Fast Kronecker Matrix-Matrix Multiplication on GPUs

Abhinav Jangda ajangda@microsoft.com Microsoft Research Redmond, Washington, USA Mohit Yadav myadav@umass.edu University of Massachusetts Amherst Amherst, Massachusetts, USA

Abstract

Kronecker Matrix-Matrix Multiplication (Kron-Matmul) is the multiplication of a matrix with the Kronecker Product of several smaller matrices. Kron-Matmul is a core operation for many scientific and machine learning computations. State-of-the-art Kron-Matmul implementations utilize existing tensor algebra operations, such as matrix multiplication, transpose, and tensor matrix multiplication. However, this design choice prevents several Kron-Matmul specific optimizations, thus, leaving significant performance on the table.

To address this issue, we present FASTKRON, an efficient technique for Kron-Matmul on single and multiple GPUs. FASTKRON is independent of linear algebra operations enabling several new optimizations for Kron-Matmul. Thus, it performs up to $40.7 \times$ and $7.85 \times$ faster than existing implementations on 1 and 16 GPUs respectively.

$\label{eq:ccs} \textit{CCS Concepts: } \bullet \textit{Computing methodologies} \rightarrow \textit{Massively parallel algorithms; Algebraic algorithms; Distributed algorithms.}$

Keywords: Graphics Processing Units, CUDA, Kronecker Product, Linear Algebra

ACM Reference Format:

Abhinav Jangda and Mohit Yadav. 2024. Fast Kronecker Matrix-Matrix Multiplication on GPUs. In *The 29th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming* (*PPoPP '24*), March 2–6, 2024, Edinburgh, United Kingdom. ACM, New York, NY, USA, 14 pages. https://doi.org/10.1145/3627535.3638489

1 Introduction

Kronecker Matrix is a widely used data format in machine learning [8, 15, 23, 29, 35, 46, 51, 52] and scientific computations [10, 18, 50]. A Kronecker Matrix is a block matrix of shape PM \times QN and is the result of the *kronecker product* of

ACM ISBN 979-8-4007-0435-2/24/03...\$15.00 https://doi.org/10.1145/3627535.3638489 two matrix factors of shape P×Q and M×N. *Kronecker Matrix-Matrix Multiplication* (Kron-Matmul) is the multiplication of an input matrix with a kronecker matrix. Kron-Matmul is the key operation for computations that represents their data as a Kronecker Matrix. For example, training Gaussian Processes (GPs), which are a class of machine learning models, involves multiplication of GPs' kernel matrix with the dataset matrix. In several state-of-the-art GPs [51, 52], the kernel matrix is a kronecker product of smaller matrix factors. Hence, training of these GPs involves Kron-Matmul of the kernel matrix with the training dataset matrix.

There are two existing algorithms for Kron-Matmul: the shuffle algorithm [11] and the fused tensor matrix multiply transpose (FTMMT) algorithm [27]. Both algorithms run for multiple iterations, where at each iteration, algorithms multiply the input matrix with a factor to generate an intermediate matrix. This intermediate of an iteration is the input for the next iteration. These algorithms differ in how they perform the multiplication in each iteration. The shuffle algorithm uses a series of reshape, matrix multiplication, and transpose operations. Existing single-GPU implementations of the shuffle algorithm, GPvTorch [16] and PvKronecker [7], use NVIDIA cuBLAS [1] for matrix multiplication and generate optimized transpose kernels. Similarly, the multi-GPU implementation in Cyclops Tensor Framework (CTF) [42] uses distributed matrix multiplication and transpose. The FTMMT algorithm represents the input matrix as a multi-dimensional tensor and fuses the transpose and multiplication using linear algebra engines, such as COGENT [25] and cuTensor [2] for single-GPU and DISTAL [53] for multi-GPU. Hence, both algorithms use existing linear algebra operations.

However, due to this design choice, these implementations miss several Kron-Matmul specific optimizations leading to the following three inefficiencies. First, the transpose in the shuffle algorithm is significantly expensive. Our experiments found that the transpose step can take up to 80% of the total execution time on both single- and multi-GPU executions. Second, linear algebra GPU kernels are not optimized for Kron-Matmul. For example, the matrix multiplication in the shuffle algorithm is performed on small and rectangular matrices, which is an inefficient case for NVIDIA cuBLAS. Moreover, when caching data from the global to shared memory, COGENT performs a high number of shared memory bank conflicts. Third, linear algebra operations used in both algorithms require full intermediates in the global memory at each iteration, leading to unnecessary loads and stores from

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org. PPoPP '24, March 2–6, 2024, Edinburgh, United Kingdom

 $[\]circledast$ 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

the global memory. This process can be optimized by fusing multiple linear algebra operations in a single kernel. However, existing implementations cannot perform this fusion because they use general linear algebra operations instead of Kron-Matmul specific operations. Similarly, multi-GPU implementations communicate per GPU intermediate at every iteration, leading to high communication volume.

In this paper, we present FASTKRON to address above limitations. FASTKRON'S Kron-Matmul algorithm is not based on existing linear algebra operations and thus, enables new optimizations for single- and multi-GPU scenarios. FASTKRON'S algorithm divides rows of the input matrix into slices of size equal to the factor's column and multiplies each slice with all columns of the factor (Section 3). Then the algorithm stores consecutive elements in the intermediate matrix as the multiplication of consecutive slices with same column. Thus, our algorithm writes output elements at the correct index, removing the need for memory shuffle operations like transpose and reshape. FASTKRON'S GPU implementation contains a novel tiling methodology that assigns multiple slices and columns to each thread (Section 4). The implementation caches inputs in the shared memory while minimizing the shared memory bank conflicts and performing coalesced global memory accesses (Section 4.1). The algorithm also enables us to fuse multiplications with multiple factors in a single GPU kernel by storing intermediates in the shared memory, leading to reduced global memory accesses (Section 4.2). Furthermore, FASTKRON'S multi-GPU implementation minimizes the communication volume by performing multiple local multiplications on each GPU before communicating the intermediate of the last local multiplication (Section 5).

FASTKRON provides significant performance speedup over state-of-the-art single and multi-GPU Kron-Matmul implementations. On an NVIDIA Tesla V100 GPU, FASTKRON provides up to 40.7× speedup over GPyTorch [16], 6.40× over COGENT [25], and 5.41× over cuTensor (Section 6.2.3). On a system with 16 NVIDIA Tesla V100 GPUs, FASTKRON performs 7.85× better than CTF [42] and 5.33× better than DIS-TAL [53] (Section 6.3). We also integrated FASTKRON into GPyTorch to reduce the training time of several Gaussian Process techniques by up to 6.20×. FASTKRON is publicly available at https://github.com/abhijangda/fastkron.

2 Kronecker Matrix-Matrix Multiplication

This section presents existing algorithms for Kronecker Matrix-Matrix Multiplication and their limitations.

A Kronecker Matrix of $G_{P^1P^2 \times Q^1Q^2}$ is the result of the kronecker product of two matrices $F^1_{P^1 \times O^1}$ and $F^2_{P^2 \times O^2}$, such that:

$$\mathbf{G}_{\mathbf{P}^{1}\mathbf{P}^{2}\times\mathbf{Q}^{1}\mathbf{Q}^{2}} = \mathbf{F}_{\mathbf{P}^{1}\times\mathbf{Q}^{1}}^{1} \otimes \mathbf{F}_{\mathbf{P}^{2}\times\mathbf{Q}^{2}}^{2} = \begin{bmatrix} f_{11}^{1}\mathbf{F}^{2} & \dots & f_{1\mathbf{Q}^{1}}^{1}\mathbf{F}^{2} \\ \vdots & \vdots & \vdots \\ f_{\mathbf{P}^{1}1}^{1}\mathbf{F}^{2} & \dots & f_{\mathbf{P}^{1}\mathbf{Q}^{1}}^{1}\mathbf{F}^{2} \end{bmatrix}$$

We refer to matrices $\mathbf{F}_{P^1 \times Q^1}^1$ and $\mathbf{F}_{P^2 \times Q^2}^2$ as *Kronecker factors* of $\mathbf{G}_{P^1P^2 \times Q^1Q^2}$. We refer to the operation of multiplying the Kronecker product of N factors $\mathbf{F}_{P^i \times Q^i}^i$ with $\mathbf{X}_{M \times \prod_i P^i}$ to compute $\mathbf{Y}_{M \times \prod_i Q^i}$ as *Kron-Matmul*. In this section, for simplicity we consider all factors are of the same shape $P \times Q$. A naive algorithm for Kron-Matmul computes the Kronecker matrix and then matrix multiply (Matmul) it with \mathbf{X} . However, this algorithm results in a high complexity of $O(MP^NQ^N)$. We now present two state-of-the-art algorithms for Kron-Matmul that are faster than the naive algorithm.

2.1 The Shuffle Algorithm

The *shuffle algorithm* [11] avoids computing the Kronecker matrix. The shuffle algorithm runs for N iterations from N to 1, with each iteration performing three steps. An iteration *i* generates an intermediate matrix $\mathbf{Y}_{M \times Q^{N+1-i}P^{i-1}}^{i}$, which is also the input for the next iteration. Figure 1 shows Kron-Matmul of $\mathbf{X}_{M \times P^2}$ with two factors $\mathbf{F}_{P \times Q}^1$ and $\mathbf{F}_{P \times Q}^2$. The first iteration multiplies \mathbf{X} with the last factor \mathbf{F}^2 . First, step (a) reshapes $\mathbf{X}_{M \times P^2}$ to $\mathbf{X}_{MP \times P}$ and then multiply $\mathbf{X}_{MP \times P}$ with $\mathbf{F}_{P \times Q}^2$ to obtain $\mathbf{Y}_{M \times P \times Q}^2$ (Figure 1a). Then, step (b) reshapes $\mathbf{Y}_{M \times P \times Q}^2$ to $\mathbf{Y}_{M \times P \times Q}^2$ and transposes the last two dimensions of $\mathbf{Y}_{M \times QP}^2$ (Figure 1b). Finally, step (c) reshapes $\mathbf{Y}_{M \times Q \times P}^2$ to $\mathbf{Y}_{M \times QP}^2$ and $\mathbf{F}_{P \times Q}^1$ to get the final result of Kron-Matmul, $\mathbf{Y}_{M \times QP}^1$. The shuffle algorithm performs O(MP $\sum_{i=1}^N Q^{N-i}P^i)$ computations, which is better than the naive approach.

