, i.e., for an edge $\{u, v\}$ in the input data, we add non-zeros (u, v) and (v, u) to the matrix.

The instances we use demonstrate significant imbalance without remapping. To avoid load imbalance, we randomly permute input indices before constructing each matrix. Given our 2D representation (and the memory layouts of our competitors), random permutation provides an adequate distribution of the input data. The same I/O and same permutation method is used for our code and for our competitors. We do not measure I/O times in any of our experiments.

B. Performance of Data Structures for Dynamic Distributed Graphs

In our first experiments, we evaluate the performance and parallel scalability of our dynamic distributed matrix data structure and our redistribution algorithm for updates.

a) *Construction*: We measure the time that it takes to construct the adjacency matrix of each input graph. We

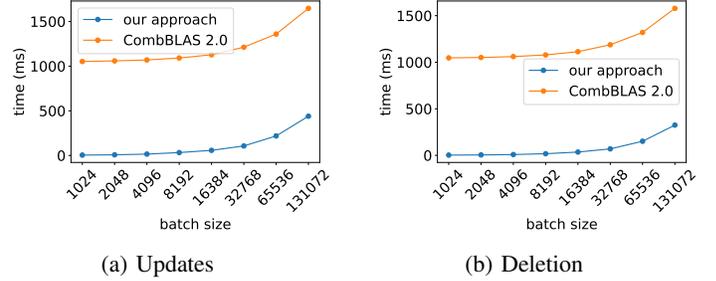


Figure 5: Mean update and deletion performance on real-world graphs. CTF and PETSc are not plotted, but at least $59.8\times$ and $477.11\times$ slower than our algorithm for updates. For deletions, CTF is at least $101.43\times$ slower than our algorithm (PETSc does not support efficient deletions).

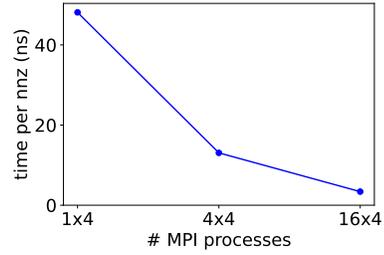


Figure 6: Weak scalability of insertions on real-world graphs, for different numbers of compute nodes. Each MPI process performs 1.3M insertions.

use our dynamic matrix data structure for this purpose; for CombBLAS, we use its default DCSC data structure, PETSc uses its parallel compressed sparse row implementation.

Figure 2 depicts the results of this experiment. We report the performance relative to CombBLAS. Our redistribution algorithm outperforms all state-of-the-art competitors. Our code is between $1.68\times$ and $2.59\times$ faster than CombBLAS (the best competitor); both CTF and PETSc are slower than both CombBLAS and our code on every instance. This difference in performance is due to two reasons: first, our dynamic matrix data structure allows for fast insertion and exploits shared-memory parallelism effectively. Secondly, our redistribution of non-zeros is faster than the competitor’s (which consists of a comparison sort and a global ALLTOALL in the case of CombBLAS). In fact, even if we construct a DCSR (i.e., the same type of data structure that CombBLAS uses) instead of

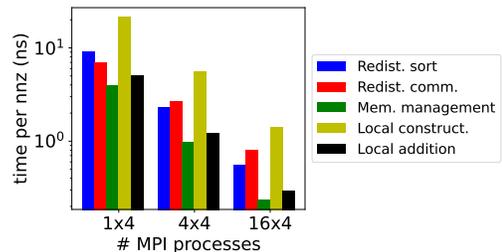
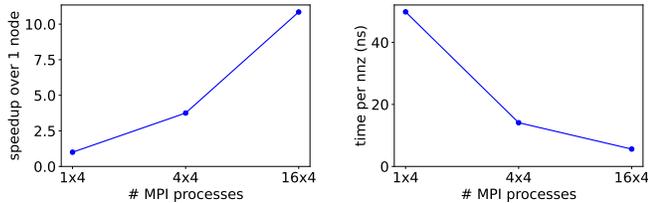


Figure 7: Breakdown of insertion running time on real-world graphs, for different numbers of compute nodes.



