

Copernicus: Characterizing the Performance Implications of Compression Formats Used in Sparse Workloads

Bahar Asgari
Georgia Institute of Technology
bahar.asgari@gatech.edu

Ramyad Hadidi
Georgia Institute of Technology
rhadidi@gatech.edu

Joshua Dierberger
Georgia Institute of Technology
jdierberger3@gatech.edu

Charlotte Steinichen
Georgia Institute of Technology
csteinichen3@gatech.edu

Hyesoon Kim
Georgia Institute of Technology
hyesoon@cc.gatech.edu

ABSTRACT

Sparse matrices are the key ingredients of several application domains, from scientific computation to machine learning. The primary challenge with sparse matrices has been efficiently storing and transferring data, for which many sparse formats have been proposed to significantly eliminate zero entries. Such formats, essentially designed to optimize memory footprint, may not be as successful in performing faster processing. In other words, although they allow faster data transfer and improve memory bandwidth utilization – the classic challenge of sparse problems – their decompression mechanism can potentially create a *computation* bottleneck. Not only is this challenge not resolved, but also it becomes more serious with the advent of domain-specific architectures (DSAs), as they intend to more aggressively improve performance. The performance implications of using various formats along with DSAs, however, has not been extensively studied by prior work. To fill this gap of knowledge, we characterize the impact of using seven frequently used sparse formats on performance, based on a DSA for sparse matrix-vector multiplication (SpMV), implemented on an FPGA using high-level synthesis (HLS) tools, a growing and popular method for developing DSAs. Seeking a fair comparison, we tailor and optimize the HLS implementation of *decompression* for each format. We thoroughly explore diverse metrics, including decompression overhead, latency, balance ratio, throughput, memory bandwidth utilization, resource utilization, and power consumption, on a variety of real-world and synthetic sparse workloads.

KEYWORDS

sparse formats, compression, decompression, domain-specific architecture, characterization

1 INTRODUCTION & MOTIVATION

Sparse matrices, first distinguished as the main workloads in scientific computation, have become an essential component in many other computation domains such as neural networks, recommendation systems, and graph analytics. Since the primary issue with sparse matrices is storing the enormous amount of non-necessary zero elements, several compression formats have been proposed to efficiently store sparse matrices. While some formats target generality, others are tailored for particular patterns of sparseness (e.g., diagonal matrices) to be more effective in saving them with

a minimum storage overhead. Such optimizations for sparse problems mainly focus only on the *storage* overhead in isolation without involving other essential performance metrics such as latency, throughput, and power efficiency. That said, a slow decompression (relative to data transfer) can even surpass the overhead of processing all zero entries in the original dense matrix format; thus, sparse formats may not necessarily guarantee fast execution. This occurs because common sparse formats are often tailored to *the distribution of data*, not *the underlying mechanism of computation*.

The challenges associated with using sparse formats not only are not resolved with the advent of domain-specific architectures (DSAs) for sparse problems, but also gain more importance. DSAs seem to soon become the main platform of sparse computations by approaching the end of Moore’s law, proven by the tremendous number of recent studies [1–11], to more efficiently accelerate the execution of sparse problems. Prior studies [6, 12–17] have also demonstrated the importance of fast compression/decompression in on-demand applications such as in the inference of neural networks. Regardless of this ongoing research, no study has shed light on the performance implications of using the variety of sparse formats. Particularly, even though prior work has studied the performance implications of software implementations of sparse formats [18–24], the *hardware* implementation of these formats on FPGAs has not been extensively characterized.

To design efficient hardware for processing sparse problems, understanding the impact of the *decompression* mechanism on the performance (measured by various metrics) is crucial. This paper endeavors to address this knowledge gap by performing a thorough characterization on **performance** implications of compression formats **used** in sparse workloads (Copernicus[†]). We study the diverse metrics, including decompression overhead, latency, throughput, memory bandwidth utilization, resource utilization, and power consumption, of seven frequently used sparse formats on a variety of real-world and synthetic sparse workloads using a hardware platform implemented on a Xilinx FPGA using Vivado high-level synthesis (HLS), a growing and popular method for designing DSAs and quickly prototyping them on FPGAs. We stream compressed sparse data from memory and process them. In an ideal (fast and power-efficient) case, we expect processing to be done at the same pace as receiving data. However, the decompression mechanism can disturb the balance. To fairly evaluate and compare the different

[†]Copernicus (/kouˈpɜːrnikəs/) is a binary star system including five planets.




Figure 1: Sparse formats: (a) The original sparse matrix in dense format, (b) CSR including offsets to indicate the number of non-zero entries per row, indices, indicating column indices of non-zero entries, and the values itself. CSC follows the same rule as CSR; (c) BCSR including offsets to indicate the number of non-zero 4×4 blocks per row, indices, indicating the index of the first column of non-zero blocks, and the values, indicating the flattened values of non-zero blocks; (d) COO including a series of (row, column, value) tuples for non-zero values; (e) DOK, which is similar to COO; (f) LIL, which pushes all the non-zero entries to top and saves the row indices; (g) ELL, which is similar to LIL but pushes the non-zero entries to left and also uses a padding; and (h) DIA, which saves the non-zero diagonals by adding the diagonal numbers as a header to each diagonal.

decompression mechanisms, for each format, we tailor and optimize the HLS implementation of the decompression. Our suggested HLS implementations and the related performance implications are also applicable to hardware accelerators that directly perform computations on compressed data because they must also reconstruct the location of each non-zero entry.

In summary, Copernicus makes the following contributions:

- It is the first work that investigates sparse formats from the *hardware point of view* for accelerating sparse problems.
- It thoroughly characterizes diverse metrics for seven frequently used sparse formats, applied on *different partition sizes* of matrices from on-demand applications.
- It provides hints to architects to mindfully choose appropriate sparse formats and, more importantly, indicates which parameters must be tuned in ideal hardware to achieve the desired performance or to optimize for a particular metric.
- It suggests an optimized pipeline architecture for processing sparse matrices and an HLS-based implementation on an FPGA that can be used as a building block in further accelerators for sparse problems.

2 SPARSE FORMATS

Compressed Sparse Row/Column (CSR/CSC): The CSR/CSC sparse format sequentially stores values in row/column order in a values array while similarly storing their column-index/row-index in a indices array. Another array, offsets, stores index pointers or range for constructing rows/columns. To do so, the adjacent pair of this array [start:stop] represents a slice from the two first arrays. Figure 1b shows an example of CSR. For an $n \times n$ matrix, the length of offsets is n (usually $n + 1$, but the first element can store absolute value to reduce the size) and the *maximum*[†] length of values and indices is n^2 .

Block CSR/CSC (BCSR/BCSC): The block(-wise) compressed sparse row/column (BCSR/BCSC) [25] sparse format is similar to CSR/CSC, but arrays are stored based on the same-shaped blocks (sub-matrices) rather than on the original matrix. This allows block-wise formats to better deal with large matrices. Figure 1c illustrates an example of BCSR for block sizes of 4×4 , the block size we choose in all our experiments as well. For an $n \times n$ matrix and $b \times b$ blocks, the length of offsets is n/b and the *maximum* length of values and indices are n^2 and $(n/b)^2$.

Coordinate (COO) & Dictionary of Keys (DOK): The COO sparse format simply stores a series of tuples, including the row index, column index, and value for each of the non-zero entries. For an $n \times n$ matrix, the *maximum* length of tuples is $3n^2$. The DOK format is similar to the COO format except that it stores coordinate-data information as key-value pairs. DOK uses hash tables to store a value with the key of (row index, column index). Figure 1d and e depict an example of COO and DOK, respectively.

