Accelerating Reduction and Scan Using Tensor Core Units

Abdul Dakkak, Cheng Li
University of Illinois Urbana-Champaign
Urbana, Illinois
{dakkak,cli99}@illinois.edu

Jinjun Xiong
IBM T. J. Watson Research Center
Yorktown Heights, New York
jinjun@us.ibm.com

Isaac Gelado
NVIDIA Corporation
Santa Clara, California
igelado@nvidia.com

Wen-mei Hwu
University of Illinois Urbana-Champaign
Urbana, Illinois
w-hwu@illinois.edu

ABSTRACT

Driven by deep learning, there has been a surge of specialized processors for matrix multiplication, referred to as Tensor Core Units (TCUs). These TCUs are capable of performing matrix multiplications on small matrices (usually 4×4 or 16×16) to accelerate HPC and deep learning workloads. Although TCUs are prevalent and promise increase in performance and/or energy efficiency, they suffer from over specialization as only matrix multiplication on small matrices is supported. In this paper we express both reduction and scan in terms of matrix multiplication operations and map them onto TCUs. To our knowledge, this paper is the first to try to broaden the class of algorithms expressible as TCU operations and is the first to show benefits of this mapping in terms of: program simplicity, efficiency, and performance. We implemented the reduction and scan algorithms using NVIDIA's V100 TCUs and achieved $89\% - 98\%$ of peak memory copy bandwidth. Our results are orders of magnitude faster (up to $100\times$ for reduction and $3\times$ for scan) than state-of-the-art methods for small segment sizes (common in HPC and deep learning applications). Our implementation achieves this speedup while decreasing the power consumption by up to 22% for reduction and 16% for scan.

1 INTRODUCTION

Deep learning's reliance on matrix-multiplication (GEMM) for compute has driven both research and industry to develop matrix-multiplication accelerator hardware — collectively called Tensor Core Units (TCUs) in this paper. TCUs are designed to accelerate Multilayer Perceptrons (MLP), Convolutional Neural Networks (CNN), and Recurrent Neural Networks (RNN) or Deep
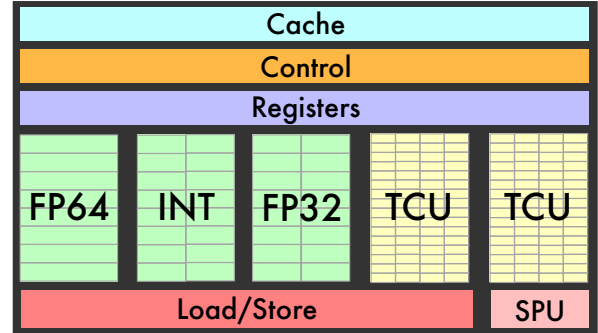
Figure 1: Each subcore (processing block) in the NVIDIA Tesla V100 PCI-E architecture contains 2 TCUs. In total, 640 TCUs are available — achieving a theoretical peek of 113 TFLOPS.

Neural Network (DNN) in general. TCUs come under the guise of different marketing terms, be it NVIDIA's Tensor Cores [55], Google's Tensor Processing Unit [19], Intel's DLBoost [69], Apple A11's Neural Engine [3], Tesla's HW3, or ARM's ML Processor [4]. They vary in the underlying hardware implementation [15, 27, 63, 71], and are prevalent [18, 55, 58] in both cloud and edge devices.

To show the theoretical benefits of TCUs, consider the NVIDIA Volta V100 GPUs architecture. Using V100 Tensor Cores, one achieves a $8\times$ throughput increase per Streaming Multiprocessors (SM) over previous Pascal GP100 generation. This throughput increase is because each V100 SM is capable of performing 1024 half precision operations per cycle using the TCUs whereas the GP100 SM is capable of performing 128 half precision operations per cycle without the TCUs. The throughput increase is enabled by the fact that the V100 dedicates a large chip area of the SM subcore to TCUs (Figure 1).

Although TCUs are prevalent and promise increase in performance and/or energy efficiency and are heavily used within supercomputers [32, 57] to achieve exascale performance, they suffer from over specialization. Currently, no algorithm other than GEMM utilizes the NVIDIA TCUs. This results in idle TCUs, low chip utilization, and limits TCUs applicability to specialized libraries or narrow application domains.

The objective of the paper is to expand the class of algorithms that can execute on TCUs— enabling the TCUs to be used within a wider range of non-GEMM algorithms. We choose reduction and scan, since a large body of work [7, 9, 36] has shown that they are key

primitives for data parallel implementations of radix sort, quicksort, lexical analysis, stream compaction, and polynomial evaluation. In this paper, we formulate a mapping of reduction or scan onto TCUs. We then introduce algorithms for cache- (warp-), processing element (PE)/core- (block-), and device- (grid-) level reduction and scan and show their performance on NVIDIA TCUs. We separate our algorithm description from implementation, making the algorithms, motivation, methods, and observations generally applicable to a broader range of TCUs and numerical precision agnostic. While the formulation is the main objective of the paper, we show that an implementation of our algorithms on NVIDIA V100 is either order of magnitude faster or rival the fastest GPU implementation, with much lower programming complexity. The key contributions of the paper are:

(1) We show how to use TCUs to compute both reduction and scan. We believe we are the first to formulate these algorithms in terms of TCU operations in a manner that is independent to the underlying TCU architecture.

(2) We implement our algorithms onto NVIDIA V100 GPUs and show orders of magnitude speedup over state-of-art algorithms for small segment sizes. Small segements are common in mathematics (e.g. evaluating polynomials), scientific applications (e.g. finite difference), and machine learning (e.g. batch norm) applications. For large segments, we are comparable to the fastest algorithms and achieve $89 - 98\%$ of theoretical peak memory copy bandwidth.

(3) We show that our implementation is up to 22% more power efficient and decreases the utilization of general purpose ALUs.

(4) We describe the current usage and programmability of the NVIDIA TensorCore and evaluate GEMM on the TCUs using cuBLAS [46], CUTLASS [49] and the CUDA TCU API.

This paper is divided as follows: we first describe the NVIDIA TCUs and show the performance of GEMM and GEMV computation in Section 2. In Section 3, we give a background of reduction and scan and show the TCU algorithms for reduction (Section 4) and scan (Section 5). We then compare our implementation against state-of-the-art in Section 6. Section 7 describes the related work, before we conclude in Section 8.

