Hypergraph Partitioning for Sparse Matrix-Matrix Multiplication

GREY BALLARD, Sandia National Laboratories, gmballa@sandia.gov
ALEX DRUINSKY, Lawrence Berkeley National Laboratory, adruinsky@lbl.gov
NICHOLAS KNIGHT, New York University, nknight@nyu.edu
ODED SCHWARTZ, Hebrew University, odedsc@cs.huji.ac.il

We propose a fine-grained hypergraph model for sparse matrix-matrix multiplication (SpGEMM), a key computational kernel in scientific computing and data analysis whose performance is often communication bound. This model correctly describes both the interprocessor communication volume along a critical path in a parallel computation and also the volume of data moving through the memory hierarchy in a sequential computation. We show that identifying a communication-optimal algorithm for particular input matrices is equivalent to solving a hypergraph partitioning problem. Our approach is sparsity dependent, meaning that we seek the best algorithm for the given input matrices.

In addition to our (3D) fine-grained model, we also propose coarse-grained 1D and 2D models that correspond to simpler SpGEMM algorithms. We explore the relations between our models theoretically, and we study their performance experimentally in the context of three applications that use SpGEMM as a key computation. For each application, we find that at least one coarse-grained model is as communication efficient as the fine-grained model. We also observe that different applications have affinities for different algorithms.

Our results demonstrate that hypergraphs are an accurate model for reasoning about the communication costs of SpGEMM as well as a practical tool for exploring the SpGEMM algorithm design space.

1. INTRODUCTION

Sparse matrix-matrix multiplication (SpGEMM) is a fundamental computation in scientific computing and data analysis. It is a key component in applications ranging from linear solvers [Brezina and Vassilevski 2011; Yamazaki and Li 2011] and graph algorithms [Rabin and Vazirani 1989; Azad et al. 2015b] to Kohn-Sham theory in computational chemistry [Borštnik et al. 2014]. Unlike dense matrix-matrix multiplication, SpGEMM is an irregular computation, and its performance is typically communication bound. Previous research on SpGEMM algorithms has focused on communication costs, both interprocessor communication and data movement within the memory hierarchy [Buluç and Gilbert 2012; Ballard et al. 2013; Azad et al. 2015a].

We focus on communication costs and argue in this paper that the amount of communication required for sequential and parallel SpGEMM algorithms depends strongly on the matrices' nonzero structures. That is, algorithms that are communication efficient for SpGEMMs within one application are not necessarily the most efficient in the context of another application. The simplest SpGEMM algorithms to design, analyze, and implement are based on dense matrix multiplication algorithms: these algorithms are parametrized only by the matrix dimensions and differ from dense matrix-matrix multiplication in that they use sparse data structures and avoid operating on and communicating zeros. We refer to these algorithms as "sparsity independent" [Ballard et al. 2013, Def. 2.5]; two examples are the row-wise algorithm with block partitioning [Faisal et al. 2014] and Sparse SUMMA (SpSUMMA) [Buluç and Gilbert 2012].

One of the main goals in this paper is to explore sparsity-dependent SpGEMM algorithms and determine how communication efficient an algorithm can be if it is able to inspect the input matrices' nonzero structures and determine the data distributions and parallelization of arithmetic operations accordingly. Currently this approach is only theoretical in general: in practice, the cost of nonzero structure inspection can exceed that of even an inefficient SpGEMM. However, practical, application-specific

algorithmic insight can be gained from the results of this paper and our techniques can be used for other applications not considered here.

We use hypergraphs to model arithmetic operations and dependencies and hypergraph partitioning to model parallelizing and scheduling operations. Hypergraphs have previously been used to model sparse matrix-vector multiplication (SpMV) (see, *e.g.*, [Vastenhouw and Bisseling 2005; Çatalyürek et al. 2010]), and hypergraph partitioning software has been developed to help design algorithms for SpMV and other computations [Çatalyürek and Aykanat 1999; Boman et al. 2007]. They have also been used in the context of SpGEMM [Akbudak and Aykanat 2014; Ballard et al. 2015a]; in this paper we present a hypergraph model that generalizes previous models for SpMV and SpGEMM. We discuss these and other related works in Sec. 2.

We present a general, fine-grained hypergraph model for SpGEMM in Sec. 3 and use it to prove sparsity-dependent communication lower bounds for both parallel (Sec. 4.1) and sequential (Sec. 4.2) algorithms. We reduce identifying a communication-optimal algorithm for given input matrices to solving a hypergraph partitioning problem. We also compare with previous lower bounds, showing that ours are more general and can be tighter.

In addition to our general SpGEMM model, we also consider several restricted classes of algorithms. In Sec. 5 we define a framework for coarsening the fine-grained model to obtain simpler hypergraphs that maintain correct modeling of communication costs. In particular, we consider the classification of 1D, 2D, and 3D SpGEMM algorithms [Ballard et al. 2013] and examine the relationships among the seven subclasses of algorithms that naturally arise from this classification. The class of 1D algorithms includes row-wise and outer-product algorithms, and the class of 2D algorithms includes SpSUMMA. Finding communication-optimal algorithms within each subclass also corresponds to solving simpler hypergraph partitioning problems.

Our experimental results, presented in Sec. 6, use existing hypergraph partitioning software to compare the communication costs of algorithms from each subclass for a variety of SpGEMMs, with input matrices coming from three particular applications. We use representative examples from each application to explore the algorithmic space of SpGEMM; our results can guide algorithmic design choices that are application specific. In particular, our empirical results lead to different conclusions for each of the three applications we consider. For example, we find that certain 1D algorithms, although simple, are sufficient in the context of an algebraic multigrid application [Brezina and Vassilevski 2011] but that 2D and 3D algorithms are much more communication efficient than 1D algorithms in the context of Markov clustering applied to social-network matrices [Satuluri and Parthasarathy 2009]. We discuss overall conclusions in more detail in Sec. 7.

As is the case for other irregular computations, there exists a tradeoff between simplicity and efficiency for SpGEMM algorithms and implementations. Using hypergraphs to model SpGEMM communication costs, we provide a mechanism for quantifying the relative efficiency of algorithms of varying complexity. We illuminate this tradeoff in the context of several applications, providing practical insight into the design of algorithms and software.

2. RELATED WORK

We base our communication models and analysis on classical results for dense matrix multiplication. Hong and Kung [1981] used a two-level sequential memory model and established communication lower bounds for dense matrix multiplication and other computations using a graph-theoretic approach. Irony, Toledo, and Tiskin [2004] used a distributed-memory parallel model and established communication lower bounds

for dense matrix multiplication using a geometric approach. See Ballard et al.'s survey [2014] for more details on these models and results.

Ballard et al. [2011; 2012] proved communication lower bounds, in both sequential and parallel models, for a general set of matrix computations that includes SpGEMM. These lower bounds are parameterized by the number of multiplications involved in the multiplication of the specific input matrices, and they are not tight in general. There has been other theoretical work in deriving bounds on the communication costs of SpGEMM in the sequential two-level memory model. Pagh and Stöckel [2014] proved matching worst-case upper and lower bounds based on the number of nonzeros in the matrices. Greiner [2012, Ch. 6] used a different approach to establish other worst-case lower and upper bounds that apply to a restricted set of matrices and algorithms. We compare our lower bounds with previous work in more detail in Secs. 4.1 and 4.2.

On the practical side of sequential SpGEMM algorithms, Gustavson [1978] proposed the first algorithm that exploits sparsity in the inputs and outputs, and Davis [2006] implemented a variant of that algorithm in the general-purpose CSPARSE library. Neither the algorithm nor implementation are designed to specifically reduce communication costs specifically, but the approach is typically efficient in terms of memory footprint and operation count. Buluç and Gilbert [2008] proposed an alternative sparse data structure and sequential algorithm that are more memory and operation-count efficient than Gustavson's algorithm for "hypersparse" matrices, which can arise in the context of parallel distribution of sparse matrices.

There exists more variation in parallel SpGEMM algorithms. We use the 1D/2D/3D classification of SpGEMM algorithms defined in [Ballard et al. 2013]. They surveyed many of the previous algorithms in the literature and theoretically analyzed their communication costs for a particular class of input matrices (Erdős-Rényi random matrices), comparing with expectation-based communication lower bounds. Buluç and Gilbert [2012] proposed a general-purpose 2D algorithm called Sparse SUMMA, which uses random permutations to achieve load balance, and they analyzed its communication cost for general inputs. Ballard, Siefert, and Hu [2015b] considered the communication costs of 1D algorithms in an algebraic multigrid application.

As a tool for minimizing communication, hypergraph partitioning has gained popularity in the parallel computing community in large part due to its application SpMV. While there is a vast literature on hypergraph partitioning for SpMV, we highlight here only a small sample from which we borrow notation. Çatalyürek and Aykanat [1999] introduced column-net and row-net models for 1D SpMV algorithms, and in a subsequent paper [2001] presented the fine-grain model for 2D SpMV. Çatalyürek, Aykanat, and Uçar [2010] summarize and compare models for 1D and 2D SpMV algorithms (see also [Vastenhouw and Bisseling 2005]). Akbudak, Kayaaslan, and Aykanat [2013] also used hypergraphs to improve cache locality of SpMV on a single processor.

Hypergraph partitioning has been used in the context of SpGEMM as well. Krishnamoorthy et al. [2006] introduced a hypergraph model for a general class of computations including SpGEMM, and described a heuristic for scheduling out-of-core algorithms with the goal of minimizing disk I/O. Akbudak and Aykanat [2014] introduced a hypergraph model for parallel outer-product algorithms that represents both the computation and the data involved in SpGEMM (see Ex. 5.2); they also implemented and benchmarked the performance based on the partitions. Our earlier conference paper [2015a] introduced a slightly simpler version of the SpGEMM hypergraph model given by Def. 3.1; that paper also presented initial results concerning the application of algebraic multigrid (see Sec. 6.1).

Fig. 1: Notation for a particular SpGEMM instance.

3. HYPERGRAPH MODEL FOR SPGEMM

3.1. Notation and Assumptions

The notation introduced in this section will be used throughout the rest of the paper.

Let $\mathbb{N} = \{0, 1, \ldots\}$; for any $n \in \mathbb{N}$, $[n] = \{1, \ldots, n\}$, with $[0] = \emptyset$. For any $n \in \mathbb{N}$ and set X, an n-way partition of X, denoted $\{X_1, \ldots, X_n\}$, is a function $[n] \ni i \mapsto X_i \subseteq X$ such that $\bigcup_{i \in [n]} X_i = X$ and $\bigcup_{i \neq j \in [n]} X_i \cap X_j = \emptyset$. Another partition of X, $\{Y_1, \ldots, Y_m\}$, refines $\{X_1, \ldots, X_n\}$ if for every $j \in [m]$ there exists an $i \in [n]$ such that $Y_j \subseteq X_i$.

Let $\mathbf{A}$ and $\mathbf{B}$ be I-by-K and K-by-J matrices with entries from a set X, which contains an element 0 and is closed under two binary operations called addition (commutative and associative with identity element 0) and multiplication (with absorbing element 0). For example, X could be a semiring.

Matrix multiplication is the function $(\mathbf{A}, \mathbf{B}) \mapsto \mathbf{C} = \mathbf{A} \cdot \mathbf{B}$, where $\mathbf{C}$ is an I-by-J matrix over X defined entrywise by $c_{ij} = \sum_{k \in [K]} a_{ik} b_{kj}$. We let $S_\mathbf{A} \subseteq [I] \times [K]$, $S_\mathbf{B} \subseteq [K] \times [J]$, and $S_\mathbf{C} \subseteq [I] \times [J]$ denote the nonzero structures of $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$. In this work, we study *SpGEMM algorithms*, the class of algorithms that evaluate and sum all *nontrivial* multiplications $a_{ik} b_{kj}$, where both $a_{ik} \neq 0$ and $b_{kj} \neq 0$, and thus depend only on $S_\mathbf{A}$ and $S_\mathbf{B}$. We do not consider algorithms that exploit additional structure on X or more general relations on the entries of $\mathbf{A}$ and $\mathbf{B}$: in particular, we ignore numerical cancellation, so $S_\mathbf{A}$ and $S_\mathbf{B}$ induce $S_\mathbf{C}$. Fig. 1 illustrates our notation for a particular SpGEMM instance; under our assumptions we do not distinguish the nonzero values of the matrices. We maintain the nonzero structure of the example in each of Figs. 1 to 4.

Given additional structure on X, there may exist many algorithms that do not simply compute the sums of products. For example, if X is a ring, then additive inverses can be exploited to obtain Strassen's algorithm, where the multiplicands do not generally equal the input matrix entries. As another example, given known relations on the entries of $\mathbf{A}$ and $\mathbf{B}$, one may be able to avoid evaluating some nontrivial multiplications.

To simplify the presentation, we assume that neither $\mathbf{A}$ or $\mathbf{B}$ have any zero rows or columns, so every nonzero input matrix entry participates in at least one nontrivial multiplication. While this can always be enforced, doing so may incur a preprocessing cost.

Fig. 2: Geometric view of the SpGEMM instance in Fig. 1. correspond to the iteration space of matrix multiplication. Each sub-cube represents a multiplication, which is nontrivial if its projections onto the **A** and **B** faces both correspond to nonzero values. The projections of nontrivial multiplications onto the **C** face define the nonzero structure of **C**.

Our hypergraph terminology borrows from Çatalyürek and Aykanat [1999]. A *hypergraph* $\mathcal{H}$ is a generalization of a graph; it consists of a set of *vertices* $\mathcal{V}$ and a set of *nets* $\mathcal{N}$, where each net $n \in \mathcal{N}$ is a subset of the vertices, $n \subseteq \mathcal{V}$. We use the term *pin* to refer to a vertex in a particular net. Each vertex $v \in \mathcal{V}$ may have a *weight* w associated with it, and each net $n \in \mathcal{N}$ may have a *cost* c associated with it. Weights and costs are typically scalar-valued but can be vector-valued; we will use vector-valued weights.

3.2. Fine-Grained Hypergraph Model

In this section we define our most general hypergraph model for SpGEMM. It may be helpful to visualize an instance of SpGEMM geometrically, as the three-dimensional "iteration space" of matrix multiplication: Fig. 2 illustrates this geometric perspective. The set of IKJ multiplications are arranged as a set of cubes so that nontrivial multiplications are distinguished by having nonzero projections on the **A**- and **B**-faces. Nonzero entries of **C** are determined by projecting the nontrivial multiplications onto the **C**-face.

We now define our hypergraph model and connect its definition to its geometric interpretation.

Definition 3.1. Consider matrices **A**, **B**, and **C** = **A** · **B**, continuing notation. Define the *fine-grained SpGEMM hypergraph* $\mathcal{H}(\mathbf{A}, \mathbf{B}) = (\mathcal{V}, \mathcal{N})$, where $\mathcal{V} = \mathcal{V}^{\mathrm{m}} \cup \mathcal{V}^{\mathrm{nz}}$, with

$$\mathcal{V}^{\mathrm{nz}} = \mathcal{V}^{\mathbf{A}} \cup \mathcal{V}^{\mathbf{B}} \cup \mathcal{V}^{\mathbf{C}},$$
$$\mathcal{V}^{\mathrm{m}} = \{v_{ikj} : (i,k) \in S_{\mathbf{A}} \wedge (k,j) \in S_{\mathbf{B}}\},$$
$$\mathcal{V}^{\mathbf{A}} = \{v_{ik}^{\mathbf{A}} : (i,k) \in S_{\mathbf{A}}\},$$
$$\mathcal{V}^{\mathbf{B}} = \{v_{kj}^{\mathbf{B}} : (k,j) \in S_{\mathbf{B}}\}, \text{ and}$$
$$\mathcal{V}^{\mathbf{C}} = \{v_{ij}^{\mathbf{C}} : (i,j) \in S_{\mathbf{C}}\},$$

and $\mathcal{N} = \mathcal{N}^{\mathbf{A}} \cup \mathcal{N}^{\mathbf{B}} \cup \mathcal{N}^{\mathbf{C}}$, with

$$\mathcal{N}^{\mathbf{A}} = \{n_{ik}^{\mathbf{A}} : (i,k) \in S_{\mathbf{A}}\},$$
$$\mathcal{N}^{\mathbf{B}} = \{n_{kj}^{\mathbf{B}} : (k,j) \in S_{\mathbf{B}}\}, \text{ and}$$
$$\mathcal{N}^{\mathbf{C}} = \{n_{ij}^{\mathbf{C}} : (i,j) \in S_{\mathbf{C}}\}.$$

Net membership is defined such that each $v_{ik}^{\mathbf{A}} \in n_{ik}^{\mathbf{A}}$, each $v_{kj}^{\mathbf{B}} \in n_{kj}^{\mathbf{B}}$, each $v_{ij}^{\mathbf{C}} \in n_{ij}^{\mathbf{C}}$, and each $v_{ikj} \in n_{ik}^{\mathbf{A}} \cap n_{kj}^{\mathbf{B}} \cap n_{ij}^{\mathbf{C}}$. Equivalently, pins are defined such that for each $(i,k) \in S_{\mathbf{A}}$,

$$n_{ik}^{\mathbf{A}} = \{v_{ikj} : (k,j) \in S_{\mathbf{B}}\} \cup \{v_{ik}^{\mathbf{A}}\},$$

for each $(j,k) \in S_{\mathbf{B}}$,

$$n_{jk}^{\mathbf{B}} = \{v_{ikj} : (i,k) \in S_{\mathbf{A}}\} \cup \{v_{kj}^{\mathbf{B}}\},$$

and for each $(i,j) \in S_{\mathbf{C}}$,

$$n_{ij}^{\mathbf{C}} = \{v_{ikj} : (i,k) \in S_{\mathbf{A}} \wedge (k,j) \in S_{\mathbf{B}}\} \cup \{v_{ij}^{\mathbf{C}}\}.$$

The vertices $\mathcal{V}$ have two types of *weights*, corresponding to computation and memory. For all $v_{ikj} \in \mathcal{V}^{\mathbf{m}}$,

$$w_{\mathbf{comp}}(v_{ikj}) = 1,$$
$$w_{\mathbf{mem}}(v_{ikj}) = 0.$$

For all $v_{ik}^{\mathbf{A}}, v_{kj}^{\mathbf{B}}, v_{ij}^{\mathbf{C}} \in \mathcal{V}^{\mathbf{m}}$,

$$w_{\mathbf{comp}}(v_{ik}^{\mathbf{A}}) = w_{\mathbf{comp}}(v_{kj}^{\mathbf{B}}) = w_{\mathbf{comp}}(v_{ij}^{\mathbf{C}}) = 0,$$
$$w_{\mathbf{mem}}(v_{ik}^{\mathbf{A}}) = w_{\mathbf{mem}}(v_{kj}^{\mathbf{B}}) = w_{\mathbf{mem}}(v_{ij}^{\mathbf{C}}) = 1.$$

The *cost* of each net $n \in \mathcal{N}$ is $c(n) = 1$.