List of List or Linked list (LIL): The LIL [26] sparse format stores one list of non-zero elements per row/column. Each element in the lists stores the column/row indices of that row/column, indices, and their value, values. Figure 1f presents an example LIL, which compresses the rows and preserves the columns (this is our assumption for LIL in Copernicus). For an $n \times n$ matrix, the *maximum* length of values and indices is n , with n list in total.

Ellpack (ELL) & Sliced ELL (SELL): In the ELL [27] sparse format, non-zero elements are extracted similarly to those of the LIL format, with their column indices and their values. However, they are stored in column-major format with the addition of explicit zero paddings to hold the data for the longest row. This format is ideal for SIMD units since the widths of all values and indices are the same. A sliced ELL (SELL) sparse format first slices the dense matrix row-wise in chunks, and then applies ELL on each chunk. Hence, it reduces the overhead of zero paddings for larger matrices. Figure 1g shows an example of ELL with a padding width of three. In Copernicus, we set this width to six. For an $n \times n$ matrix, the *maximum* length of values and indices is n (longest possible row). The width in ELL is n , and that in SELL varies based on the pattern of data. Variants of ELL formats such as ELL+COO, Jagged Diagonal Storage (JDS) [28], and SELL-C- σ [29] are also popular. ELL+COO mixes ELL and COO formats to reduce the width of long rows. The JDS format sorts the rows in ELL from longest to shortest (for vector machines). SELL-C- σ is a variant of JDS that only sorts rows within a window of σ .

Diagonal (DIA): The DIA [30] sparse format operates by specifying a diagonal number (0 for the main diagonal, negative/positive for diagonals which start on a lower/higher row/column) followed by the values that fall on the diagonal, diagonals. Figure 1h illustrates an example of DIA. For an $n \times n$ matrix, the *maximum* number of non-zero diagonals is $2n - 1$ and the *maximum* length of a diagonal is $n + 1$ (the additional element contains the diagonal number).

[†]Note that these worst-case scenarios are used for on-chip memory allocation. The storage overhead is still defined by the number of non-zero entries.

3 SPARSE WORKLOADS

Sparse matrices are the main data structure in several application domains, from scientific computations to graph analytics. This section first introduces various instances of sparse matrices and lists our *real-world* workloads. Then, it reviews the common regular structures of sparseness and introduces our *synthetic* workloads to evaluate a wider range of sparse applications. Finally, this section explains the common computation of several sparse problems, on which we build our hardware platform for evaluation.

3.1 Sparse Matrices in Various Domains

The main ingredient of different fields of *scientific computations* is modeling physical phenomena, such as sound, heat, elasticity, fluid dynamics, and quantum mechanics. Such models often use partial differential equations (PDEs). To use digital computers for solving PDEs, they are discretized into a 3D grid. Discretization is used to convert PDEs to a linear system of algebraic equations, $Ax = b$, in which A is the coefficient matrix. Since not all the points in a 3D grid are used in the discretization of a phenomenon, *the coefficient matrix A is sparse*. Sparse matrices also exist in *graph analytics*. A common approach for representing graphs is to use an adjacency matrix, each entry of which represents an edge in the graph. As in many applications, not all the nodes in a graph are connected; the graphs are sparse and so are their equivalent *adjacency matrices* containing several zeros. As a representative for various size sparse matrices in the domain of scientific computation and graph analytics, with different density, we obtained the matrices listed in Table 1 from SuiteSparse [31] matrix collection.

The third group of common sparse problems is the applications of machine learning, including the *inference of neural networks* and the *recommendation systems*. Since after training, close-to-zero values are assigned to several model parameters, a common practice is to prune those values. Pruning results in sparse *weight matrices*. Sparse matrices in neural networks are often not as sparse as extremely sparse matrices in the first two application domains. Besides, in neural networks, sparseness is more random and varies case by case with the algorithm of pruning and the other design factors such as the desired accuracy. The recommendation system models are the other instance of sparse problems. They include an embedding table followed by a neural network. Although the embedding tables are dense, accesses to them are random and sparse [32]. To evaluate the whole range of the variety of randomness in sparse machine learning applications, Copernicus covers the synthetic matrices introduced in the following.

3.2 Synthetic Sparse Matrices

Besides the real-world matrices listed in Table 1, our workloads consist of two groups of synthetic sparse matrices. The first group includes *randomly generated* sparse matrices, the density of which varies from 0.0001 to 0.5. We generate the denser random matrices (i.e., the density of 0.1 to 0.5) as a representation for those in machine learning applications. On the other hand, the more sparse random matrices (i.e., density between 0.0001 to 0.01) represent scientific and graph applications with no particular structure. The second group of our synthetic sparse matrices denotes the common structure in sparse matrices: *diagonal and band matrices*. A band

Table 1: Matrices from SuiteSparse [31] matrix collection.

ID	Name	Dim.(M) ¹	NNZ(M) ²	Kind
2C	2cubes_sphere	0.101	1.647	Electromagnetics Problem
FR	Freescape2	2.9	14.3	Circuit Sim. Matrix
RE	N_reactome	0.016	0.043	Biochemical Network
AM	amazon0601	0.4	3.3	Directed Graph
DW	dwt_918	0.000918	0.0073	Structural Problem
EO	europe_osm	50.9	108	Undirected Graph
FL	flickr	0.82	9.8	Directed Graph
HC	hcircuit	0.1	0.51	Circuit Sim. Problem
HU	hugebubbles	18.3	54.9	Undirected Graph
KR	kron_g500-logn21	2	182	Undirected Multigraph
RL	rail582	0.056	0.4	Linear Prog. Problem
RJ	rajat31	4.6	20.3	Circuit Sim. Problem
RO	roadNet-TX	1.3	3.8	Undirected Graph
RC	road_central	14	33.8	Undirected Graph
LJ	soc-LiveJournal1	4.8	68.9	Directed Graph
TH	thermomech_dK	0.2	2.8	Thermal Problem
WE	wb-edu	9.8	57.1	Directed Graph
WG	web-Google	0.91	5.1	Directed Graph
WT	wiki-Talk	2.3	5	Directed Graph
WI	wikipedia	3.5	45	Directed Graph

¹ Dim.: dimension or the number of columns/rows of a square matrix.

² NNZ: the number of non-zero entries.

matrix is a sparse matrix, the non-zero entries of which are confined to a diagonal band, including the main diagonal and more than one diagonal on each side. The width of a band matrix is the number k such that $a_{i,j} = 0$ if $|i - j| > k/2$. We generate and evaluate band matrices of size 8000 with widths of 2, 4, 16, 32, and 64. Numerical problems in higher dimensions often lead to band matrices (e.g., a PDE on a square domain). A type of band matrices consisting of only the main diagonal (i.e., $k = 1$) is called a diagonal matrix. Diagonal matrices also occur in many fields of linear algebra.

3.3 Common Computation in Sparse Problems

This section shows that sparse matrix-vector multiplication (SpMV) is the key sparse kernel in all of the three aforementioned domains of sparse problems. Designing DSAs for accelerating SpMV has been the focus of several recent studies [7, 9, 10, 17, 33–36]

Scientific Computations: The coefficient sparse matrix A in a linear system of algebraic equations is very large for two or higher dimensional problems (e.g., elliptic, parabolic, or hyperbolic PDEs). While for small A s, solving PDEs is doable by direct methods – those based on matrix inversion – they become less efficient for larger A s. Such systems of linear equations with a large symmetric positive-definite matrix A can be solved by iterative algorithms such as conjugate gradient (CG) methods. In both cases (i.e., direct and iterative methods) for solving PDEs, the key sparse kernel is *SpMV*. Symmetric Gauss-Seidel iteration [37] used in the CG algorithm or the lower-upper (LU) decomposition used for matrix inversion are examples that contain *SpMV*.