2 TENSOR CORES UNITS (TCUS)

A marquee feature of NVIDIA's GPUs (Volta's Tesla V100 and Turning's TU102 architectures) and Google's TPUs are their TCUs— a programmable matrix multiply and accumulate hardware units, called Tensor Cores by NVIDIA and matrix-multiply-units (MXUs) by Google [1]. While there are other competing TCU implementations, both NVIDIA Tensor Cores and Google's TPU are by far the most popular. At a high level, their functionality and architectural design are similar. They both subdivide the device into cores, with each having multiple processing block (or subcores) and TCUs. Figure 1 illustrates a subcore in an NVIDIA SM, with the V100 containing 80 SMs and each having 4 subcores. In turn, each subcore contains two Tensor Cores — for a total of 640 Tensor Cores and achieve a $12\times$ throughput improvement over previous generation Tesla P100 [54]. Google's TPUv3 device, on the other hand, has 8 cores — 4 chips each with 2 cores — with each core having 2 MXUs.

Since Google TPUs currently can only be used within Google Cloud using the XLA compiler [33] and the NVIDIA V100 TCUs are widely available and are installed in supercomputers [32, 57], this section will only describe the TCU usage and results for NVIDIA V100. Similar analysis can be performed for other TCUs.

Each NVIDIA V100 Tensor Core provides a $4 \times 4 \times 4$ tensor processing array capable of performing the operation $D = A \cdot B + C$ within a cycle, where A, B, C and D are 4×4 matrices. The inputs A and B must be in half precision format while the accumulators, C and D, can be either single or half precision. Each Tensor Core can perform 64 $(4 \times 4 \times 4)$ FMA operations per cycle. Therefore, using the TCU each SM can perform 1024 $(64 \times 2 \times 8)$ floating point operations per cycle, since each FMA consists of two floating point operations and each SM contains 8 Tensor Cores. This is an $8\times$ SM throughput increase compared to Pascal for floating point operations [54]. This section first describes the current usage of the NVIDIA Tensor Cores, then details the current NVIDIA Tensor Cores API, and presents evaluation results to motivate our work.

2.1 Current Library Usage

Currently, Tensor Cores have only been used to accelerate GEMM operations, most prominently through NVIDIA's CUDA libraries (such as cuBLAS [46] and cuDNN [48]). These libraries require users to opt-in to use the Tensor Cores to accelerate GEMM computation.NVIDIA also provides the CUTLASS (CUDA Templates for Linear Algebra Subroutines) [49] library, which is a C++ templated library that provides building block primitives to write high performance GEMM-like kernels. Deep learning frameworks such as NVCaffe [25], Caffe2 [8], MXNet [43], PyTorch [59], TensorFlow [67], and TensorRT [56] leverage these NVIDIA libraries for DNN training [50] and inference acceleration.

2.2 Programming Interface

Aside from the libraries, NVIDIA also provides a CUDA C++ Warp Matrix Multiply and Accumulate (WMMA) [47] API to program the Tensor Cores directly. The current WMMA API provides warp-level matrix operations for matrix load (`load_matrix_sync`), matrix store (`store_matrix_sync`), and matrix multiply and accumulate (`mma_sync`). These APIs operate on a special data type `fragment`, which holds a matrix tile in thread-local registers. A helper function to broadcast a scalar constant into a fragment (`fill_fragment`) is provided as well. No API currently exists for calling TCU operations at sub warp level — neither in the IR nor in the PTX [52, 53].

The `load_matrix_sync` function distributes values of the matrix across the warp lanes. Threads within a warp utilize multiple Tensor Cores concurrently to perform the `mma_sync` operation — collaborating to compute the $D_{M \times N} = A_{M \times K} \cdot B_{K \times N} + C_{M \times N}$, with M, N, K denoting the matrix dimensions. The API imposes limitations on the dimensions — requiring the shape $\langle M, N, K \rangle$ to be either $\langle 16, 16, 16 \rangle$, $\langle 32, 8, 16 \rangle$, or $\langle 8, 32, 16 \rangle$.

Listing 1 shows a CUDA kernel that computes a $\langle 16, 16, 16 \rangle$ matrix multiplication within a warp using the WMMA API. Lines 4–6 declare the matrix fragments. The API supports 3 kinds of matrices — `matrix_a` (A), `matrix_b` (B), and `accumulator` (C or D) — with

---

[1]We will use TCU and Tensor Core interchangeably in this paper.
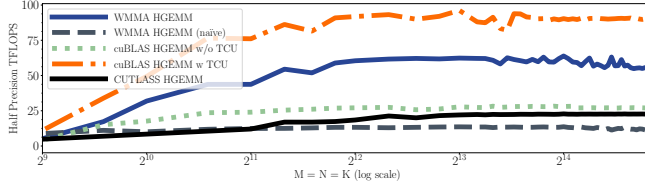
(a) GEMM with half precision input and half precision output.
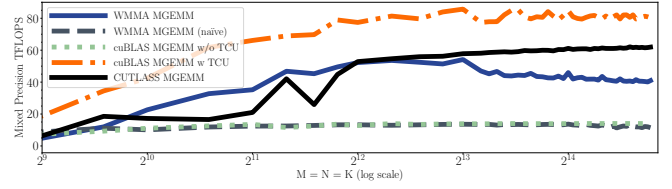
(b) Mixed precision GEMM with half precision input and single precision output.

Figure 2: General matrix-matrix multiplication (GEMM) performance using Tensor Cores for both half- (2a) and mixed- (2b) precision on a V100 PCI-E GPU with a clock frequency of 1380 MHz and a 113 TFLOPS peek performance. The inputs are square matrices with variable $\langle M, N, K \rangle$ dimensions. The optimized and naïve WMMA GEMM algorithms are described in the text.

```
1  #include <mma.h>
2  using namespace nvcuda::wmma;
3  __global__ void dot_wmma_16x16(half *a, half *b, half *c) {
4      fragment<matrix_a, 16, 16, 16, half, col_major> a_frag;
5      fragment<matrix_b, 16, 16, 16, half, row_major> b_frag;
6      fragment<accumulator, 16, 16, 16, half> c_frag;
7      load_matrix_sync(a_frag, a, /* row stride */ 16);
8      load_matrix_sync(b_frag, b, /* row stride */ 16);
9      fill_fragment(c_frag, 0.0f);
10     mma_sync(c_frag, a_frag, b_frag, c_frag);
11     store_matrix_sync(c, c_frag, 16, row_major);
12 }
```
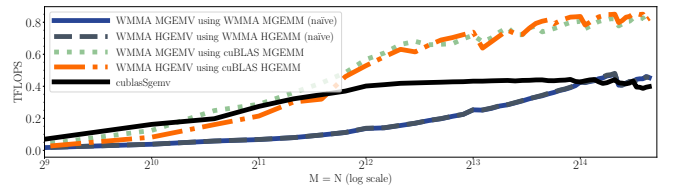
Listing 1: A simple CUDA kernel performing $\langle 16, 16, 16 \rangle$ matrix multiplication ($C = A \cdot B + C$) in half precision using the CUDA WMMA API.