In summary, each vertex v of $\mathcal{H}(\mathbf{A}, \mathbf{B})$ corresponds either to a nontrivial multiplication or to a nonzero entry of $\mathbf{A}$, $\mathbf{B}$, or $\mathbf{C}$, and each net corresponds to a nonzero of $\mathbf{A}$, $\mathbf{B}$, or $\mathbf{C}$. Each *multiplication vertex* $v \in \mathcal{V}^{\mathbf{m}}$ is a member of the three nets corresponding to nonzeros involved in the multiplication; each *nonzero vertex* $v \in \mathcal{V}^{\mathbf{nz}}$ is a member of its corresponding net. Conversely, each net corresponds to a nonzero value and includes as members its corresponding nonzero vertex as well all associated multiplication vertices.

In the geometric context of Fig. 2, we can think of each multiplication vertex v_{ikj} as a point in the three-dimensional iteration space, and we can think of each nonzero vertex $v_{ik}^{\mathbf{A}}$, $v_{kj}^{\mathbf{B}}$, or $v_{ij}^{\mathbf{C}}$ as a point on one of the (two-dimensional) faces of the iteration space. Likewise, a net $n_{ik}^{\mathbf{A}}$, $n_{kj}^{\mathbf{B}}$, or $n_{ij}^{\mathbf{C}}$ corresponds to a (one-dimensional) fiber perpendicular to one of the faces, and its pins are the vertices that lie along the fiber.

We provide two other visualizations of the hypergraph defined by the SpGEMM instance given in Fig. 1. Fig. 3 depicts the hypergraph as a bipartite graph, where hypergraph vertices are denoted by circles, hypergraph nets are denoted by squares, and net membership is defined by edges in the graph. Fig. 4 shows the incidence matrix of the hypergraph, where rows correspond to vertices, columns correspond to nets, and net membership is defined by the nonzero pattern.

This hypergraph is nearly the same as the one proposed in our earlier paper [Ballard et al. 2015a, Definition 1]. The difference here is the introduction of nonzero vertices and the inclusion of each in its corresponding nonzero's net. The purpose of adding these vertices is to enable enforcing memory balance in the parallel case (see Sec. 4.1), and for relating the hypergraph model to a computation DAG model in the sequential case (see Sec. 4.2).
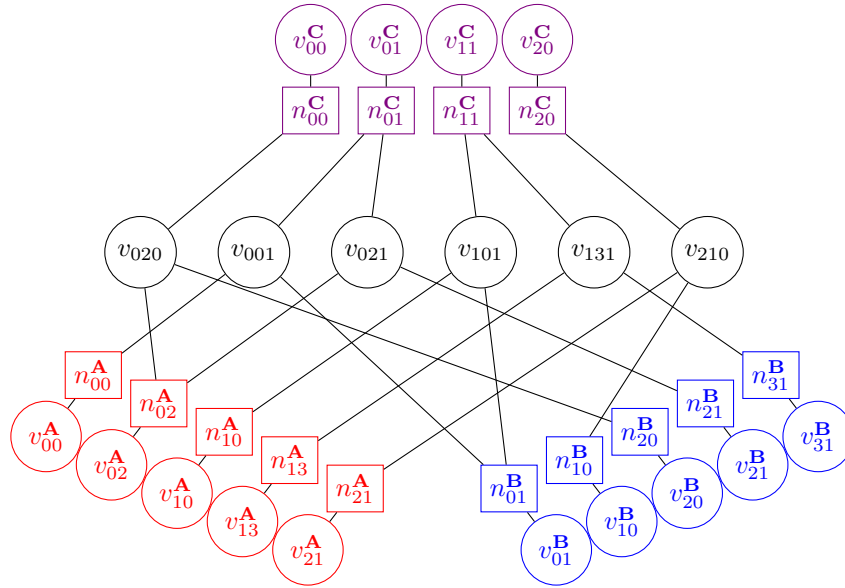
Fig. 3: Hypergraph of the SpGEMM instance in Fig. 1. Circles correspond to vertices and squares correspond to nets.

	$n_{00}^{\mathbf{A}}$	$n_{02}^{\mathbf{A}}$	$n_{10}^{\mathbf{A}}$	$n_{13}^{\mathbf{A}}$	$n_{21}^{\mathbf{A}}$	$n_{01}^{\mathbf{B}}$	$n_{10}^{\mathbf{B}}$	$n_{20}^{\mathbf{B}}$	$n_{21}^{\mathbf{B}}$	$n_{31}^{\mathbf{B}}$	$n_{00}^{\mathbf{C}}$	$n_{01}^{\mathbf{C}}$	$n_{11}^{\mathbf{C}}$	$n_{20}^{\mathbf{C}}$
v_{020}		∗						∗			∗			
v_{001}	∗					∗						∗		
v_{021}		∗							∗			∗		
v_{101}			∗			∗							∗	
v_{131}				∗						∗			∗	
v_{210}					∗		∗							∗

Fig. 4: Submatrix of the incidence matrix of the hypergraph of the SpGEMM instance in Fig. 1. Rows correspond to multiplication vertices and columns correspond to nets. The rows of the incidence matrix not shown here correspond to nonzero vertices; they form a diagonal submatrix.

4. COMMUNICATION LOWER BOUNDS

4.1. Parallel Lower Bound

We consider performing SpGEMM with input matrices $\mathbf{A}$ and $\mathbf{B}$ on a parallel machine with p processors with disjoint memories. Let $\mathcal{H} = \mathcal{H}(\mathbf{A}, \mathbf{B})$. A *parallelization* is a p-way partition of $\mathcal{V}^{\mathrm{m}}$ (assigning multiplications to processors) and a *data distribution* is a p-way partition of $\mathcal{V}^{\mathrm{nz}}$ (assigning nonzeros to processors). Thus, a partition of $\mathcal{V}$ defines both a parallelization and a data distribution. We define a *parallel SpGEMM algo-*

rithm by not only a parallelization and data distribution, but also by a communication pattern and each processor's local computation algorithm(s). So, a single partition of $\mathcal{V}$ corresponds to a family of parallel SpGEMM algorithms, and each parallel SpGEMM algorithm corresponds to a single partition of $\mathcal{V}$. For the rest of Sec. 4.1, "algorithm" means "parallel SpGEMM algorithm".

Given a partition of $\mathcal{V}$, the communication for any corresponding algorithm comprises two phases: the *expand* phase, where the processors exchange nonzero entries of $\mathbf{A}$ and $\mathbf{B}$ (initially distributed according to the partitions of $\mathcal{V}^{\mathbf{A}}$ and $\mathcal{V}^{\mathbf{B}}$) in order to perform their multiplications (assigned according to the partition of $\mathcal{V}^{\mathbf{m}}$), and the *fold* phase, where the processors communicate to reduce partial sums for nonzero entries of $\mathbf{C}$ (finally distributed according to the partition of $\mathcal{V}^{\mathbf{C}}$).

We now bound below the parallel communication costs of an algorithm in terms of the cuts of $\mathcal{H}$ induced by that algorithm's associated partition of $\mathcal{V}$.

Definition 4.1. Given a partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_p\}$ of $\mathcal{V}$, for each $i \in [p]$, Q_i denotes the subset of $\mathcal{N}$ whose elements (nets) have nonempty intersections with both $\mathcal{V}_i$ and $\mathcal{V} \setminus \mathcal{V}_i$.

LEMMA 4.2. *Given a partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_p\}$ of $\mathcal{V}$, for any associated algorithm, the number of words each processor $i \in [p]$ sends or receives is at least $|Q_i|$, and the critical-path communication cost is at least $\max_{i \in [p]} |Q_i|$.*

PROOF. For each processor $i \in [p]$, for each net in Q_i, processor i must either receive or send the corresponding nonzero, since at most one processor owns each nonzero at the start and end of the computation. The bound on the critical-path communication cost is obtained by maximizing over $i \in [p]$, as each processor can send only one word at a time along the critical path [Ballard et al. 2014]. □

Furthermore, for each partition, there exists an algorithm that attains these lower bounds, within constant factors in the cases of the per-processor costs and within a logarithmic factor in the case of the critical-path cost.

LEMMA 4.3. *Given a partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_p\}$, there exists an associated algorithm such that the number of words each processor $i \in [p]$ sends or receives is $O(|Q_i|)$, and the critical-path communication cost is $O(\log p \cdot \max_{i \in [p]} |Q_i|)$.*

PROOF. The basic idea of the algorithm is as follows. For the expand phase, every input nonzero that corresponds to a cut net will be sent to the processors whose parts $\mathcal{V}_i$ intersect the net via a binary-tree broadcast. For the fold phase, every output nonzero that corresponds to a cut net will be reduced to the processor whose part includes the nonzero vertex via a binary-tree reduction.

In more detail, consider first the expand phase. All of the broadcasts are performed synchronously: in the first step, each root sends its nonzero to two other processors; in the second step, every processor that receives a nonzero from a root sends it to two other processors in the second step; and so on. Since at most p processors can be involved in any one broadcast, the expand phase requires at most $O(\log p)$ steps. Processor i receives each of its required nonzeros at most once and sends each at most twice, for a total (per-processor) cost of $O(|Q_i|)$. Furthermore, at each expand step, processor i performs one receive and two sends in each of a subset of $|Q_i|$ broadcast trees, so each expand step involves $O(\max_{i \in [p]} |Q_i|)$ sends or receives along any critical path. Analysis of the fold phase is symmetric, using binary-tree reductions. □

For efficient algorithms, we constrain the partitions to enforce load balance of memory and computation. We define sets of load-balanced partitions in terms of two parameters:

Definition 4.4. For any $\delta, \epsilon \in [0, p-1]$, let $\Pi_{\delta,\epsilon}$ be the set of all partitions $\{\mathcal{V}_1, \ldots, \mathcal{V}_p\}$ of $\mathcal{V}$ where

$$w_{\mathrm{mem}}(\mathcal{V}_i) \leq (1+\delta)\frac{|\mathcal{V}^{\mathrm{nz}}|}{p} \qquad \text{and} \qquad w_{\mathrm{comp}}(\mathcal{V}_i) \leq (1+\epsilon)\frac{|\mathcal{V}^{\mathrm{m}}|}{p}$$

for each $i \in [p]$. We say an algorithm with partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_p\} \in \Pi_{\delta,\epsilon}$ is (δ, ϵ)-*load balanced*.

Given $\Pi_{\delta,\epsilon}$, if $\delta = 0$ then the memory requirements are perfectly load-balanced, and if $\epsilon = 0$ then the computations (multiplications) are perfectly load-balanced. If $\delta = \epsilon = p - 1$ then $\Pi_{\delta,\epsilon}$ includes the trivial partition wherein no interprocessor communication is required, since one processor stores all three matrices and performs the whole computation.

Finally, we state our parallel communication lower bound:

THEOREM 4.5. *The critical-path communication cost of an (δ, ϵ)-load balanced parallel SpGEMM algorithm is at least*

$$\min_{\{\mathcal{V}_1, \ldots, \mathcal{V}_p\} \in \Pi_{\delta,\epsilon}} \max_{i \in [p]} |Q_i|.$$

This lower bound is tight up to at most a logarithmic factor in the number of processors.

PROOF. The lower bound follows directly from Lem. 4.2 and Def. 4.4. The tightness of the lower bound follows from applying Lem. 4.3 to any optimal partition. □

We compare Thm. 4.5 with the distributed-memory parallel communication lower bounds for SpGEMM derived by Ballard et al. [2011; 2012]. In our notation, with the additional assumption that each processor is restricted to using M words of its local memory, the critical-path communication cost is bounded below by a function contained in the union of the sets

$$\Omega\left(\frac{|\mathcal{V}^{\mathrm{m}}|}{pM^{1/2}} - \alpha M\right) \qquad \text{and} \qquad \Omega\left(\frac{|\mathcal{V}^{\mathrm{m}}|^{2/3}}{p^{2/3}} - \beta\frac{|\mathcal{V}^{\mathrm{nz}}|}{p}\right) \tag{1}$$

for some $\alpha, \beta > 0$. These two asymptotic lower bounds are called memory-dependent and memory-independent, respectively; the memory-independent bound requires an additional asymptotic assumption that the computational load balance parameter $\epsilon = O(1)$; we refer to [Ballard et al. 2011; Ballard et al. 2012] for details regarding these asymptotic expressions, like the constants α, β, and others suppressed by asymptotic notation, and a discussion of when one lower bound dominates the other.

The combined lower bound eq. (1) is tight (within constant factors) when **A** and **B** are dense matrices [Irony et al. 2004], but asymptotically loose in expectation when $I = J = K = n$ and each entry of **A** and **B** is nonzero with probability d/n for any d in the intersection of the sets $\omega(1)$, $o(\sqrt{M})$, and $o(\sqrt{n})$ [Ballard et al. 2013]. On the other hand, by considering the nonzero structures of **A** and **B**, the conclusion of Lem. 4.2 is always tight, within a logarithmic factor.

4.2. Sequential Lower Bound

Now we turn to the communication costs of SpGEMM algorithms executed on a sequential machine with a two-level memory. In this computation model, communication is data movement between a fast memory of M-word capacity, $M \geq 3$, and a slow memory of unbounded capacity; operations are only performed on data in fast memory. Communication cost is defined as the total number of loads and stores, i.e., words moved from slow to fast and from fast to slow memory. We will derive communication

lower bounds in terms of the SpGEMM hypergraph model (see Sec. 3.2. For the rest of Sec. 4.2, "algorithm" means "sequential SpGEMM algorithm".

We follow Hong and Kung's approach [1981, Sec. 3] for deriving lower bounds on communication in this sequential model, as well as their graph-theoretic SpGEMM model [1981, Sec. 6]. Hong and Kung model algorithms schematically as uninterpreted directed acyclic graphs (DAGs): the SpGEMM hypergraph model in Def. 3.1 is closely related to Hong and Kung's DAG model of SpGEMM. In particular, each SpGEMM hypergraph $\mathcal{H}(\mathbf{A}, \mathbf{B})$ represents a family of SpGEMM DAGs $\mathcal{G}(\mathbf{A}, \mathbf{B})$ whose members differ only in the order in which they sum the (nontrivial) multiplications. After formalizing this relationship, we exploit it to approximate and simplify Hong and Kung's communication lower bounds framework using hypergraphs.

We define a *summation tree* as any DAG with exactly one zero-outdegree vertex, called the *root*; the zero-indegree vertices are called the *leaves*. The members of $\mathcal{G}(\mathbf{A}, \mathbf{B})$ are parameterized by a function that associates each $(i, j) \in S_{\mathbf{C}}$ with a summation tree T_{ij} whose leaves are $\{v_{ikj} : (\exists k)\, v_{ikj} \in \mathcal{V}\}$ and whose non-leaves are disjoint from $\mathcal{V}$ and the other trees. Having associated disjoint summation trees with $S_{\mathbf{C}}$ in this manner, the generic member $(V, E) \in \mathcal{G}(\mathbf{A}, \mathbf{B})$ is defined by

$$V = \mathcal{V}^{\mathbf{A}} \cup \mathcal{V}^{\mathbf{B}} \cup \{V(T_{ij}) : (i, j) \in S_{\mathbf{C}}\}$$

$$E = \quad \{(v_{ik}^{\mathbf{A}}, v_{ikj}) : (\exists i, k, j)\, v_{ik}^{\mathbf{A}}, v_{ikj} \in \mathcal{V}\}$$

$$\cup \; \{(v_{ik}^{\mathbf{B}}, v_{ikj}) : (\exists i, k, j)\, v_{kj}^{\mathbf{B}}, v_{ikj} \in \mathcal{V}\}$$

$$\cup \; \{E(T_{ij}) : (i, j) \in S_{\mathbf{C}}\}.$$

Of the vertices V, we distinguish the *inputs* I, which have no predecessors, and the *outputs* O, which have no successors; by construction, these two sets are disjoint. We remark that $\mathcal{G}(\mathbf{A}, \mathbf{B})$ can be specialized to obtain Hong and Kung's SpGEMM DAGs, by restricting each T_{ij} to be a binary tree with at least two leaves; this excludes, e.g., multiplication of diagonal matrices.

The key to Hong and Kung's communication analysis is the set of *S-partitions* of a DAG (V, E). A partition $\{V_1, \ldots, V_h\}$ of V is an S-partition if additionally

(1) Each V_i has a *dominator set* of size at most S — a dominator set of V_i contains a vertex on every path from I to V_i);
(2) Each V_i has a *minimum set* of size at most S — the minimum set of V_i is the set of all elements of V_i with no successors in V_i); and
(3) There is no cyclic dependence among $V_1, \ldots, V_h$ — where V_i depends on V_j if $(V_i \times V_j) \cap E \neq \emptyset$.

We exploit $\mathcal{G}(\mathbf{A}, \mathbf{B})$'s members' common structure to derive lower bounds on dominator and minimum set sizes.