Graph Analytics: Graph algorithms, such as breadth-first search, single-source shortest path, and PageRank, which traverse vertices and edges of a graph to compute some properties based on the

graph relationships, can be implemented as a sparse matrix-vector operation. Such an implementation is compatible with the vertex-centric programming model [38], which divides a graph algorithm into two main phases. In the first phase, a row of the adjacency matrix is processed using a vector-vector operation between the row of the matrix, and a property vector varied based on the algorithm type. Then, the output vector from the first phase is reduced by some form of reduction operations (e.g., summation). The combination of the two phases creates a dot-product, and combining the dot-products for processing the whole graph leads to an *SpMV*.

Machine Learning: Machine learning applications consist of SpMV or sparse matrix-matrix multiplication, both of which rely on the same underlying dot-product engine. For instance, in convolution neural networks, the computation of convolutional layers can be viewed as a matrix-matrix multiplication. In other words, convolving a 3D input with a given number of filters can be represented as an equivalent matrix-matrix multiplication that multiplies the 2D flatten weight matrix by the input matrix, which is originally two-dimensional. In recommendation systems, the sparse embedding-table look-ups end up as a reduction operation (e.g., summation) that can also be implemented using a dot-product engine, initially designed for *SpMV*.

4 EXPERIMENTAL SETUP

4.1 Platform and System Overview

We describe our hardware platform in C++ and use this implementation for generating RTL to be executed on an FPGA. To generate RTL (in Verilog), we use the high-level synthesis (HLS) tool of Vivado and related #pragmas as hints to detail the architecture. After simulations in Vivado_HLS, we use Vivado to synthesise and implement our hardware on an xq7z020 FPGA from Zynq-7000 family connected to a DDR3 memory. The clock frequency is set to 250 MHz. We verify the functionality of the RTL using synthetic data as a C++ testbench for C/RTL co-simulation. The inputs to our hardware platform are sparse workloads (SuitSparse and synthetic sparse matrices). We use Matlab to preprocess the workloads and compress them in target formats, studied in Copernicus.

Although all sparse formats eliminate transferring a large portion of zero entries, applying them directly on the large original matrices is not efficient. For instance, formats such as CSR transfer one index for all rows, even for all-zero rows. Therefore, their overhead grows with the size (i.e., dimension) of the matrix. Therefore, as compressing and transferring large units of data (e.g., the entire original matrix) is not beneficial for some of the formats, a common efficient practice is to apply the compression on the smaller partitions of the original matrix, which also allows coarse-grained parallel processing of partitions. Additionally, since some sparse matrices in real-world applications capture a certain level of spatial locality, by using partitioning, we can eliminate transferring and processing the all-zero partitions. Following this common practice, Copernicus also applies all the compression formats *only on the non-zero partitions* of large matrices. The size of these partitions is one of our hyperparameters.

4.2 Metrics & Hyperparameters

Copernicus evaluates the performance implications of sparse formats using the metrics introduced in the following. First, we define σ as a metric to measure the latency **overhead of decompression**:

$$\sigma = \frac{T_{decomp} + nnz_rows \times T_{dot}}{p \times T_{dot}} \quad (1)$$

in which T_{decomp} indicates decompression latency, which consists of latency for BRAM accesses and logic, nnz_rows are non-zero rows for which we must perform a dot-product, each taking T_{dot} , and p is the partition size. As a result, for the dense format, $\sigma = 1$. Besides σ , we measure the breakdowns of latency: (i) **memory latency**, the time to transfer a compressed partition (data and meta-data) to FPGA and buffer it in the BRAM; and (ii) **computation latency** consisting of decompression, dot-product, and necessary BRAM accesses. Furthermore, seeking an appropriate sparse format for achieving balanced streaming, we proposed evaluating a **balance ratio**, which we define as the average ratio of memory latency to compute latency for all non-zero partitions. The balanced ratio of perfectly balanced streaming would be one. An imbalance streaming leads to idle computation or pauses in data transfer.

As another important factor in the execution of sparse problems, we evaluate **throughput**, defined as bytes processed per second, which reflects the bubbles in the streaming pipeline caused by imbalance streaming (balance ratio $\neq 1$). Besides throughput, we compare the **memory-bandwidth utilization**, the ratio of useful data over all transmitted data (i.e., useful data plus metadata). Our other metrics for the full design-space exploration are **resource utilization** and dynamic/static **power consumption**. By resource utilization, we mean the percentage of FPGA resources used by all components (entire Figure 2). We evaluate the aforementioned metrics while varying the hyperparameters including practical partition sizes of 8, 16, and 32, and the width of 2, 4, 8, 16, 32, and 64 for band matrices to represent realistic matrices, and the density of random matrices from 0.0001 to 0.5.

5 COPERNICUS

5.1 Architecture

This section proposes our hardware architecture for processing sparse matrices (i.e., executing SpMV). To be effective in accelerating any problem, two ingredients are key: parallel computation and fast accesses to memory. Parallelizing the computation of SpMV is relatively straightforward and can be implemented in various ways. For instance, the vector operand can be multiplied with the non-zero rows of the sparse matrix operand in parallel, or the columns of the matrix and the vector operand can be partitioned into a few chunks and be processed concurrently. On top of these coarse-grained approaches, element-wise fine-grained parallelism can always be applied to each of the dot products, the fundamental operations in SpMV. For instance, the intuitive way to perform parallel dot products is by implementing a fixed-size compute logic (i.e., a multiplier array attached to an adder tree) and using it on the chunks of the operands of the original SpMV. We build our hardware platform (shown in Figure 2) based on the building blocks of such a fine-grained-parallel dot product engine (Figure 2, ⑨)




Figure 2: The architecture of our evaluation platform: Streaming the compressed partitions of a sparse matrix from the memory to FPGA through AXI stream interfaces and processing them (i.e., SpMV) in a pipeline. The *decompression* component varies based on sparse format.

and investigate how to effectively integrate it with stream memory accesses as the second fundamental ingredient to achieve high performance. Instances of this architecture can be aggregated for implementing coarse-grain parallelism.

The components of our evaluation architecture to stream the compressed partitions and process them in an FPGA include (i) the global memory, (ii) an AXI streaming (AXIS) interface through which the partitions and the vector operand of SpMV are transferred from the memory to FPGA, and (iii) the high-level three-stage pipeline (Figure 2, ①) implemented in the FPGA that receives the partitions in the memory-read stage, processes them in the compute stage (i.e., SpMV block in Figure 2), and streams the partial output vector back to the memory in the memory-write stage. The input buffer contains a partition compressed in a particular format (e.g., it contains values, offsets, and column indices for CSR). The following explains the details of the compute stage of this high-level pipeline, which itself comprises a two-stage pipeline, including decompress and dot-product stages. The block of SpMV iteratively creates dense non-zero rows in the first stage (Figure 2, ②). Then, the second stage (Figure 2, ③) performs a dot-product between the result of the first stage and the vector operand of SpMV. We implement the dot-product as an array of multipliers connected to a balanced adder tree. Since the output of the SpMV is a vector (not necessarily sparse), we do not include a recompression stage in hardware.