Limitations State-of-the-art GPU based implementations of the shuffle algorithm, GPyTorch [16] and PyKronecker [7], uses NVIDIA cuBLAS MatMul in step (a) and an efficient GPU kernel for transposing two inner dimensions of a 3-D tensor for step (b). However, this transpose of the 3-D tensor cannot be fused with the Matmul. Since all transpose steps in the algorithm performs $O(M \sum_{i=1}^{N} Q^{N-i}P^i)$ memory accesses, both implementations spends up to 80% of the total time in the transpose. Moreover, the cuBLAS Matmul is not optimized for multiplying a large skinny and small matrix.

2.2 Fused Tensor-Matrix Multiply Transpose

We can avoid the expensive transpose operation by representing the input $\mathbf{X}_{M \times P^N}$ and intermediates $\mathbf{Y}_{M \times Q^{N+1-i}p^{i-1}}^{i}$ as 3 dimensional tensor and fusing the transpose with the computation using tensor multiplication engines [25, 32]. This algorithm, which we call as Fused Tensor-Matrix Multiply Transpose (FTMMT) algorithm [27] works as follows. The algorithm goes from N to 1 iterations and in an iteration *i* multiplies the tensor with the *i*th factor and transposes the last two dimensions then reshape to the columns of next factor. Consider Kron-Matmul of $\mathbf{X}_{M \times P^3}$ with factors $\mathbf{F}_{P \times Q}^1$, $\mathbf{F}_{P \times Q}^2$, and $\mathbf{F}_{P \times Q}^3$. In the first iteration, we reshape $\mathbf{X}_{M \times P^3}$ to 3-D tensor $\mathbf{X}_{M \times P^2 \times P}$. Then, we multiply the last dimension of

Fast Kronecker Matrix-Matrix Multiplication on GPUs



Figure 1. First iteration of the shuffle algorithm for Kron-Matmul of $X_{2\times4}$ and $F_{2\times2}^1 \otimes F_{2\times2}^2$. Reshape transforms shape of a tensor to other shape. Transpose exchanges the elements of two dimensions of a multi-dimensional tensor.



Figure 2. First iteration of the FASTKRON Kron-Matmul algorithm of $\mathbf{X}_{2\times4}$ with $\mathbf{F}_{2\times2}^1 \otimes \mathbf{F}_{2\times2}^2$. Elements of \mathbf{Y}^2 with the same color are generated by a column of \mathbf{F}^2 with the same color. The result of first iteration, \mathbf{Y}^2 , is same as in Figure 1.

 $X_{M\times P^2\times P}$ with the last factor $F^3_{P\times Q}$, to obtain $Y^3_{M\times P^2\times Q}$. Finally, we transpose the second and last dimensions of $Y^3_{M\times P^2\times Q}$ to $Y^3_{M\times Q\times P^2}$ and reshape to $Y^3_{M\times QP\times P}$. The second iteration multiplies $Y^3_{M\times QP\times P}$ with $F^2_{P\times Q}$ and transposes the first and last dims of the result and reshape to $Y^2_{M\times Q^2\times P}$. Similarly, the third iteration multiplies with $F^1_{P\times Q}$ to obtain $Y^1_{M\times Q^2\times Q}$ and reshape to $Y^1_{M\times Q^3}$.

Limitations Although existing single- and multi-GPU tensor multiplication systems, COGENT [25], cuTensor [2], and DISTAL [53], improves over the shuffle algorithm, they do not execute each iteration efficiently and do not optimize across iterations. For example, COGENT caches data in fast memories using the standard approach, i.e., cache contiguous P elements of the last dimension from the shared memory to P registers of consecutive threads. This approach leads to shared memory bank conflicts because every P element lies in the same shared memory bank when P is a multiple of the number of banks. Moreover, these systems store the output intermediate of the current iteration in the global memory and load the intermediate in the next iteration, leading to high memory accesses and communication volume.

In summary, existing implementations perform high memory accesses and high communication volume because they do not optimize for Kron-Matmul.

3 The FASTKRON Algorithm

This section presents a novel Kron-Matmul algorithm, which enables us to develop new optimizations for Kron-Matmul.

Algorithm 1 is FASTKRON's algorithm for Kron-Matmul. For brevity, we present the algorithm for the common case in our evaluation dataset, where all Kronecker factors are of the same shape. However, it is straightforward to generalize this algorithm to factors of different shape. The algorithm works as follows. First, the algorithm allocates two intermediate matrices and set the number of cols of input matrix for the first iteration (line 3–5). These intermediates are swapped after every iteration. The algorithm starts the multiplication from the last factor (lines 6–17). For each factor, the algorithm first computes number of cols of the output intermediate (line 8). Then the algorithm performs *Sliced Multiply* for each row of **X** with all columns of **F**ⁱ, where it divides the row into slices of size P (line 10) and multiplies each slice with each column of **F**ⁱ (line 14). Then, the algorithm writes the result of each slice and column in **Y**¹ (line 15). Finally, the algorithm returns the final result (line 18). The algorithm performs O(MP $\sum_{i=1}^{N} Q^{N-i}P^i$) computations and O(M $\sum_{i=1}^{N} Q^{N-i}P^i$) memory accesses. Hence, the ratio of computations to memory accesses is P.

Figure 2 shows an example of Kron-Matmul of $X_{2\times4}$ and $F_{2\times2}^1 \otimes F_{2\times2}^2$ using FASTKRON's algorithm. First, both columns of F^2 are multiplied with each slice of each row of X (Figure 2a). Two elements generated by sliced multiplication of the first column of F^2 are stored as the first two elements of the first row of Y^2 . Then, elements generated by sliced multiplication of the second column of F^2 are stored as the third and fourth elements of the first row of Y^2 . Finally, the intermediate result with the current factor, Y^2 , is used as the input matrix for the next factor.

Comparison with Existing Algorithms There is a key difference that separates FASTKRON's algorithm and MatMul. In MatMul, consecutive elements in a row of the output are the result of multiplication of consecutive columns of the second matrix with the same row of the first matrix. However, in FASTKRON's algorithm, consecutive elements are obtained by multiplying consecutive slices of rows with the same column of the factor. Thus, FASTKRON's algorithm

Algorithm	17	Гhe Fa	STKRON	Kron-l	Matmul	algorithm
-----------	----	--------	--------	--------	--------	-----------

1: Input: Matrix $\mathbf{X}_{M \times P^N}$ and N Kronecker Factors $\mathbf{F}_{P \times O}^1$ 2: **Output**: Result of Kron-Matmul of **X** with all N \mathbf{F}^{i} s 3: \mathbf{Y}^1 and \mathbf{Y}^2 are new matrices of size $M \times \max_{f=0}^N (Q^{N-f}P^f)$ 4: $Y^1 = X$ \triangleright Copy X to \mathbf{Y}^1 5: $K = P^N$ ▶ Input Intermediate No. of Cols 6: for $f = N \rightarrow 1$ do for $i = 1 \rightarrow M$ do 7: $L = (K \div P) \times Q$ ▶ Output Intermediate No. of Cols 8: for $j = 1 \rightarrow L$ do ▶ Sliced-Multiply **X** row and **F**^t col 9: rowSlice = $(j \times P) \mod K$ 10: $kCol = (j \div P^{N-1}) \mod P$ 11: acc = 012: for $k = 1 \rightarrow P$ do ▶ Sliced Multiply Accumulate 13: acc += $\mathbf{Y}^{0}[i]$ [rowSlice + k] × $\mathbf{F}^{f}[k]$ [kCol] 14: $Y^{1}[i][j] = acc$ 15: $Y^1, Y^2 = Y^2, Y^1$ ▶ Swap intermediates 16: K = L17: 18: return Y¹

writes output elements at the right index, removing any need for extra memory operations, like transpose.

4 FASTKRON'S CUDA Implementation

This section presents FASTKRON'S CUDA kernel with an efficient shared memory caching technique (Section 4.1) and fusion of multiple sliced multiplications (Section 4.2).

FASTKRON provides Python and C++ APIs for Kron-Matmul for several data types, including float and double. All the API functions call into a type generic C++ implementation of Algorithm 1. The implementation executes the loop of lines 6–17 of Algorithm 1 and returns the output. The implementation invokes SLICEDMULTIPLYKERNEL (in Figure 3) to sliced-multiply $X_{M\times K}$ and $F_{P\times Q}$ to generate $Y_{M\times \frac{KQ}{P}}$. The kernel takes global memory addresses for X, F, and Y, along with their shapes. Each thread block of the kernel performs the following steps:

- 1. Load slices of rows of X and cols of F into shared memory.
- 2. Load part of slices from the shared memory to registers.
- 3. Perform sliced multiply accumulate on register buffers to compute multiple elements of **Y**.
- 4. When elements of **Y** are computed, transfer elements from registers to global memory.