(a) Strong scalability, 2^{30} insertions in total. (b) Weak scalability, 2^{28} insertions per MPI process.

Figure 8: Parallel scalability of insertions on synthetic R-MAT graphs, for different numbers of compute nodes.

a dynamic matrix, we are still on average $1.15\times$ faster than CombBLAS, although CombBLAS constructs its DCSC data structure faster on some instances.

b) Insertions: We now compare the performance of inserting new non-zeros into the adjacency matrix, for our implementation versus the state-of-the-art competitors. In this experiment, we insert half of the non-zeros initially; this does not contribute to the running time we measure. Afterwards, we insert randomly chosen non-zeros from the remaining half into the already existing matrix, by first constructing an update matrix (in hypersparse layout for the competitors that support it) and adding the update matrix to the adjacency matrix (which uses a dynamic storage for our implementation). Insertions are performed in batches of various sizes. The batch size denotes the number of insertions that *each* MPI process performs – i.e., the full update matrix has (batch size \times p)-many non-zeros. In our experiments, this results in between 65K to 8.3M non-zeros in each update matrix. We perform 10 batches per instance (such that up to 83M entries are inserted into the already existing matrix).

Results are depicted in Figure 4. Since all competitors have to use a static matrix data structure that they have to rebuild after insertions, we outperform the competitors. In particular, we outperform CombBLAS by $3.63\times$ (for batch size 131072) to $227.68\times$ (for batch size 1024). As expected, our speedup over CombBLAS decreases with the batch size; as update matrices become denser, the cost of rebuilding the output matrix amortizes more effectively. Compared to CTF, we are always at least $55.15\times$ faster, and compared to PETSc, we are at least $460.83\times$ faster.

c) Updates and Deletions: To evaluate updates and deletions of the matrix, we proceed similarly to the insertion experiment. However, for update and deletion experiments, we insert the full adjacency matrix initially (and only draw non-zeros for the update matrix from existing non-zeros of the adjacency matrix). We note that PETSc does not support an efficient way to mask non-zeros in matrices; thus, we do not compare against PETSc for deletions.

The results are visualized in Figure 5a and 5b. Similarly to the insertion case, we are $3.75\times$ to $263.57\times$ faster for updates and $4.86\times$ to $393.85\times$ faster for deletions than CombBLAS. Likewise, our algorithm performs updates at least $59.8\times$ faster than CTF and at least $477.11\times$ faster than PETSc. For deletions, we are always at least $101.43\times$ faster than CTF.

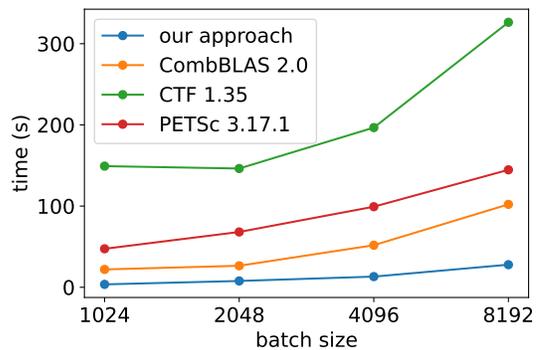


Figure 9: Mean performance of dynamic SpGEMM, **algebraic case**, on real-world graphs. Batch sizes are *per MPI process*, i.e., update matrices have up to 524K non-zeros.

d) Parallel Scalability of Graph Updates: We evaluate the scalability of insertions across different numbers of compute nodes. We vary the number of compute nodes between 1, 4 and 16 (such that the process grid is square), while keeping the OpenMP thread count constant at 6, and the number of MPI processes per node constant at 4. Insertions are performed in the same way as in the previous section, with the batch size fixed to 131072. This can be seen as a weak scaling experiment since nnz/p remains constant, where nnz denotes the number of non-zeros in the update matrix.