We use this architecture (Figure 2) as the baseline and tailor the decompression stage for each format (Section 5.2) and then compare the outcomes. The *decompression* (Figure 2, ②) is the bottleneck-prone stage and the focus of Copernicus. In our platform, the width of the dot-product engine (Figure 2, ③) is the same as the width of the partitions. The *partition density* and, more specifically the *row density* (i.e., the density of non-zero rows), defines the computation utilization of the dot-product engine at run time. Besides, the *number of non-zero rows* in the partitions determines the utilization of the inner pipeline. Figure 3 illustrates the three mentioned factors for three partition sizes: 8, 16, and 32. These raw statistics can be interpreted along with the evaluation results reported in Section 5.

5.2 Optimized Decompression in HLS

Here, we discuss our tailored HLS implementation of the decompression mechanism (Figure 2, ②) for the seven target sparse formats.

CSR: As Listing 1 shows, since CSR uses three vectors (i.e., offsets, column indices, and values) to represent a sparse matrix, for decompressing a non-zero row, we need to first access the offsets




Figure 3: Density and spatial locality: the average percentage of (a) non-zero values in partitions, (b) non-zero values in non-zero rows, and (c) non-zero rows in partitions for the sparse matrices listed in Table 1.

(Listing 1, line 7), and then, based on numVal, we can read as many column indices and values as required (Listing 1, line 10). As a result, decompression from the CSR format is likely to be compute-bound because of the overhead of one extra access to BRAM. Additionally, to retrieve the column indices and values, *since accesses to the BRAM blocks are sequential*, we do not know in advance which elements of column indices and values are going to be accessed. Thus, we cannot partition and allocate those two vectors across the blocks of BRAM to guarantee parallel accesses. Because of the sequential accesses in a non-zero row, the latency of decompressing a row depends on the number of non-zero elements in that row. By assuming that we stream the offsets and column indices using two streamlines in parallel, the one with more non-zero elements (longer one) defines the latency of memory access. To reduce the negative impact of the accesses to offsets on performance, we pipeline this progress to concurrently create non-zero rows (if more than one).

```

1 function decompressCSR(A, readInx, oldInx)
2   // A in CSR: offsets[OFFSET_LENGTH]
3   //           colInx[COL_INX_LENGTH]
4   //           values[VAL_LENGTH]
5   // readInx: current row index
6   // oldInx: last row index
7   numVal = offsets[readInx] - offsets[readInx-1]
8   for i=0 to numVal:
9     #pragma HLS pipeline
10    drow[colInx[oldInx+i]] = values[oldInx+i]
11 return drow

```

Listing 1: CSR-decompression HLS pseudo code

BCSR: As Listing 2 shows, the decompression of BCSR is similar to that of CSR, whereas instead of individual non-zero elements, the non-zero blocks are processed. To initiate the accesses to column indices and values, one extra access to the offsets is required per each row of blocks (Listing 2, line 9). The advantage of BCSR over CSR is that we can distribute the values and column indices over BRAM blocks and access their elements in parallel, for which, as

lines 1 and 2 in Listing 2 show, we completely partition the values and colInx across their second dimension before calling the decompression function. This allows us to unroll the for loop (line 12), the iterations of which access different BRAM blocks in parallel. The downsides of BCSR, however, are (i) the overhead of transferring zero elements in the non-zero blocks and (ii) processing all the rows in the non-zero blocks whether they (the rows) are all zero or not. The latency of decompressing the blocks in a row depends on the number of non-zero *blocks* in that row. On the other hand, since the values has the longest length (compared to offsets and colInx), transferring values defines the memory latency.

```

1 #pragma HLS array_partition variable=values dim=2
2 #pragma HLS array_partition variable=colInx dim=2
3 function decompressBCSR(A, readInx, oldInx)
4 // A in BCSR: offsets[OFFSET_LENGTH]
5 // colInx[COL_INX_LENGTH]
6 // values[VAL_LENGTH][VAL_WIDTH]
7 // readInx: current row index
8 // oldInx: last row index
9 numBlocks = offsets[readInx] - offsets[readInx-1]
10 for i=0 to numBlocks:
11     for j=0 to VAL_WIDTH:
12         #pragma HLS unroll
13         drows[j / BLOCK_LENGTH][colInx[oldInx + i]
14         + j mod BLOCK_LENGTH] = values[oldInx + i][j]
15 return drows

```

Listing 2: BCSR-decompression HLS pseudo code

CSC: Listing 3 shows the pseudo code of CSC decompression, in which the *columns* are compressed. On the contrary, the hardware requires *rows* of the matrix for performing SpMV. Because of this mismatch, the decompression mechanism must iteratively traverse all the columns of the matrix to find the values corresponding to the current row (Listing 3, line 12). Although this mismatch makes the decompression inefficient, we still include this extreme case in our evaluation to explore how much performance is hurt if the format and the hardware are not aligned.

```

1 function decompressCSC(A, readInx)
2 // A in CSC: offsets[OFFSET_LENGTH]
3 // rowInx[ROW_INX_LENGTH]
4 // values[VAL_LENGTH]
5 // readInx: current row index
6 numVal = offsets[colInx] - offsets[Inx-1]
7 for i=0 to CSC_ROW_INX_LENGTH
8     && colInx < CSC_OFFSETS_LENGTH:
9         startInx = i
10        #pragma HLS pipeline
11        while i < startInx + numVal:
12            if rowInx[i] == read_inx:
13                drow[colInx-1] = values[i]
14                break
15            i++
16            i = offsets[colInx++]
17 return drow

```

Listing 3: CSC-decompression HLS pseudo code

LIL: The LIL decompression (Listing 4) avoids extra accesses to BRAM and enables *deterministic* parallel accesses to the BRAM blocks for decompressing the non-zero rows of a sparse matrix. Since the columns of values and indices can always be accessed in parallel, we partition both of them (Listing 4, lines 1 and 2) before calling decompression. As a result, no extra read access is required

to determine the number of next read accesses. Thus, the latency of processing a matrix depends on the number of non-zero rows. In other words, creating a non-zero row consists of the latency of one BRAM access (since the accesses are parallel) plus the latency for creating the input of the dot product, which is done by a simpler logic compared to those of CSR or BCSR. To recognize the end of the non-zero rows, one additional BRAM access is required. Memory latency for LIL is defined by the number of non-zero rows, the size of rows, and transferring one additional row for indicating the end of non-zero rows.

```

1 #pragma HLS array_partition variable=values dim=2
2 #pragma HLS array_partition variable=Inx dim=2
3 function decompressLIL(A, readInx[], oldInx)
4 // A in LIL: values[HEIGHT][WIDTH]
5 // Inx[HEIGHT][WIDTH]
6 // readInx: current row index
7 // oldInx: last row index
8 minInx = inf
9 for i=0 to WIDTH:
10    #pragma HLS pipeline
11    if readInx[i]<HEIGHT && Inx[readInx[i]][i]<minInx:
12        minInx = Inx[readInx[i]][i]
13    for i=0 to WIDTH:
14        #pragma HLS unroll
15        if Inx[readInx[i]][i] == minInx:
16            drow[i] = values[readInx[i]][i]
17            readInx[i]++
18 return drow