The above workflow is similar to NVIDIA CUTLASS [3] and BLIS [48] but with differences in the shared memory caching and element to thread assignment optimized for Kron-Matmul. We now explain above steps in Figure 3 using an example workflow for $\mathbf{Y}_{2\times512}$, $\mathbf{X}_{2\times512}$, and $\mathbf{F}_{8\times8}$ in Figure 4. **Thread Block Tiles** Each thread block sliced multiply $\{\mathsf{T}_{M},\mathsf{T}_{K}\}$ block of $\mathbf{X}_{M\times K}$ with T_Q cols of $\mathbf{F}_{P\times Q}$ to produce a block of $\{\mathsf{T}_{M}, \left(\frac{\mathsf{T}_{K}}{P} \times \mathsf{T}_Q\right)\}$ of \mathbf{Y} . Thus, the kernel is invoked with

1 Ks = (Slices = (T_K/P))*T_P; 2 shared T Xs[T_M][Ks], Fs[T_P][T₀]; 3 register T Yr[T_M][R_K][R₀]={0}; 4 //Compute Element Locations in Y $5 yQ = (tid.x / Slices) * R_0;$ 6 yK = (tid.x % Slices) * R_K; 7 for $(t_P = 0; t_P < P; t_P += T_P)$ /*Step 1: Global to Shared Memory*/ 8 9 ShiftGToS(Xg, Xs, K, T_M, T_P, Ks, R_K); 10 DirectGToS(Fg, Fs, Q, T_P, T₀); 11 syncthreads(); 12 for $(r_P = 0; r_P < T_P; r_P += R_P)$ { register T Xr[T_M][R_K][R_P], Fr[R_P][R₀]; 13 /*Step 2: Shared to Registers*/ 14 15 ShiftSToR(Xs, Xr, r_P, yK, R_K, R_P, T_M, T_P); 16 DirectSToR(Fs, Fr, r_P, T_P, T₀, R₀); /*Step 3: Multiply Accumulate*/ 17 for $(m = 0; m < T_M; m++)$ for $(k = 0; k < R_K; k++)$ 18 for $(q = 0; q < R_Q; q++)$ for $(p = 0; p < R_P; p++)$ 19 Yr[m][k][q] += Xr[m][k][p] * Fr[p][q]; 20 21 } syncthreads(); } 22 /*Step 4: Registers to Global Memory*/ 23 for(r = 0; r < T_M ; r++) 24 for (b = 0; b < R_0 ; b++) for (e = 0; e < R_K ; e++){ 25 $yRow = r + bid.x*T_M;$ $yCol = (yQ + b) * (T_K/P) + yK + e;$ 26 27 $yCol = (yCol / (T_K/P)) * (K/P) +$ 28 bid.y * (T_K/P) + yCol % (T_K/P) ;

Y[yRow * K + yCol] = Yr[r][b][e];}

29

Figure 3. FASTKRON'S SLICEDMULTIPLYKERNEL for $X_{M \times K}$ and $F_{P \times Q}$ to compute $Y_{M \times \frac{KQ}{P}}$. Shift* and Direct* transfers data from global/shared memory to shared memory/registers.

 $\begin{cases} \frac{M}{T_M}, \frac{K}{T_K}, \frac{Q}{T_Q} \end{cases} \text{ threadblocks. Each thread produces } R_K \times R_Q \text{ elements for each } T_M \text{ row. Thus, each thread block contains } \\ \frac{T_K \div P}{R_K} \times \frac{T_Q}{R_Q} \text{ threads. In Figure 4a, each threadblock sliced-multiply } T_P \times T_K = 512 \text{ block of } \mathbf{X}_{2 \times 512} \text{ with } T_Q = 2 \text{ cols of } \\ \mathbf{F}_{8 \times 8} \text{ to produce } \frac{512}{8} \times 2 = 128 \text{ elements of a row of } \mathbf{Y}_{2 \times 512}. \end{cases}$

Global to Shared Memory Each thread block caches T_P elements of slices of rows of **X** and cols of **F** in the shared memory (lines 9–10 in Figure 3). The main loop of the kernel iterates over all T_P tiles and multiplies T_P elements of slices and cols to produce elements of **Y** (lines 7–21). In Figure 4a, the thread block caches $T_P = 4$ elements of each slice and col in the shared memory. Section 4.1 presents our efficient shared memory access approach to minimize bank conflicts.

Shared Memory to Registers A thread loads R_K slices of **X** and R_Q cols of **F** from the shared memory to registers (lines 15–16 in Figure 3). These slices and columns are multiplied to compute $R_K \times R_Q$ elements of **Y**. Therefore, higher values of R_K increases the reuse of cols of **F** and higher values of R_0 imply higher reuse of slices of **X**. Figure 4b shows that

Fast Kronecker Matrix-Matrix Multiplication on GPUs



(a) Thread block 0 is assigned to 1^{st} row of **X** and 2 cols of **F** to produce $\frac{512}{8} \times 2 = 128$ elements of **Y**. The thread block load 4 elements of all 64 slices into Xs and of columns into Fs, and multiplies each slice and column to produce partial 128 elements of **Y**. Then, the thread block moves to the next 4 elements, and updates partial elements to get its final elements.



(b) Thread 0 is assigned to first 2 slices of Xs and all 2 cols of Fs to produce 4 elements of Yr. A thread loads 2 elements of both slices in Xr and of cols in Fr, and multiplies each slice and column to produce 4 partial elements. Then, the thread moves to next 2 elements and updates partial elements to get final elements. Elements for col 1 are stored at index 0 and for col 2 are stored at index 64 of **Y**.

Figure 4. FASTKRON's tiling to sliced-multiply $X_{2\times512}$ and $F_{8\times8}$ to produce $Y_{2\times512}$ with $T_M = 1$, $T_K = 512$, $T_Q = 2$, $T_P = 4$, $R_P = 2$, $R_Q = 2$, $R_K = 2$. There are $\frac{512}{8} = 64$ slices for each **X** row. The CUDA kernel is invoked with $\left\{\frac{2}{1}, \frac{512}{512}, \frac{8}{2}\right\}$ threadblocks. Xs and Fs are shared memory buffers. Xr, Fr, and Yr are register buffers.

each thread loads $R_P = 2$ elements of $R_K = 2$ slices and of $R_Q = 2$ columns to register buffers (Xr and Fr).

Elements to Thread Mapping Each thread computes $R_K \times R_Q$ elements of **Y** stored in registers. The thread controls the computation intensity using R_P , i.e., each thread loads and multiplies R_P elements of R_K slices of **X** and R_Q cols of **F** (lines 12–21 in Figure 3). In Figure 4b, with $R_K = R_Q = 2$, two slices are multiplied with two cols to obtain 4 elements of Y.

Registers to Global Memory After computing its elements, each thread stores these elements to **Y** (line 23–29 in Figure 3). Since consecutive elements of **Y** are obtained by multiplying slices of **X** with the same col of **F**, all R_K elements are stored consecutively. Moreover, a group of R_K elements for a col *c* are stored at the address starting from $c \times \frac{K}{P}$. Therefore, the thread computes its index in **Y** (lines 5–6 in Figure 3) and write elements. In Figure 4b, the $R_K = 2$ elements for the first col of **F** are stored at indices starting from 0, while the elements for the second col are stored at indices $1 \times \frac{512}{8} = 64$.

4.1 Efficient Data Movement

This section describes FASTKRON'S *shift caching* that minimizes shared memory bank conflicts in our algorithm.

The standard approach, which we call *direct caching*, transfers data from the global to shared memory by assigning consecutive threads to contiguous elements and loads contiguous elements from the shared memory to registers of consecutive threads. This method is used in Matmul and tensor contractions of CUTLASS [3] and COGENT [25]. However, using this method, in our case, leads to high shared memory bank conflicts. In Figure 4b, as $T_P = 4$, thread 0 load elements 0–3 of slice 0 stored in Xs[0] to Xs[3], and thread 1 load elements 0–3 of slice 2 stored in Xs[8] to Xs[11]. If

the number of banks, i.e., warpSize, is 4, then all elements at Xs[0], Xs[8], ..., Xs[64] fall into the same bank leading to 4 conflicts for every read.

FASTKRON'S shift caching performs coalesced global memory accesses and minimizes shared memory bank conflicts to $\left[\frac{WarpSize}{T_P}\right]$. Figure 5 shows the implementation of the shift caching. When loading from the global to shared memory ShiftGToS finds the slice index for each element and shifts the element forward by the ratio of the slice index and the number of slices per thread (R_K) (line 4). Consequently, when loading elements from the shared memory to registers, ShiftSToR takes the starting slice index of the thread, divides it by the number of slices per thread to obtain the shift, and shifts the element back in the thread's register tile (lines 10-13). In Figure 4b, thread 1 stores slice 2 to the shared memory by shifting elements forward by $\frac{2}{R_{\nu}} = 1$, i.e., elements 0-2 are stored at Xs[9]-X[11], and element 3 is stored at Xs[8]. Similarly, when storing slice 4, thread 2 shift all elements forward by $\frac{4}{R_{K}} = 2$ and stores element 0 of this slice at Xs[14]. When loading to registers, thread 1 loads element 0 of slice 2 from Xs[9] while thread 2 loads element 0 of slice 4 from Xs[14]. If warpSize is 4, Xs[9] and Xs[14] fall in different banks, avoiding any bank conflicts. Section 6.2.2 shows the effectiveness of the shift method in reducing bank conflicts over the direct method in COGENT [25] and cuTensor [2].

4.2 Fusing Consecutive Sliced Multiplications

This section describes FASTKRON's fusion mechanism that perform multiple sliced multiplications in a single kernel leading to significant decrease in global memory accesses. PPoPP '24, March 2-6, 2024, Edinburgh, United Kingdom

```
1 ShiftGToS(Xg, Xs, K, T_N, T_P, Ks, R_K){
2
    for (m = 0; m < T_M; m++)
3
    for (k = tid; k < Ks; k += bdim) {
 4
     elem = k%T_P; slice = k/T_P; shift = slice/R_K;
5
     col = slice *T_P + (elem + shift) %T_P;
     Xs[m][col] = Xg[(m + bid.x) * K + k]; \}
6
7
8 ShiftSToR(Xs, Xr, r<sub>P</sub>, yK, R<sub>K</sub>, R<sub>P</sub>, T<sub>M</sub>, T<sub>P</sub>){
9
    for (m = 0; m < T_M; m++) for (q = 0; q < R_0; q++)
10
     slice = (yK + q); shift = (yK / R_K)%T_P;
11
     for (p = 0; p < R_P; p++) {
12
      elem = r_P + p; round = (elem + shift)%T_P;
      Xr[m][q][p] = Xs[m][slice*T_P + round];}
13
```

Figure 5. FASTKRON'S shift caching method. ShiftGToS caches from global to shared memory. ShiftSToR caches from shared memory to registers.