Figure 6 depicts the results of this experiment. In particular, the time per non-zero decreases with increasing numbers of compute nodes, indicating that our algorithm does not yet hit a scalability bottleneck on our cluster. Figure 7 shows that all steps of our algorithm scale well with the number of compute nodes. A large fraction of the running time consists of local operations as opposed to communication.

e) Scalability on Synthetic Graphs: In Figure 8, we present strong scaling and weak scaling results for insertions into our data structure on synthetic R-MAT graphs. We use the same R-MAT parameters as the Graph500 benchmark. In the strong scaling experiment, each MPI process generates $2^{30}/p$ non-zeros according to the R-MAT model. In the weak scaling experiment, each process generates 2^{28} non-zeros. As in our experiments on real-world graphs, we apply a (global) permutation to the row/column indices to ensure that load is evenly balanced across the process grid. We use a batch size of 131072 entries and insert all entries into a dynamic matrix. Our algorithm scales well with increasing numbers of MPI processes: for 16 compute nodes, we achieve a strong speedup of $10.85\times$ over a single compute node. Likewise, in the weak scaling model, our time per non-zero drastically decreases with the number of compute nodes, indicating that our algorithm does not hit an efficiency bottleneck in the configurations that we tested.

C. Performance of Dynamic SpGEMM

We perform experiments on both dynamic SpGEMM algorithms that are presented in Section V.

a) Algebraic case: This experiment repeatedly computes $C' = A'B$, where in each iteration we insert non-zeros

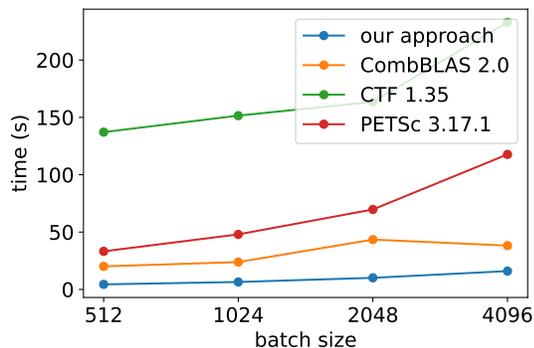


Figure 10: Mean performance of dynamic SpGEMM, **general case**, on real-world graphs. Batch sizes are *per MPI process*, i.e., update matrices have up to 262K non-zeros.

into A' , while B is static. In particular, we initialize A' to an empty matrix (i.e., the zero matrix) for each group. B is initialized to the full adjacency matrix of the graph. Initialization of B is not included in the measured runtime. Afterwards, insertions are performed in batches, with batch sizes between 1024 and 8192 vertices per MPI process. These insertions equate to the expression $A' = A + A^*$, where A^* is the matrix of insertions. We draw the insertions from the adjacency matrix of the graph (i.e., drawing all possible insertions would result in the computation of A^2). Each MPI process draws insertions individually, independently, and uniformly at random. Furthermore, the method (and random seed) to draw non-zeros is the same for our competitors and for our approach. We perform 10 batches per instance, therefore the total number of non-zeros in the left-hand side varies between 655K and 5.2M (= number of batches per instance \times batch size \times number of MPI processes). Batch sizes of over 8192 produced result matrices that did not fit in the RAM of our compute cluster, both for the competitors and for our algorithm, hence they were excluded. We perform the multiplication in the $(+, \cdot)$ semiring, so we can utilize the algorithm for algebraic updates presented in Section V-A. As previously demonstrated, our algorithm reduces $C' = A'B$ to $C' = C + A*B$, provided B is static. Our competitors compute $A*B$ using their distributed SpGEMM algorithms and add the result to C .