```

Listing 4: LIL-decompression HLS pseudo code

ELL: Similar to LIL, the representation of matrix A in ELL (Listing 5) includes values and indices that can be accessed in parallel and thus are partitioned and distributed over BRAM blocks (Listing 5, lines 1 and 2), which subsequently allows unrolling the for loop for parallel processing (line 7). The difference between LIL and ELL, however, is the direction of compression. Although the direction of compression in ELL enables a simple assignment shown in line 8, it prevents skipping the all-zero rows that in turn can cause a performance drop. Since we completely unroll the for loop (line 7), reducing ELL_MAX_COMP_ROW_LENGTH in the ELL implementation and using optimizations such as ELL-COO only impact the resource utilization of FPGA, not the performance.

```

1 #pragma HLS array_partition variable=values dim=2
2 #pragma HLS array_partition variable=indices dim=2
3 function decompressELL(A)
4 // A in ELL: values[ELL_MAX_COMP_ROW_LENGTH]
5 // indices[ELL_MAX_COMP_ROW_LENGTH]
6 for i=0 to ELL_MAX_COMP_ROW_LENGTH:
7     #pragma HLS unroll
8     drow[indices[i]] = values[i]
9 return drow

```

Listing 5: ELL-decompression HLS pseudo code

COO: As Listing 6 shows, since COO saves tuples for representing a matrix, its decompression mechanism is pretty straightforward, including a simple assignment, shown in Listing 6, line 7. The downside of COO, however, is that we do not know in advance how many elements exist in each row. Thus, we cannot partition and allocate the vector of tuples across the blocks of BRAM to guarantee parallel accesses. For the same reason, we pipeline the for loop (line 5) rather than unrolling it. The same procedure is also applicable to DOK.




Figure 4: Decompression overhead for SuiteSparse: comparing the latency overhead, σ (lower is better), of seven sparse formats for partition size of 16×16 . A darker color indicates less sparsity (i.e., higher density).

```

1 function decompressCOO(A, row)
2   // A in COO: tuples[COO_NUM_TUPLES][3]
3   // row: the current row
4   for i=0 to COO_NUM_TUPLES:
5     #pragma HLS pipeline
6     if (tuples[i][0]!=inf && tuples[i][0]==row):
7       drow[tuples[i][1]] = tuples[i][2];
8   return drow

```

Listing 6: COO-decompression HLS pseudo code

DIA: Line 7 of Listing 7 shows the pseudo code for the decompression mechanism of DIA, the most domain-specific format studied in Copernicus. DIA saves matrix A as diags, a two-dimensional matrix including all non-zero diagonals of matrix A. The first element of each row of diags indicates the diagonal number. To decompress the rows of matrix A, the decompression function traverses all rows of diags to find the elements corresponding to the current row. To this end, we use two helper functions, DiaInxForRow (line 1) and IsRowOnDiagonal (line 4). As the decompression mechanism suggests, although in terms of memory footprint DIA should be beneficial for diagonal matrices, its decompression mechanism is not quite compatible with even a simple computation such as a fine-grained parallel SpMV, which subsequently requires rows of the matrix. Such an overhead worsens when non-zero elements are scattered over multiple diagonals but do not completely fill them.

```

1 function DiaInxForRow(row, d)
2   return (row + d < row) ? row + d : row
3
4 function IsRowOnDiagonal(row, d)
5   return d <= WIDTH - 1 - row && d >= -row
6
7 function decompressDIA(A, row)
8   // A in DIA:
9   // diags[NUM_DIAGONALS][MAX_DIAGONAL_LEN]
10  // row: the current row
11  for i=0 to NUM_DIAGONALS:
12    #pragma HLS pipeline
13    d = diags[i][0]
14    if (!IsRowOnDiagonal(row, d)) continue
15    // column index is the row + d
16    // + 1 since 0 element is the diagonal number
17    drow[row+d] = diags[i][DiaInxForRow(row,d)+1]
18  return drow

```

Listing 7: DIA-decompression HLS pseudo code

6 CHARACTERIZATION

6.1 The Overhead of Decompression

This section explores σ , the latency overhead of decomposition (Figure 2, ①) for SuiteSparse, random, and structured band matrices in Figures 4, 5, and 6, respectively. In Figures 4, the bars lower than




Figure 5: Decompression overhead for random matrices: comparing the σ (lower is better) of seven sparse formats for 16×16 partitions when density varies from 0.0001 to 0.5.

one illustrate faster computation than the baseline dense format. The overhead of the dense baseline, for which $\sigma = 1$, is computing and transferring zero entries, and the overhead of all sparse formats (only for computations) goes to the decompression. As Figures 4 shows, the overhead of sparse formats can, in some cases, exceed that of the dense format. The worst-case scenario of decompression occurs with the CSC format because the orientation of data is opposite to that of the computation mechanism in hardware.

From Figures 4, we do not observe any relationship between the density (darkness of the bars in Figures 4) and σ in highly sparse matrices. Thus, Figures 5 clarifies such a relationship for a wider range of density based on our randomly generated synthetic workloads. Likewise, Figure 6 shows the latency of band (and diagonal) matrices when the width increases. As the two figures illustrate, although the σ of all formats increase with density and width of band matrices, it more dramatically increases for COO, CSR, and CSC. Besides, the time to reconstruct the rows of a matrix from a column-oriented compression format (i.e., CSC) leads to up to




Figure 6: Decompression overhead for band matrices: comparing the latency overhead, σ , of seven sparse formats for partition size of 16×16 when the width varies from 1 to 64.




Figure 7: Decompression overhead for various partition sizes: comparing the average σ (lower is better) of seven sparse formats for three types of workloads (SuiteSparse, random, band) and partition sizes of 8, 16, and 32.

21 \times and 30 \times slower computation than if we were to process all zero entries of the dense format, respectively, for random and band matrices. In such cases, preprocessing the sparse data to a format compatible with a hardware accelerator is highly suggested.

Figures 7 illustrates the impact of partition size on σ . In all workloads (i.e., SuiteSparse, random, and structured), the computation latency of ELL is proportional to that of the dense format and does not change with the pattern of sparsity. This is because, in ELL, we are still processing a whole non-zero matrix regardless of its individual entries. However, since the length of these new squares (in our case, six) is smaller than that of the original dense partitions (8, 16, or 32), the computation latency of ELL decreases as the partition size increases. Seeking a relatively generic sparse format that can provide moderate computation latency for random and structured matrices, BCSR could be a fair option. However, it is not as good for random matrices when the partition size increases. This is because of the additional dot products that must be done per each non-zero block regardless of the individual values of the entries.

6.2 Latency & Balance Ratio

Since memory accesses and computation are pipelined, the sum of their *maximum for each partition* defines the total latency. Therefore, the latency overhead (discussed in Section 6.1), which stems from only the computation, does not provide any information about which one (computation or memory) defines the total latency. Details in that regard are discussed in this section. Figure 8 shows both memory and compute latency and thus implicitly shows the balance ratio: points below the *balance* line have balance ratio < 1 .

Since only non-zero entries of non-zero partitions are transmitted, the latency to transmit data and metadata (i.e., memory latency) for all sparse formats is much lower than that for the dense format, as expected. The computation latency of sparse formats, however, is not always lower than for the dense format. While some of the

frequently used formats such as CSR, CSC, and DIA have been successful in lowering the memory latency, their computation latency is higher than the baseline, which negates their benefits. Other formats such as LIL and ELL also do well in reducing memory latency (i.e., data-transfer time) but have similar computations as the baseline. For instance, the computation latency of LIL is defined by the longest column; hence, in some cases, it is the same as or more than that of dense format. For ELL, on the other hand, when the width of the ELL matrix is slightly smaller than the width of the original partition (e.g., the 8 \times 8 case), the computation latency of ELL is just slightly higher than dense format because of the overhead of decomposition, even though it is small. Similar to random and structured matrices, the CSC format is the slowest for SuiteSparse workloads with up to 27 \times higher latency compared to the dense baseline. All in all, Figure 8 suggests that in terms of latency, COO or BCSR could be appropriate candidates (even though in some cases they perform as good as the dense format) to be used for diverse matrices from scientific and graph applications.