The linear algebra kernels used by existing Kron-Matmul implementations require inputs in the global memory. Thus, these implementations store the output intermediate of each multiplication in the global memory and load the intermediate again for the next multiplication. Since FASTKRON is independent of linear algebra kernels, it can fuse consecutive sliced multiplications and store intermediates in the shared memory, thereby, avoiding expensive global memory accesses. We now discuss the working of fused kernel.

The fused kernel sliced multiplies all cols of N_{fused} factors with T_K elements and stores the intermediate in shared memory. Hence, the algorithm using the fused kernel runs for $\lceil \frac{N}{N_{fued}} \rceil$ iterations. After every sliced multiply the number of shared memory elements that are contiguous in the full global memory intermediate reduces by the factor of P. Figure 6 shows that the fused kernel for Kron-Matmul of $X_{1 \times 256}$ and $F_{4 \times 4}$ with $T_{K} = 128$, sliced multiplies $N_{\text{fused}} = 2$ factors to generate T_K elements in the shared memory for each factor. After the first sliced multiply, there are $T_Q^1 = 4$ sets of $\frac{T_{K}}{p^{1}}$ = 32 elements of shared memory with a stride of 32 in the global intermediate. After the second multiply, we get $T_Q^2 = 16$ sets of $\frac{T_K}{P^2} = 8$ contiguous elements with a stride of 8 in the global intermediate. In general, after the i^{th} sliced multiply, there are $T_Q{}^i$ sets of $\frac{T_K}{P^i}$ contiguous elements with a stride of $\frac{T_{K}}{p_{i}}$ in the global intermediate. Thus, the fused kernel can compute a maximum of $N_{fused} = \lfloor \log_P T_K \rfloor$ consecutive sliced multiplications. Moreover, the fusion is valid when all elements of all slices can be stored in the shared memory, i.e., $T_P = P$. Our experiments found this is true for $P \le 32$ and $Q \leq 32$.

After the last multiply, the fused kernel transfers elements from the shared to global memory using StoreFusedShMem function (Figure 7). Consider storing element 41 of the shared memory to global memory in Figure 6. The function first Abhinav Jangda and Mohit Yadav



Figure 6. Workflow of the fused kernel for Kron-Matmul of $\mathbf{X}_{1\times 256}$ with 4 factors $\mathbf{F}_{4\times 4}$ and $\mathsf{T}_{\mathsf{K}} = 128$ by the first thread block. The kernel fuses $N_{\text{fused}} = 2$ sliced multiplications (max is 3). A thread block sliced multiply all 4 cols of \mathbf{F}^4 and \mathbf{F}^3 with T_{K} shared memory elements and store them in another shared memory buffer. After each multiplication, the corresponding indices in global intermediate (\mathbf{Y}^1) are shown. Finally, the output of 2nd sliced multiply is written to \mathbf{Y}^1 .

computes (i) the number of slices in global and shared memory, i.e., 64 and 32 in our example, and (ii) the slices of fusion of N_{fused} factors in the global and shared memory, i.e., 16 and 8 in our example (lines 2–3). The function now iterates on all T_K elements of each row in the shared memory and store them to global memory using the below steps: 1) Compute the slice of the element in the shared memory tile and scale it to the global memory (line 7), i.e., 64 in our example; 2) Compute the fused slice index in the shared memory tile and scale to the global memory (line 9), i.e., 16 in our example; 3) Compute the element index within the fused slice (line 12), i.e., 1 in our example; 4) Finally, store the element at the sum of the above indices (line 15), i.e., at 81. Section 6.2.2 shows that fusion is a key optimization for small P.

4.3 Autotuning Kernel Parameters

This section describes FASTKRON's autotuning mechanism to find efficient tile sizes for any Kron-Matmul shape.

Kron-Matmul can be performed on matrices of diverse shapes. However, there is no single set of tile size parameter values that performs efficiently for all shapes. Therefore, FASTKRON performs auto-tuning over a range of tile size parameter values for the given shape. The auto-tuning phase considers all combinations of following values of tile size parameters till the maximum shared memory usage and registers per thread is reached:

Thread Block Tile Sizes (i) $T_K \in \text{all multiples of P till K, (ii)}$ $T_P \in \text{all factors of P, (iii)}$ $T_Q \in \text{all factors of Q, and (iv)}$ even values of T_M until the number of threadblocks executing in parallel by all SMs reaches a maximum value. Fast Kronecker Matrix-Matrix Multiplication on GPUs

```
1 StoreFusedShMem(Xg, Xs, K, P, Q, T<sub>m</sub>, T<sub>K</sub>, N<sub>fused</sub>) {
2
   XgSlics = K/P; XsSlics = T_K/P;
   XgFuseSlics=K/P**N<sub>fused</sub>; XsFuseSlics=T<sub>K</sub>/P**N<sub>fused</sub>;
3
 4
    for(e = tid; e < T_m * T_K; e += bdim) {
5
     m = e/T_K; c = e%T_K
     //Scale Shared Mem Slice Idx to Global Mem Idx
6
     slice = (c/XsSlics)*XgSlics;
7
     //Scale Shared Fused Slice to Global Mem
8
9
     fusedSlice = ((c%XsSlics)/XsFuseSlics) *
10
                      XgFuseSlics;
11
     //Elem Idx in Fused Slice
12
     elem = bid.y * XsFuseSlics + c % XsFuseSlics;
     //Column index in Global Memory
13
     col = slice + fusedSlice + elem;
14
     Yg[(m+bid.x*T_M) * K/P*Q + col] = Xs[m*T_M + c];
15
```

Figure 7. Write output of N_{fused} fused sliced multiply kernels from shared to global memory. $X \star Y$ represents X^{Y} .

Thread Tile Sizes (i) $R_P \in all$ factors of T_P , (ii) $R_Q \in all$ factors of T_Q , and (iv) $R_K \in \text{all factors of } \frac{T_K}{T_R}$.

This narrowing down of tile size choices based on the available resource usage and the occupancy significantly decreases the search space of all choices. FASTKRON compile CUDA kernels for all combinations of the above tile sizes in parallel and find the kernel with the least execution time.

Distributed Kron-Matmul 5

Existing distributed implementations of the shuffle algorithm in CTF [42] and the FTMMT algorithm using DISTAL [53], executes each iteration by (i) dividing the input among all GPUs, (ii) performing multiplications on each GPU to generate its local intermediate, and (iii) communicate intermediates among all GPUs to obtain a globally distributed intermediate as the output of iteration. This section presents FASTKRON's distributed Kron-Matmul that minimizes the communication by (i) performing multiple N_{local} sliced multiplications on each GPU to generate intermediates local to each GPU and (ii) communicating local intermediates to obtain the globally distributed output intermediate of N_{local} iterations.

FASTKRON performs distributed Kron-Matmul on a homogenous 2D grid of $\{G_M, G_K\}$ GPUs by dividing the computation into a block of size $\frac{M}{G_M} \times \frac{K}{G_K}$ per GPU. Since each factor $\mathbf{F}_{P \times Q}$ is much smaller than $\mathbf{X}_{M \times K}$, FASTKRON requires that all factors are accessible on all GPUs. Algorithm 2 is FASTKRON'S distributed Kron-Matmul algorithm, which is executed by each GPU. The algorithm assumes that all factors are of the same shape, but it is straightforward to support the general case. First, the algorithm computes GPU block size, allocates intermediate matrices on each GPU, and computes N_{local} (lines 2–4). Then each GPU can perform $N_{local} = \log_{P} \frac{K}{G_{V}}$ local sliced multiplications before communicating local intermediates to obtain the globally distributed

Algorithm 2 Multi-GPU Kron-Matmul using $G_M \times G_K$ GPUs

- 1: Current GPU $\left\{g_M,g_K\right\}$ with in the grid of $\left\{G_M,G_K\right\}$ GPUs
- 2: $TG_M, TG_K = \frac{M}{G_M}, \frac{K}{G_K} \longrightarrow Tile comput$ 3: \mathbf{Y}^1 and \mathbf{Y}^2 are new matrices of size $TG_M \times TG_K$ ▶ Tile computed by each GPU
- 4: $N_{local} = \lfloor \log_P TG_K \rfloor$
- 5: $\mathbf{Y}^1 = \mathbf{X}[\mathbf{g}_M \times \mathsf{TG}_M : (\mathbf{g}_M + 1) \times \mathsf{TG}_M][\mathbf{g}_K \times \mathsf{TG}_K : (\mathbf{g}_K + 1) \times \mathsf{TG}_K]$
- 6: **for** $f = N \rightarrow 1$ with step N_{local} **do**
- $\textbf{for } b = 0 \rightarrow N_{local} 1 \textbf{ do} \quad \triangleright \text{ Do } N_{local} \text{ multiplies per GPU}$ 7:
- $\mathbf{Y}^2 = \text{SLICEDMULTIPLYKERNEL}(\mathbf{Y}^1_{\text{TG}_M \times \text{TG}_K}, \mathbf{F}^{f-b}_{P \times O})$ 8:
- $Y^1, Y^2 = Y^2, Y^1$ 9:
- $\label{eq:fordst} \textbf{for} \ \texttt{dst} = 1 \longrightarrow G_K \ \textbf{do} \quad \triangleright \ \texttt{Share result among GPUs with the}$ 10: same g_M
- 11: if dst = g_K then
- 12: **for** src = $1 \rightarrow G_K$ **do**
- 13:
- $$\begin{split} \mathbf{Y}_{src}^{1} &= \mathbf{Y}^{1}[1:\mathsf{TG}_{\mathsf{M}}][src \times \frac{\mathsf{TG}_{\mathsf{K}}}{\mathsf{G}_{\mathsf{K}}}:(src+1) \times \frac{\mathsf{TG}_{\mathsf{K}}}{\mathsf{G}_{\mathsf{K}}}] \\ \mathbf{Y}_{dst}^{2} &= \mathbf{Y}^{2}[1:\mathsf{TG}_{\mathsf{M}}][dst \times \frac{\mathsf{TG}_{\mathsf{K}}}{\mathsf{G}_{\mathsf{K}}}:(dst+1) \times \frac{\mathsf{TG}_{\mathsf{K}}}{\mathsf{G}_{\mathsf{K}}}] \\ \text{if } src \neq dst \text{ then} \end{split}$$
 14:
- 15:
- $\mathbf{Y}_{\mathrm{src}}^1 = \operatorname{Recv}(\mathbf{g}_{\mathrm{M}}, \operatorname{src})$ 16:
- $STOREGPUTILE(\mathbf{Y}_{dst}^2, \, \mathbf{Y}_{src}^1, \, K, \, P, \, Q, \, \mathsf{TG}_{\mathsf{M}}, \, \mathsf{TG}_{\mathsf{K}}, \, N_{local})$ 17:
- 18:
- $\mathbf{Y}_{dst}^{1} = \mathbf{Y}^{1} [1: \mathsf{TG}_{\mathsf{M}}] [dst \times \frac{\mathsf{TG}_{\mathsf{K}}}{\mathsf{G}_{\mathsf{K}}} : (dst+1) \times \frac{\mathsf{TG}_{\mathsf{K}}}{\mathsf{G}_{\mathsf{K}}}]$ 19:
- $Send(\mathbf{Y}_{dst}^1, \mathbf{g}_M, dst)$ 20:
- 21: $\mathbf{Y}^1, \mathbf{Y}^2 = \mathbf{Y}^2, \mathbf{Y}^1$
- 22: return Y¹