The experiment (results shown in Figure 9) demonstrates that our algorithm is $3.41\times$ (for a batch size of 8192) to $6.18\times$ (for batch size 1024) faster than CombBLAS (which is the best competitor in this experiment). We are also at least $11.73\times$ faster than CTF and at least $5.2\times$ faster than PETSc. As expected, the speedup decreases for increasing batch sizes. In particular, for large batch sizes, update matrices are not hypersparse on all instances anymore. In these cases, our algorithm is expected to perform worse than SUMMA due to its more complicated communication pattern.

b) General case: This experiment is performed using the same setup as before, but using the general dynamic SpGEMM algorithm described in Section V-B instead. We use a $(\min, +)$ semiring to differentiate from the algebraic case. To perform an equivalent operation using our competitors,

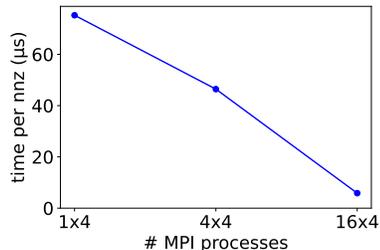


Figure 11: Weak scalability of dynamic SpGEMM (algebraic case) on real-world graphs, for different numbers of compute nodes. This experiment uses 81920 non-zeros per MPI process.

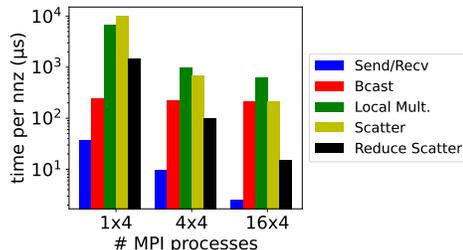


Figure 12: Breakdown of dynamic SpGEMM (algebraic case) running time on real-world graphs, for different numbers of compute nodes.

it is no longer enough to compute just $A*B$, we have to recompute $A'B$ from scratch (since insertions into the matrix are incompatible with the \min operator that the semiring uses for addition). Unlike our other competitors, PETSc does not support arbitrary semirings; thus, we continue to use the $(+, \cdot)$ semiring for PETSc in this experiment.

The results for this experiment are depicted in Figure 10. Our findings show that – depending on the batch size – our algorithm is $2.39\times$ to $4.57\times$ faster than CombBLAS, the best competitor in this experiment. Additionally, we are always at least $14.58\times$ faster than CTF and at least $6.9\times$ faster than PETSc. While our competitors communicate all non-zeros of the left-hand side, our algorithm can avoid this using our Bloom filter for entries that do not contribute to the result. As the matrix becomes denser, the probability that a non-zero does not contribute to the result decreases. Hence for larger batch sizes it is expected to be more efficient to simply transfer all non-zeros, because the overhead of a Bloom filter is less economical.

c) Parallel Scalability of Dynamic SpGEMM: To test the scalability of our algorithm, we perform the SpGEMM experiment with algebraic updates on a varying number of MPI processes. We change the number of compute nodes between 1, 4 and 16, while keeping the OpenMP thread count constant at 6, and the number of MPI processes per node constant at 4, as with previous experiments. We excluded instances PayDomain, wikipedia, webbase, uk2005, friendster, and twitter for these experiments as the number of non-zeros generated by the SpGEMM did not fit into the RAM of a single compute node. The batch size in this experiment is fixed to 8192.

Figure 11 shows the relative runtime for different numbers of MPI processes. The time that our algorithm takes per non-zero decreases with the number of MPI processes, indicating that we do not yet hit a performance bottleneck on our cluster. The findings in Figure 11 are supported by Figure 12. Local multiplication, reduce/scatter and initial send/receive rounds scale well with the number of compute nodes. However, broadcasting matrices takes a larger fraction of the overall running time for higher numbers of compute nodes (as expected).

VIII. CONCLUSIONS

In this paper we proposed a data structure for dynamic sparse graphs/matrices distributed over MPI processes. This data structure allows for fast updates and redistribution. With this data structure and an adapted communication mechanism, we designed a dynamic SpGEMM algorithm that usually performs several times faster in practice than the static state of the art. Future work could investigate replacing the SUMMA algorithm in the 3D SpGEMM by Azad et al. [21] by our algorithm to obtain a new dynamic one with further improved communication volume.

Acknowledgments: This work is partially supported by German Research Foundation (DFG) grant GR 5745/1-1 (DyANE) and DFG grant ME 3619/4-1 (ALMACOM).