LEMMA 4.6. *Consider any* $(V, E) \in \mathcal{G}(\mathbf{A}, \mathbf{B})$ *and any* $U \subseteq V$. *Let* D_U *be any dominator set of* U *and let* M_U *be the minimum set of* U.

$$|D_U| \geq \max \left(|\{(i, k) : (\exists j)\, v_{ikj} \in U\}|, |\{(k, j) : (\exists i)\, v_{ikj} \in U\}|\right)$$

$$|M_U| \geq |\{(i, j) : (\exists k)\, v_{ikj} \in U\}|$$

PROOF. There exists a set of disjoint paths, each starting from a distinct element of $\{v_{ik}^{\mathbf{A}} : (\exists j)\, v_{ikj} \in U\} \subset I$ and ending in U. Similarly, there exists a set of disjoint paths, each starting from a distinct element of $\{v_{kj}^{\mathbf{B}} : (\exists i)\, v_{ikj} \in U\} \subset I$ and ending in U. By definition, D_U must contain at least one vertex from every path in each of the two sets, so the first lower bound follows. Now observe that for each pair (i, j) such that $v_{ikj} \in U$ for some $k \in [K]$, $T_{ij} \cap U$ is nonempty and, since the summation trees are disjoint, has

at least one element with no successors in U. Thus each such T_{ij} contributes at least one element to M_U, establishing the second lower bound. $\square$

We now relate the DAGs $\mathcal{G}(\mathbf{A}, \mathbf{B})$ and the hypergraph $\mathcal{H}(\mathbf{A}, \mathbf{B})$.

Definition 4.7. Consider any partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_h\}$ of $\mathcal{V}$. For each part $\mathcal{V}_i$, let $W_i^{\mathbf{A}}$, $W_i^{\mathbf{B}}$, and $W_i^{\mathbf{C}}$ be the subsets of $\mathcal{N}^{\mathbf{A}}$, $\mathcal{N}^{\mathbf{B}}$, and $\mathcal{N}^{\mathbf{C}}$, resp., having nonempty intersections with $\mathcal{V}_i$.

LEMMA 4.8. *Consider any $G \in \mathcal{G}(\mathbf{A}, \mathbf{B})$ and $S \in \{1, 2, \ldots\}$. For any S-partition $\{V_1, \ldots, V_h\}$ of G, there exists a partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_h\}$ of $\mathcal{V}$ with $|W_i^{\mathbf{A}}|, |W_i^{\mathbf{B}}|, |W_i^{\mathbf{C}}| \le S$ for each $i \in [h]$.*

PROOF. We construct $\{\mathcal{V}_1, \ldots, \mathcal{V}_h\}$ from $\{V_1, \ldots, V_h\}$ as follows. Define a partition $\{V_1', \ldots, V_h'\}$ of $\mathcal{V}^{\mathbf{m}}$ by intersecting each subset $V_1, \ldots, V_h$ with $\mathcal{V}^{\mathbf{m}}$. Every element of $\mathcal{N}$ has a nontrivial intersection with at least one subset $V_1', \ldots, V_h'$. Moreover, applying Lem. 4.6 to each subset V_i', we have that $|W_i'^{\mathbf{A}}|, |W_i'^{\mathbf{B}}|, |W_i'^{\mathbf{C}}| \le S$. Now add each $v_{ik}^{\mathbf{A}} \in \mathcal{V}^{\mathbf{A}}$ to an arbitrarily chosen V_i' that intersects the corresponding net $n_{ik}^{\mathbf{A}}$, and likewise for each $v_{kj}^{\mathbf{B}} \in \mathcal{V}^{\mathbf{B}}$ and each $v_{ij}^{\mathbf{C}} \in \mathcal{V}^{\mathbf{C}}$. In this manner, we obtain a partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_h\}$ of $\mathcal{V}$. Moreover, augmenting $V_1', \ldots, V_h'$ with elements of $\mathcal{V}^{\mathbf{nz}}$ in this manner does not change the intersected nets, i.e., $W_i^{\mathbf{A}} = W_i'^{\mathbf{A}}$, $W_i^{\mathbf{B}} = W_i'^{\mathbf{B}}$, and $W_i^{\mathbf{C}} = W_i'^{\mathbf{C}}$ for each $(\mathcal{V}_i, V_i')$ pair. $\square$

LEMMA 4.9. *Consider any $S \in \{1, 2, \ldots\}$ and partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_h\}$ of $\mathcal{V}$ with $|W_i^{\mathbf{A}}|, |W_i^{\mathbf{B}}|, |W_i^{\mathbf{C}}| \le S$ for each $i \in [h]$. There exists an algorithm with communication cost $O(Mh(S/M + 1)^3)$.*

PROOF. For each $i \in [h]$, partition $W_i^{\mathbf{A}}, W_i^{\mathbf{B}}, W_i^{\mathbf{C}}$ so that each part has size $m = \lfloor M/3 \rfloor$ except possibly for one smaller part. These three partitions naturally induce a partition $\{\mathcal{U}_1, \ldots, \mathcal{U}_g\}$ of $\mathcal{V}$, whose parts are called *blocks*. Observe that $\{\mathcal{U}_1, \ldots, \mathcal{U}_g\}$ refines $\{\mathcal{V}_1, \ldots, \mathcal{V}_h\}$; in particular, each part $\mathcal{V}_i$ is partitioned into to at most $\lceil S/m \rceil^3$ blocks. Thus, the total number of blocks $g \le h \lceil S/m \rceil^3$.

The algorithm considers $\mathcal{U}_j$ for each $j \in [g]$:

(1) Load $\mathbf{A}$- and $\mathbf{B}$-matrix entries associated with $W_j^{\mathbf{A}}$ and $W_j^{\mathbf{B}}$, as well as any partial sums of $\mathbf{C}$-matrix entries $W_j^{\mathbf{C}}$ previously computed — at most $3m$ loads.
(2) Perform all possible nontrivial multiplications and additions — this can be done with no data movement by processing each $\mathbf{C}$-matrix entry in sequence.
(3) Store the updated $\mathbf{C}$-matrix entries — at most m stores.

There are at most $4mg \le 4Mh(5S/M + 1)^3/3$ loads and stores in total. $\square$

Finally, we apply these results within Hong and Kung's communication lower bounds framework.

THEOREM 4.10. *The communication cost of an algorithm is at least $M(h-1)$, where h is the minimum cardinality of a partition $\{\mathcal{V}_1, \ldots, \mathcal{V}_h\}$ of $\mathcal{V}$ with $|W_i^{\mathbf{A}}|, |W_i^{\mathbf{B}}|, |W_i^{\mathbf{C}}| \le 2M$ for each $i \in [h]$. When $h > 1$, this lower bound is tight up to a constant factor.*

PROOF. Any algorithm can be represented in Hong and Kung's model by some $G \in \mathcal{G}(\mathbf{A}, \mathbf{B})$. By Hung and Kung's argument [1981, Lem. 3.1], the communication cost is at least $M(P(2M) - 1)$, where $P(S)$ denotes the minimum cardinality of an S-partition of G. We apply Lem. 4.8 to bound $P(2M)$ below by h. Tightness follows from applying Lem. 4.9 with $S = 2M$. $\square$

Thm. 4.10 simplifies Hong and Kung's key lemma by requiring a single hypergraph partitioning problem, vs. $|\mathcal{G}(\mathbf{A},\mathbf{B})|$ graph partitioning problems. Additionally, Hong and Kung did not address attainability of their key lemma.

We note that another of Hong and Kung's results [1981, Thm. 6.1] yields the lower bound $\Omega(|\mathcal{V}^{\mathbf{m}}|/M^{1/2})$, a sequential analogue of the memory-dependent bound in eq. (1). While this lower bound is attainable in the cases where $\mathbf{A}$ and $\mathbf{B}$ are both dense (see, e.g., [Ballard et al. 2011]), it is asymptotically loose in many sparse cases. For example, if $I = J = K = n$ and $S_{\mathbf{A}} = S_{\mathbf{B}} = \{(i,i) : i \in [n]\}$, then any algorithm requires moving at least $3|\mathcal{V}^{\mathbf{m}}|$ words between the two levels, the nonzero entries of $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$. However, a trivial lower bound of $|\mathcal{V}^{\mathrm{nz}}|$ also exists, under the assumption that fast memory must be empty before and after the computation. It remains open to show whether or not the combination of this bound and the memory-dependent one is asymptotically loose.

We mention some other communication bounds that have appeared previously. Pagh and Stöckel [2014] proved matching worst-case upper and lower bounds based on the number of nonzeros in the input and output matrices. That is, given a number of nonzeros, they showed that there exist input matrices that require a certain amount of communication; they also gave an algorithm that never requires more than a constant factor times this worst-case cost. Greiner [2012, Ch. 6] considered restricted classes of algorithms and matrices and used a different approach to establish other worst-case lower and upper bounds, which are not always tight. In contrast to these results, Thm. 4.10 is input specific rather than worst case.

5. RESTRICTING SPGEMM ALGORITHMS

Each SpGEMM hypergraph (defined in Sec. 3.2) models a class of SpGEMM algorithms via partitions of its vertices. This hypergraph model can be simplified in the study of subclasses of algorithms and matrix nonzero structures. Here, we only consider the distributed-memory parallel case, as described in Sec. 4.1.

In Sec. 5.1, we describe how to simplify SpGEMM hypergraphs using a vertex coarsening technique while still correctly modeling algorithmic costs. We use vertex coarsening to model subclasses of algorithms, restricting the parallelization (Sec. 5.2), the data distribution (Sec. 5.3), and both (Sec. 5.4). We can also use vertex coarsening when modeling subclasses of matrix nonzero structures, as we demonstrate in Sec. 5.5 in the case of *sparse matrix-vector multiplication* (SpMV). Lastly, in Sec. 5.6, we show how vertex coarsening can also be used to model SpGEMM-like algorithms that violate two of the assumptions in Sec. 3.1, exploiting relations between the input matrix entries (Sec. 5.6.1) and only computing a subset of the output matrix entries (Sec. 5.6.2).

5.1. Vertex Coarsening

We can restrict the class of algorithms modeled by an SpGEMM hypergraph by forcing certain subsets of vertices to be *monochrome*, or all assigned to the same part, in the partition. For example, we can enforce that a certain subset of multiplications will be performed by the same processor or that a particular multiplication is performed by the processor that owns a corresponding nonzero entry. To implement such restrictions in an SpGEMM hypergraph, we *coarsen* vertices, choosing a new vertex to represent each subset of *constituent* vertices that must be monochrome. We will use the symbol "$\sim$" to mean "represents" in this sense.

After vertex coarsening, in order to model communication costs correctly, we define net membership as follows: a coarsened vertex is a member of a net if any of its constituent vertices was a member of that net. This is because the computation corresponding to the coarsened vertex requires the matrix entries associated with any nets that contained a constituent vertex.

In order to model computational and memory costs correctly, the weights of a coarsened vertex equal the sum of the constituent vertices' weights. This is because the processor assigned the coarsened vertex will perform the computation corresponding to all of the multiplication constituent vertices and will own all of the nonzero constituent vertices.

After coarsening the vertices and updating net memberships, there may be multiple *coalesced* nets that contain identical sets of vertices, and there may be *singleton* nets that contain only a single vertex. To reduce the number of nets, we can combine coalesced nets into one coarsened net; the cost of the coarsened net equals the sum of the coalesced nets' costs. Because singleton nets cannot be cut, we can omit them, a further simplification.

5.2. Restricted Parallelizations

We first apply vertex coarsening to the multiplication vertices $\mathcal{V}^m \subset \mathcal{V}$ of an SpGEMM hypergraph $(\mathcal{V}, \mathcal{N}) = \mathcal{H}(\mathbf{A}, \mathbf{B})$, representing parallelizing computation across processors at a coarser granularity than individual multiplications. We represent each coarsened set by a new vertex $\hat{v}$, and let $\hat{\mathcal{V}}^m$ denote the set of these new $\hat{v}$ vertices. In this section, we leave the remaining vertices $\mathcal{V}^{nz}$ untouched, so $\hat{\mathcal{V}} = \hat{\mathcal{V}}^m \cup \mathcal{V}^{nz}$; we will address restricting data distribution subsequently in Secs. 5.3 and 5.4. Net memberships are updated in terms of the $\hat{v}$ vertices as described in Sec. 5.1, i.e., $\mathcal{N}$ induces a set $\hat{\mathcal{N}}$ of nets on $\hat{\mathcal{V}}$, thus obtaining the simplified hypergraph $(\hat{\mathcal{V}}, \hat{\mathcal{N}})$. Since each net contains a distinct nonzero vertex, no nets are coalesced/combined, so $|\hat{\mathcal{N}}| = |\mathcal{N}|$ and the nets' costs remain unit. Additionally, the nonzero vertices are untouched, so the only weights that change are the computational weights of elements $\hat{v}$ of $\hat{\mathcal{V}}^m$: each $w_{\text{comp}}(\hat{v})$ equals the number of multiplication vertices it combined.

The simplified hypergraph has fewer vertices than the original: $\mathcal{V}^m$ is replaced by $\hat{\mathcal{V}}^m$, whose cardinality is that of the given partition of $\mathcal{V}^m$. Additionally, while the number of nets is unchanged, the number of pins generally decreases.

In the remainder of Sec. 5.2, we will study two types of coarsenings of $\mathcal{V}^m$: by slice and by fiber. Fig. 5 shows visualizations of each type of coarsening.

Slice-wise coarsenings, which model what we call *1D parallelizations*, coarsen the computational cube by slices parallel to one of the three coordinate planes. Coarsening by slices means we identify all v_{ikj} with the same i, j, or k index — the choice of index gives rise to three 1D models, called *row-wise*, *column-wise*, and *outer-product*:

$$\begin{aligned}
\text{row-wise} \quad & \hat{\mathcal{V}}^m = \{\hat{v}_i : i \in [I]\}, \quad \hat{v}_i \sim \{v_{its} : (\exists t, s)\, v_{its} \in \mathcal{V}\} \\
\text{column-wise} \quad & \hat{\mathcal{V}}^m = \{\hat{v}_j : i \in [J]\}, \quad \hat{v}_j \sim \{v_{rtj} : (\exists r, t)\, v_{rtj} \in \mathcal{V}\} \\
\text{outer-product} \quad & \hat{\mathcal{V}}^m = \{\hat{v}_k : k \in [K]\}, \quad \hat{v}_k \sim \{v_{rks} : (\exists r, s)\, v_{rks} \in \mathcal{V}\}.
\end{aligned}$$

Fiber-wise coarsenings, which model what we call *2D parallelizations*, coarsen the computational cube by fibers parallel to one of the three coordinate axes. Coarsening by fibers means we identify all v_{ikj} with the same pair (i, k), (k, j), or (i, j), of indices — this choice gives rise to three 2D models, called *monochrome-$\mathbf{A}$*, *monochrome-$\mathbf{B}$*, and *monochrome-$\mathbf{C}$*:

$$\begin{aligned}
\text{monochrome-}\mathbf{A} \quad & \hat{\mathcal{V}}^m = \{\hat{v}_{ik} : (i, k) \in S_{\mathbf{A}}\}, \quad \hat{v}_{ik} \sim \{v_{iks} : (\exists s)\, v_{iks} \in \mathcal{V}\} \\
\text{monochrome-}\mathbf{B} \quad & \hat{\mathcal{V}}^m = \{\hat{v}_{kj} : (k, j) \in S_{\mathbf{B}}\}, \quad \hat{v}_{kj} \sim \{v_{rkj} : (\exists r)\, v_{rkj} \in \mathcal{V}\} \\
\text{monochrome-}\mathbf{C} \quad & \hat{\mathcal{V}}^m = \{\hat{v}_{ij} : (i, j) \in S_{\mathbf{C}}\}, \quad \hat{v}_{ij} \sim \{v_{itj} : (\exists t)\, v_{itj} \in \mathcal{V}\}.
\end{aligned}$$

Each model is identified with a set of parallelizations. In a row-wise parallelization, every $\mathbf{B}$-slice is monochrome, and the parallelization corresponds to a partition of the row-wise model; column-wise and outer-product parallelizations are defined similarly.
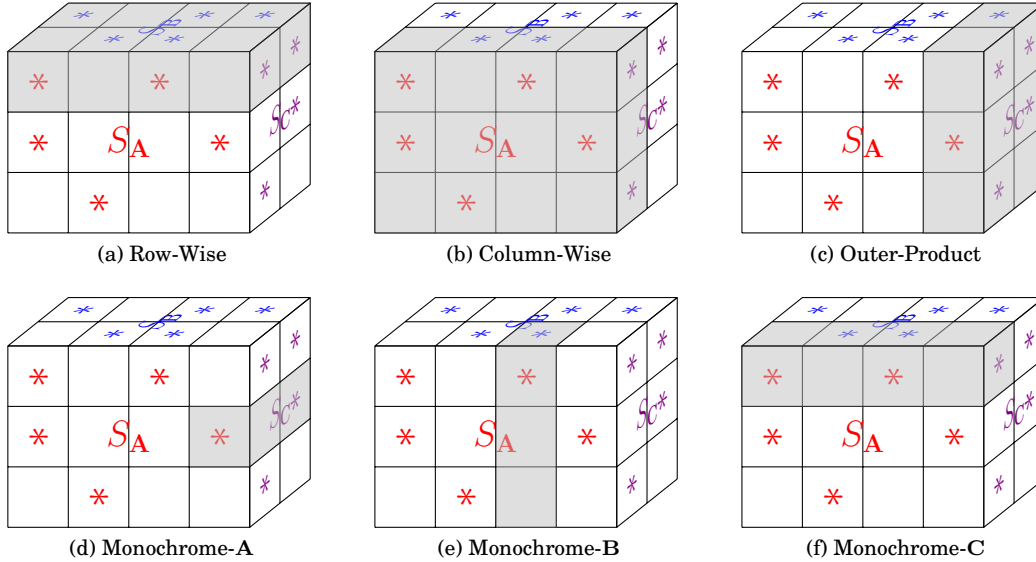
Fig. 5: Visualizations of restricted parallelizations. The top row (Figs. 5a to 5c) corresponds to slice-wise coarsenings or 1D parallelizations; the bottom row (Figs. 5d to 5f) corresponds to fiber-wise coarsenings or 2D parallelizations. In each figure, a particular coarsened vertex is highlighted as a shaded set of cubes; the coarsened vertex corresponds to all multiplication vertices in the shaded set.