As Figure 8 shows, for all types of matrices (i.e., SuiteSparse, random, and band), the balance ratio of dense format is higher than most of the sparse formats. This is because the zero entries impact both memory and computation latency. In fact, the balance ratio of dense format is closer to one (the perfect case) – but it moves toward a memory-bound as partition size (indicated by marker size) increases. In some formats, such as LIL, increasing the partition size helps to achieve a better balance, while in some others, such as ELL, it is the opposite. In random and structured band sparse matrices (Figures 8b and 8c), higher density and/or larger bandwidth of band matrices leads to a memory bottleneck for BCSR, LIL, and DIA. In such cases, if adding more memory bandwidth to the system is possible, using BCSR or LIL for less sparse application (e.g., for the inference of neural networks) or using DIA for applications with diagonal/band matrices is suggested. Otherwise, COO seems to offer a reasonable balance for various densities as well as the varieties of band matrices. The same hypotheses for balance ratio are also applicable for the more diverse SuiteSparse workloads, as shown in Figure 8a.

6.3 Throughput & Bandwidth Utilization

This section studies throughput and memory bandwidth utilization. First, Figure 9 explores the relationship between throughput and the total time to process an 8000 \times 8000 matrix. The following parameters contribute to throughput: (i) the total processed data consisting of data and metadata and (ii) the total time to process, which is the maximum of memory latency (data-transfer time) and




Figure 8: Balance ratio: the relationship between the memory and compute latency for various partition sizes indicated by the size of markers for (a) SuiteSparse, (b) random workloads, and (c) band matrices. The blue line indicates balance ratio = 1.




Figure 9: Throughput vs. latency: comparing the throughput and the average time to apply SpMV on an 8000×8000 : (a) CSR, (b) BCSR, (c) COO, (d) LIL, (e) ELL, (f) DIA, and (g) CSC. Thicker lines indicate larger partition sizes.




Figure 10: Memory bandwidth utilization for random matrices: comparing seven sparse formats for partition size of 16×16 when density varies from 0.0001 to 0.5.

computation latency for each partition. As a result, in a sparse format such as ELL in which both total latency and data grow with the same pace, throughput does not change with latency (this is also the case for the dense baseline). For all formats but ELL, throughput increases with latency and then reaches a maximum. As Figure 9 suggests, BCSR, LIL, and DIA reach a higher throughput compared to the other four formats. Besides, for all formats but CSC, increasing partition size, shown by the thickness of lines in Figure 9, results in higher throughput because both latency and data decrease as the partition size increases. Since throughput does not reflect the impact of transmitting and computing *useful* data, we study throughput along with the utilization of memory bandwidth.

Memory bandwidth utilization and its relationship with density, width of band matrices, and partition size are shown in Figures 10, 11, and 12, respectively. As they all indicate, the memory bandwidth utilization of COO is always 0.3 since it always transmits two indices per one non-zero entry. Besides, as Figure 11 indicates, the memory bandwidth utilization of DIA for diagonal matrices is close to one – the slight difference occurs because of saving the diagonal number




Figure 11: Memory bandwidth utilization for band matrices: comparing seven sparse formats for partition size of 16×16 when the width varies from 1 to 64.




Figure 12: Memory bandwidth utilization for various partition sizes: comparing the average memory bandwidth utilization (higher is better) on SuiteSparse, random, and band workloads for partition sizes of 8, 16, and 32.

for the main diagonal. As partition size grows, this memory bandwidth utilization approaches full utilization. However, for other band matrices, we see that the DIA format does not offer better memory bandwidth compared to *more generic* formats such as COO, ELL, or LIL, among which LIL is a better candidate to cover more extreme sparseness as well as a wider variety of random matrices (Figure 10) while offering a better balance ratio at larger partitions compared to COO and ELL. Finally, Figure 12 demonstrates that, as expected, for all formats but COO, the memory bandwidth utilization of denser matrices (density > 0.1) and structured ones is higher than that of extremely sparse matrices (e.g., SuiteSparse).

6.4 Resource Utilization & Power Consumption

Table 2 compares FPGA resource utilization and the dynamic power consumption. We must dedicate enough BRAM blocks to envision the worst-case scenarios even though they occur rarely. The other factor impacting the BRAM utilization is the degree of parallelism. To enable parallelism, we partition the matrices and distribute them to BRAM blocks. Because of these two factors, we see that CSR and CSC utilized the lowest number of BRAM blocks, whereas BCSR utilizes the same blocks as the dense implementation does. In some cases, such as ELL, smaller partitions (i.e., 8 and 16) use more flip flops (FFs) compared to larger partitions (i.e., 32). This is because, in a small partition size, the buffering is automatically implemented using FFs rather than BRAM blocks. It is also demonstrated by the fewer BRAM blocks utilized by the 8×8 ELL.

The dynamic power consumption listed in Table 2 suggests that while larger partition sizes cause higher power consumption in some formats (i.e., CSR, BCSR, COO, and LIL), for the others (i.e., dense, CSC, ELL, and DIA), the maximum power is consumed at the 16×16 partition size and the minimum case may occur at 8×8 (e.g., dense and CSC) or at 32×32 (e.g., ELL and DIA). To clarify, Figure 13a, 13b, and 13c illustrate the dynamic power consumed by logic, BRAM, and signals, respectively. As Figure 13 shows, the power consumption of logic always increases or stays steady as partition size increases, while that of BRAM may decrease (e.g., dense


Table 2: Resource utilization and the total dynamic power consumption for three partition sizes (8, 16, and 32).

part. size:	BRAM_18K			FF ($\times 1000$)			LUT ($\times 1000$)			DY Power(W)		
	8	16	32	8	16	32	8	16	32	8	16	32
DENSE	8	16	32	1.5	1.9	4.3	0.7	0.7	1.2	0.02	0.08	0.03
CSR	2	2	8	0.7	0.8	3.8	0.9	0.9	1.1	0.04	0.04	0.07
BCSR	8	16	32	1.6	2.4	4.4	1.2	1.4	2.2	0.05	0.06	0.06
CSC	1	1	9	0.9	1	2.7	1	1.2	1.1	0.01	0.05	0.03
LIL	4	4	6	2.9	5.8	9.1	1.6	2.7	4.8	0.05	0.08	0.07
ELL	1	7	9	2	3.2	0.9	0.9	1	0.8	0.06	0.10	0.06
COO	3	3	8	1.8	1.3	3.2	1.2	2.5	5.4	0.02	0.04	0.04
DIA	3	3	11	2.2	5	9.2	1.5	2.8	4.6	0.07	0.12	0.05
Total	140			106.4			53.2			N/A		

and BCSR). Therefore, comparing Figure 13 against the dynamic power listed in Table 2 indicates that the trend of overall dynamic power consumption partially depends on BRAM, but more generally follows the same trend as the power consumption of signals (Figure 13c). By evaluating total latency and power consumption together, we see that for SuiteSparse matrices, not only does COO consume the least dynamic power, but also it is the fastest in terms of total latency. However, if achieving high throughput at lower power is the goal, BCSR is a better fit. On the other hand, for structured matrices, LIL and ELL are the fastest in terms of latency and throughput, among which ELL performs better for band matrices with wider bandwidths and consumes less power. The static power consumption of dense, CSR, BCSR, LIL, and ELL is 0.121W and that of CSC, COO, and DIA is 0.103W. The static energy, which depends on time, can be an issue for those slower sparse formats that require less amount of dynamic energy.