Figure 3: General matrix-vector multiplication (GEMV) performance using Tensor Cores on a V100 PCI-E GPU. GEMV can be implemented in terms of a GEMM (with dimensions $\langle M, N, 16 \rangle$) or calling the GEMV method in CUBLAS (which currently does not support half precision).

each having their own internal data layout [2] as well as loading, storing, and computing semantics. Users specify both the data type and the $\langle M, N, K \rangle$ shape of the fragments. For both the *A* and *B* kinds, users specify whether the matrix is in column- or row-major order. Users also specify the stride between rows and load the data from either shared or global memory (Lines 7–8). Line 9 initializes the `matrix_c` elements to zero by broadcasting the scalar value 0 into the fragment. Once the data is loaded, users perform the matrix multiplication operation (Line 10) and store the results (Line 11).

The kernel in Listing 1 can be generalized to implement GEMM for arbitrary matrix dimensions in a manner similar to tiling matrix multiplication. For example, a naive implementation (referred to as `WMMA HGEMM (naïve)`) assigns a strip of 16 rows from matrix *A* and a strip of 16 columns from matrix *B* columns to each warp to compute a 16×16 tile of the output *C*. Each warp iterates through the *A* rows and *B* columns by loading 16×16 tiles of *A* and *B* from global memory into the fragments using `load_matrix_sync`, then performing `mma_sync`, and repeats. After all rows of *A* and columns of *B* have been consumed, the warp uses `store_matrix_sync` to store the accumulated *C* values into global memory. An optimized implementation (referred to as `WMMA HGEMM`) utilizes persistent threads where each thread block collaboratively loads multiple tiles of matrix *A* and *B* into shared memory (to facilitate tile re-use). The tiles are then loaded into fragments and the `mma_sync` operation is performed.

2.3 GEMM Evaluation

To show the TCU performance, we evaluate GEMM using Tensor Cores on an NVIDIA Tesla V100 PCI-E GPU with CUDA 9.2.88 through cuBLAS, CUTLASS (version 0.1.1), and hand written kernels using the WMMA API (Figure 2).

For half precision GEMM (HGEMM), shown in Figure 2a, cuBLAS HGEMM with Tensor Cores achieves a maximum performance of 96.3 TFLOPS — approximately 85% the peak performance — and over 3.4× that of cuBLAS without the use of TCUs. For mixed precision GEMM (MGEMM), shown in Figure 2b, a maximum performance of 85.8 TFLOPS is achieved on NVIDIA TCUs using cuBLAS, approximately 76% the peak performance, for a 6.2× speedup over cuBLAS without Tensor Cores (the degradation of performance compared to HGEMM is due to output bytes count being twice as large). CUTLAS MGEMM is more performant than HGEMM, this is due to compiler and hardware optimizations for mixed precision that are absent from half precision [62].

2.4 GEMV Evaluation

The order of magnitude speedup of GEMM with TCU raises the question: can we formulate other algorithms in terms of matrix multiplication and also benefit from the TCU? The most obvious algorithm is matrix-vector multiplication (GEMV). We implement HGEMV (half precision GEMV) and MGEMV (mixed-precision GEMV) using cuBLAS HGEMM or MGEMM with dimension $\langle M, N, K = 16 \rangle$.This method wastes at least $15N$ memory loads and performs $15MN$ extra

---

[2]The mapping between individual matrix elements to their residing thread(s) is purposely opaque [47] and undocumented. We discuss how we alleviate some of the constraints in Section 6.1.
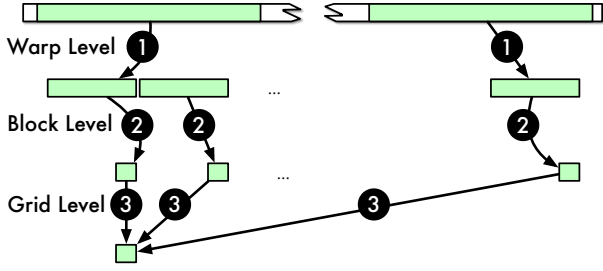
Figure 4: The reduction algorithm is ❶ composed of warp-level reduction that reduces each segment and is used to ❷ implement block-level reduction that further reduces each segment of partially reduced values. The partially reduced values are reduced across the grid ❸ to perform full reduction.

flops. We evaluate our implementations against cuBLAS SGEMV, since half precision GEMV is not present within cuBLAS.

Figure 3 shows that even when accounting for both resource and computation waste, HGEMV, implemented using cuBLAS HGEMM with Tensor Cores, outperforms cuBLAS SGEMV by at least $2\times$ and saturates at 900 GFLOPS due to the HBM2 global memory bandwidth. Naïve [3] HGEMV and MGEMV are super imposed atop each other since the overhead of using mixed-precision is dwarfed by the inefficient memory access. Both naïve versions still outperform cuBLAS' SGEMV for large inputs.

The GEMV evaluation shows that the performance of matrix multiplication on NVIDIA TCUs is high enough to tolerate resource and computation waste in algorithms. Driven by this observation, we examine how to formulate two widely used primitives — reduction and scan — to utilize TCUs.

3 REDUCTION AND SCAN ON GPUS

We start by defining reduction and scan. Reduction (also called *fold* or *total*) of a vector $A = [a_1, a_2, \ldots, a_n]$ is defined by its sum $\Sigma_{i=1}^{n} a_i$. Segmented reduction is defined as reductions on subsets of the input vector. In a regular segmented reduction, all segments are the same size[4]. The *scan* operation (also called *prefix sum*) for the same vector A is defined by the vector $[a_1, a_1 + a_2, \ldots, \Sigma_{i=1}^{n} a_i]$. Segmented scan is defined similarly to segmented reduction.

```
1   __device__ half warp_reduce(half val) {
2     for (int offset=WARP_SIZE/2; offset>0; offset/=2)
3       val += __shfl_down_sync(0xFFFFFFFFU, val, mask);
4     return val; }
5   __device__ half warp_scan(half val) {
6     for (int offset=1; offset<WARP_SIZE; offset*=2) {
7       auto n = __shfl_up_sync(0xFFFFFFFFU, val, mask);
8       if (laneid >= offset) val += n; }
9     return val; }
```

Listing 2: NVIDIA's recommended warp-level reduction and scan implementations utilizing shuffle instructions.

---

[3]Note that one implicitly performs tiling when utilizing the WMMA API.
[4]Irregular segmented reduction is implemented in terms of regular segmented reduction by padding the input.