REFERENCES

- [1] K. Stüben, “A review of algebraic multigrid,” in *Numerical Analysis: Historical Developments in the 20th Century*, C. Brezinski and L. Wuytack, Eds. Amsterdam: Elsevier, 2001, pp. 331–359. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/B978044450617750015X>
- [2] J. Kepner, S. Alford, V. Gadepally, M. Jones, L. Milechin, R. A. Robinett, and S. Samsi, “Sparse deep neural network graph challenge,” in *2019 IEEE High Performance Extreme Computing Conference, HPEC 2019, Waltham, MA, USA, September 24-26, 2019*. IEEE, 2019, pp. 1–7. [Online]. Available: <https://doi.org/10.1109/HPEC.2019.8916336>
- [3] S. Van Dongen, “Graph clustering via a discrete uncoupling process,” *SIAM Journal on Matrix Analysis and Applications*, vol. 30, no. 1, pp. 121–141, 2008. [Online]. Available: <https://doi.org/10.1137/040608635>
- [4] A. Yaşar, S. Rajamanickam, M. Wolf, J. Berry, and Ü. V. Çatalyürek, “Fast triangle counting using cilk,” in *2018 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE, 2018, pp. 1–7.
- [5] J. Kepner and J. R. Gilbert, Eds., *Graph Algorithms in the Language of Linear Algebra*, ser. Software, environments, tools. SIAM, 2011, vol. 22.
- [6] J. Kepner, P. Aaltonen, D. A. Bader, A. Buluç, F. Franchetti, J. R. Gilbert, D. Hutchison, M. Kumar, A. Lumsdaine, H. Meyerhenke, S. McMillan, C. Yang, J. D. Owens, M. Zalewski, T. G. Mattson, and J. E. Moreira, “Mathematical foundations of the graphblas,” in *2016 IEEE High Performance Extreme Computing Conference, HPEC 2016, Waltham, MA, USA, September 13-15, 2016*. IEEE, 2016, pp. 1–9. [Online]. Available: <https://doi.org/10.1109/HPEC.2016.7761646>
- [7] J. Gao, W. Ji, Z. Tan, and Y. Zhao, “A systematic survey of general sparse matrix-matrix multiplication,” *CoRR*, vol. abs/2002.11273, 2020. [Online]. Available: <https://arxiv.org/abs/2002.11273>
- [8] T. G. Mattson, C. Yang, S. McMillan, A. Buluç, and J. E. Moreira, “Graphblas C API: ideas for future versions of the specification,” in *2017 IEEE High Performance Extreme Computing Conference, HPEC 2017, Waltham, MA, USA, September 12-14, 2017*. IEEE, 2017, pp. 1–6. [Online]. Available: <https://doi.org/10.1109/HPEC.2017.8091095>
- [9] M. Winter, D. Mlakar, R. Zayer, H. Seidel, and M. Steinberger, “Adaptive sparse matrix-matrix multiplication on the GPU,” in *Proceedings of the 24th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP 2019, Washington, DC, USA, February 16-20, 2019*, J. K. Hollingsworth and I. Keidar, Eds. ACM, 2019, pp. 68–81. [Online]. Available: <https://doi.org/10.1145/3293883.3295701>
- [10] Y. Nagasaka, S. Matsuoka, A. Azad, and A. Buluç, “High-performance sparse matrix-matrix products on intel KNL and multicore architectures,” in *The 47th International Conference on Parallel Processing, ICPP 2018, Workshop Proceedings, Eugene, OR, USA, August 13-16, 2018*. ACM, 2018, pp. 34:1–34:10. [Online]. Available: <https://doi.org/10.1145/3229710.3229720>
- [11] Z. Gu, J. Moreira, D. Edelsohn, and A. Azad, “Bandwidth optimized parallel algorithms for sparse matrix-matrix multiplication using propagation blocking,” in *SPAA '20: 32nd ACM Symposium on Parallelism in Algorithms and Architectures, Virtual Event, USA, July 15-17, 2020*, C. Scheideler and M. Spear, Eds. ACM, 2020, pp. 293–303. [Online]. Available: <https://doi.org/10.1145/3350755.3400216>
- [12] J. A. Ellis and S. Rajamanickam, “Scalable inference for sparse deep neural networks using kokkos kernels,” in *2019 IEEE High Performance Extreme Computing Conference, HPEC 2019, Waltham, MA, USA, September 24-26, 2019*. IEEE, 2019, pp. 1–7. [Online]. Available: <https://doi.org/10.1109/HPEC.2019.8916378>