In a monochrome-**A** parallelization, every **A**-fiber is monochrome, and the parallelization corresponds to a partition of the monochrome-**A** model; monochrome-**B** and -**C** parallelizations are defined similarly. A general parallelization is called fine-grained.

We now discuss relationships among these seven classes of parallelizations, as illustrated in Fig. 6. Let $\mathcal{F}$ denote the set of all (fine-grained) parallelizations. Let $\mathcal{R}, \mathcal{L}, \mathcal{U}, \mathcal{A}, \mathcal{B}, \mathcal{C} \subseteq \mathcal{F}$ denote the row-wise, column-wise, outer-product, monochrome-**A**, -**B**, and -**C** parallelizations, respectively. Fig. 6 depicts the inclusions that hold among these 7 sets. Notice in particular that these 7 sets define a 13-way partition of $\mathcal{F}$, whose parts are listed in the first column of Tab. I; we now verify that these 13 parts indeed form a partition of $\mathcal{F}$. Observe that $\mathcal{R} \subseteq \mathcal{A} \cap \mathcal{C}$, $\mathcal{L} \subseteq \mathcal{B} \cap \mathcal{C}$, and $\mathcal{U} \subseteq \mathcal{A} \cap \mathcal{B}$. Actually, $\mathcal{U} = \mathcal{A} \cap \mathcal{B}$. To see the converse inclusion, consider any $P \in \mathcal{A} \cap \mathcal{B}$, and observe that every **C**-slice must be monochrome in P; this is because **C**-slices are Cartesian products of their projections onto the **A**- and **B**-faces, but a similar property does not generally hold for the **A**- and **B**-slices. The preceding inclusions suffice to demonstrate correctness of the 13-way partition of $\mathcal{F}$ in Fig. 6.

Fig. 6: Venn diagram showing relationships among restricted parallelizations.

Tab. I uses the following four SpGEMM instances to show that each of the 13 parts is not empty.

$$\begin{bmatrix} * & * \\ * & * \end{bmatrix} = \begin{bmatrix} * & * \\ * & * \end{bmatrix} \begin{bmatrix} * & * \\ * & * \end{bmatrix} \tag{2}$$

$$\begin{bmatrix} * & * \\ * & * \end{bmatrix} = \begin{bmatrix} * & \\ & * \end{bmatrix} \begin{bmatrix} * & * \\ * & * \end{bmatrix} \tag{3}$$

$$\begin{bmatrix} * & * \\ * & * \end{bmatrix} = \begin{bmatrix} * & * \\ * & * \end{bmatrix} \begin{bmatrix} * & \\ & * \end{bmatrix} \tag{4}$$

$$\begin{bmatrix} * & * \\ * & * \end{bmatrix} = \begin{bmatrix} * & & * & \\ & * & & * \end{bmatrix} \begin{bmatrix} * \\ * \\ & * \\ & * \end{bmatrix} \tag{5}$$

5.3. Restricted Data Distributions

Now we apply vertex coarsening to the nonzero vertices $\mathcal{V}^{\mathrm{nz}} \subset \mathcal{V}$ of an SpGEMM hypergraph $(\mathcal{V}, \mathcal{N}) = \mathcal{H}(\mathbf{A}, \mathbf{B})$, representing distributing data across processors at a coarser granularity than individual matrix entries. We represent each coarsened set by a new vertex $\hat{v}$, and let $\hat{\mathcal{V}}^{\mathrm{nz}}$ be the set of these new $\hat{v}$ vertices. In this section, we leave the remaining vertices $\mathcal{V}^{\mathrm{m}}$ untouched, so $\hat{\mathcal{V}} = \mathcal{V}^{\mathrm{m}} \cup \hat{\mathcal{V}}^{\mathrm{nz}}$; we addressed parallelizations previously in Sec. 5.2 and will synthesize the two approaches subsequently in Sec. 5.4. Net memberships are updated in terms of the $\hat{v}$ vertices, i.e., $\mathcal{N}$ induces a set $\hat{\mathcal{N}}$ of nets on $\hat{\mathcal{V}}$, thus obtaining the simplified hypergraph $(\hat{\mathcal{V}}, \hat{\mathcal{N}})$. Now, it is possible that some nets are coalesced, and thus are combined in $\hat{\mathcal{N}}$; the cost of each net in $\hat{\mathcal{N}}$ is the number of coalesced nets it combines. Additionally, the multiplication vertices are untouched, so the only weights that change are the memory weights of elements $\hat{v}$ of $\hat{\mathcal{V}}^{\mathrm{nz}}$: each $w_{\mathrm{mem}}(\hat{v})$ equals the number of nonzero vertices it combines.

Table I: Examples showing that each of the 13 parts in Fig. 6 is nonempty. The "finest" parallelization is where each nontrivial multiplication is assigned to a distinct processor. A parallelization "by fiber" means that each fiber (of a given type) is assigned to a distinct processor. A parallelization "by slice" means that each slice (of a given type) is assigned to a distinct processor. The "coarsest" parallelization is where all nontrivial multiplications are assigned to the same processor.

Part	SpGEMM Instance	Parallelization
$\mathcal{F} \setminus (\mathcal{A} \cup \mathcal{B} \cup \mathcal{C})$	eq. (2)	finest
$\mathcal{A} \setminus (\mathcal{B} \cup \mathcal{C})$	eq. (2)	by $\mathbf{A}$-fiber
$\mathcal{B} \setminus (\mathcal{A} \cup \mathcal{C})$	eq. (2)	by $\mathbf{B}$-fiber
$\mathcal{C} \setminus (\mathcal{A} \cup \mathcal{B})$	eq. (2)	by $\mathbf{C}$-fiber
$((\mathcal{B} \cap \mathcal{C}) \setminus \mathcal{A}) \cap \mathcal{L}$	eq. (2)	by $\mathbf{A}$-slice
$((\mathcal{A} \cap \mathcal{C}) \setminus \mathcal{B}) \cap \mathcal{R}$	eq. (2)	by $\mathbf{B}$-slice
$(\mathcal{A} \cap \mathcal{B}) \setminus \mathcal{C}$	eq. (2)	by $\mathbf{C}$-slice
$\mathcal{A} \cap \mathcal{B} \cap \mathcal{C} \cap \mathcal{R} \cap \mathcal{L}$	eq. (2)	coarsest
$((\mathcal{B} \cap \mathcal{C}) \setminus \mathcal{A}) \setminus \mathcal{L}$	eq. (3)	finest
$(\mathcal{A} \cap \mathcal{B} \cap \mathcal{C} \cap \mathcal{R}) \setminus \mathcal{L}$	eq. (3)	by $\mathbf{A}$-fiber
$((\mathcal{A} \cap \mathcal{C}) \setminus \mathcal{B}) \setminus \mathcal{R}$	eq. (4)	finest
$(\mathcal{A} \cap \mathcal{B} \cap \mathcal{C} \cap \mathcal{L}) \setminus \mathcal{R}$	eq. (4)	by $\mathbf{B}$-fiber
$(\mathcal{A} \cap \mathcal{B} \cap \mathcal{C}) \setminus (\mathcal{R} \cup \mathcal{L})$	eq. (5)	finest

The simplified hypergraph has fewer vertices than the original: $\mathcal{V}^{\mathrm{nz}}$ is replaced by $\hat{\mathcal{V}}^{\mathrm{nz}}$, whose cardinality is that of the given partition of $\mathcal{V}^{\mathrm{nz}}$.

In the remainder of Sec. 5.2, we will study two classes of data distributions, *row-wise* and *column-wise*, meaning the matrix nonzeros are distributed by rows or columns. We coarsen $\mathcal{V}^{\mathbf{A}}$, $\mathcal{V}^{\mathbf{B}}$, and $\mathcal{V}^{\mathbf{C}}$ separately, leading to sets of coarsened vertices $\hat{\mathcal{V}}^{\mathbf{A}}$, $\hat{\mathcal{V}}^{\mathbf{B}}$, and $\hat{\mathcal{V}}^{\mathbf{C}}$:

$$
\begin{aligned}
\mathbf{A}\ \text{row-wise} \quad & \hat{\mathcal{V}}^{\mathbf{A}} = \{\hat{v}_i^{\mathbf{A}} : i \in [I]\}, \quad \hat{v}_i^{\mathbf{A}} \sim \{v_{it}^{\mathbf{A}} : (\exists t)\, v_{it}^{\mathbf{A}} \in \mathcal{V}\} \\
\mathbf{A}\ \text{column-wise} \quad & \hat{\mathcal{V}}^{\mathbf{A}} = \{\hat{v}_k^{\mathbf{A}} : k \in [K]\}, \quad \hat{v}_k^{\mathbf{A}} \sim \{v_{rk}^{\mathbf{A}} : (\exists r)\, v_{rk}^{\mathbf{A}} \in \mathcal{V}\} \\
\mathbf{B}\ \text{row-wise} \quad & \hat{\mathcal{V}}^{\mathbf{B}} = \{\hat{v}_k^{\mathbf{B}} : k \in [K]\}, \quad \hat{v}_k^{\mathbf{B}} \sim \{v_{ks}^{\mathbf{B}} : (\exists s)\, v_{ks}^{\mathbf{B}} \in \mathcal{V}\} \\
\mathbf{B}\ \text{column-wise} \quad & \hat{\mathcal{V}}^{\mathbf{B}} = \{\hat{v}_j^{\mathbf{B}} : j \in [J]\}, \quad \hat{v}_j^{\mathbf{B}} \sim \{v_{tj}^{\mathbf{B}} : (\exists t)\, v_{tj}^{\mathbf{B}} \in \mathcal{V}\} \\
\mathbf{C}\ \text{row-wise} \quad & \hat{\mathcal{V}}^{\mathbf{C}} = \{\hat{v}_i^{\mathbf{C}} : i \in [I]\}, \quad \hat{v}_i^{\mathbf{C}} \sim \{v_{is}^{\mathbf{C}} : (\exists s)\, v_{is}^{\mathbf{C}} \in \mathcal{V}\} \\
\mathbf{C}\ \text{column-wise} \quad & \hat{\mathcal{V}}^{\mathbf{C}} = \{\hat{v}_j^{\mathbf{C}} : j \in [J]\}, \quad \hat{v}_j^{\mathbf{C}} \sim \{v_{rj}^{\mathbf{C}} : (\exists r)\, v_{rj}^{\mathbf{C}} \in \mathcal{V}\}.
\end{aligned}
$$

The unrestricted choices $\hat{\mathcal{V}}^{\mathbf{A}} = \mathcal{V}^{\mathbf{A}}$, $\hat{\mathcal{V}}^{\mathbf{B}} = \mathcal{V}^{\mathbf{B}}$, and $\hat{\mathcal{V}}^{\mathbf{C}} = \mathcal{V}^{\mathbf{C}}$, are called *fine-grained*. Having picked a distribution for each of $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$, we assemble $\hat{\mathcal{V}}^{\mathrm{nz}} = \hat{\mathcal{V}}^{\mathbf{A}} \cup \hat{\mathcal{V}}^{\mathcal{V}} \cup \hat{\mathcal{V}}^{\mathbf{C}}$.

We remark here that some choices of these distributions naturally correspond to the 1D parallelization schemes. That is, a row-wise parallelization corresponds to a matching distribution of the rows of $\mathbf{A}$ and $\mathbf{C}$, a column-wise parallelization corresponds to a matching distribution of the columns of $\mathbf{B}$ and $\mathbf{C}$, and an outer-product parallelization corresponds to a matching distribution of the columns of $\mathbf{A}$ and rows of $\mathbf{B}$. We revisit this connection next, in Sec. 5.4, when we consider pairing parallelizations and data distributions.

5.4. Restricted Algorithms

In Secs. 5.2 and 5.3, we considered vertex coarsening using partitions of $\mathcal{V}^{\mathrm{m}}$ or $\mathcal{V}^{\mathrm{nz}}$. Now we extend this idea to partitions of $\mathcal{V}$ whose parts may combine vertices from both

$\mathcal{V}^{\mathrm{m}}$ and $\mathcal{V}^{\mathrm{nz}}$. This models the algorithmic constraint where a processor who performs a certain set of multiplications must also store a certain set of matrix entries.

Revisiting the parallelizations in Sec. 5.2, when coarsening $\mathcal{V}^{\mathrm{m}}$ by fiber, there is a natural choice of a nonzero vertex to associate with each fiber; likewise, when coarsening $\mathcal{V}^{\mathrm{m}}$ by slice, there is a natural set of nonzero vertices to associate with each slice. More generally, if all the multiplication vertices associated with a particular set of nonzero vertices are to be coarsened, it is natural to include those nonzero vertices. (However, this assignment does impose additional constraints on load balance, which may be unnatural in certain applications.) An expression (xyz), where $x, y, z \in \{\mathrm{r,c,f,R,C,F}\}$, indicates the data distributions (row-wise, column-wise, or fine-grained) used for $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$. Capital letters indicate that the data distribution is a natural choice, in the sense just described, and the corresponding nonzero vertices have been absorbed into coarsened vertices. We highlight the following coarsenings:

row-wise (RfR): $\quad \hat{\mathcal{V}} = \bigcup \left\{ \begin{array}{c} \{\hat{v}_i : i \in [I]\}, \\ \mathcal{V}^{\mathbf{B}} \end{array} \right\}, \quad \hat{v}_i \sim \bigcup \left\{ \begin{array}{c} \{v_{its} : (\exists t, s)\, v_{its} \in \mathcal{V}\}, \\ \{v_{it}^{\mathbf{A}} : (\exists t)\, v_{it}^{\mathbf{A}} \in \mathcal{V}\}, \\ \{v_{is}^{\mathbf{C}} : (\exists s)\, v_{is}^{\mathbf{C}} \in \mathcal{V}\} \end{array} \right\}$

column-wise (fCC): $\quad \hat{\mathcal{V}} = \bigcup \left\{ \begin{array}{c} \{\hat{v}_j : j \in [J]\}, \\ \mathcal{V}^{\mathbf{A}} \end{array} \right\}, \quad \hat{v}_j \sim \bigcup \left\{ \begin{array}{c} \{v_{rtj} : (\exists r, t)\, v_{rtj} \in \mathcal{V}\}, \\ \{v_{tj}^{\mathbf{B}} : (\exists t)\, v_{tj}^{\mathbf{B}} \in \mathcal{V}\}, \\ \{v_{rj}^{\mathbf{C}} : (\exists r)\, v_{rj}^{\mathbf{C}} \in \mathcal{V}\} \end{array} \right\}$

outer-product (CRf): $\quad \hat{\mathcal{V}} = \bigcup \left\{ \begin{array}{c} \{\hat{v}_k : k \in [K]\}, \\ \mathcal{V}^{\mathbf{C}} \end{array} \right\}, \quad \hat{v}_k \sim \bigcup \left\{ \begin{array}{c} \{v_{rks} : (\exists r, s)\, v_{rks} \in \mathcal{V}\}, \\ \{v_{rk}^{\mathbf{A}} : (\exists r)\, v_{rk}^{\mathbf{A}} \in \mathcal{V}\}, \\ \{v_{ks}^{\mathbf{B}} : (\exists s)\, v_{ks}^{\mathbf{B}} \in \mathcal{V}\} \end{array} \right\}$

monochrome-A (Fff): $\quad \hat{\mathcal{V}} = \bigcup \left\{ \begin{array}{c} \{\hat{v}_{ik} : (i, k) \in S_{\mathbf{A}}\}, \\ \mathcal{V}^{\mathbf{B}}, \mathcal{V}^{\mathbf{C}} \end{array} \right\}, \quad \hat{v}_{ik} \sim \bigcup \left\{ \begin{array}{c} \{v_{iks} : (\exists s)\, v_{iks} \in \mathcal{V}\}, \\ \{v_{ik}^{\mathbf{A}}\} \end{array} \right\}$

monochrome-B (fFf): $\quad \hat{\mathcal{V}} = \bigcup \left\{ \begin{array}{c} \{\hat{v}_{kj} : (k, j) \in S_{\mathbf{B}}\}, \\ \mathcal{V}^{\mathbf{A}}, \mathcal{V}^{\mathbf{C}} \end{array} \right\}, \quad \hat{v}_{kj} \sim \bigcup \left\{ \begin{array}{c} \{v_{rkj} : (\exists r)\, v_{rkj} \in \mathcal{V}\}, \\ \{v_{kj}^{\mathbf{B}}\} \end{array} \right\}$

monochrome-C (ffF): $\quad \hat{\mathcal{V}} = \bigcup \left\{ \begin{array}{c} \{\hat{v}_{ij} : (i, j) \in S_{\mathbf{C}}\}, \\ \mathcal{V}^{\mathbf{A}}, \mathcal{V}^{\mathbf{B}} \end{array} \right\}, \quad \hat{v}_{ij} \sim \bigcup \left\{ \begin{array}{c} \{v_{itj} : (\exists t)\, v_{itj} \in \mathcal{V}\}, \\ \{v_{ij}^{\mathbf{C}}\} \end{array} \right\}$

The coarsened vertices now have both nonzero computational and memory weights, indicating the number of multiplication and nonzero vertices, respectively, that they combined. Note that the nets corresponding to the coarsened nonzero vertices become singletons, and are therefore omitted from the simplified hypergraph — in the six cases above, the nets induced by $\mathcal{N}^{\mathbf{A}} \cup \mathcal{N}^{\mathbf{C}}$, $\mathcal{N}^{\mathbf{B}} \cup \mathcal{N}^{\mathbf{C}}$, $\mathcal{N}^{\mathbf{A}} \cup \mathcal{N}^{\mathbf{B}}$, $\mathcal{N}^{\mathbf{A}}$, $\mathcal{N}^{\mathbf{B}}$, and $\mathcal{N}^{\mathbf{C}}$, respectively, are omitted from $\hat{\mathcal{N}}$.