3.1 State-of-the-art Implementations

For GPUs, state of the art libraries [1, 12, 38] implement both reduction and scan in terms of warp-, block-, and device-level collectives, as illustrated in Figure 4. The warp-level are commonly implemented using shuffle instructions [28], shown in Listing 2, which allows threads within a warp to share values via registers without synchronization or using shared memory. Shuffle instructions can be a bottleneck due to their limited throughput, however. For example, on the NVIDIA Volta architecture only 32 warp shuffle operations can be performed within a clock cycle per SM.

4 TCU REDUCTION ALGORITHM

Intuitively, reduction can be implemented using TCUs by representing it as a special case of matrix multiplication, since

$$Reduction[a_1, a_2, \ldots, a_n] = \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix} \cdot \begin{pmatrix} a_1 & a_2 & \cdots & a_n \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix}^T = \begin{pmatrix} \Sigma_{i=1}^{n} a_i & 0 & \cdots & 0 \\ 0 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 0 \end{pmatrix}$$

The challenge is to map generic input sizes onto the fixed matrix dimensions supported by the TCUs. For simplicity, this paper will assume that the TCU supports only matrices with 16×16 dimension. Other hardware may require other dimensions and those can be used without modifying the core idea of the algorithms. The algorithms are also presented in a precision agnostic way.

We use $Reduction_K$ to represent a K regular segmented reduction — partial reductions of the input uniformly partitioned into K element subsets. We will use P to denote the matrix which has ones for the first row and zero otherwise (i.e. $p_{r,c} = \begin{cases} 1 & \text{if } r = 0 \\ 0 & \text{if } r \neq 0 \end{cases}$), and the notation $\mathbf{X}$ for a matrix where all elements are the constant value X.

To make our formulation non-NVIDIA WMMA API specific, we present our algorithms in an API neutral way. In the following sections, we use **LoadTile** in place of the load_matrix_sync which takes a memory address, a matrix layout (default is row-major), and stride (default is 16) as input. We abstract store_matrix_sync to make it addressable as if it were a linear array. We will also use the notation $A \cdot B + C$ to denote the mma_sync operation. This paper however uses the standard CUDA terminology for warp, block, and grid to explain the algorithms, since no other standard nomenclature exists. The warp, block and device used in this paper correspond to the three memory hierarchy levels: L-Cache, PE/core, and device.

4.1 L-Cache (Warp)-level Reduction

We introduce warp-level reduction first, since it is the building block for both block- and grid-level reductions. We formulate reduction using TCUs for segment sizes 16, 256, and multiples of 16 and 256. Support for arbitrary segment sizes can be supported either by padding the input with zeros or by masking the P matrix. We find that padding introduces minimal overhead and is required in some cases to maintain the memory alignment imposed by the TCU APIs.

Segment Size 16*:* The $Reduction_{16}$ algorithm, shown in Algorithm 1 and Figure 5, performs warp-level reduction on 256 elements which represent 16 segments of size 16. On Line 3 in Algorithm 1 or Step ❶ in Figure 5, the data is loaded from memory into a column-major order fragment (matrix A). Each row is then reduced using
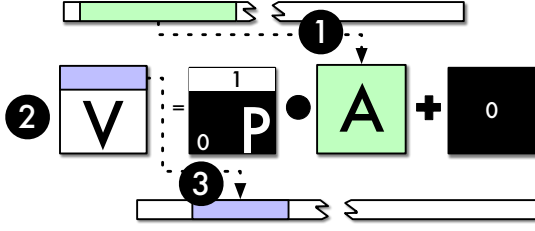
Figure 5: The $Reduction_{16}$ algorithm **❶** **each warp loads** 256 **elements into the matrix** A **in column major order from the input vector,** **❷** **performs the TCU operation where the** P **matrix has ones for the first row, and then** **❸** **the result, which is in the first row of** V**, is stored into the output vector.**

$V = P \cdot A$ (Line 4 or Step **❷**). The result — first row of V — is stored in the output memory (Line 5 or Step **❸**).

Algorithm 1 The $Reduction_{16}$ algorithm.

1: Initialize P matrix.
2: $idx \leftarrow$ **global offset**
3: $A \leftarrow$ **LoadTile** $\left(in\left[idx \ldots idx + 256\right], \text{"colmajor"}\right)$
4: $V \leftarrow P \cdot A + \underline{\mathbf{0}}$
5: **if** $laneIdx < 16$ **then** $out\left[\frac{idx}{16} + laneIdx\right] \leftarrow V\left[laneIdx\right]$

Segment Size 256: For handling segments of size 256, one follows a pattern similar to $Reduction_{16}$. The algorithm is shown in Algorithm 2 and is a single iteration of the algorithm illustrated in Figure 6. First, all 256 elements are loaded onto the TCU (Line 3). The rows are reduced using the same procedure as $Reduction_{16}$ (Line 2-4) the resulting columns are reduced using P^T (Line 5) before we store the scalar result (Line 6) into memory.

Algorithm 2 The $Reduction_{256}$ algorithm.

1: Initialize P matrix
2: $idx \leftarrow$ **global offset**
3: $A \leftarrow$ **LoadTile** $\left(in\left[idx \ldots idx + 256\right], \text{"colmajor"}\right)$
4: $V \leftarrow P \cdot A + \underline{\mathbf{0}}$
5: $V \leftarrow V \cdot P^T + \underline{\mathbf{0}}$
6: **if** $laneIdx = 0$ **then** $out\left[\frac{idx}{256}\right] \leftarrow V\left[0\right]$

Segment Size Multiples of 256: With the above $Reduction_{16}$ and $Reduction_{256}$ warp-level primitives, we can express segments that are multiples of either 16 (denoted by $16N$) or 256 (denoted by $256N$). We will first look at the $256N$ algorithm, since it will be used for representing the $16N$ algorithm.

A naïve way is to implement the $256N$ segmented reduction as N-repeated applications of the $Reduction_{256}$, shown in Figure 6. While this is correct, it is work inefficient — wasting one matrix multiplication for each iteration. Instead of performing two reductions in each iteration, we can implement a work efficient $256N$ segmented reduction by first reducing each row of the 16×16 matrix ($Reduction_{16}$)

in each iteration and then using the row of reduced values as an accumulator. In the final iteration, the final row is reduced into a scalar. Figure 7 illustrates the work-efficient algorithm.

Segment Size Multiples of 16: Similar to $Reduction_{256}$, segmented reduction where the segment size is multiples of 16 ($16N$) can be performed in two ways. The first is a strided segmented reduction, shown in Figure 8 (for the case where $N = 2$). During each iteration i, a warp loads 16 segments (each of length 16) into the matrix A with a stride of $16N$ (Steps **❶** and **❹**), i.e., the beginning of each 16-element segment is 16N elements away from the beginning of the next segment in the original input vector. The 16 columns of A are then reduced and accumulated into the first row of V (Steps **❷** and **❺**). This repeats for N iterations. This method is simple, works for arbitrary multiple of 16, and leverages GPU cache for small N. For large N this method suffers from bad locality.