- [13] Z. Xie, G. Tan, W. Liu, and N. Sun, “A pattern-based spgemm library for multi-core and many-core architectures,” *IEEE Trans. Parallel Distributed Syst.*, vol. 33, no. 1, pp. 159–175, 2022. [Online]. Available: <https://doi.org/10.1109/TPDS.2021.3090328>
- [14] A. Azad, O. Selvitopi, M. T. Hussain, J. Gilbert, and A. Buluç, “Combinatorial blas 2.0: Scaling combinatorial algorithms on distributed-memory systems,” *IEEE Trans. Parallel and Distributed Systems*, pp. 1–1, 2021.
- [15] T. Trilinos Project Team, *The Trilinos Project Website*.
- [16] E. Solomonik, D. Matthews, J. R. Hammond, J. F. Stanton, and J. Demmel, “A massively parallel tensor contraction framework for coupled-cluster computations,” *Journal of Parallel and Distributed Computing*, vol. 74, no. 12, pp. 3176–3190, 2014, domain-Specific Languages and High-Level Frameworks for High-Performance Computing. [Online]. Available: <https://www.sciencedirect.com/science/article/pii/S074373151400104X>
- [17] T. Brandes, E. Schricker, and T. Soddemann, “The LAMA approach for writing portable applications on heterogeneous architectures,” in *Scientific Computing and Algorithms in Industrial Simulations, Projects and Products of Fraunhofer SCAI*, M. Griebel, A. Schüller, and M. A. Schweitzer, Eds. Springer, 2017, pp. 181–198. [Online]. Available: https://doi.org/10.1007/978-3-319-62458-7_9
- [18] S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith, “Efficient management of parallelism in object oriented numerical software libraries,” in *Modern Software Tools in Scientific Computing*, E. Arge, A. M. Bruaset, and H. P. Langtangen, Eds. Birkhäuser Press, 1997, pp. 163–202.
- [19] A. Buluç, J. R. Gilbert, and V. B. Shah, “Implementing sparse matrices for graph algorithms,” in *Graph Algorithms in the Language of Linear Algebra*, ser. Software, environments, tools, J. Kepner and J. R. Gilbert, Eds. SIAM, 2011, vol. 22, pp. 287–313. [Online]. Available: <https://doi.org/10.1137/1.9780898719918.ch13>
- [20] A. Buluç and J. R. Gilbert, “New ideas in sparse matrix matrix multiplication,” in *Graph Algorithms in the Language of Linear Algebra*, ser. Software, environments, tools, J. Kepner and J. R. Gilbert, Eds. SIAM, 2011, vol. 22, pp. 315–337. [Online]. Available: <https://doi.org/10.1137/1.9780898719918.ch14>
- [21] A. Azad, G. Ballard, A. Buluç, J. Demmel, L. Grigori, O. Schwartz, S. Toledo, and S. Williams, “Exploiting multiple levels of parallelism in sparse matrix-matrix multiplication,” *SIAM Journal on Scientific Computing*, vol. 38, no. 6, pp. C624–C651, 2016. [Online]. Available: <https://doi.org/10.1137/15M104253X>
- [22] A. Buluç and K. Madduri, “Graph partitioning for scalable distributed graph computations,” in *Graph Partitioning and Graph Clustering, 10th DIMACS Implementation Challenge Workshop, Georgia Institute of Technology, Atlanta, GA, USA, February 13-14, 2012. Proceedings*, ser. Contemporary Mathematics, D. A. Bader, H. Meyerhenke, P. Sanders, and D. Wagner, Eds., vol. 588. American Mathematical Society, 2012, pp. 83–102. [Online]. Available: <http://www.ams.org/books/conm/588/11709>
- [23] H. Jin, N. Wang, D. Yu, Q.-S. Hua, X. Shi, and X. Xie, “Core maintenance in dynamic graphs: A parallel approach based on matching,” *IEEE Transactions on Parallel and Distributed Systems*, vol. 29, no. 11, pp. 2416–2428, 2018.
- [24] E. Bergamini, H. Meyerhenke, M. Ortman, and A. Slobbe, “Faster betweenness centrality updates in evolving networks,” in *16th International Symposium on Experimental Algorithms, SEA 2017, June 21-23, 2017, London, UK*, ser. LIPIcs, C. S. Iliopoulos, S. P. Pissis, S. J. Puglisi, and R. Raman, Eds., vol. 75. Schloss Dagstuhl