Further simplifications are possible, but perhaps less natural. For instance, in the row-wise and monochrome-A cases, it is arguably natural to coarsen $\mathcal{V}^{\mathbf{B}}$ along rows of $\mathbf{B}$: every entry in a row of $\mathbf{B}$ is involved in multiplications with the same entries of $\mathbf{A}$. After this coarsening, the nets induced by $\mathcal{N}^{\mathbf{B}}$ are coalesced along each row of $\mathbf{B}$, and so are combined into a single net whose cost equals the number of coalesced nets (also the sum of their costs). Similarly, in the column-wise and monochrome-B cases, it is arguably natural to coarsen $\mathcal{V}^{\mathbf{A}}$ along columns of $\mathbf{A}$: every entry in a column of $\mathbf{A}$ multiplies the same entries of $\mathbf{B}$; the nets induced by $\mathcal{N}^{\mathbf{A}}$ are coalesced/combined along column of $\mathbf{A}$ as before.

These (arguably) natural simplifications don't always coarsen the nonzeros of all three matrices. In particular, in the outer-product, monochrome-A, and monochrome-

B cases, there is no natural choice of rows or columns of C, and in the monochrome-C case, there is no natural choice of rows or columns of **A** or **B**. Of course, nonzero vertex coarsening may be dictated by the application. For example, **C** may need to be distributed in a certain way for use after the SpGEMM. Or, in the row-wise case, if **A** and **C** are to be distributed row-wise, then additionally distributing **B** row-wise may be simplest from a software-design standpoint (and similarly distributing **A** column-wise in the column-wise case).

We have presented several restricted classes of algorithms that enable simplifications to the SpGEMM hypergraph. For concreteness, we conclude with four self-contained examples, stated similarly to Def. 3.1. (We continue our simplifying assumption that **A** and **B** have no zero rows or columns.)

Example 5.1 (*Row-wise (RrR)*). Our first example pairs a row-wise parallelization with row-wise distributions of **A**, **B**, and **C**, where the distributions of **A** and **C** are matched with the parallelization while the distribution of **B** is not.

$$
\mathcal{V} = \{v_i : i \in [I]\} \cup \{v_k^{\mathbf{B}} : k \in [K]\},
$$
$$
\mathcal{N} = \{n_k^{\mathbf{B}} = \{v_i : (\exists i)\,(i,k) \in S_{\mathbf{A}}\} \cup \{v_k^{\mathbf{B}}\} : k \in [K]\},
$$
$$
w_{\mathrm{comp}}(v_i) = |\{(i,k,j) : (\exists k)\,((i,k) \in S_{\mathbf{A}} \wedge (\exists j)\,(k,j) \in S_{\mathbf{B}})\}|,
$$
$$
w_{\mathrm{mem}}(v_i) = |\{(i,k) : (\exists k)\,(i,k) \in S_{\mathbf{A}}\}| + |\{(i,j) : (\exists j)\,(i,j) \in S_{\mathbf{C}}\}|,
$$
$$
w_{\mathrm{comp}}(v_k^{\mathbf{B}}) = 0, \quad w_{\mathrm{mem}}(v_k^{\mathbf{B}}) = c(n_k^{\mathbf{B}}) = |\{(k,j) : (\exists j)\,(k,j) \in S_{\mathbf{B}}\}|.
$$

We have that $|\mathcal{V}| = I + K$ and $|\mathcal{N}| = K$, where each net has between 2 and $I + 1$ pins. Note that computing the weights $v_{\mathrm{mem}}(v_i)$ requires determining $S_{\mathbf{C}}$.

Example 5.2 (*Outer-product (CRf)*). Our second example pairs an outer-product parallelization with a column-wise distribution of **A**, a row-wise distribution of **B**, and a fine-grained distribution of **C**, where the distributions of **A** and **B** are matched with the parallelization while the distribution of **C** is not.

$$
\mathcal{V} = \{v_k : k \in [K]\} \cup \{v_{ij}^{\mathbf{C}} : (i,j) \in S_{\mathbf{C}}\},
$$
$$
\mathcal{N} = \{n_{ij}^{\mathbf{C}} = \{v_k : (\exists k)\,(i,k) \in S_{\mathbf{A}} \wedge (k,j) \in S_{\mathbf{B}}\} \cup \{v_{ij}^{\mathbf{C}}\} : (i,j) \in S_{\mathbf{C}}\},
$$
$$
w_{\mathrm{comp}}(v_k) = |\{i : (\exists k)\,(i,k) \in S_{\mathbf{A}}\}| \cdot |\{j : (\exists k)\,(k,j) \in S_{\mathbf{B}}\}|,
$$
$$
w_{\mathrm{mem}}(v_k) = |\{i : (\exists k)\,(i,k) \in S_{\mathbf{A}}\}| + |\{j : (\exists k)\,(k,j) \in S_{\mathbf{B}}\}|,
$$
$$
w_{\mathrm{comp}}(v_{ij}^{\mathbf{C}}) = 0, \quad w_{\mathrm{mem}}(v_{ij}^{\mathbf{C}}) = c(n_{ij}^{\mathbf{C}}) = 1.
$$

We have that $|\mathcal{V}| = K + |S_{\mathbf{C}}|$ and $|\mathcal{N}| = S_{\mathbf{C}}$, where each net has between 2 and $K + 1$ pins. Note that constructing $\mathcal{V}$ and $\mathcal{N}$ requires determining $S_{\mathbf{C}}$. A closely related hypergraph model for this example and two specializations (CRr and CRc) was previously studied [Akbudak and Aykanat 2014].

Example 5.3 (*Monochrome-**A** (Frf)*). Our third example pairs a monochrome-**A** parallelization with a row-wise distribution of **B** and fine-grained distributions of **A** and **C**, where the distribution of **A** is matched with the parallelization while the dis-

tributions of $\mathbf{B}$ and $\mathbf{C}$ are not.

$$\mathcal{V} = \{v_{ik} : (i,k) \in S_{\mathbf{A}}\} \cup \{v_k^{\mathbf{B}} : k \in [K]\} \cup \{v_{ij}^{\mathbf{C}} : (i,j) \in S_{\mathbf{C}}\},$$

$$\mathcal{N} = \bigcup \left\{ \begin{array}{c} \{n_k^{\mathbf{B}} = \{v_{ik} : (\exists i)\,(i,k) \in S_{\mathbf{A}}\} \cup \{v_k^{\mathbf{B}}\} : k \in [K]\}, \\ \{n_{ij}^{\mathbf{C}} = \{v_{ik} : (\exists k)\,(i,k) \in S_{\mathbf{A}} \wedge (k,j) \in S_{\mathbf{B}}\} \cup \{v_{ij}^{\mathbf{C}}\} : (i,j) \in S_{\mathbf{C}}\} \end{array} \right\},$$

$$w_{\mathrm{comp}}(v_{ik}) = w_{\mathrm{mem}}(v_k^{\mathbf{B}}) = c(n_k^{\mathbf{B}}) = |\{j : (\exists j)\,(k,j) \in S_{\mathbf{B}}\}|,$$

$$w_{\mathrm{mem}}(v_{ik}) = w_{\mathrm{mem}}(v_{ij}^{\mathbf{C}}) = c(n_{ij}^{\mathbf{C}}) = 1, \qquad w_{\mathrm{comp}}(v_k^{\mathbf{B}}) = w_{\mathrm{comp}}(v_{ij}^{\mathbf{C}}) = 0.$$

We have that $|\mathcal{V}| = |S_{\mathbf{A}}| + K + |S_{\mathbf{C}}|$ and $|\mathcal{N}| = K + |S_{\mathbf{C}}|$, where each of the K nets $n_k^{\mathbf{B}}$ has between 2 and $I+1$ pins, and each of the $|S_{\mathbf{C}}|$ nets $n_{ij}^{\mathbf{C}}$ has between 2 and $K+1$ pins. Note that constructing $\mathcal{V}$ and $\mathcal{N}$ requires determining $S_{\mathbf{C}}$.

Example 5.4 (*Monochrome-$\mathbf{C}$ (ffF)*).   Our fourth example pairs a monochrome-$\mathbf{C}$ parallelization with fine-grained distributions of $\mathbf{A}$, $\mathbf{B}$, and $\mathbf{C}$, where the distribution of $\mathbf{C}$ is matched with the parallelization while the distributions of $\mathbf{A}$ and $\mathbf{B}$ are not.

$$\mathcal{V} = \{v_{ij} : (i,j) \in S_{\mathbf{C}}\} \cup \{v_{ik}^{\mathbf{A}} : (i,k) \in S_{\mathbf{A}}\} \cup \{v_{kj}^{\mathbf{B}} : (k,j) \in S_{\mathbf{B}}\},$$

$$\mathcal{N} = \bigcup \left\{ \begin{array}{c} \{n_{ik}^{\mathbf{A}} = \{v_{ij} : (\exists j)\,(k,j) \in S_{\mathbf{B}}\} \cup \{v_{ik}^{\mathbf{A}}\} : (i,k) \in S_{\mathbf{A}}\}, \\ \{n_{kj}^{\mathbf{B}} = \{v_{ij} : (\exists i)\,(i,k) \in S_{\mathbf{A}}\} \cup \{v_{kj}^{\mathbf{B}}\} : (k,j) \in S_{\mathbf{B}}\} \end{array} \right\},$$

$$w_{\mathrm{comp}}(v_{ij}) = |\{k : (\exists k)\,(i,k) \in S_{\mathbf{A}} \wedge (k,j) \in S_{\mathbf{B}}\}|,$$

$$w_{\mathrm{mem}}(v_{ij}) = w_{\mathrm{mem}}(v_{ik}^{\mathbf{A}}) = w_{\mathrm{mem}}(v_{kj}^{\mathbf{B}}) = c(n_{ik}^{\mathbf{A}}) = c(n_{kj}^{\mathbf{B}}) = 1,$$

$$w_{\mathrm{comp}}(v_{ik}^{\mathbf{A}}) = w_{\mathrm{comp}}(v_{kj}^{\mathbf{B}}) = 0.$$

We have that $|\mathcal{V}| = |S_{\mathbf{A}}| + |S_{\mathbf{B}}| + |S_{\mathbf{C}}|$ and $|\mathcal{N}| = |S_{\mathbf{A}}| + |S_{\mathbf{B}}|$, where each of the $|S_{\mathbf{A}}|$ nets $n_{ik}^{\mathbf{A}}$ has between 2 and $J+1$ pins, and each of the $|S_{\mathbf{B}}|$ nets $n_{kj}^{\mathbf{B}}$ has between 2 and $I+1$ pins. Note that constructing $\mathcal{V}$ requires determining $S_{\mathbf{C}}$.

In each of these four examples, the simplified hypergraph is potentially much smaller than the original (fine-grained) hypergraph in terms of the numbers of vertices, nets, or pins. Our experiments in Sec. 6 using the PaToH hypergraph partitioning software demonstrated orders-of-magnitude improvements in runtime when partitioning simplified (vs. fine-grained) hypergraphs.

We caution that constructing the simplified hypergraphs in these four examples requires determining $S_{\mathbf{C}}$, which can be as expensive as computing $\mathbf{C}$. This cost can be avoided, e.g., in the row-wise (RrR) case (Ex. 5.1), if we omit memory weights.

5.5. Sparse Matrix-Vector Multiplication (SpMV)

In the case of multiplying a sparse matrix by a dense vector (SpMV), our hypergraph model simplifies in several ways. If the input matrix $\mathbf{B}$ is a dense vector of length K, then the output matrix $\mathbf{C}$ is a dense vector of length I, since $\mathbf{A}$ has a nonzero in every row under the assumptions of Sec. 3.1. In this case, there is only one multiplication performed for each nonzero of $\mathbf{A}$, so the multiplication vertices of $\mathcal{V}$ can be indexed by only two parameters, letting us write $\mathcal{V}^{\mathrm{m}} = \{v_{ik} : (i,k) \in S_{\mathbf{A}}\}$. Furthermore, the $\mathbf{B}$ and $\mathbf{C}$ nonzero vertices can be indexed by only one parameter and simplify to $\mathcal{V}^{\mathbf{B}} = \{v_k^{\mathbf{B}} : k \in [K]\}$ and $\mathcal{V}^{\mathbf{C}} = \{v_i^{\mathbf{C}} : i \in [I]\}$.

In addition to the simplifications based on the inputs of SpMV, we can also apply vertex coarsening, restricting the set of algorithms, to reproduce the "fine-grain" model for SpMV of Çatalyürek and Aykanat [2001]; this model assumes that $I = K$, i.e., $\mathbf{A}$ is a square matrix. There are three steps in the following derivation.

First, because each nonzero of $\mathbf{A}$ is involved in exactly one multiplication, we force v_{ik} and $v_{ik}^{\mathbf{A}}$ to be monochrome for each $(i, k) \in S_{\mathbf{A}}$. (This corresponds to monochrome-$\mathbf{A}$ (Fff), discussed in Sec. 5.4.) We use the notation $\hat{v}_{ik}$ to refer to the coarsened vertex, which has weights $w_{\mathrm{comp}}(\hat{v}_{ik}) = 1$ and $w_{\mathrm{mem}}(\hat{v}_{ik}) = 1$. Note that this coarsening implies that all nets in $\mathcal{N}^{\mathbf{A}}$ have exactly one vertex, and thus we omit them.

Second, we apply further coarsening to attain a symmetric partitioning of the input and output vector entries. We do this to satisfy the "consistency condition" [Çatalyürek and Aykanat 2001], motivated by the problem of performing repeated SpMV operations with the same matrix. For every $i \in [I] = [K]$, we force $v_i^{\mathbf{B}}$, $v_i^{\mathbf{C}}$, and $\hat{v}_{ii}$ (if it exists) to be monochrome. This corresponds to assigning vector entries to the processor that owns the corresponding diagonal entry of the matrix. We maintain the notation $\hat{v}_{ii}$ to reference the coarsened vertex, which has weights $w_{\mathrm{comp}}(\hat{v}_{ii}) = 1$ and $w_{\mathrm{mem}}(\hat{v}_{ii}) = 3$, if $(i, i) \in S_{\mathbf{A}}$, and $w_{\mathrm{comp}}(\hat{v}_{ii}) = 0$ and $w_{\mathrm{mem}}(\hat{v}_{ii}) = 2$, if $(i, i) \notin S_{\mathbf{A}}$.

Third, because the fine-grain SpMV model does not explicitly enforce memory balance, we drop the memory balance constraint on partitions by setting $\delta = p - 1$. Even though we removed the memory balance constraints, computational balance in this model guarantees a balanced partition of the nonzeros of $\mathbf{A}$. However, the memory allocated to the input and output vectors is not accounted for, and so the overall memory imbalance can potentially be higher than the computational imbalance.

The preceding derivation reproduces the "fine-grain" model for SpMV [Çatalyürek and Aykanat 2001], which consists of a vertex for every nonzero in the matrix (and a zero-weight "dummy" vertex for every zero diagonal entry) and a net for each row and each column. Similar simplifications to 1D SpGEMM algorithms (see Sec. 5.4) yield row-wise and column-wise SpMV algorithms [Çatalyürek and Aykanat 1999]. In particular, the row-wise (RrR) hypergraph (Ex. 5.1) is identical to the "column-net" SpMV hypergraph (modeling a row-wise algorithm), except for the presence of the B-nonzero vertices and the memory weights. Similarly, the outer-product (CRf) hypergraph (Ex. 5.2) is identical to the "row-net" SpMV hypergraph (modeling a column-wise algorithm), except for the presence of the C-nonzero vertices and the memory weights. On the other hand, the column-wise SpGEMM parallelizations applied to SpMV have no analogue in Çatalyürek and Aykanat's models [1999; 2001] since there is no parallelism. Additionally, the monochrome-A, -B, and -C SpGEMM parallelizations, in the case of SpMV, correspond to the aforementioned "fine-grain", "row-net", and "column-net" SpMV hypergraphs.

5.6. Generalizing SpGEMM Algorithms

In this section, we explore two directions in which the class of SpGEMM algorithms (defined in Sec. 3) can be generalized. The first generalization exploits relations among the input matrices' entries to reduce algorithmic costs — recall that Sec. 3.1 assumed $\mathbf{A}$- and $\mathbf{B}$-entries are unrelated. A common application is exploiting symmetry of $\mathbf{A}$, $\mathbf{B}$, or $\mathbf{C}$. The second generalization considers the case where only a subset of the output entries are desired, a task called *masked* SpGEMM — recall that Sec. 3.1 assumed all nonzero C-entries are computed.

We show that both classes of algorithms can be modeled by straightforward simplifications to the SpGEMM hypergraph (defined in Sec. 3.2). In particular, we apply vertex coarsening similarly to as in Sec. 5.1, except here we adjust vertex weights and net costs in a simpler manner. The hypergraph-based communication bounds in Secs. 4.1 and 4.2, as well as the restrictions studied in the previous subsections of Sec. 5, extend to both these classes of algorithms with only minor changes required to address the vertex coarsening.

5.6.1. Exploiting Input Matrix Relations. We consider modifying SpGEMM algorithms to exploit known equivalence relations on the nonzero entries of $\mathbf{A}$ and $\mathbf{B}$ to reduce algorithmic costs. For simplicity, we will consider equality relations: e.g., if $\mathbf{A} = \mathbf{A}^T$, then nonzeros in the upper triangle of $\mathbf{A}$ equal the corresponding entries in the lower triangle of $\mathbf{A}$. However, the following approach extends to cases like $\mathbf{A} = -\mathbf{A}^T$, or more generally to the case where each equivalence class is assigned a single value and each nonzero in that equivalence class is a function of that value.