Algorithm 3 makes better use of cache locality and reduces uncoalesced memory accesses. The algorithm implements $Reduction_{16N}$ in terms of $Reduction_{256N}$ for $N > 256$. The left over, $Reduction_{(N\%16) \times 16}$, can be implemented using the strided segmented $16N$ reduction method.

4.2 PE/Core (Block)-level Reduction

When the segment size is large, collaborative reduction within a block becomes profitable. We follow standard practice [38] to implement block-level reduction, but differ in that we still use the TCU to perform reduction on the partially reduced values within a block.
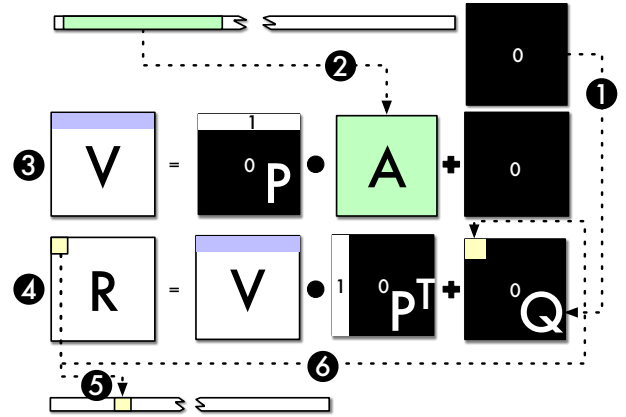
Figure 6: The work-inefficient $Reduction_{256N}$ algorithm **❶** **initializes the** Q **matrix with all zeros and** **❷** **loads the** 256 **input elements into a matrix** A **in column major order.** **❸** **A dot product** $V = P \cdot A + \underline{\mathbf{0}}$ **where the** P **matrix has the first row as ones and the rest of the values are zeros is performed to reduce each row into a scalar.** **❹** **the dot product** $R = V \cdot P^T + Q$ **reduces the first row into a scalar.** **❺** **If the segmented reduction size is equal to the matrix size (i.e.** $N = 1$**) or for the last iteration, then the first element of the** R **matrix is stored in the output array, otherwise** **❻** **the first element of** R **is used as the first element of the** Q **matrix and the procedure is iterated starting from step** **❷**.
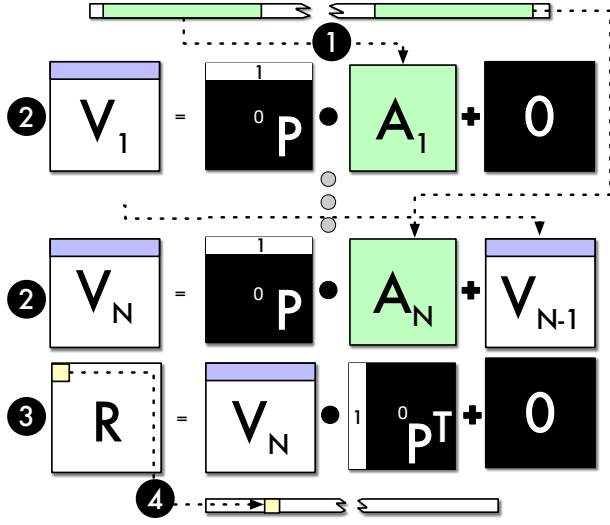
Figure 7: The work-efficient $Reduction_{256N}$ algorithm ❶ loads 256 **input elements into matrix** A_i **in each iteration. It then** ❷ **performs a matrix multiplication** $V_i = P \cdot A_i + V_{i-1}$ **for** i **between** 1 **and** N **with** $V_0 = \underline{0}$**. The final vector is reduced** ❸ **by performing the** $R = V_N \cdot P^T$ **operation and the** ❹ **result stored as output.**
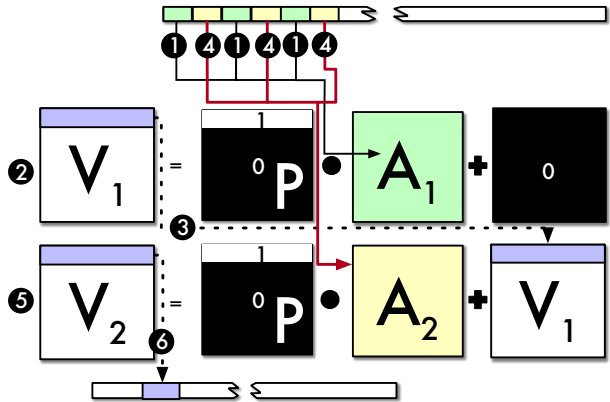

Figure 8: A strided $Reduction_{16N}$ algorithm for $N = 2$ ❶ loads 256 **elements where the stride between each row is** 16N**.** ❷ **We then perform the matrix multiplication** $V_1 = P \cdot A_1$ **and** ❸ **use the** V_1 **matrix as an accumulator for the next iteration where** ❹ **we again load the next** 256 **elements with the leading dimension set to** 16N**. The** ❺ **matrix multiplication** $V_2 = P \cdot A_2 + V_1$ **is performed and** ❻ **the first row is stored in the output vector.**

Algorithm 4 shows how warp-level reduction is used to implement the block-level $Reduction_{256N}$ kernel.

4.3 Device (Grid)-level Reduction

When the segment size is very large a grid-level reduction might be needed. A naïve grid-level reduction for a list of length N involves two kernel launches. The first kernel launch performs a segmented

---

Algorithm 3 The Coalesced $Reduction_{16N}$ algorithm.

1: Initialize P matrix
2: $V \leftarrow \underline{0}$
3: $gidx \leftarrow$ **global offset**
4: $numSegs \leftarrow \lfloor \frac{16N}{256} \rfloor$ ▷ Number of 256 segments
5: **for** $i \leftarrow 0; i < numSegs; i \leftarrow i + 1$ **do**
6:      $idx \leftarrow gidx + 256i$
7:      $A \leftarrow$ **LoadTile** $\left(in \left[idx \ldots idx + 256 \right], \text{"colmajor"} \right)$
8:      $V \leftarrow P \cdot A + V$
9:      $\ldots$ ▷ Reduce rest of segments using $StridedReduction_{16}$
10: **if** $laneIdx < 16$ **then** $out \left[\frac{gidx}{16N} + laneIdx \right] \leftarrow V \left[laneIdx \right]$

Algorithm 4 The Block-level $Reduction_{256N}$ algorithm.