7 RELATED WORK

Proposing DSAs for SpMV in various application domains have been the topic of several recent studies [7, 9, 10, 17, 33–36], in which sparse formats play a key role. While SIGMA [33], Two-Step [36], and ExTensor [9] use popular sparse formats such as CSR, CSC, and COO, others tailor a sparse format to their DSA [7, 34] or propose new formats [10, 17, 35] as follows: SpArch [34] proposes a condensed format, which is equivalent to storing a matrix in the CSR format and fetching the elements with the same index for all rows; Alrescha [7] modifies BCSR so that data follow a desired order dictated by DSA; Tensaurus [35] proposes compressed interleaved sparse slice (CISS) format, which allows accessing sparse data in a vectorized and streaming manner; SparTen [17] suggests a bitmask representation, in which a sparse tensor is a two tuple of a bit mask, called SparseMap, and a set of non-zero values; and SMASH [10] presents a hierarchical bitmap compression that uses a non-zero values array (NZA) for holding the values.

**Figure 13: Dynamic power consumption: breakdown into (a) logic, (b) BRAM, and (c) signals.**

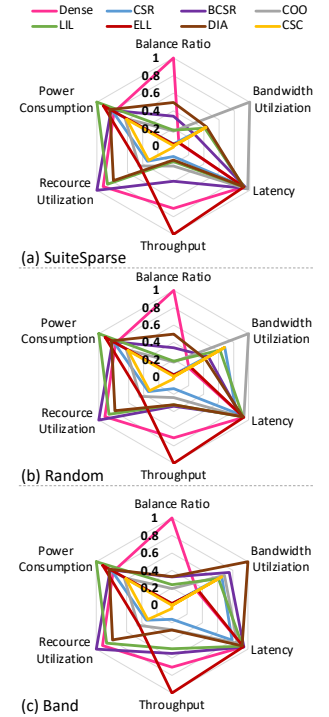
8 INSIGHTS & CONCLUSIONS

Copernicus investigated the performance implication of HLS-based FPGA implementations of seven sparse formats. We hope our results lead architects to knowingly choose the required sparse format and tailor their FPGA designs for sparse applications. Figure 14 summarizes all the six metrics for three group of workloads by normalizing each metric to its maximum achieved number so that "1" represents the best case and "0" represents the worst case. Here, we overview some of the key insights:

- Unlike a common belief, the memory bandwidth is not always the bottleneck; hence the performance sparse problems cannot always be improved by simply adding more memory bandwidth to the system. Thus, when using a format such as CSR to efficiently use storage, a lower-bandwidth low-cost memory is sufficient. Otherwise, the implementation of the computations must be further improved (if possible).

- Although in scientific computation and graph analytics (Figure 14a), the common patterns of sparse matrices are diagonal and band, our study shows that a non-specialized format such as COO performs faster and better utilizes the memory bandwidth compared to a specialized format such as DIA. This is because of the compatibility of more generic formats with a generic hardware for common computations. Besides, a generic format better tolerate the variations in the distribution of non-zero entries. If power consumption and FPGA resource utilization must also be considered, LIL or BCSR are other candidates.

- For structured band matrices (Figure 14c), a pattern-specific format such as DIA, near-perfectly utilizes the memory bandwidth and does it better as the partition size increases. However, to allow such utilization to effectively impact the other performance metrics, the computation engine must also be tailored to the format if DIA must be used in a particular application. Otherwise, the mismatch would create a computation bottleneck.
- For less sparse (density > 0.1) applications such as the inference of neural network, optimizations beyond simple partitioning of size 8x8 or at most 16x16 hurt the performance even though it might help reduce the memory footprint (possibly, not too much). Extracting the non-zero partitions from the neural network can be done with the aid of structure pruning schemes [3, 39–42].

**Figure 14: Comparing sparse formats for: (a) SuiteSparse, (b) random, and (c) band matrices. 1 and 0 show best and worst, respectively.**