intermediate (line 7). Now the layout of each column of local intermediate is such that (i) a column stores G_K parts of size $\frac{K+G_K}{G_K}$, where the *i*th local part is stored on the *i*th GPU of the globally distributed intermediate and (ii) a column contains $\frac{K+G_K}{p^N_{local}}$ elements that lie apart by the same value in the global intermediate. Figure 8 shows an example of this layout on 4 GPUs. The algorithm relocates elements on each GPU by sharing these parts among all GPUs (line 20 and 16) and storing the received elements to correct place using STOREGPUTILE function (line 17), which is similar to STORE-FUSEDSHMEM. Therefore, the algorithm communicates exactly $G_M \times \left(\frac{N \times TG_M \times (K - TG_K)}{\log_P TG_K}\right)$ number of values. FASTKRON uses NVIDIA NCCL [4] for SEND and RECV. If all NVIDIA GPUs in the same g_M supports Point-to-Point accesses, FASTKRON implement lines 10-20 in a single CUDA kernel, which is more efficient than NCCL.

Similar to SUMMA [47], we divide the process grid into $\{\sqrt{G}, \sqrt{G}\}$, where G is the number of GPUs. If G is not a perfect square, we set the grid to $\{2^{\lceil \log_2 \sqrt{G} \rceil}, 2^{\lfloor \log_2 \sqrt{G} \rfloor}\}$. Although complex partitioning approaches for distributed Matmul exists [13, 26, 41], our experiments found that our partitioning approach performs well.



Figure 8. The element distribution of local intermediates of all 4 GPUs for Kron-Matmul on $X_{1\times 256}$ with 4 factors $F_{4\times 4}$ with $\{G_M, G_K\} = \{1, 4\}$. Each GPU computes $N_{local} = 2$ (max value is 3) sliced multiplications locally on the block size $\frac{M}{G_M} \times \frac{K}{G_K} = 1 \times 64$. After $N_{local} = 2$ multiplications, each local intermediate stores $\frac{K+G_K}{G_K} = 16$ elements for all 4 GPUs and $\frac{K+G_K}{G_K} = 4$ contiguous elements with stride of 4 in the full distributed intermediate.



Figure 9. Performance of GPyTorch, COGENT, FASTKRON, and FASTKRON without fusion for float with M = 1024, P = 8 to 128, and the two largest allocatable values of P^{N} .

6 Evaluation

In this section, we evaluate FASTKRON against state-of-theart implementations of the shuffle algorithm and the FTMMT algorithm on diverse Kron-Matmul sizes.

Experimental Setup We run our experiments on a single NVIDIA DGX-2 machine, which contains dual 24-core Intel Xeon CPUs and 16 NVIDIA Tesla V100 GPUs connected using NVLINK 2. Each Tesla V100 GPU contains 32 GB of global memory, and provides 15.7 Tera Floating Point Operations per Second (TFLOPS) for float and 7.8 TFLOPS for double. We use CUDA 12.2 on Ubuntu 22.04 and report the average TFLOPS of 100 runs after a warmup of 10 runs.

6.1 Autotuning Time

We perform a search over the tile size parameters for each Kron-Matmul problem size to obtain the best performing kernel parameters (Section 4.3). The autotuner generates upto 10,000 configurations for each problem size. By compiling kernels in parallel, the auto-tuner takes less than 2 minutes to find the fastest kernel.

6.2 Single GPU Evaluation

We first evaluate FASTKRON on a single GPU using microbenchmarks and then on a real world dataset.

6.2.1 Evaluation Systems. We perform experiments on the following systems:

GPyTorch [16] and PyKronecker [7] are two state-of-the-art GPU based implementations of the shuffle algorithm. Since both implementations call into NVIDIA cuBLAS for Matmul and generate identical CUDA kernels for the transpose, both implementations perform within 10% of each other. Thus, we use GPyTorch 1.11 as the baseline for the shuffle algorithm.

COGENT [25] is a state-of-the-art GPU code generator for tensor contractions. It fuses the transpose with the multiplication and generates optimized code and tile sizes for the FTMMT algorithm for the given Kron-Matmul shape.

cuTensor [2] is a state-of-the-art library for tensor contractions by NVIDIA. It fuses transpose with multiplications and autotunes at runtime over several tile sizes for the FTMMT algorithm. We use cuTensor as baseline because we found it performs as good as manually tuned CUTLASS [3].

FASTKRON with all optimizations.

FASTKRON-wo-Fuse is FASTKRON without the fusion of consecutive sliced multiplications (Section 4.2).

6.2.2 Microbenchmarks. We now present results for Kron-Matmul of $X_{M \times P^N}$ with N factors $F_{P \times P}$. Figure 9 shows the performance of each system with M=1024, the power of two values of P, and several values of N. The performance improves with the increase in P and N because (i) the ratio of computation to memory accesses, P increases, and (ii) large N increases the amount of parallelism. For the largest size, FASTKRON achieves 87% of the maximum FLOPS of the GPU.

Impact of Fusion The fusion optimization improves the performance by $2.20 \times$ for 8^5 to $1.15 \times$ for 32^3 . Since the shared memory is limited, the number of fused sliced multiplications decreases with an increase in P, leading to less improvement

Fast Kronecker Matrix-Matrix Multiplication on GPUs

Р	Ν	GPy	Torch (m	s)	COGENT	FastKron
		Matmul Trans. Total		(ms)	(ms)	
8	6	26	45	71.0	36.4	5.76
16	5	64	169	238	104	29.7
32	4	44	159	203	64.4	38.8
64	3	8.7	36	45.7	14.8	8.74

Table 1. Execution time of GPyTorch's Matmul and transpose, COGENT, and FASTKRON for float with M = 1024 and largest values of P and N on a 32GB GPU.

Р	N	COGEN	$MT(\times 10^7)$	FastKi	$ron(\times 10^7)$	Reduction in	
		Loads	Stores	Loads	Stores	Loads	Stores
8	6	6.93	1.06	2.24	1.04	3.10	1.02
16	5	27.8	6.29	11.9	2.48	2.33	2.54
32	4	27.7	10.4	20.2	3.32	1.37	3.13
64	3	6.85	4.71	3.97	1.48	1.72	3.18

Table 2. Shared memory load and store transactions in FASTKRON and COGENT, and reduction in transactions for M = 1024 and diverse values of P and N.

			FastKron		JEN I	GPyTorch	
		Float	Double	Float	Double	Float	Double
8	8	3.90	1.80	0.67	0.26	0.26	0.13
16	6	6.17	3.20	1.98	0.91	0.46	0.21
32	5	7.75	3.88	5.38	2.26	1.36	0.64
64	4	11.0	5.40	7.98	3.40	2.70	1.29

Table 3. Achieved TFLOPS of GPyTorch, COGENT, and FASTKRON for float and double with M = 16 and largest P^N .

with increasing P. The auto-tuner finds $T_P = 32$ for $P \ge 64$, hence, fusion is not applied to $P \ge 64$.

Speedup over GPyTorch FASTKRON provides a speedup of 7.62× for 8⁵ to 3.11× for 128³ over GPyTorch because: (i) FASTKRON avoids the transpose step of GPyTorch. Table 1 shows that the transpose step takes up to 80% of the total time of GPyTorch; and (ii) FASTKRON performs better than cuBLAS Matmul for small and rectangular shapes. Table 1 shows that FASTKRON is 0.95×-4.51× faster than cuBLAS.

Speedup over COGENT and cuTensor Both COGENT and cuTensor provides similar performance. FASTKRON provides a speedup of $6.40 \times$ for 8^5 to $1.47 \times$ for 128^3 over COGENT. Similarly, FASTKRON provides a speedup of $3.32 \times$ for 8^5 to $1.2 \times$ for 128^3 with maximum speedup of $5.41 \times$ at 16^4 . This speedup is because: (i) FASTKRON fuse consecutive sliced multiplications and store intermediates in the shared memory; and (ii) the shift caching method decreases shared memory bank conflicts over COGENT and cuTensor, which uses the direct caching method. For example, Table 2 shows that FASTKRON generates up to $3.10 \times$ less load and $3.18 \times$ less store transactions than COGENT.