1: $wpb \leftarrow$ **warps per block**
2: $prtls \leftarrow$ **alloc shared mem** $\left[wpb \right]$
3: $partial \leftarrow Reduction_{256 \frac{N}{wpb}} \left(in \right)$
4: **if** $laneIdx = 0$ **then** $prtls \left[warpIdx \right] \leftarrow partial$
5: **sync threads**
6: **if** $warpIdx = 0$ **then** $out \left[blockIdx \right] \leftarrow Reduction_{wpb} \left(prtls \right)$

---

reduction with the output stored in a list of partials. A second kernel then reduces the partials into a scalar. Although this algorithm is naïve, its performance is on par with the fastest algorithm.

5 TCU SCAN ALGORITHM

It might be less intuitive to represent scan as matrix multiplication. For a vector V of 256 elements, we can store it in row-major order within a 16×16 matrix A — with $a_{i,j} = V \left[16 \left(j - 1 \right) + i \right]$.

$$A = \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,16} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,16} \\ \vdots & \vdots & \ddots & \vdots \\ a_{16,1} & a_{16,2} & \cdots & a_{16,16} \end{pmatrix}$$

We notice that a row-wise scan can be obtained by multiplying the matrix A with an upper diagonal matrix — with the values of the upper diagonals being 1 and the rest 0.

$$RowScan \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,16} \\ a_{2,1} & a_{2,2} & \cdots & a_{2,16} \\ \vdots & \vdots & \ddots & \vdots \\ a_{16,1} & a_{16,2} & \cdots & a_{16,16} \end{pmatrix} = A \cdot U = A \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 0 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix} = \begin{pmatrix} a_{1,1} & \cdots & \Sigma_{i=1}^{16} a_{1,i} \\ a_{2,1} & \cdots & \Sigma_{i=1}^{16} a_{2,i} \\ \vdots & \ddots & \vdots \\ a_{16,1} & \cdots & \Sigma_{i=1}^{16} a_{16,i} \end{pmatrix}$$

Similarly, to get the scan of each column one can use a lower diagonal matrix. We use a strict lower diagonal, i.e. the diagonal is 0, to get an exclusive scan of each column.

$$ExclusiveColumnScan \begin{pmatrix} a_{1,1} & a_{1,2} & \cdots & a_{1,16} \\ \vdots & \vdots & \ddots & \vdots \\ a_{16,1} & a_{16,2} & \cdots & a_{16,16} \end{pmatrix} = L \cdot A =$$

$$\begin{pmatrix} 0 & 0 & \cdots & 0 \\ 1 & 0 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 0 \end{pmatrix} \cdot A = \begin{pmatrix} 0 & 0 & \cdots & 0 \\ a_{1,1} & a_{1,2} & \cdots & a_{1,16} \\ a_{1,1}+a_{2,1} & a_{1,2}+a_{2,2} & \cdots & a_{1,16}+a_{2,16} \\ \vdots & \vdots & \ddots & \vdots \\ \Sigma_{j=1}^{15} a_{j,1} & \Sigma_{j=1}^{15} a_{j,2} & \cdots & \Sigma_{j=1}^{15} a_{j,16} \end{pmatrix}$$

We then use the $L \cdot A$ matrix to create a G matrix where each element $G_{j,i}$ is the reduction of the j^{th} row of $L \cdot A$. That is, all elements in the j^{th} row of G are of the same value — the sum of all elements preceding the j^{th} row of A, i.e. $G_{j,i} = \Sigma_{k=1}^{j-1} \Sigma_{i=1}^{16} A_{k,i}$. The G matrix can be generated by multiplying $L \cdot A$ with a matrix with

all element values set to 1. We then add G to the $A \cdot U$ matrix to generate the scan of V — which is read in linear row-major order.

$$Scan(V) = L \cdot A \cdot \begin{pmatrix} 1 & 1 & \cdots & 1 \\ 1 & 1 & \cdots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \cdots & 1 \end{pmatrix} + A \cdot U = G + A \cdot U =$$

$$\begin{pmatrix} a_{1,1} & a_{1,1}+a_{1,2} & \cdots & \sum_{i=1}^{16} a_{1,i} \\ a_{2,1}+\sum_{i=1}^{16} a_{1,i} & a_{2,1}+a_{2,2}+\sum_{i=1}^{16} a_{1,i} & \cdots & \sum_{j=1}^{2}\sum_{i=1}^{16} a_{j,i} \\ \vdots & \vdots & \ddots & \vdots \\ a_{16,1}+\sum_{j=1}^{15}\sum_{i=1}^{16} a_{j,i} & a_{16,1}+a_{16,2}+\sum_{j=1}^{15}\sum_{i=1}^{16} a_{j,i} & \cdots & \sum_{j=1}^{16}\sum_{i=1}^{16} a_{j,i} \end{pmatrix}$$

Throughout this section we will use U to represent the upper diagonal matrix where the upper diagonal values are one, and use L to represent the strict lower diagonal matrix where the values below the lower diagonal are one — i.e. $U_{r,c} = \begin{cases} 1 & \text{if } r >= c \\ 0 & \text{if } r < c \end{cases}$

and $L_{r,c} = \begin{cases} 1 & \text{if } r < c \\ 0 & \text{if } r >= c \end{cases}$.

5.1 L-Cache (Warp)-level Scan

With the above derivation, we follow a similar structure to Section 4: first introducing warp-level primitives before presenting the block- and grid-level primitives. We write $Scan_K$ to represent a K regular segmented scan. Since the process of building warp-level, block-level, and grid-level scans from $Scan_K$ is very similar to that of reduction, we will only highlight the key differences.

Segment Size 16: Is the *RowScan* equation above and is illustrated in Figure 9 as steps **1**, **2**, and **3**.

Segment Size 256: Is implemented using 3 matrix multiplications shown in Figure 9 and presented mathematically above.

Segment Size Multiples of 16: Is similar to strided $16N$ reduction, with the key difference being that we broadcast the last column rather than the reduced scalar value and is shown in Algorithm 5.
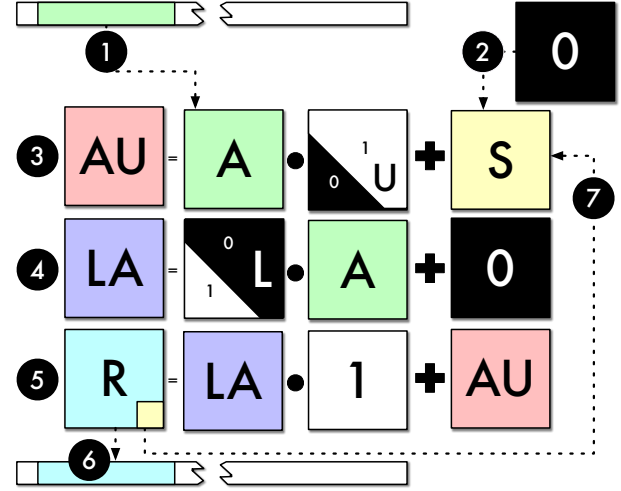

Algorithm 5 The $Scan_{16N}$ algorithm.