Consider any equality relation on the nonzeros of $\mathbf{A}$ and $\mathbf{B}$. We first consider SpGEMM-like algorithms that store exactly one copy of repeated $\mathbf{A}$- and $\mathbf{B}$-entries, but still perform all nontrivial multiplications and compute all nonzero $\mathbf{C}$-entries, regardless of whether these values also include repetitions. Such algorithms can be naturally modeled by altering the SpGEMM hypergraph $(\mathcal{V}, \mathcal{N}) = \mathcal{H}(\mathbf{A}, \mathbf{B})$ according to a given partition of $\mathcal{V}^{\mathbf{A}} \cup \mathcal{V}^{\mathbf{B}}$, where each part is a set of nonzero entries that must have the same value. For example, if an algorithm only inputs the unique entries of $\mathbf{A} = \mathbf{A}^T$, then the partition groups each pair of distinct vertices $v_{ik}^{\mathbf{A}}$ and $v_{ki}^{\mathbf{A}}$ while keeping the nonzero vertices associated with the diagonal of $\mathbf{A}$, as well as all of $\mathcal{V}^{\mathbf{B}}$, in singleton parts. Now coarsen these vertex sets, following the approach in Sec. 5.1 except setting the memory costs of the coarsened vertices to 1, rather than the number of coarsened vertices. Since only nonzero vertices are combined, no nets are coalesced by this coarsening.

We next extend the preceding class of SpGEMM-like algorithms to avoid performing redundant nontrivial multiplications and computing redundant $\mathbf{C}$-entries. We say that two nontrivial multiplications are redundant if their left operands are equal and their right operands are equal. For example, if $\mathbf{A} = \mathbf{A}^T$ and $\mathbf{B} = \mathbf{B}^T$, then any nontrivial multiplication $a_{lm}b_{ml}$ must equal $a_{ml}b_{lm}$. More generally, if there exist (i, k, j) and (r, t, s) such that $a_{ik} = a_{rt}$ and $b_{kj} = b_{ts}$, then $a_{ik}b_{kj} = a_{rt}b_{ts}$ — only one of these two multiplicationsies needs to be performed. We say that two $\mathbf{C}$-entries are redundant if their summands can be matched as pairs of redundant nontrivial multiplications. For example, if all entries of $\mathbf{A}$ and $\mathbf{B}$ are equal to the same nonzero value, then all $\mathbf{C}$-entries are equal. More generally, if there exist (i, j) and (r, s) in $S_{\mathbf{C}}$ and a bijection

$$\phi \colon \mathcal{K} = \{k : (\exists k)\,((i, k) \in S_{\mathbf{A}} \wedge (k, j) \in S_{\mathbf{B}})\} \to \{t : (\exists t)\,((r, t) \in S_{\mathbf{A}} \wedge (t, s) \in S_{\mathbf{B}})\}$$

such that $a_{ik} = a_{r\phi(k)}$ and $b_{kj} = b_{\phi(k)s}$ for each $k \in \mathcal{K}$, then $c_{ij} = c_{rs}$, and only one of the two needs to be computed. We further alter the SpGEMM hypergraph to exploit these savings as follows. Let the notation $u \equiv v$ assert that $u, v \in \mathcal{V}$ are contained in the same part in some fixed partition of $\mathcal{V}$. We extend the partition of $\mathcal{V}^{\mathbf{A}} \cup \mathcal{V}^{\mathbf{B}}$ to all of $\mathcal{V}$ in two steps, partitioning $\mathcal{V}^{\mathbf{m}}$ and then $\mathcal{V}^{\mathbf{C}}$. First, for each pair $v_{ikj}, v_{rts} \in \mathcal{V}^{\mathbf{m}}$, we specify that $v_{ikj} \equiv v_{rts}$ if $v_{ik}^{\mathbf{A}} \equiv v_{rt}^{\mathbf{A}}$ and $v_{kj}^{\mathbf{B}} \equiv v_{ts}^{\mathbf{B}}$. Second, for each $v_{ij}^{\mathbf{C}}, v_{rs}^{\mathbf{C}} \in \mathcal{V}^{\mathbf{C}}$, we specify that $v_{ij}^{\mathbf{C}} \equiv v_{rs}^{\mathbf{C}}$ if $n_{ij}^{\mathbf{C}}$ and $n_{rs}^{\mathbf{C}}$ intersect the same parts of $\mathcal{V}^{\mathbf{m}}$. We then coarsen $\mathcal{V}$ according to this partition: we follow the approach in Sec. 5.1 except setting both the memory costs of the coarsened nonzero vertices and the computation costs of the coarsened multiplication vertices to 1, rather than to the numbers of coarsened nonzero and multiplication vertices. Now, it is possible that some nets are coalesced: coalesced nets can be combined without increasing net costs since only one nonzero needs to be stored/sent/received.

Lastly, we next extend the preceding class of SpGEMM-like algorithms to exploit commutative multiplication, which is not guaranteed in our model (see Sec. 3). For example, if $\mathbf{B} = \mathbf{A}$, then $a_{lm}b_{ml} = a_{ml}b_{lm}$; or, if $\mathbf{B} = \mathbf{A}^T$, then $\mathbf{C} = \mathbf{C}^T$; neither of these redundancies necessarily occur without commutative multiplication. We model algorithms that avoid this larger class of redundant multiplications (and $\mathbf{C}$-entries) by

augmenting the preceding construction in just one place: for each pair $v_{ikj}, v_{rts} \in \mathcal{V}^{\mathbf{m}}$, we additionally specify that $v_{ikj} \equiv v_{rts}$ if $v_{ik}^{\mathbf{A}} \equiv v_{rt}^{\mathbf{B}}$ and $v_{kj}^{\mathbf{A}} \equiv v_{ts}^{\mathbf{B}}$.

5.6.2. Masked SpGEMM. We consider modifying SpGEMM algorithms to only compute a subset of the entries of $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$; this is known as *masked* SpGEMM (see, e.g., [Azad et al. 2015b]), where $S \subseteq S_{\mathbf{C}}$ indexes the desired subset of C-entries and $S_{\mathbf{C}} \setminus S$ is called the *mask*. That is, only for each $(i, j) \in S$ are c_{ij} and its associated nontrivial multiplications $a_{ik}b_{kj}$ computed. We develop a hypergraph model starting with the usual SpGEMM hypergraph, $(\mathcal{V}, \mathcal{N}) = \mathcal{H}(\mathbf{A}, \mathbf{B})$. We remove from $\mathcal{N}^{\mathbf{C}}$ each net $n_{ij}^{\mathbf{C}}$ where $(i, j) \notin S$, as well as that net's elements (vertices) from $\mathcal{V}$. In particular, for each $(i, j) \in S_{\mathbf{C}} \setminus S$, this removes from $\mathcal{V}^{\mathbf{C}}$ the vertex $v_{ij}^{\mathbf{C}}$ and from $\mathcal{V}^{\mathbf{m}}$ each vertex v_{ikj} with $k \in [K]$.

Depending on $S_{\mathbf{A}}$, $S_{\mathbf{B}}$, and S, it is possible that some entries of $\mathbf{A}$ and $\mathbf{B}$ are not involved in any nontrivial multiplications after masking. Avoiding this triviality motivated our simplifying assumption in Sec. 3 that $\mathbf{A}$ and $\mathbf{B}$ have no zero rows or columns. In terms of the hypergraph, this manifests as elements of $\mathcal{N}^{\mathbf{A}}$ and $\mathcal{N}^{\mathbf{B}}$ becoming singletons, containing only their associated nonzero vertices $v_{ik}^{\mathbf{A}}$ or $v_{kj}^{\mathbf{B}}$. To model algorithms that reduce memory costs by not storing the associated nonzero entries, these nets and vertices can be removed from $\mathcal{N}^{\mathbf{A}}$ or $\mathcal{N}^{\mathbf{B}}$, and $\mathcal{V}^{\mathbf{A}}$ or $\mathcal{V}^{\mathbf{B}}$, respectively.

Next, we consider a generalization of masked SpGEMM. When the underlying set X has a right multiplicative identity 1, masking can be viewed as computing $(\mathbf{A} \cdot \mathbf{B}) \odot \mathbf{M}$, where $\odot$ denotes the Hadamard (or entrywise) matrix product and where the $\{0, 1\}$-valued matrix $\mathbf{M}$ has the same dimensions as $\mathbf{C} = \mathbf{A} \cdot \mathbf{B}$ and $S_{\mathbf{M}} = S$. To generalize, consider computing $\mathbf{C} = (\mathbf{A} \cdot \mathbf{B}) \odot \mathbf{M}$, where $\mathbf{M}$ is a general sparse matrix with the same dimensions as $\mathbf{C}$. Let $(\mathcal{V}, \mathcal{N})$ denote the hypergraph for masked SpGEMM with $S = S_{\mathbf{M}}$, as in the preceding paragraphs. Suppose the algorithmic constraint that any processor who stores some c_{ij} also must both store m_{ij} and perform the multiplication involving m_{ij}. To model this, for each of the (unmasked) nonzero vertices $v_{ij}^{\mathbf{C}}$, we set its memory cost to 2 and its computation cost to 1. It is possible to relax this algorithmic constraint by introducing additional multiplication vertices v_{ij}' for the Hadamard product, additional nonzero vertices $v_{ij}^{\mathbf{M}}$ for the M-entries, and additional nets $n_{ij}^{\mathbf{M}} = \{v_{ij}', v_{ij}^{\mathbf{M}}, v_{ik}^{\mathbf{C}}\}$.

6. EXPERIMENTAL RESULTS

We now compare our fine-grained hypergraph model from Sec. 3.2 with the six restricted parallelizations developed in Sec. 5.2, in the context of three applications: algebraic multigrid (AMG, Sec. 6.1), linear programming (LP, Sec. 6.2), and Markov clustering (MCL, Sec. 6.3). A list of the SpGEMMs that we study in our experiments is shown in Tab. II. Our experiments consider only the distributed-memory parallel case, where we expect communication cost to be most closely correlated with execution time. Because hypergraph partitioning is NP-hard (in general), hypergraph partitioners use heuristics to approximate optimal partitions; results we present are not guaranteed to be optimal. However, it remains open whether partitioning SpGEMM hypergraph instances are NP-hard.

The hypergraph partitioner we use is PaToH (version 3.2) [Çatalyürek and Aykanat 1999]. PaToH minimizes the "connectivity metric," defined as the sum, over all nets, of the product of each net's cost and its number of incident parts minus one. The communication costs shown below represent the maximum, over all parts, of the sum of each part's non-internal incident nets' costs (matching Lem. 4.2). This cost is most closely aligned with the parallel running time

Table II: The SpGEMMs that we studied in our experiments and their parameters. The model AMG problem is represented by the rows 27-AP and 27-PTAP, corresponding to the $A{\cdot}P$ and $P^T{\cdot}(AP)$ SpGEMMs, respectively, and the SA-ρAMGe problem is similarly represented by the rows SA-AP and SA-PTAP. The AMG experiments were carried out in a weak scaling regime and are represented in this table by the largest instances. The LP experiments have one input matrix and compute $\mathbf{C} = \mathbf{A} \cdot \mathbf{A}^T$, and the MCL experiments have one input matrix and compute $\mathbf{C} = \mathbf{A} \cdot \mathbf{A}$. In addition to the the dimensions of the SpGEMM, we provide the average number of nonzeros per row of each matrix and the ratio of nontrivial multiplications to output nonzeros.

	Name	I	K	J	$|S_\mathbf{A}|/I$	$|S_\mathbf{B}|/K$	$|S_\mathbf{C}|/I$	$|\mathcal{V}^\mathbf{m}|/|S_\mathbf{C}|$
AMG (Sec. 6.1)	27-AP	970,299	970,299	35,937	26.5	4.5	12.1	9.9
	27-PTAP	35,937	970,299	35,937	4.5	12.1	25.4	49.0
	SA-AP	1,088,640	1,088,640	31,496	26.4	20.1	38.5	13.9
	SA-PTAP	31,496	1,088,640	31,496	696.3	38.5	216.4	139.3
LP (Sec. 6.2)	fome21	67,748	216,350	67,748	6.9	2.2	9.5	1.6
	pds80	129,181	434,580	129,181	7.2	2.1	9.7	1.6
	pds100	156,243	514,577	156,243	7.0	2.1	9.4	1.6
	cont11l	1,468,599	1,961,394	1,468,599	3.7	2.7	12.3	1.5
	sgpf5y6	246,077	312,540	246,077	3.4	2.7	11.3	1.2
MCL (Sec. 6.3)	biogrid11	5,853	5,853	5,853	21.5	21.5	2,105.7	1.6
	dip	5,051	5,051	5,051	8.7	8.7	200.9	1.6
	wiphi	5,955	5,955	5,955	8.4	8.4	85.6	1.5
	dblp	425,957	425,957	425,957	4.9	4.9	64.8	1.7
	enron	36,692	36,692	36,692	10.0	10.0	831.0	1.7
	facebook	4,039	4,039	4,039	43.7	43.7	717.1	6.5
	roadnetca	1,971,281	1,971,281	1,971,281	2.8	2.8	6.5	1.4

devoted to communicating words of data, called the "critical-path bandwidth cost" by Ballard et al. [2014].

We do not constrain memory load balance (i.e., $\delta = p - 1$) — and so we omit $\mathcal{V}^\mathrm{nz}$ from the hypergraphs — but we do constrain computation load balance, $\epsilon = 0.01$. In a few cases, our partitioner failed to produce such a balanced partition. We discuss these cases below.

We performed our experiments on two machines at the National Energy Research Supercomputing Center in Berkeley, California. The experiments with the AMG model problem were carried out on a node (1 TB memory) of the machine Carver; all other experiments were carried out on a node (128 GB memory) of the machine Cori.

Building the fine-grained hypergraph and partitioning it is as expensive as performing the associated SpGEMM, and furthermore our partitioner is sequential. Although we are not focused on partitioning times in this study, fast partitioning is important for a user who seeks to find the best SpGEMM strategy for a particular problem. In our experiments, the partitioning times varied from a few seconds up to 5 hours, while the time to build the hypergraph was negligible in comparison. We observed that the partitioning time varied with the dimensions of the matrices (smaller was faster), the hypergraph model (1D was fastest while fine-grained was slowest), and on whether the matrices have a regular nonzero structure (the 27-point stencil matrices in Sec. 6.1 were quicker to partition, while the social-network matrices in Sec. 6.3 were slower).
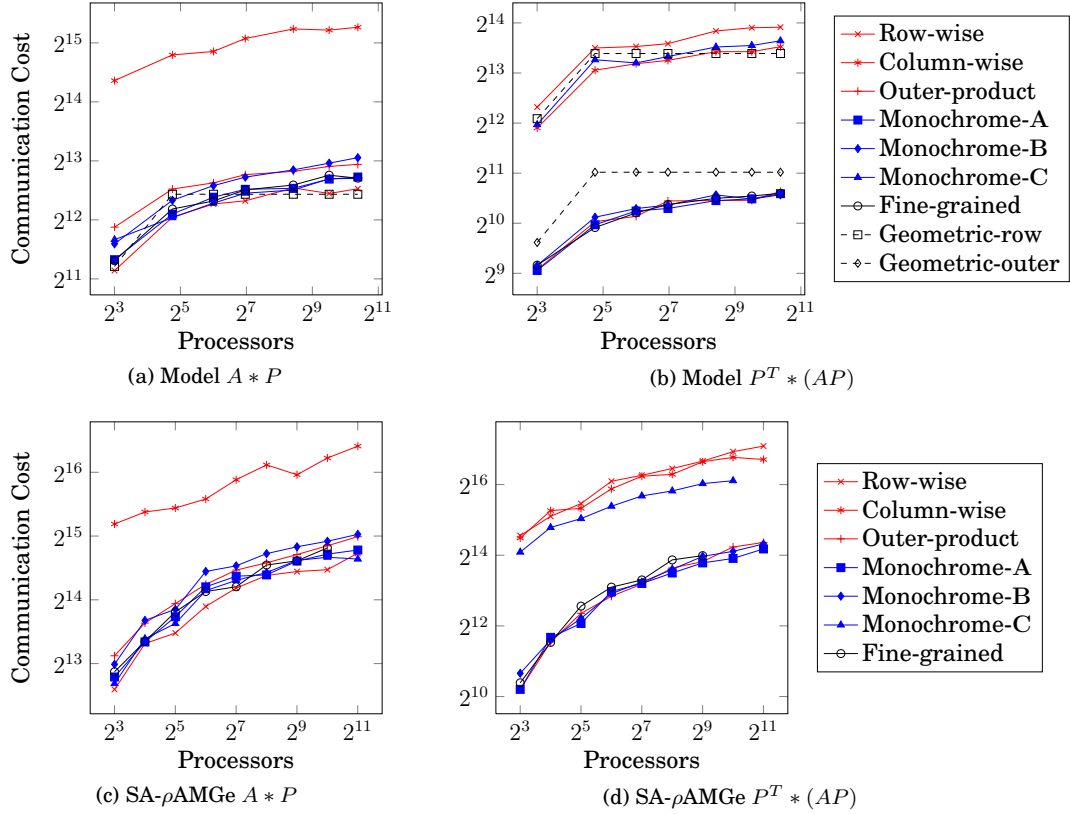
Fig. 7: Communication costs of various parallelizations for performing the top-level triple product in the setup phase of algebraic multigrid (AMG — Sec. 6.1).

6.1. Application: Algebraic Multigrid

First, we consider the SpGEMMs that arise in the context of algebraic multigrid (AMG). AMG methods are linear solvers typically applied to systems of linear equations arising from the discretization of partial differential equations (PDEs). AMG methods operate in two stages: setup and solve. The setup stage consists of forming a "grid hierarchy," which is a sequence of matrices that represent the PDE at progressively coarser levels of discretization.