ID	0	940	$(\mathbf{p}^{N_i} \dots \mathbf{p}^{N_i})$
ID	Source	{Mi}	$\{P_i : X \subseteq Q_i\}$
		20	$2^7 \times 2^7$
1.5	LOTM and DNINI [92]	{20, 50}	$2^9 \times 2^9$
1-5	LSTM and KINN [25]	20	$2^{10} \times 2^{10}$
		1	$2^{11} \times 2^{11}$
		10	$52 \times 50, 65 \times 20$
6-8	ML Compression [46]	50	$32 \times 8, 64 \times 128$
		10	$52 \times 65, 50 \times 20$
0.10	II-DA [10]	{4, 8, 16, 20}	$2^9 \times 2^9$
9-10	HYPA [10]	{4, 8, 16, 20}	$8^3 \times 8^3$
		1024	$3^7 \times 3^7$
17-19	Graphs [29]	1024	$4^7 \times 4^7$
		1024	$6^7 \times 6^7$
20-21	Biology [18]	1	$5^3 \times 5^3, 2 \times 2$
20-21	biology [18]	1	$5^2 \times 5^2, 2 \times 2, 25 \times 25$
		1526	$4^{6} \times 4^{6}$
22-24	Drug-Targets [50]	156	$8^3 \times 8^3$
		2967	$4^7 \times 4^7$
		16	$8^8 \times 8^8$
25_20	CD [8 15 25 51 50]	16	$16^{6} \times 16^{6}$
25-20	GI [0, 13, 33, 31, 32]	16	$32^{6} \times 32^{6}$
		16	$64^3 \times 64^3$

 $\begin{array}{l} \textbf{Table 4. Real world Kron-Matmul sizes. The first column is the id for each size and the second column is the source of these sizes. The third column <math display="inline">\{M_i\}$ contains one or more values of M for the same size of factors. The final column $\{P_i^{N_i} \times Q_i^{N_i}\}$ represents N_i consecutive factors of the shape $P_i \times Q_i. \end{array}$

Small M and Double Type Table 3 shows the performance for float and double types with M = 16. FASTKRON provides up to 13.4× speedup for float and up to 15.24× for double over GPyTorch. FASTKRON also runs up to 5.82× and 6.92× faster than COGENT for float and double respectively.

In summary, FASTKRON provides significant speedup over baselines for a diverse mix of matrix sizes and data type.

6.2.3 Real World Dataset. We now perform experiments on the real world Kron-Matmul sizes used in machine learning compression [23, 46], scientific computations [10], graph computations [29], computational biology [18], drugs [50], and gaussian process kernels [8, 15, 35, 51, 52]. Table 4 shows that our dataset contains 28 diverse cases with odd and nonpower of two values of M, factors with distinct and odd sizes, and N from 2 to 11. Figure 10 shows that FASTKRON performs $5.70 \times -40.7 \times$ faster than GPyTorch, $1.43 \times -8.14 \times$ faster over COGENT, and $1.55 \times -6.45 \times$ faster than cuTensor on the real world dataset.

6.3 Multiple GPUs Evaluation

We now evaluate the multi-GPU performance of FASTKRON on 16 NVIDIA Tesla V100 GPUs. In this experiment, we allocate all factors on all GPUs and compare FASTKRON



Figure 10. Speedup of FASTKRON over GPyTorch and CO-GENT on real-world Kron-Matmul sizes. The x-axis shows id of each problem size of Table 4.



Figure 11. Weak scaling of FASTKRON, CTF, and DISTAL on 1 to 16 GPUs with increasing M for P = 64, N = 4 (left) and P = 128, N = 4 (right) with float type.

against the following state-of-the-art distributed tensor algebra frameworks:

Cyclops Tensor Framework (CTF) [42] implements distributed tensor matrix multiply as a series of distributed Matmuls and transposes. Thus, our implementation of the distributed version of the shuffle algorithm in CTF uses a series of distributed tensor matrix multiplies.

DISTAL [53] allows a user to manually specify a distributed schedule for the given tensor algebra computation. We implemented each iteration of the FTMMT algorithm in DISTAL. Our schedule of the algorithm follows the same distribution as FASTKRON, i.e., divide M and K by \sqrt{G} . However, it is not possible to specify our distributed Kron-Matmul in DISTAL because DISTAL communicates intermediate after every sliced multiplication.

6.3.1 Results. Figure 11 shows the weak scaling (memory per GPU remains constant) performance of all systems with increasing M for P = 64, N = 4 and P = 128, N = 4. We chose these values of P and N because they provide maximum FLOPs per GPU. FASTKRON provides speedup of 7.85× over CTF and 5.33× over DISTAL at 16 GPUs. Moreover, for

Dataset	$\mathbf{P}^{\mathbf{N}}$	Spee	dup on	1 GPU	Speedup on 16 GPUs			
		SKI	SKIP	LOVE	SKI	SKIP	LOVE	
autompg	8 ⁷	$1.1 \times$	1.1×	1.2×	1.3×	1.3×	1.5×	
kin40k	8 ⁸	$1.5 \times$	1.3×	$1.2 \times$	3.1×	1.8×	1.6×	
airfoil	16 ⁵	1.1×	1.1×	1.3×	1.2×	1.2×	1.5×	
yacht	16^{6}	$1.8 \times$	1.7×	1.9×	3.8×	3.3×	5.2×	
servo	32^{4}	1.1×	1.1×	1.2×	1.3×	1.2×	1.5×	
airfoil	32^{5}	$1.8 \times$	1.8×	$1.8 \times$	6.2×	4.9×	5.0×	
3droad	64^{3}	1.1×	1.1×	1.2×	1.2×	1.2×	1.1×	
servo	64^{4}	$2.1 \times$	2.0×	2.2×	4.5×	3.8×	5.4×	

Table 5. Speedups in training GPs on real world datasets after

 integrating FASTKRON in GPyTorch over vanilla GPyTorch.

16 GPUs, FASTKRON reaches 69% of the maximum FLOPs. DISTAL performs better than CTF because CTF performs distributed transposes, which DISTAL avoids. FASTKRON performs better than both DISTAL and CTF because FASTKRON minimizes the communication volume by performing multiple sliced multiplications on each GPU before communicating their intermediates to obtain the full intermediate. Thus, FASTKRON is an efficient distributed Kron-Matmul engine.

6.4 Case Study: Fast Training of Gaussian Processes

Gaussian Processes (GPs) are a class of machine learning models that provide predictions with uncertainty and interpretability [37]. GPs contain a kernel matrix **K** and represent the training dataset of M points as a vector **V** of length M. The training process of GPs computes $(\mathbf{K})^{-1}\mathbf{V}$ [38, 52], which can be expensive for large values of M. Structured Kernel Interpolation (SKI) [51, 52] is a GP that interpolates the kernel matrix as $\mathbf{W}(\mathbf{K}^1 \otimes \mathbf{K}^2 \dots \mathbf{K}^N)\mathbf{W}^T$, where $\mathbf{W}_{M \times P^N}$ is an interpolation weight matrix and $\mathbf{K}^i_{P \times P}$ is a Kronecker kernel matrix. Using the conjugate gradient algorithm, the inverse computation is done using a series of Kron-Matmuls of **V** and $\otimes_i \mathbf{K}^i$. Thus, Kron-Matmul is a key process in training SKI and its variants, SKIP [15] and LOVE [35]. Also, large values of P and N improve the accuracy of GPs.

We integrated FASTKRON in GPyTorch to accelerate Kron-Matmul and use it to evaluate the reduction in training time of SKI, SKIP, and LOVE on the UCI dataset [5] with 150 to 3×10^5 points. We set the conjugate gradient method to consider 16 samples, i.e., M = 16, and runs for 10 iterations in each epoch. These datasets and parameters have been used in prior works [15, 35, 51, 52]. We perform experiments on the highest value of P that can be allocated in the GPU memory for each dataset. Table 5 shows that integrating FASTKRON in GPyTorch provides a speedup of up to 1.95× on a single GPU and up to $6.20 \times$ on 16 GPUs when training GPs. GPyTorch does not support multi-GPU execution for these GPs and thus executes several other operations on a single GPU, leading to a speedup increase of up to 3.33× with 16 GPUs over 1 GPU. Thus, integrating FASTKRON enables faster training of GPs on larger kernel matrices.

7 Related Work

Tensor Contractions The traditional way to execute tensor contractions first transpose input tensors to a valid Matmul, and then transpose the output to required tensor. Since these transposes are expensive, multiple works have developed efficient transpose routines for CPUs [44, 45], and GPUs [20, 30, 44, 49]. TTC [44] and TTLG [49] are compilers for transpose routines. TAL_SH [30] uses the state-of-the-art cuTT [20] library for efficient transpose on GPUs. Several works avoid the transpose and directly perform tensor contraction [25, 31, 32, 40, 43]. TBLIS [32] fuses transpose with BLIS [48] Matmul kernels on CPUs. GETT [43] uses a highly tuned macro kernel where its operands reside in the cache hierarchy. CUTLASS [3] and cuTensor [2] extends GETT approach to GPUs. Nelson et. al [33] and Patabandi et. al. [34] uses machine learning to tune tile size parameters of a tensor algebra GPU kernel. COGENT [25] improves over these approaches by generating a specialized kernel and tile sizes for tensor contractions for GPUs. Thus, these works can efficiently execute each multiplication with a Kronecker factor in the FTMMT algorithm. However, unlike FASTKRON, these works do not optimize for memory accesses across multiplications in the algorithm. Kim et. al. [24] improves tensor contractions for coupled cluster methods in quantum chemistry by fusing multiple contractions. However, their approach performs transpose in shared memory and these tensor contractions are different from contractions in Kron-Matmul.

Kronecker Matrix-Matrix Multiplication GPytorch [16] and PyKronecker [7] are two state-of-the-art single-GPU implementations for the shuffle algorithm [11]. Dayar and Orhan [12] presents an improvement to the shuffle algorithm for Kronecker matrix-vector products. Fackler [27] proposed an algorithm similar to FTMMT that avoids the transpose by representing the input matrix as a tensor. We use CO-GENT [24, 25] as a baseline for this algorithm. FASTKRON improves over these implementations by avoiding transpose and using optimizations, such as, shift caching and fusion of iterations. Moreover, FASTKRON provides a distributed algorithm while above systems are only for single node.

Optimizing Small and Skinny Matmul Many works has optimized Matrix Multiplication and Matrix-Vector Multiplication computations on small and skinny matrices on GPUs [6, 19, 39]. He et al. [19] proposes an optimal warp allocation strategy for matrix-vector multiplication. KBLAS [6] uses double-buffering to overlap data motion with computation to optimize matrix-vector multiplication. TSM2X [39] optimizes GEMM of rectangular matrices with small matrices. These techniques can improve the Matmul part in the shuffle algorithm but will still suffer from high transpose cost. However, FASTKRON avoids transpose operations and also provides an efficient multi-GPU execution.