1: Initialize U matrix.
2: $gidx \leftarrow$ **global offset**
3: $S \leftarrow \mathbf{0}$
4: $lid \leftarrow laneIdx$
5: **for** $i \leftarrow 0; i < N; i \leftarrow i+1$ **do**
6:     $idx \leftarrow gidx + 16i$
7:     $A \leftarrow$ **LoadTile** $(in\,[idx \ldots idx+256N]\,, stride = 16N)$
8:     $R \leftarrow A \cdot U + S$
9:     $S \leftarrow$ **Broadcast** $(\textbf{LastColumn}\,(R))$
10:     **if** $lid < 16$ **then**
11:        $oi \leftarrow idx + lid * 16N$
12:        $out\,[oi \ldots oi+16] \leftarrow R\,[16lid \ldots 16lid+16]$

Segment Size Multiples of 256: Only a small modification to $Scan_{256}$ is needed to implement $Scan_{256N}$ and is illustrated in Figure 9 and Algorithm 6. Line 11 in Algorithm 6 shows that we keep track of the sum (last element of the R matrix) and broadcast it to the S matrix after each iteration. The S matrix is then used when performing subsequent iterations.

Figure 9: The $Scan_{256N}$ algorithm **1** loads 256 elements from the input vector into a matrix A and **2** initializes the S matrix to $\mathbf{0}$. The **3** $AU = A \cdot U + S$ and **4** $LA = L \cdot A + \mathbf{0}$ matrix multiplications are performed to compute the prefix sum of each row and column. **5** A row wise reduction is performed on the LA and added to the AU matrix. **6** The result R is stored in the output vector. **7** If the segment size is a multiple of 256, then the last element of R (position $16, 16$) is broadcasted into the S matrix and the procedure is repeated.

Algorithm 6 The $Scan_{256N}$ algorithm.

1: Initialize U and L matrices.
2: $gidx \leftarrow$ **global offset**
3: $S \leftarrow \mathbf{0}$
4: **for** $i \leftarrow 0; i < N; i \leftarrow i+1$ **do**
5:     $idx \leftarrow gidx + 256i$
6:     $A \leftarrow$ **LoadTile** $(in\,[idx \ldots idx+256])$
7:     $AU \leftarrow A \cdot U + S$
8:     $LA \leftarrow L \cdot A + \mathbf{0}$
9:     $R \leftarrow LA \cdot \mathbf{1} + AU$
10:     $out\,[idx \ldots idx+256] \leftarrow R$
11:     $S \leftarrow$ **Broadcast** $(R\,[255])$

5.2 PE/Core (Block)-level Scan

Algorithm 7 shows how to perform the scan at the block level. It first computes the segmented scan using the warp primitives (Line 8-13), stores the reduced values into a partials list (Line 16), performs a scan on the partial list (Line 17), and adds the values to the intermediate results to get the output (Line 19-23).

Algorithm 7 also exercises the TCU to perform the scan on the partially reduced values across tiles. On Line 16 we use the offset of the last row (240) and 256 as the leading dimension when loading the tile. This loads the last row of R across tiles into E. Line 17 then

Algorithm 7 The Block-level $Scan_{256N}$ algorithm.

1: Initialize U and L matrices.
2: $gidx \leftarrow$ **global offset**
3: $wpb \leftarrow$ **warps per block**   ▷ Assumed to be less than 16
4: $sout \leftarrow$ **alloc shared mem** $[256 \times 16]$
5: $prtls \leftarrow$ **alloc shared mem** $[16]$ ▷ Partial sums
6: $S \leftarrow \underline{\mathbf{0}}$
7: **for** $i \leftarrow 0; i < N; i \leftarrow i + warpsPerBlock$ **do**
8:     $idx \leftarrow gidx + 256\,(i + warpIdx)$
9:     $A \leftarrow$ **LoadTile** $(in\,[idx \ldots idx + 256])$
10:     $AU \leftarrow A \cdot U + S$
11:     $LA \leftarrow L \cdot A + \underline{\mathbf{0}}$
12:     $R \leftarrow LA \cdot \underline{\mathbf{1}} + AU$
13:     $sout\,[256warpIdx \ldots 256warpIdx + 256] \leftarrow R$
14:     **sync threads**
15:     **if** $warpIdx = 0$ **then**
16:         $E \leftarrow$ **LoadTile** $(sout\,[240 \ldots 4096]\,, stride = 256)$
17:         $prtls \leftarrow LastColumnScan_{16}\,(E)$ ▷ Exclusive scan
18:     **sync threads**
19:     **for** $j \leftarrow 1; j \leq 256; j \leftarrow j + warpSize$ **do**
20:         $it \leftarrow j + laneIdx$
21:         $val \leftarrow sout\,[256warpIdx + it] + prtls\,[warpIdx]$
22:         $out\,[idx + it] \leftarrow val$
23:     $S \leftarrow$ **Broadcast** $(prtls\,[15])$

performs an exclusive scan on the last column of the E and stores the results into the list of partials[5].

5.3 Device (Grid)-level Scan

Similar to reduction, the segmented scan is used as a building block for the grid-level scan. The grid-level scan uses a text book implementation, scan-then-propagate strategy, and involves 3 kernel calls. The first kernel uses segmented scan and produces partially reduced values for each block. The second kernel performs a scan on the partially reduced values. The third kernel then uniformly adds the partially scanned values to their corresponding segments.

6 EVALUATION

We implemented the algorithms presented in Sections 4 and 5 using NVIDIA's WMMA API. The code (available at https://github.com/c3sr/tcu_scope) is implemented as a C++ header library with an API similar to CUB's — providing functions such as SegmentedReduce, Reduce, SegmentedScan, and Scan. We employ auto-tuning to select the ideal algorithm, number of warps (or independent TCU operations) per block, coarsening factor (the number segments to perform within a warp), and block dimensions for the user-provided segment size.

We evaluate our implementation on an Intel Xeon E5-2698 with CentOS 4.3, CUDA Driver 396.26, and CUDA Runtime 9.2.88 installed. We use the Tesla V100-PCIE GPU with 16GB of GPU HBM2 memory and a theoretical peak bandwidth of $900GBs$ or 450

---

[5] The implementation of $LastColumnScan_{16}$ is performed by loading the last column values into the first row and performing an TCU version of the exclusive scan algorithm. Formulating the intermediate operation this way is needed to adhere to the CUDA WMMA API's byte alignment constraint for loading fragments.

---

billion half precision elements per second. All the results below show the throughput of the algorithms in terms of billions of half precision elements per second.