REFERENCES

- [1] Colin Yu Lin, Ngai Wong, and Hayden Kwok-Hay So. Design space exploration for sparse matrix-matrix multiplication on fpgas. *International Journal of Circuit Theory and Applications*, 41(2):205–219, 2013.
- [2] Qiuling Zhu, Tobias Graf, H Ekin Sumbul, Larry Pileggi, and Franz Franchetti. Accelerating sparse matrix-matrix multiplication with 3d-stacked logic-in-memory hardware. In *High Performance Extreme Computing Conference (HPEC), 2013 IEEE*, pages 1–6. IEEE, 2013.
- [3] Bahar Asgari, Ramyad Hadidi, Hyesoon Kim, and Sudhakar Yalamanchili. Eridanus: Efficiently running inference of dnn using systolic arrays. *IEEE Micro*, 39(5):46–54, 2019.
- [4] Asit K Mishra, Eriko Nurvitadhi, Ganesh Venkatesh, Jonathan Pearce, and Debbie Marr. Fine-grained accelerators for sparse machine learning workloads. In *Design Automation Conference (ASP-DAC), 2017 22nd Asia and South Pacific*, pages 635–640. IEEE, 2017.
- [5] Eriko Nurvitadhi, Asit Mishra, and Debbie Marr. A sparse matrix vector multiply accelerator for support vector machine. In *Proceedings of the 2015 International Conference on Compilers, Architecture and Synthesis for Embedded Systems*, pages 109–116. IEEE Press, 2015.
- [6] Udit Gupta, Brandon Reagen, Lillian Pentecost, Marco Donato, Thierry Tambe, Alexander M Rush, Gu-Yeon Wei, and David Brooks. Masr: A modular accelerator for sparse rns. In *2019 28th International Conference on Parallel Architectures and Compilation Techniques (PACT)*, pages 1–14. IEEE, 2019.
- [7] Bahar Asgari, Ramyad Hadidi, Tushar Krishna, Hyesoon Kim, and Sudhakar Yalamanchili. Alrescha: A lightweight reconfigurable sparse-computation accelerator. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 249–260. IEEE, 2020.
- [8] Subhankar Pal, Jonathan Beaumont, Dong-Hyeon Park, Apurva Amarnath, Siying Feng, Chaitali Chakrabarti, Hun-Seok Kim, David Blaauw, Trevor Mudge, and Ronald Dreslinski. Outerspace: An outer product based sparse matrix multiplication accelerator. In *HPCA*, pages 724–736. IEEE, 2018.
- [9] Kartik Hegde, Hadi Asghari-Moghaddam, Michael Pellauer, Neal Crago, Aamer Jaleel, Edgar Solomonik, Joel Emer, and Christopher W Fletcher. Extensor: An accelerator for sparse tensor algebra. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 319–333, 2019.
- [10] Konstantinos Kanellopoulos, Nandita Vijaykumar, Christina Giannoula, Roknoddin Azizi, Skanda Koppula, Nika Mansouri Ghiasi, Taha Shahroodi, Juan Gomez Luna, and Onur Mutlu. Smash: Co-designing software compression and hardware-accelerated indexing for efficient sparse matrix operations. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 600–614, 2019.
- [11] Vidushi Dadu, Jian Weng, Sihao Liu, and Tony Nowatzki. Towards general purpose acceleration by exploiting common data-dependence forms. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 924–939, 2019.
- [12] R David Evans, Lufei Liu, and Tor M Aamodt. Jpeg-act: Accelerating deep learning via transform-based lossy compression.
- [13] Animesh Jain, Amar Phanishayee, Jason Mars, Lingjia Tang, and Gennady Pekhimenko. Gist: Efficient data encoding for deep neural network training. In *2018 ACM/IEEE 45th Annual International Symposium on Computer Architecture (ISCA)*, pages 776–789. IEEE, 2018.
- [14] Song Han, Xingyu Liu, Huizi Mao, Jing Pu, Ardavan Pedram, Mark A Horowitz, and William J Dally. Eie: efficient inference engine on compressed deep neural network. *ACM SIGARCH Computer Architecture News*, 44(3):243–254, 2016.
- [15] Angshuman Parashar, Minsoo Rhu, Anurag Mukkara, Antonio Puglielli, Rangharajan Venkatesan, Bruce Khailany, Joel Emer, Stephen W Keckler, and William J Dally. Scnn: An accelerator for compressed-sparse convolutional neural networks. *ACM SIGARCH Computer Architecture News*, 45(2):27–40, 2017.
- [16] Song Han, Junlong Kang, Huizi Mao, Yiming Hu, Xin Li, Yubin Li, Dongliang Xie, Hong Luo, Song Yao, Yu Wang, et al. Ese: Efficient speech recognition engine with sparse lstm on fpga. In *Proceedings of the 2017 ACM/SIGDA International Symposium on Field-Programmable Gate Arrays*, pages 75–84, 2017.
- [17] Ashish Gondimalla, Noah Chesnut, Mithuna Thottethodi, and TN Vijaykumar. Sparten: A sparse tensor accelerator for convolutional neural networks. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 151–165, 2019.
- [18] Guy E Blelloch, Michael A Heroux, and Marco Zagha. Segmented operations for sparse matrix computation on vector multiprocessors. Technical report, CARNEGIE-MELLON UNIV PITTSBURGH PA SCHOOL OF COMPUTER SCIENCE, 1993.
- [19] Dahai Guo, William Gropp, and Luke N Olson. A hybrid format for better performance of sparse matrix-vector multiplication on a gpu. *The International Journal of High Performance Computing Applications*, 30(1):103–120, 2016.
- [20] A Cristiano I Malossi, Yves Ineichen, Constantine Bekas, Alessandro Curioni, and Enrique S Quintana-Orti. Performance and energy-aware characterization of the sparse matrix-vector multiplication on multithreaded architectures. In *2014 43rd International Conference on Parallel Processing Workshops*, pages 139–148. IEEE, 2014.
- [21] Georgios Goumas, Kornilios Kourtis, Nikos Anastopoulos, Vasileios Karakasis, and Nectarios Koziris. Understanding the performance of sparse matrix-vector multiplication. In *16th Euromicro Conference on Parallel, Distributed and Network-Based Processing (PDP 2008)*, pages 283–292. IEEE, 2008.
- [22] Vasileios Karakasis, Georgios Goumas, and Nectarios Koziris. Exploring the performance-energy tradeoffs in sparse matrix-vector multiplication. In *Workshop on emerging supercomputing technologies (WEST)—ICS*, volume 11, 2011.
- [23] Samuel Williams, Leonid Oliker, Richard Vuduc, John Shalf, Katherine Yelick, and James Demmel. Optimization of sparse matrix-vector multiplication on emerging multicore platforms. In *SC’07: Proceedings of the 2007 ACM/IEEE Conference on Supercomputing*, pages 1–12. IEEE, 2007.
- [24] Naser Sedaghati, Te Mu, Louis-Noel Pouchet, Srinivasan Parthasarathy, and P Sadayappan. Automatic selection of sparse matrix representation on gpus. In *Proceedings of the 29th ACM on International Conference on Supercomputing*, pages 99–108, 2015.
- [25] Richard W Vuduc and Hyun-Jin Moon. Fast sparse matrix-vector multiplication by exploiting variable block structure. In *International Conference on High Performance Computing and Communications*, pages 807–816. Springer, 2005.
- [26] SciPy. List-of-list sparse matrix, 2020. [Online; accessed April-2020].
- [27] David R Kincaid, Thomas C Oppe, and David M Young. Itpackv 2d user’s guide. Technical report, Texas Univ., Austin, TX (USA). Center for Numerical Analysis, 1989.
- [28] Youcef Saad. Numerical solution of large nonsymmetric eigenvalue problems. *Computer Physics Communications*, 53(1-3):71–90, 1989.
- [29] Moritz Kreutzer, Georg Hager, Gerhard Wellein, Holger Fehske, and Alan R Bishop. A unified sparse matrix data format for efficient general sparse matrix-vector multiplication on modern processors with wide simd units. *SIAM Journal on Scientific Computing*, 36(5):C401–C423, 2014.
- [30] Yousef Saad. *Iterative methods for sparse linear systems*, volume 82. siam, 2003.
- [31] Timothy A Davis and Yifan Hu. The university of florida sparse matrix collection. *ACM Transactions on Mathematical Software (TOMS)*, 38(1):1, 2011.
- [32] Maxim Naumov, Dheevatsa Mudigere, Hao-Jun Michael Shi, Jianyu Huang, Narayanan Sundaraman, Jongsoo Park, Xiaodong Wang, Udit Gupta, Carole-Jean Wu, Alisson G Azzolini, et al. Deep learning recommendation model for personalization and recommendation systems. *arXiv preprint arXiv:1906.00091*, 2019.
- [33] Eric Qin, Ananda Samajdar, Hyounkjun Kwon, Vineet Nadella, Sudarshan Srinivasan, Dipankar Das, Bharat Kaul, and Tushar Krishna. Sigma: A sparse and irregular gemm accelerator with flexible interconnects for dnn training. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 58–70. IEEE, 2020.
- [34] Zhekai Zhang, Hanrui Wang, Song Han, and William J Dally. Sparch: Efficient architecture for sparse matrix multiplication. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 261–274. IEEE, 2020.
- [35] Nitish Srivastava, Hanchen Jin, Shaden Smith, Hongbo Rong, David Albonese, and Zhiru Zhang. Tensaurus: A versatile accelerator for mixed sparse-dense tensor computations. In *2020 IEEE International Symposium on High Performance Computer Architecture (HPCA)*, pages 689–702. IEEE, 2020.
- [36] Fazle Sadi, Joe Sweeney, Tze Meng Low, James C Hoe, Larry Pileggi, and Franz Franchetti. Efficient spmv operation for large and highly sparse matrices using scalable multi-way merge parallelization. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture*, pages 347–358, 2019.
- [37] Gene H Golub and Charles F Van Loan. *Matrix computations*, volume 3. JHU press, 2012.
- [38] Grzegorz Malewicz, Matthew H Austern, Aart JC Bik, James C Dehnert, Ilan Horn, Naty Leizer, and Grzegorz Czajkowski. Pregel: a system for large-scale graph processing. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of data*, pages 135–146. ACM, 2010.
- [39] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. In *Advances in Neural Information Processing Systems*, pages 2074–2082, 2016.
- [40] Huizi Mao, Song Han, Jeff Pool, Wenshuo Li, Xingyu Liu, Yu Wang, and William J Dally. Exploring the regularity of sparse structure in convolutional neural networks. *arXiv preprint arXiv:1705.08922*, 2017.
- [41] Pavlo Molchanov et al. Pruning convolutional neural networks for resource efficient inference. *arXiv:1611.06440*, 2016.
- [42] HT Kung, Bradley McDanel, and Sai Qian Zhang. Packing sparse convolutional neural networks for efficient systolic array implementations: Column combining under joint optimization. In *Proceedings of the Twenty-Fourth International Conference on Architectural Support for Programming Languages and Operating Systems*, pages 821–834. ACM, 2019.