To build the grid hierarchy, we generate a sequence of "prolongators" $P_1, \ldots, P_{\ell-1}$, where ℓ is the number of levels in the hierarchy, an algorithm parameter. Prolongators are matrices with more rows than columns and the number of rows in each prolongator is equal to the number of columns in its predecessor. The grid hierarchy $A_1, \ldots, A_\ell$ is then formed by successively computing

$$A_2 = P_1^T A_1 P_1, \qquad A_3 = P_2^T A_2 P_2, \qquad \ldots, \qquad A_\ell = P_{\ell-1}^T A_{\ell-1} P_{\ell-1}, \tag{6}$$

where A_1 is the fine-grained discretization of the PDE, an algorithm input.

Typically, AMG methods are designed such that all of the matrices in eq. (6) are sparse, so each triple product is computed using two SpGEMMs: one that forms $A_k \cdot P_k$, followed by another that forms $P_k^T \cdot (A_k P_k)$ for $k \in [\ell - 1]$. In practice, the SpGEMMs in eq. (6) can take a large part of the time spent in the setup phase, and the setup phase

can take a large part of the total time. Considerable effort has been made recently to optimize these SpGEMMs (see, e.g., [Ballard et al. 2015b; Park et al. 2015]).

We consider the first level of the grid hierarchy for both a model problem and a realistic problem. In the model problem, the rows of the input matrix A_1 correspond to points of an $N \times N \times N$ regular grid, and its nonzero structure corresponds to a 27-point stencil so that each point is adjacent to its (at most) 26 nearest neighbors. The prolongator matrix P_1 is $N^3 \times (N/3)^3$ (we assume N is divisible by 3) and is defined so that $3 \times 3 \times 3$ sub-grids correspond to single points in the coarser grid, and its values are computed using the technique of "smoothed aggregation" (using damped Jacobi). This model problem is one that has been previously studied [Ballard et al. 2015a; Ballard et al. 2015b]. Because the grid is regular, we can compare our partitioning results with regular, geometric partitions. For example, the natural partition of the rows of A corresponds to assigning each processor a contiguous $(N/p^{1/3}) \times (N/p^{1/3}) \times (N/p^{1/3})$ subcube of points (assuming $N/p^{1/3}$ is an integer).

The second set of matrices comes from an oil-reservoir simulation called SPE10 [Christie and Blunt 2001]. In this problem, the Darcy equation, in dual form, is used to model fluid flow in porous media. The spatial domain is a rectangular prism, tessellated using a hexahedral mesh and discretized using linear finite elements [Barker et al. 2015; Christensen et al. 2015]. Brezina and Vassilevski [2011] solved this problem using an AMG variant called SA-ρAMGe, wherein P_1 has approximately $35\times$ more rows than columns (slightly more aggressive coarsening than the model problem) and is constructed using a polynomial smoother, giving more nonzeros. For details of the implementation, see [Kalchev 2012; Kalchev et al. 2013].

We consider both SpGEMMs for both problems, comparing the fine-grained model with all six restricted parallelizations from Sec. 5.2. In the case of the model problem, we also compare against 1D algorithms based on geometric partitions of the regular grid; these natural and efficient (but not necessarily optimal) instances can validate the quality of results obtained from the hypergraph partitioner. For both problems, we present a weak-scaling study, maintaining I/p, where $I = K$ is the number of rows and columns of A_1. The model problem ranges from dimension $18^3 = 5832$ on 8 processors up to dimension $99^3 = 970229$ on 1331 processors, and the SA-ρAMGe problem ranges from dimension 4900 on 8 processors up to dimension 1088640 on 2048 processors.

In four cases involving the SA-ρAMGe matrices, our partitioner failed with an out-of-memory error. The P_1 matrix in SA-ρAMGe has more nonzero entries than in the model problem, and so the corresponding hypergraphs and their partitioning require more memory. The SpGEMMs that failed were the largest for the fine-grained model of $A \cdot P$, the largest for the monochrome-C model of $P^T \cdot (AP)$, and the two largest for the fine-grained model of $P^T \cdot (AP)$.

The results for the model problem are given in the top row of Fig. 7. We see in Fig. 7a that for the first SpGEMM, all parallelizations except for column-wise require about the same amount of communication (to within a factor of 2). Note that the row-wise algorithm with geometric partitioning (labeled "Geometric-row") achieves the least communication on 1331 processors; this implies that PaToH is not finding an optimal partition at that scale (at least in the row-wise, monochrome-A, monochrome-C, or fine-grained cases), but we believe that the optimal partition is not much better than the geometric one. We can also compare the relative costs discovered by hypergraph partitioning with Ballard et al.'s theoretical analysis [2015b, Table 2] based on geometric partitioning. Theoretically, for geometric partitions, the row-wise algorithm requires factors of 2.0 and 4.9 less communication than outer-product and column-wise algorithms, respectively; the ratios for partitions discovered by PaToH are 1.3 and 6.7 on 1331 processors. We conclude from this data that a row-wise parallelization is the

simplest parallelization that is nearly optimal for computing $A \cdot P$; this agrees with previous results [Ballard et al. 2015a; Ballard et al. 2015b].

For the second SpGEMM, we see that the outer-product, monochrome-**A**, and monochrome-**B** parallelizations all perform as well as the fine-grained model (recall from Fig. 6 that monochrome-**A** and -**B** parallelizations refine outer product parallelizations). The other parallelizations all perform similarly to each other, and require about $10\times$ more communication than the fine-grained parallelization. We note that the outer-product algorithm based on geometric partitioning (labeled "Geometric-outer") is (slightly) outperformed by the partitions discovered by PaToH. For geometric partitions, the outer-product algorithm communicates a factor of $5\times$ less than the row-wise algorithm (matching previous theoretical analyses [Ballard et al. 2015b, Tab. 2]); using hypergraph partitioning that difference increases to $7\times$ on 1331 processors. We conclude that an outer-product parallelization is the simplest parallelization that is nearly optimal for computing $P^T \cdot (AP)$. In our earlier paper [Ballard et al. 2015a], for the second SpGEMM, we compared Geometric-row with the fine-grained model and concluded that Geometric-row was suboptimal. The present experiments add the course-grained models and Geometric-outer and draw stronger conclusions.

The results for the SA-ρAMGe problem are shown in the bottom row of Fig. 7. For both SpGEMMs, the results are qualitatively very similar to those of the model problem. Because the mesh's geometry is not available within the AMG code, we do not have parallelizations based on geometric partitions with which to compare. The conclusions are the same: the row-wise algorithm is nearly optimal for the case of $A \cdot P$, and the outer-product parallelization is nearly optimal for the case of $P^T \cdot (AP)$.

6.2. Application: Linear Programming Normal Equations

We next consider SpGEMMs within a linear programming application, studied previously by Akbudak and Aykanat [2014], wherein a linear program is solved by an interior-point method. During each iteration of this method, normal equations of the form $A \cdot D^2 \cdot A^T$ are constructed via SpGEMM, and the resulting system is solved either directly (via Cholesky factorization) or iteratively (via a Krylov subspace method). Boman, Parekh, and Phillips [2005] reported that these SpGEMMs can dominate the overall runtime. The matrix A encodes the linear constraints and remains fixed throughout the iterations, while the positive diagonal matrix D varies with each iteration. Therefore, the SpGEMMs always have the form $\mathbf{A} \cdot \mathbf{B}$, where $S_{\mathbf{B}} = S_{\mathbf{A}^T}$ and both $S_{\mathbf{A}}$ and $S_{\mathbf{B}}$ remain unchanged throughout the iterations. This means there is potential to amortize the cost of forming and partitioning the SpGEMM hypergraph.

Akbudak and Aykanat considered SpGEMM algorithms similar to what we call outer-product (CRf), (CRr), and (CRc) — see Ex. 5.2).

We consider the same five linear programming constraint matrices as Akbudak and Aykanat — fome21, pds-80, pds-100, sgpf5y6, and cont11l — all from the University of Florida collection [Davis and Hu 2011]. In each case, the SpGEMM dimensions $I = J < K$, i.e., the output matrix is smaller in dimension than the input matrices. In Fig. 8, we compare the fine-grained hypergraph model with the six restricted parallelizations from Sec. 5.2 for these five problems; note that since $S_{\mathbf{B}} = S_{\mathbf{A}^T}$, columnwise is equivalent to row-wise and monochrome-**B** is equivalent to monochrome-**A**, so we omit these curves. We perform a strong-scaling study, increasing the number of processors for fixed matrices.

In the first four problems (Figs. 8a to 8d), the fine-grained, outer-product, and monochrome-**A** parallelizations were the most communication efficient, while the row-wise and monochrome-**C** parallelizations were the least. The largest difference between row-wise and outer-product parallelizations was observed in the fourth problem (Fig. 8d), with a communication cost ratio of $23\times$ (on 128 processors). This trend is

less clear in the fifth problem (Fig. 8e), where all parallelizations' costs varied within a factor of two. For all problems, we observed the differences between parallelization costs to be much less sensitive to processor count than in the following experiment (see Fig. 9).

Since monochrome-A includes outer-product and monochrome-C includes row-wise (in the sense of Fig. 6), these results suggest that 2D parallelizations don't provide much advantage over 1D parallelizations; on the other hand, it is important to use outer-product instead of row-wise. The observation in all cases that the outer-product curve closely tracks the fine-grained curve supports Akbudak and Aykanat's decision to consider only outer-product parallelizations, at least for these problems.

6.3. Application: Markov Clustering

Van Dongen's Markov Clustering Algorithm (MCL) [2000] is an iterative scheme for determining clusters in graphs. The basic idea is to square the adjacency matrix of a graph, "inflate" and prune the result based on its values to maintain sparsity, and then iterate on the result. The computational bottleneck for the algorithm is using SpGEMM to square the sparse matrix, so in this section we explore squaring several different matrices coming from social-network and protein-interaction graphs, where Markov clustering is commonly used.

We consider only the first iteration of MCL, squaring the original adjacency matrix, as a representative example of the iteration. There have been several proposed variants of MCL, including (multi-level) regularized MCL [Satuluri and Parthasarathy 2009; Niu et al. 2014], that perform slightly different SpGEMMs. We believe that the present experiments can help inform algorithmic choices for parallelizing any clustering algorithm applied to scale-free graphs that uses SpGEMM as its computational workhorse.

Of the 11 non-synthetic data sets considered by Niu et al. [2014], we present results for the 7 matrices that are publicly available and whose fine-grained hypergraphs fit in memory on our machine. The matrices dblp, enron, and facebook are social network matrices, roadnetca is a graph representing roads and intersections, and dip, wiphi, and biogrid11 are protein-protein interaction matrices. Fig. 9 presents the results.

We consider the fine-grained hypergraph model as well as four of the restricted parallelizations: row-wise, outer-product, monochrome-A, and monochrome-B. Because we're squaring symmetric matrices in this section, column-wise and monochrome-B parallelizations are equivalent. Note that we do not exploit symmetry in these experiments. As in Sec. 6.2, we perform a strong-scaling study.

The overall conclusion from Fig. 9 is that 2D algorithms seem to require significantly less communication than 1D algorithms, particularly at high processor counts. The largest difference was for the facebook matrix, where the ratio of communication costs between monochrome-C and outer-product was $83\times$ (on 4096 processors). Furthermore, the downward trend of 2D and 3D plot lines show that per-processor communication costs decreased as the number of processors increased, indicating scalability. The 1D plot lines are mostly horizontal, implying that 1D algorithms are not scalable for these problems. The partitions of the 1D models we obtained from the partitioner violated our load-balance constraint ($\epsilon = 0.01$) and were increasingly imbalanced as we increased the number of parts. We attribute this to the presence of "heavy" vertices whose weights exceeded the maximum per-part weight prescribed by our constraint. The exception to these trends is the roadnetca matrix, which is qualitatively different from the social network and protein-protein interaction matrices (we include it for completeness).

We mention that Faisal, Parthasarathy, and Sadayappan [2014] parallelized regularized MCL within a graph framework. Their MPI-based algorithm used a 1D (row-wise)

(a) fome21
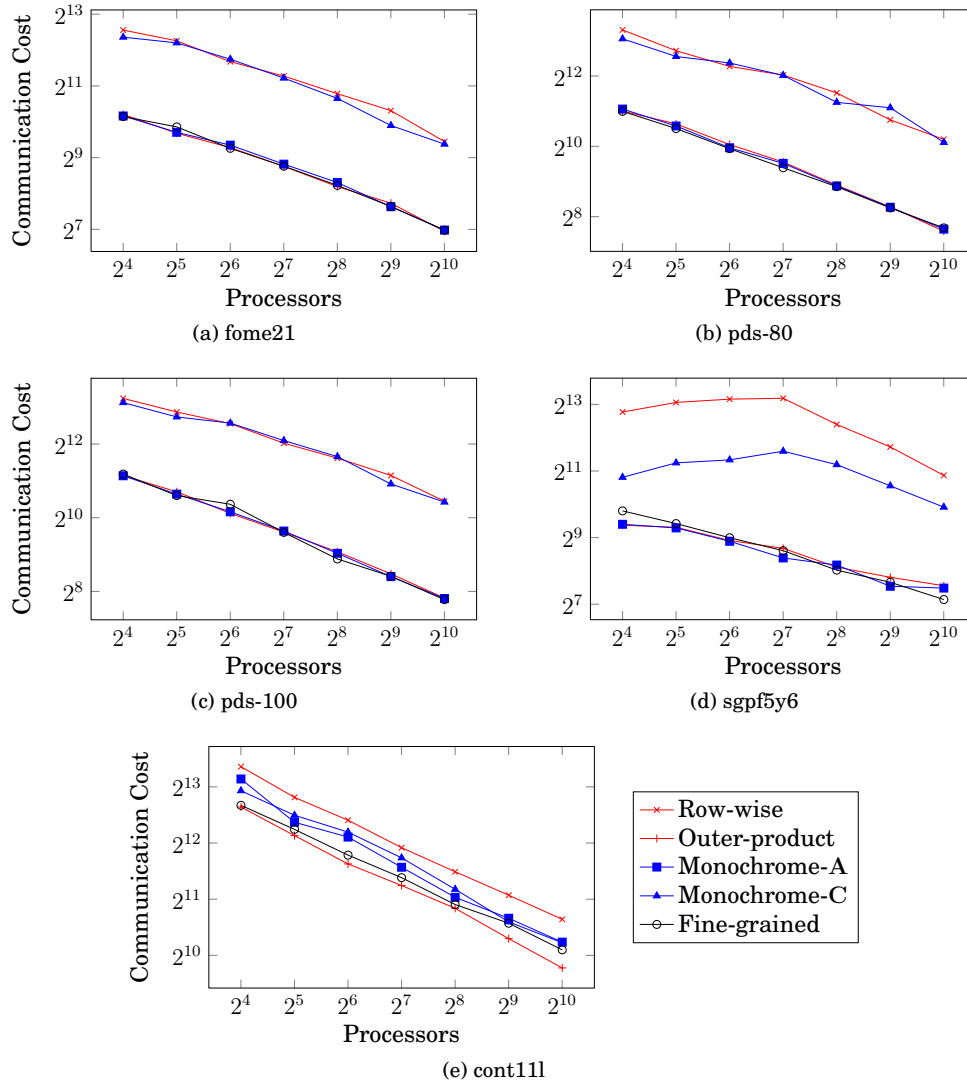
(b) pds-80

(c) pds-100

(d) sgpf5y6

(e) cont11l

Fig. 8: Communication costs of various hypergraph models for constructing the coefficient matrix for the normal equations in the context of the interior point methods for Linear Programming (LP — Sec. 6.2).

parallelization, but the data from these experiments suggests that a 2D parallelization would likely perform better for these types of matrices. This conclusion agrees with Boman, Devine, and Rajamanickam's premise [2013] in the context of SpMV: when $S_{\mathbf{A}}$ is a scale-free graph, 1D partitions of the SpMV hypergraph are often sub-optimal.