Distributed Tensor Algebra Cannon [9, 28] and SUMMA [47] are one of the first algorithms for distributed Matmul. Solomonik et. al. [41] presented a 2.5D algorithm that distributes the summation dimension and CARMA [13] is a recursive algorithm. COSMA [26] is an near I/O-optimal algorithm for distributed Matmul. Rajbhandari et. al. [36] presents a communication-optimal algorithm for distributed tensor contraction. Cyclops Tensor Framework (CTF) [42] and DISTAL [53] are two state-of-the-art distributed tensor algebra systems. CTF executes tensor contractions as a series of distributed transposes and Matmuls, while DISTAL allows fusion of the transpose with contraction to perform better than CTF. For Kron-Matmul, both approaches communicates intermediate for each iteration, while FASTKRON minimizes the communication by performing multiple sliced multiplications on each GPU and communicate intermediate of last multiplication.

Distributed Computations Distributed Halide [14] extends Halide with scheduling primitives for distributing dimensions of loops. Recent works [17, 22] supports overlapping CUDA computations with communication. However, Kron-Matmul algorithm cannot be represented in these frameworks.

8 Conclusion

In this paper, we proposed a novel algorithm for Kron-Matmul, which is not based on existing linear algebra operations. This advantage enabled us to develop new optimizations for Kron-Matmul implementations on GPUs. Experimental results demonstrates that our implementation outperforms state-ofthe-art techniques on both single and multiple GPUs.

References

- Accessed: 2022-07-30. NVIDIA cuBLAS. https://developer.nvidia.com/ cublas.
- [2] Accessed: 2023-07-30. cuTENSOR: A High-Performance CUDA Library For Tensor Primitives. https://docs.nvidia.com/cuda/cutensor/latest/ index.html.
- [3] Accessed: 2023-07-30. NVIDIA cuTLASS: CUDA Templates for Linear Algebra Subroutines. https://github.com/NVIDIA/cutlass.
- [4] Accessed: 2023-07-30. NVIDIA NCCL: Optimized primitives for collective multi-GPU communication. https://github.com/NVIDIA/nccl.
- [5] Accessed: 2023-07-30. UCI ML Dataset. https://archive.ics.uci.edu/ datasets.
- [6] Ahmad Abdelfattah, David Keyes, and Hatem Ltaief. 2016. KBLAS: An Optimized Library for Dense Matrix-Vector Multiplication on GPU Accelerators. ACM Trans. Math. Softw, Article 18 (2016), 31 pages. https://doi.org/10.1145/2818311
- [7] Edward Antonian, Gareth W. Peters, and Michael Chantler. 2023. PyKronecker: A Python Library for the Efficient Manipulation of Kronecker Products and Related Structures. *Journal of Open Source Software* 8 (2023). https://doi.org/10.21105/joss.04900
- [8] Edwin V Bonilla, Kian Chai, and Christopher Williams. 2007. Multitask Gaussian Process Prediction. In Advances in Neural Information Processing Systems, J. Platt, D. Koller, Y. Singer, and S. Roweis (Eds.). Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/ paper/2007/file/66368270ffd51418ec58bd793f2d9b1b-Paper.pdf
- [9] Lynn Elliot Cannon. 1969. A Cellular Computer to Implement the Kalman Filter Algorithm. Ph. D. Dissertation. USA.
- [10] Rong Chen Chencheng Cai and Han Xiao. 2023. Hybrid Kronecker Product Decomposition and Approximation. *Journal of Computational* and Graphical Statistics 32 (2023). https://doi.org/10.1080/10618600. 2022.2134873 arXiv:https://doi.org/10.1080/10618600.2022.2134873
- Marc Davio. 1981. Kronecker products and shuffle algebra. *IEEE Trans. Comput.* C-30 (1981). https://doi.org/10.1109/TC.1981.6312174
- [12] Tuğrul Dayar and M. Can Orhan. 2015. On Vector-Kronecker Product Multiplication with Rectangular Factors. SIAM Journal on Scientific Computing 37, 5 (2015), S526–S543. https://doi.org/10.1137/140980326 arXiv:https://doi.org/10.1137/140980326
- [13] James Demmel, David Eliahu, Armando Fox, Shoaib Kamil, Benjamin Lipshitz, Oded Schwartz, and Omer Spillinger. 2013. Communication-Optimal Parallel Recursive Rectangular Matrix Multiplication. In 2013 IEEE 27th International Symposium on Parallel and Distributed Processing. 261–272. https://doi.org/10.1109/IPDPS.2013.80
- [14] Tyler Denniston, Shoaib Kamil, and Saman Amarasinghe. 2016. Distributed Halide. In Proceedings of the 21st ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming. https: //doi.org/10.1145/2851141.2851157
- [15] Jacob Gardner, Geoff Pleiss, Ruihan Wu, Kilian Weinberger, and Andrew Wilson. 2018. Product Kernel Interpolation for Scalable Gaussian Processes. In Proceedings of the Twenty-First International Conference on Artificial Intelligence and Statistics (Proceedings of Machine Learning Research, Vol. 84). PMLR. https://proceedings.mlr.press/v84/gardner18a. html
- [16] Jacob R. Gardner, Geoff Pleiss, David Bindel, Kilian Q. Weinberger, and Andrew Gordon Wilson. 2018. GPyTorch: Blackbox Matrix-Matrix Gaussian Process Inference with GPU Acceleration. In Proceedings of the 32nd International Conference on Neural Information Processing Systems (Montréal, Canada) (NIPS'18). 11 pages.
- [17] Tobias Gysi, Jeremia Bär, and Torsten Hoefler. 2016. dCUDA: Hardware Supported Overlap of Computation and Communication. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. https://doi.org/10.1109/SC.2016.51
- [18] Saskia Haupt, Alexander Zeilmann, Aysel Ahadova, Hendrik Bläker, Magnus von Knebel Doeberitz, Matthias Kloor, and Vincent Heuveline. 2021. Mathematical modeling of multiple pathways in colorectal

carcinogenesis using dynamical systems with Kronecker structure. *PLOS Computational Biology* (2021). https://doi.org/10.1371/journal. pcbi.1008970

- [19] Guixia He, Jiaquan Gao, and Jun Wang. 2018. Efficient dense matrixvector multiplication on GPU. Concurrency and Computation: Practice and Experience 30 (2018). https://doi.org/10.1002/cpe.4705
- [20] Antti-Pekka Hynninen and Dmitry I. Lyakh. 2017. cuTT: A High-Performance Tensor Transpose Library for CUDA Compatible GPUs. *CoRR* abs/1705.01598 (2017). arXiv:1705.01598
- [21] Abhinav Jangda. 2023. (Artifact) Fast Kronecker Matrix Multiplication on GPUs. (12 2023). https://doi.org/10.6084/m9.figshare.24803229.v1
- [22] Abhinav Jangda, Jun Huang, Guodong Liu, Amir Hossein Nodehi Sabet, Saeed Maleki, Youshan Miao, Madanlal Musuvathi, Todd Mytkowicz, and Olli Saarikivi. 2022. Breaking the Computation and Communication Abstraction Barrier in Distributed Machine Learning Workloads. In Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (Lausanne, Switzerland) (ASPLOS '22). https://doi.org/10.1145/3503222.3507778
- [23] Cijo Jose, Moustapha Cissé, and François Fleuret. 2017. Kronecker Recurrent Units. CoRR abs/1705.10142 (2017). http://arxiv.org/abs/ 1705.10142
- [24] Jinsung Kim, Aravind Sukumaran-Rajam, Changwan Hong, Ajay Panyala, Rohit Kumar Srivastava, Sriram Krishnamoorthy, and P. Sadayappan. 2018. Optimizing Tensor Contractions in CCSD(T) for Efficient Execution on GPUs. In Proceedings of the 2018 International Conference on Supercomputing (Beijing, China) (ICS '18). https://doi.org/10.1145/3205289.3205296
- [25] Jinsung Kim, Aravind Sukumaran-Rajam, Vineeth Thumma, Sriram Krishnamoorthy, Ajay Panyala, Louis-Noël Pouchet, Atanas Rountev, and P. Sadayappan. 2019. A Code Generator for High-Performance Tensor Contractions on GPUs. In 2019 IEEE/ACM International Symposium on Code Generation and Optimization (CGO). 85–95. https: //doi.org/10.1109/CGO.2019.8661182
- [26] Grzegorz Kwasniewski, Marko Kabić, Maciej Besta, Joost VandeVondele, Raffaele Solcà, and Torsten Hoefler. 2019. Red-Blue Pebbling Revisited: Near Optimal Parallel Matrix-Matrix Multiplication. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis (Denver, Colorado) (SC '19). https://doi.org/10.1145/3295500.3356181
- [27] Amy N. Langville and William J. Stewart. 2004. The Kronecker Product and Stochastic Automata Networks. J. Comput. Appl. Math. 167, 2 (jun 2004). https://doi.org/10.1016/j.cam.2003.10.010
- [28] Hyuk-Jae Lee, James P. Robertson, and José A. B. Fortes. 1997. Generalized Cannon's Algorithm for Parallel Matrix Multiplication. In *Proceedings of the 11th International Conference on Supercomputing* (Vienna, Austria) (*ICS '97*). https://doi.org/10.1145/263580.263591
- [29] Jure Leskovec, Deepayan Chakrabarti, Jon Kleinberg, Christos Faloutsos, and Zoubin Ghahramani. 2010. Kronecker Graphs: An Approach to Modeling Networks. J. Mach. Learn. Res. 11 (2010).
- [30] Dmitry Lyakh. 2015. An efficient tensor transpose algorithm for multicore CPU, Intel Xeon Phi, and NVidia Tesla GPU. *Computer Physics Communications* 189 (2015). https://doi.org/10.1016/j.cpc.2014.12.013
- [31] Wenjing Ma, Sriram Krishnamoorthy, Oreste Villa, Karol Kowalski, and Gagan Agrawal. 2013. Optimizing tensor contraction expressions for hybrid CPU-GPU execution. *Cluster Computing* 16 (Mar 2013). https://doi.org/10.1007/s10586-011-0179-2
- [32] Devin A. Matthews. 2018. High-Performance Tensor Contraction without Transposition. SIAM Journal on Scientific Computing 40 (2018). https://doi.org/10.1137/16M108968X
- [33] Thomas Nelson, Axel Rivera, Prasanna Balaprakash, Mary Hall, Paul D. Hovland, Elizabeth Jessup, and Boyana Norris. 2015. Generating Efficient Tensor Contractions for GPUs. In 2015 44th International Conference on Parallel Processing. https://doi.org/10.1109/ICPP.2015.106
- [34] Tharindu R. Patabandi, Anand Venkat, Abhishek Kulkarni, Pushkar Ratnalikar, Mary Hall, and Justin Gottschlich. 2021. Predictive Data