- Leibniz-Zentrum für Informatik, 2017, pp. 23:1–23:16. [Online]. Available: <https://doi.org/10.4230/LIPIcs.SEA.2017.23>
- [25] A. van der Grinten, E. Bergamini, O. Green, D. A. Bader, and H. Meyerhenke, “Scalable katz ranking computation in large static and dynamic graphs,” in *26th Annual European Symposium on Algorithms, ESA 2018, August 20–22, 2018, Helsinki, Finland*, ser. LIPIcs, Y. Azar, H. Bast, and G. Herman, Eds., vol. 112. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018, pp. 42:1–42:14. [Online]. Available: <https://doi.org/10.4230/LIPIcs.ESA.2018.42>
- [26] E. J. Riedy, “Updating pagerank for streaming graphs,” in *2016 IEEE International Parallel and Distributed Processing Symposium Workshops, IPDPS Workshops 2016, Chicago, IL, USA, May 23–27, 2016*. IEEE Computer Society, 2016, pp. 877–884. [Online]. Available: <https://doi.org/10.1109/IPDPSW.2016.22>
- [27] A. van der Grinten, M. Predari, and F. Willich, “A fast data structure for dynamic graphs based on hash-indexed adjacency blocks,” Technical report and code available at <http://github.com/hu-macsy/dhb>, 2022.
- [28] A. Buluç and J. R. Gilbert, “On the representation and multiplication of hypersparse matrices,” in *IPDPS*. IEEE, 2008, pp. 1–11.
- [29] E. Solomonik and T. Hoefer, “Sparse tensor algebra as a parallel programming model,” *CoRR*, vol. abs/1512.00066, 2015.
- [30] M. D. Schatz, R. A. van de Geijn, and J. Poulson, “Parallel matrix multiplication: A systematic journey,” *SIAM J. Sci. Comput.*, vol. 38, no. 6, 2016. [Online]. Available: <https://doi.org/10.1137/140993478>
- [31] A. Azad, A. Buluç, and J. R. Gilbert, “Parallel triangle counting and enumeration using matrix algebra,” in *IPDPS Workshops*. IEEE Computer Society, 2015, pp. 804–811.
- [32] F. G. Gustavson, “Two fast algorithms for sparse matrices: Multiplication and permuted transposition,” *ACM Trans. Math. Softw.*, vol. 4, no. 3, p. 250–269, Sep. 1978. [Online]. Available: <https://doi.org/10.1145/355791.355796>
- [33] J. R. Gilbert, C. Moler, and R. Schreiber, “Sparse matrices in matlab: Design and implementation,” *SIAM J. Matrix Anal. Appl.*, vol. 13, pp. 333–356, 1992.
- [34] J. Leskovec and A. Krevl, “SNAP Datasets: Stanford large network dataset collection,” <http://snap.stanford.edu/data>, Jun. 2014.
- [35] R. A. Rossi and N. K. Ahmed, “The network data repository with interactive graph analytics and visualization,” in *AAAI*, 2015. [Online]. Available: <https://networkrepository.com>