7. CONCLUSION

We have proposed a framework for studying SpGEMM algorithms, called the fine-grained hypergraph model, which captures communication and computation costs, as well as memory usage (see Sec. 3). We have defined hypergraph partitioning problems
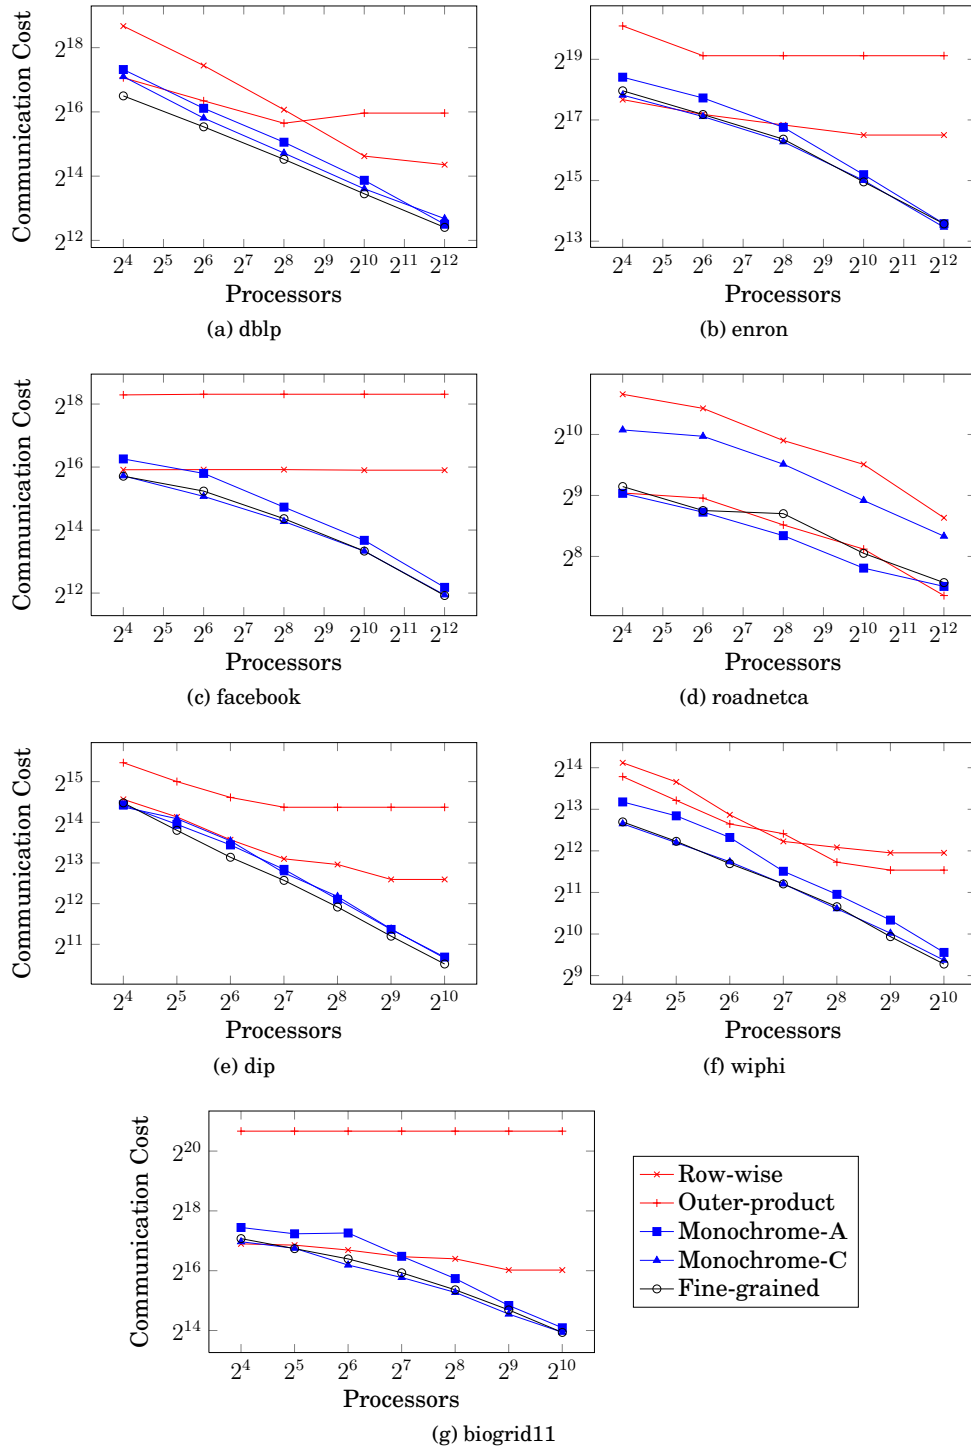
Fig. 9: Communication costs of various hypergraph models for squaring symmetric matrices in the context of the Markov Clustering Algorithm (MCL — Sec. 6.3).

whose solutions yield lower bounds on communication costs on both parallel and sequential machines, and algorithms that attain these lower bounds within constant factors (see Sec. 4). We have also shown how to simplify the fine-grained hypergraph model, applying vertex coarsening to focus on special classes of SpGEMM algorithms such as 1D and 2D classes (see Sec. 5). We applied these theoretical tools to real-world SpGEMM problems, experimentally confirming that hypergraph partitioning helps identify algorithms with reduced communication costs (see Sec. 6).

A key feature of our approach to SpGEMM analysis is that it is sparsity dependent. The fact that our communication lower bounds, parameterized by the nonzero structures of the input matrices, are tight suggests that a simpler parameterization — in terms of numbers of nonzeros, number of arithmetic operations, etc. — is unlikely. However, for specific classes of instances, simpler tight lower bounds are possible: e.g., in the case where the input matrices are both dense, tight lower bounds are parameterized by the matrix dimensions (see, e.g., [Ballard et al. 2014]).

The class of SpGEMM algorithms studied in this paper is quite general, and efficiently implementing a generic member of this class remains a practical challenge. On the other hand, implementations of 1D and 2D algorithms are much simpler, and our experiments demonstrated that in some cases these simpler approaches suffice to minimize communication. For example, in the AMG and LP applications (Secs. 6.1 and 6.2), we observed that some (but not all) of the 1D algorithms could minimize communication. On the other hand, in the MCL application (Sec. 6.3), 2D algorithms minimized communication and 1D algorithms did not, assuming that the partition of the fine-grained model is optimal, or nearly so.

We have proposed hypergraph analysis as a practical tool for SpGEMM algorithm design. We do not advocate solving hypergraph partitioning problems on the fly: this cost could dominate that of the actual SpGEMM computation. Our approach would be most practical when the preprocessing overhead can be amortized over many SpGEMMs, like when a sequence of matrices with the same (or similar) nonzero structure are to be multiplied (see, e.g., [Boman et al. 2005; Kalchev et al. 2013]).

One direction for future work is refining our parallel communication cost model to capture latency cost (number of messages) in addition to bandwidth cost (number of words). It is straightforward to derive a lower bound on latency cost by modifying Lem. 4.2 to count the number of adjacent parts instead of the number of adjacent nets. On the other hand, the algorithm in Lem. 4.3 may not attain the latency lower bound, due to sending messages of containing single words.

Another interesting extension to our model is considering parallel machines with shared memory. A promising approach is to generalize our sequential model to have multiple processors, each with their own fast memory. Of course, Hong and Kung's communication lower bound result, invoked in Thm. 4.10, must be replaced by a shared-memory extension, of which several have appeared in the literature.

Another direction for future work is considering additional coarsenings of the fine-grained model. In the case of SpMV, "regular" partitions have been proposed to limit latency costs [Vastenhouw and Bisseling 2005; Boman et al. 2013]. In the case of SpGEMM, SpSUMMA [Buluç and Gilbert 2012] corresponds to a regular partition, a special case of monochrome-C. Coarsening the fine-grained model and restricting to regular partitions yields simpler algorithms that may also exhibit reduced latency cost.

REFERENCES

K. Akbudak and C. Aykanat. 2014. Simultaneous input and output matrix partitioning for outer-product–parallel sparse matrix-matrix multiplication. *SIAM Journal on Scientific Computing* 36, 5 (2014), C568–C590. DOI:http://dx.doi.org/10.1137/13092589X

K. Akbudak, E. Kayaaslan, and C. Aykanat. 2013. Hypergraph partitioning based models and methods for exploiting cache locality in sparse matrix-vector multiplication. *SIAM Journal on Scientific Computing* 35, 3 (2013), C237–C262. DOI:http://dx.doi.org/10.1137/100813956

A. Azad, G. Ballard, A. Buluc, J. Demmel, L. Grigori, O. Schwartz, S. Toledo, and S. Williams. 2015a. *Exploiting multiple levels of parallelism in sparse matrix-matrix multiplication*. Technical Report 1510.00844. arXiv. http://arxiv.org/abs/1510.00844

A. Azad, A. Buluç, and J. Gilbert. 2015b. Parallel triangle counting and enumeration using matrix algebra. In *Proceedings of the IPDPSW, Workshop on Graph Algorithm Building Blocks (GABB '15)*. 804 – 811. DOI:http://dx.doi.org/10.1109/IPDPSW.2015.75

G. Ballard, A. Buluç, J. Demmel, L. Grigori, B. Lipshitz, O. Schwartz, and S. Toledo. 2013. Communication optimal parallel multiplication of sparse random matrices. In *Proceedings of the 25th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '13)*. ACM, 222–231. DOI:http://dx.doi.org/10.1145/2486159.2486196

G. Ballard, E. Carson, J. Demmel, M. Hoemmen, N. Knight, and O. Schwartz. 2014. Communication lower bounds and optimal algorithms for numerical linear algebra. *Acta Numerica* 23 (May 2014), 1–155. DOI:http://dx.doi.org/10.1017/S0962492914000038

G. Ballard, J. Demmel, O. Holtz, B. Lipshitz, and O. Schwartz. 2012. Brief announcement: strong scaling of matrix multiplication algorithms and memory-independent communication lower bounds. In *Proceedings of the 24th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '12)*. ACM, New York, NY, USA, 77–79. DOI:http://dx.doi.org/10.1145/2312005.2312021

G. Ballard, J. Demmel, O. Holtz, and O. Schwartz. 2011. Minimizing communication in numerical linear algebra. *SIAM Journal on Matrix Analysis and Applications* 32, 3 (September 2011), 866–901. DOI:http://dx.doi.org/10.1137/090769156

G. Ballard, A. Druinsky, N. Knight, and O. Schwartz. 2015a. Brief announcement: hypergraph partitioning for parallel sparse matrix-matrix multiplication. In *Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures (SPAA '15)*. ACM, New York, NY, USA, 86–88. DOI:http://dx.doi.org/10.1145/2755573.2755613

G. Ballard, C. Siefert, and J. Hu. 2015b. *Reducing communication costs for sparse matrix multiplication within algebraic multigrid*. Technical Report SAND2015-3275. Sandia National Laboratories. http://prod.sandia.gov/techlib/access-control.cgi/2015/153275.pdf

A. Barker, D. Kalchev, I. Mishev, P. Vassilevski, and Y. Yang. 2015. Accurate coarse-scale AMG-based finite volume reservoir simulations in highly heterogeneous media. In *SPE Reservoir Simulation Symposium*. SPE–173277–MS. DOI:http://dx.doi.org/10.2118/173277-MS

E. Boman, K. Devine, L. Fisk, R. Heaphy, B. Hendrickson, C Vaughan, Ü. Çatalyürek, D. Bozdag, W. Mitchell, and J. Teresco. 2007. *Zoltan 3.0: Parallel Partitioning, Load Balancing, and Data-Management Services; User's Guide*. Technical Report SAND2007-4748W. Sandia National Laboratories. http://www.cs.sandia.gov/Zoltan/ug_html/ug.html

E. Boman, K. Devine, and S. Rajamanickam. 2013. Scalable matrix computations on large scale-free graphs using 2D graph partitioning. In *Proceedings of the International Conference on High Performance Com-*

puting, Networking, Storage and Analysis (SC '13). ACM, New York, NY, USA, Article 50, 12 pages. DOI:http://dx.doi.org/10.1145/2503210.2503293

E. Boman, O. Parekh, and C. Phillips. 2005. *LDRD final report on massively-parallel linear programming: the parPCx system*. Technical Report SAND2004-6440. Sandia National Laboratories. http://prod.sandia.gov/techlib/access-control.cgi/2004/046440.pdf

U. Borštnik, J. VandeVondele, V. Weber, and J. Hutter. 2014. Sparse matrix multiplication: the distributed block-compressed sparse row library. *Parallel Computing* 40, 5–6 (2014), 47–58. DOI:http://dx.doi.org/10.1016/j.parco.2014.03.012

M. Brezina and P. Vassilevski. 2011. Smoothed aggregation spectral element agglomeration AMG: SA-ρAMGe. In *Proceedings of the 8th International Conference on Large-Scale Scientific Computing (LSSC) (Lecture Notes in Computer Science)*, I. Lirkov, S. Margenov, and J. Waśniewski (Eds.), Vol. 7116. Springer, 3–15. DOI:http://dx.doi.org/10.1007/978-3-642-29843-1_1

A. Buluc and J. Gilbert. 2008. On the representation and multiplication of hypersparse matrices. In *IEEE International Symposium on Parallel and Distributed Processing (IPDPS '08)*. 1–11. DOI:http://dx.doi.org/10.1109/IPDPS.2008.4536313

A. Buluç and J. Gilbert. 2012. Parallel sparse matrix-matrix multiplication and indexing: implementation and experiments. *SIAM Journal on Scientific Computing* 34, 4 (2012), C170–C191. DOI:http://dx.doi.org/10.1137/110848244

Ü. Çatalyürek and C. Aykanat. 1999. *PaToH: Partitioning Tool for Hypergraphs*. Technical Report. (Revised March 2011). http://bmi.osu.edu/umit/PaToH/manual.pdf

Ü. Çatalyürek and C. Aykanat. 2001. A fine-grain hypergraph model for 2D decomposition of sparse matrices. In *Proceedings of the 15th International Parallel and Distributed Processing Symposium (IPDPS '01)*. 118–123. DOI:http://dx.doi.org/10.1109/IPDPS.2001.925093

Ü. Çatalyürek, C. Aykanat, and B. Uçar. 2010. On two-dimensional sparse matrix partitioning: models, methods, and a recipe. *SIAM Journal on Scientific Computing* 32, 2 (2010), 656–683. DOI:http://dx.doi.org/10.1137/080737770

U. Çatalyürek and C. Aykanat. 1999. Hypergraph-partitioning-based decomposition for parallel sparse-matrix vector multiplication. *IEEE Transactions on Parallel and Distributed Systems* 10, 7 (Jul 1999), 673–693. DOI:http://dx.doi.org/10.1109/71.780863

M. Christensen, U. Villa, and P. Vassilevski. 2015. Multilevel techniques lead to accurate numerical upscaling and scalable robust solvers for reservoir simulation. In *SPE Reservoir Simulation Symposium*. SPE–173257–MS. DOI:http://dx.doi.org/10.2118/173257-MS

M. Christie and M. Blunt. 2001. Tenth SPE comparative solution project: a comparison of upscaling techniques. *SPE Reservoir Evaluation & Engineering* 4, 04 (August 2001), 308–317. DOI:http://dx.doi.org/10.2118/72469-PA

T. Davis. 2006. *Direct Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics. DOI:http://dx.doi.org/10.1137/1.9780898718881

T. Davis and Y. Hu. 2011. The University of Florida sparse matrix collection. *ACM Transactions on Mathematical Software* 38 (2011), 1:1–1:25.

S. Faisal, S. Parthasarathy, and P. Sadayappan. 2014. Global graphs: a middleware for large scale graph processing. In *2014 IEEE International Conference on Big Data (Big Data)*. 33–40. DOI:http://dx.doi.org/10.1109/BigData.2014.7004369

G. Greiner. 2012. *Sparse Matrix Computations and Their I/O Complexity*. Dissertation. Technische Universität München, München. http://mediatum.ub.tum.de?id=1113167

F. Gustavson. 1978. Two fast algorithms for sparse matrices: multiplication and permuted transposition. *ACM Transactions on Mathematical Software* 4, 3 (Sept. 1978), 250–269. DOI:http://dx.doi.org/10.1145/355791.355796

J. Hong and H. Kung. 1981. I/O complexity: the red-blue pebble game. In *Proceedings of the Thirteenth Annual ACM Symposium on Theory of Computing (STOC '81)*. ACM, 326–333. DOI:http://dx.doi.org/10.1145/800076.802486

D. Irony, S. Toledo, and A. Tiskin. 2004. Communication lower bounds for distributed-memory matrix multiplication. *Journal of Parallel and Distributed Computing* 64, 9 (2004), 1017–1026. DOI:http://dx.doi.org/10.1016/j.jpdc.2004.03.021

D. Kalchev. 2012. *Adaptive Algebraic Multigrid for Finite Element Elliptic Equations with Random Coefficients*. Master's thesis. Sofia University. https://e-reports-ext.llnl.gov/pdf/594392.pdf

D. Kalchev, C. Ketelsen, and P. Vassilevski. 2013. Two-level adaptive algebraic multigrid for a sequence of problems with slowly varying random coefficients. *SIAM Journal on Scientific Computing* 35, 6 (2013), B1215–B1234. DOI:http://dx.doi.org/10.1137/120895366

S. Krishnamoorthy, Ü. Çatalyürek, J. Nieplocha, A. Rountev, and P. Sadayappan. 2006. Hypergraph partitioning for automatic memory hierarchy management. In *Proceedings of the ACM/IEEE SC 2006 Conference (SC '06)*. 34–46. DOI:http://dx.doi.org/10.1109/SC.2006.36

Q. Niu, P.-W. Lai, S. Faisal, S. Parthasarathy, and P. Sadayappan. 2014. A fast implementation of MLR-MCL algorithm on multi-core processors. In *21st International Conference on High Performance Computing (HiPC '14)*. 1–10. DOI:http://dx.doi.org/10.1109/HiPC.2014.7116888

R. Pagh and M. Stöckel. 2014. The input/output somplexity of sparse matrix multiplication. In *Algorithms - ESA 2014*, A. Schulz and D. Wagner (Eds.). Lecture Notes in Computer Science, Vol. 8737. Springer Berlin Heidelberg, 750–761. DOI:http://dx.doi.org/10.1007/978-3-662-44777-2_62

J. Park, M. Smelyanskiy, U. Meier Yang, D. Mudigere, and P. Dubey. 2015. High-performance algebraic multigrid solver optimized for multi-core based distributed parallel systems. In *Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis (SC '15)*. Article 54, 12 pages. DOI:http://dx.doi.org/10.1145/2807591.2807603

M. Rabin and V. Vazirani. 1989. Maximum matchings in general graphs through randomization. *Journal of Algorithms* 10, 4 (1989), 557 – 567. DOI:http://dx.doi.org/10.1016/0196-6774(89)90005-9

V. Satuluri and S. Parthasarathy. 2009. Scalable graph clustering using stochastic flows: applications to community discovery. In *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD '09)*. ACM, New York, NY, USA, 737–746. DOI:http://dx.doi.org/10.1145/1557019.1557101

S. van Dongen. 2000. *Graph Clustering by Flow Simulation*. Ph.D. Dissertation. University of Utrecht. http://www.library.uu.nl/digiarchief/dip/diss/1895620/full.pdf

B. Vastenhouw and R. Bisseling. 2005. A two-dimensional data distribution method for parallel sparse matrix-vector multiplication. *SIAM Review* 47, 1 (2005), 67–95. DOI:http://dx.doi.org/10.1137/S0036144502409019

I. Yamazaki and X. Li. 2011. On techniques to improve robustness and scalability of a parallel hybrid linear solver. In *9th International Conference on High Performance Computing for Computational Science (VECPAR '10)*, José M. Laginha M. Palma, Michel Daydé, Osni Marques, and João Correia Lopes (Eds.). Springer Berlin Heidelberg, 421–434. DOI:http://dx.doi.org/10.1007/978-3-642-19328-6_38