Locality Optimization for Higher-Order Tensor Computations. In Proceedings of the 5th ACM SIGPLAN International Symposium on Machine Programming (Virtual, Canada) (MAPS 2021). https://doi.org/10.1145/ 3460945.3464955

- [35] Geoff Pleiss, Jacob Gardner, Kilian Weinberger, and Andrew Gordon Wilson. 2018. Constant-Time Predictive Distributions for Gaussian Processes. In Proceedings of the 35th International Conference on Machine Learning (Proceedings of Machine Learning Research). PMLR. https://proceedings.mlr.press/v80/pleiss18a.html
- [36] Samyam Rajbhandari, Akshay Nikam, Pai-Wei Lai, Kevin Stock, Sriram Krishnamoorthy, and P. Sadayappan. 2014. A Communication-Optimal Framework for Contracting Distributed Tensors. In SC '14: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis. https://doi.org/10.1109/SC.2014.36
- [37] Carl Edward Rasmussen and Christopher K. I. Williams. 2005. Gaussian Processes for Machine Learning. The MIT Press. https://doi.org/10. 7551/mitpress/3206.001.0001
- [38] Carl Edward Rasmussen and Christopher K. I. Williams. 2006. *Gaussian* processes for machine learning. MIT Press.
- [39] Cody Rivera, Jieyang Chen, Nan Xiong, Jing Zhang, Shuaiwen Leon Song, and Dingwen Tao. 2021. TSM2X: High-performance tall-andskinny matrix-matrix multiplication on GPUs. *J. Parallel and Distrib. Comput.* 151 (2021). https://doi.org/10.1016/j.jpdc.2021.02.013
- [40] Yang Shi, U. N. Niranjan, Animashree Anandkumar, and Cris Cecka. 2016. Tensor Contractions with Extended BLAS Kernels on CPU and GPU. In 2016 IEEE 23rd International Conference on High Performance Computing (HiPC). https://doi.org/10.1109/HiPC.2016.031
- [41] Edgar Solomonik and James Demmel. 2011. Communication-Optimal Parallel 2.5D Matrix Multiplication and LU Factorization Algorithms. In *Euro-Par 2011 Parallel Processing*. https://doi.org/10.1007/978-3-642-23397-5_10
- [42] Edgar Solomonik, Devin Matthews, Jeff R. Hammond, John F. Stanton, and James Demmel. 2014. A Massively Parallel Tensor Contraction Framework for Coupled-Cluster Computations. *J. Parallel Distrib. Comput.* 74 (2014). https://doi.org/10.1016/j.jpdc.2014.06.002
- [43] Paul Springer and Paolo Bientinesi. 2018. Design of a High-Performance GEMM-like Tensor-Tensor Multiplication. ACM Trans. Math. Softw. 44 (2018). https://doi.org/10.1145/3157733
- [44] Paul Springer, Aravind Sankaran, and Paolo Bientinesi. 2016. TTC: A Tensor Transposition Compiler for Multiple Architectures. In Proceedings of the 3rd ACM SIGPLAN International Workshop on Libraries,

Languages, and Compilers for Array Programming (Santa Barbara, CA, USA). https://doi.org/10.1145/2935323.2935328

- [45] Paul Springer, Tong Su, and Paolo Bientinesi. 2017. HPTT: A High-Performance Tensor Transposition C++ Library. In Proceedings of the 4th ACM SIGPLAN International Workshop on Libraries, Languages, and Compilers for Array Programming (Barcelona, Spain) (ARRAY 2017). https://doi.org/10.1145/3091966.3091968
- [46] Urmish Thakker, Paul N. Whatmough, Zhi Gang Liu, Matthew Mattina, and Jesse G. Beu. 2021. Doping: A technique for Extreme Compression of LSTM Models using Sparse Structured Additive Matrices. In *Proceedings of Machine Learning and Systems 2021, MLSys 2021*, Alex Smola, Alex Dimakis, and Ion Stoica (Eds.). mlsys.org.
- [47] R. A. Van De Geijn and J. Watts. 1997. SUMMA: scalable universal matrix multiplication algorithm. *Concurrency: Practice and Experience* 9 (1997). https://doi.org/10.1002/(SICI)1096-9128(199704)9:4<255::AID-CPE250>3.0.CO;2-2
- [48] Field G. Van Zee and Robert A. van de Geijn. 2015. BLIS: A Framework for Rapidly Instantiating BLAS Functionality. ACM Trans. Math. Software 41 (2015). https://doi.acm.org/10.1145/2764454
- [49] Jyothi Vedurada, Arjun Suresh, Aravind Sukumaran Rajam, Jinsung Kim, Changwan Hong, Ajay Panyala, Sriram Krishnamoorthy, V. Krishna Nandivada, Rohit Kumar Srivastava, and P. Sadayappan. 2018. TTLG - An Efficient Tensor Transposition Library for GPUs. In 2018 IEEE International Parallel and Distributed Processing Symposium (IPDPS). https://doi.org/10.1109/IPDPS.2018.00067
- [50] Markus Viljanen, Antti Airola, and Tapio Pahikkala. 2022. Generalized vec trick for fast learning of pairwise kernel models. *Machine Learning* 111 (2022). https://doi.org/10.1007/s10994-021-06127-y
- [51] Andrew Gordon Wilson, Christoph Dann, and Hannes Nickisch. 2015. Thoughts on Massively Scalable Gaussian Processes. arXiv:1511.01870 [cs.LG]
- [52] Andrew Gordon Wilson and Hannes Nickisch. 2015. Kernel Interpolation for Scalable Structured Gaussian Processes (KISS-GP). In Proceedings of the 32nd International Conference on International Conference on Machine Learning - Volume 37 (Lille, France) (ICML'15).
- [53] Rohan Yadav, Alex Aiken, and Fredrik Kjolstad. 2022. DISTAL: The Distributed Tensor Algebra Compiler. In Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (San Diego, CA, USA) (PLDI 2022). https://doi. org/10.1145/3519939.3523437

A Artifact Appendix

The artifact [21] contains implementation of FASTKRON and scripts to reproduce our key results. The artifact provides a Dockerfile, which contains all prerequisites installed. Latest source code is available at https://github.com/abhijangda/fastkron.

A.1 Hardware

FASTKRON supports both systems with a single NVIDIA GPU and multiple NVIDIA GPUs. In our experiments we use a DGX-2 machine with 16 NVIDIA Tesla V100 GPUs connected using NVLINK 2.

A.2 Docker Container

Download the artifact zip file from [21], unzip it, and create the container.

unzip fastkron-ppopp-24-ae.zip cd fastkron-ae docker build -t fastkron-ppopp-24-ae . docker run -it --gpus all fastkron-ppopp-24-ae

Check if PyTorch supports CUDA:

```
python
>>> import torch
>>> torch.cuda.is_available()
True
```

A.3 Getting Started

We will now build FastKron and execute tests. In the container, the FastKron directory is available at /fastkron and the benchmark infrastructure is in /fastkron-benchmarks. **Setup CMake** Setup CMake inside FastKron Directory

```
mkdir /fastKron/build
cd /fastKron/build
cmake ..
```

Single GPU Test We can execute one of the single GPU tests as below:

```
make gen-single-gpu-kernels
make run-single-gpu-no-fusion-tests -j
```

Multi GPU Test We can execute one of the multi GPU tests as below:

```
make gen-multi-gpu-tests-kernel
make run-multi-gpu-nccl-no-fusion-tests -j
```

Execute all Tests (Optional) We can execute all tests from the FastKron directory

```
cd /fastkron
python tests/run-tests.py
```

If all above tests run fine and do not give any error then we have successfully setup the benchmarking.

A.4 Step by Step Instructions

We will now reproduce results in Figure 9, Table 3, Figure 10, Figure 11, and Table 5. These commands generate figures as PDF in the benchmarks directory and table as CSV in the benchmarks directory.

Change to the benchmark directory:

cd /fastkron-benchmarks

Figure 9 [Time 30 mins] Generate Figure-9.pdf in the benchmarks directory by executing:

```
python run_benchmarks.py -fk-dir /fastkron\
    -fk-bench-dir /fastkron-benchmarks \
    -bench Figure-9
make Figure-9.pdf
```

Table 3 [Time 15 mins] Generate Table-3-float.csv for Float type and Table3-double.csv for Double type in the benchmarks directory by executing:

python run_benchmarks.py -fk-dir /fastkron\
 -fk-bench-dir /fastkron-benchmarks \
 -bench Table-3

Figure 10 [Time 40 mins] Generate Figure-10.pdf in the benchmarks directory by executing:

```
python run_benchmarks.py -fk-dir /fastkron\
    -fk-bench-dir /fastkron-benchmarks \
    -bench Figure-10
make Figure-10.pdf
```

Figure 11 [Time 40 mins] Generate Figure-11-64.pdf and Figure-11-128.pdf:

```
python run_benchmarks.py -fk-dir /fastkron\
    -fk-bench-dir /fastkron-benchmarks \
    -bench Figure-11
make Figure-11-64.pdf Figure-11-128.pdf
```

 Table 5 [Time 30 mins] Generate Table-5.csv by executing:

python gps-Table-5.py ./uci 10