6.1 Relaxing the WMMA API Constraints

Constraints arise when using the current WMMA API for non-GEMM computation. These limitations would not exist if one is to perform just GEMM computation. The constraints observed were:

(1) Loads or stores must be performed at fragment granularity.
(2) Loading and storing fragments can only be performed using global or shared memory; constant memory cannot be used.
(3) The memory layout for the matrix kinds are not the same and their is no API to perform casts between them.

We address these limitations in different ways within our implementation. For (1) and (2) we use knowledge about the internal layout of the fragment [26] and implemented WMMA API enhancements tailored to our usage. Listing 3 shows an example of our API extensions for operating on partial fragments.

```
using frag_b = fragment<matrix_b, 16, 16, 16, half, row_major>;
__device__ int matrix_b_get_row_idx() {
    const int laneIdx = threadIdx.x % warpSize;
    return laneIdx&0x10 >> 2 + laneIdx&0x0B;
}
__device__ void matrix_b_set_upper_triangular(frag_b &f) {
#pragma unroll
    for (int ii = 0; ii < f.num_elements; ii++)
        f.x[ii] = matrix_b_get_row_idx() < ii ? 0.0f : 1.0f; }
__device__ void matrix_b_get_first_column(half* out, frag_b f) {
    const int laneid = threadIdx.x % warpSize;
    if (laneid & 0x04) return ; // avoid redundant writes
    out[matrix_b_get_row_idx()] = f.x[0];
}
```

Listing 3: The WMMA API can only perform load/store from shared or global memory and lacks the ability to fill an TCU fragment from constant memory or operate on sub-fragments. This code shows how we enhance the NVIDIA WMMA API, using knowledge of the fragment layout, to create an upper triangular matrix and get the first column of a fragment for the `matrix_b` fragment kind.

Although we can use the layout information to shuffle registers to address (3), we opt instead to express the cast in terms of load/store APIs available through the WMMA API. For example, to cast a matrix in the `matrix_a` format to `matrix_b` format, we first store the matrix into shared memory and then perform a load from memory to `matrix_b`. Using our API extensions for fragment layout information requires less block synchronization — which increases the performance of our implementation by up to 5%. Since relying on fragment layout information is not portable, we omit these results.

6.2 Optimizing CUB for Half Precision

CUB is a C++ template library that contains multiple algorithms for the collectives. The CUB library contains the fastest [40, 41] implementation for the reduction and scan collectives and is used by libraries such as Thrust [5] as well as most deep learning frameworks [8, 42, 43, 59, 67]. We compare against the latest release of CUB [38] (version 1.8) and evaluate against different parameters of the collectives. As an optimization, warp-level shuffle reduction and scan are implemented in PTX within CUB for integer, float, and double data types, since NVCC is currently unable to use the
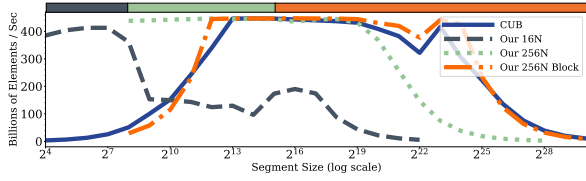
Figure 10: We evaluate the segmented reduction for the algorithms presented on different segment sizes (between 16 and 2^{30}) for a fixed 2^{30} element list. Through a combination of the algorithms presented, for the range between 16 and 2^{24} we are able to achieve throughput within 90% and 98% of ideal throughput (the theoretical peak is 450 billion half precision elements per second). The bar on top of the figure shows the best performing algorithm for each range of segment sizes.

shuffle instruction's predicate to guard against invalid peers [28, 44]. We observerved that CUB does not contain these shuffle-based optimizations for half precision. To make the evaluations fair and meaningful, we implement these optimization for the half precision data type in CUB. The modified CUB is used for the evaluation to provide a more aggressive base of comparison.

6.3 Warp- and Block-level Reduction and Scan

Theoretically (on V100) our warp-level TCU reduction algorithms require less than one fourth of the cycles of the warp-level shuffle reduction. For example, consider performing a warp-level $Reduction_{256}$: the warp-level reduction shown in Listing 2 requires 8 iterations of 32 element reduction to reduce each segment. The total cycles is therefore 256, since each shuffle instruction and addition takes 4 cycles. Our algorithm performs the reduction using two matrix multiplications or 64 cycles — since each TCU WMMA matrix multiplication requires 32 cycles. However, reduction is known to be memory bound, with the ideal performance bounded by memory copy speed.

We evaluate the TCU segmented reduction algorithms against `cub::DeviceSegmentedReduce::Sum` by fixing the number of input elements and varying the segment size (Figure 10). When the segment size is less than 256, the $16N$ algorithm is used. The $16N$ algorithm's performance degrades for large segment sizes due to its strided access pattern resulting in uncoalesced global memory access. When the segment size is larger than 256, the $256N$ algorithm is used, but again suffers from performance degradation after segment size 2^{15} due to low occupancy. When the segment size is large (greater than 2^{15}) the block-level $256N$ reduction is used. Figure 10 shows that our TCU implementation achieves more than 90% of the peak throughput for variable segment size and is always better than CUB.

When the segment size is large and the number of segments is small, the performance of both CUB and our implementation drops. Since each segment gets mapped onto a block, a small number of segments causes some SMs to be idle. For example when segment size is 2^{25}, both CUB and our implementation achieve an occupancy of around 0.25 and SM efficiency of around 40%. A better strategy for these cases would be to assign multiple thread blocks to

collaboratively reduce each segment when the size of the segments is very large. This optimization can be achieved using CUDA 9's cooperative groups [45], but is outside the focus of this paper.

Our TCU implementation largely outperforms CUB's device-wide segmented reduction for different segment size. Through profiling, we identified the major predictors of performance to be, in the order of importance, the number of half-precision floating-point instructions (`inst_fp_16` in the NVProf [51] metrics), warp instructions (`inst_inter_thread_communication`), and integer instructions (`inst_integer`). We consistently find that our implementation's half-precision instructions is approximately equal to the number of total elements (2^{31}) while CUB's is much larger. Moreover, CUB requires large number of integer and warp shuffle instructions while our implementation uses no warp shuffle instructions and a smaller number of integer instructions. This contributes to the $100\times$ speedup for segment size 16.

We examined the power consumption by measuring the average power draw within the execution phase of the kernel using NVProf. Based on these measurements, we find that our implementation consumes $7.4 - 22.3\%$ less power compared to CUB across different segment sizes. Again, this is because of the efficient use of the $FP16$ and INT ALUs as well as better SM and DRAM utilization. We note that our algorithm leaves the general purpose ALUs idle, allowing less contention on these units.

CUB provides a `cub::WarpReduce`, applicable for segment sizes 16 and 32, to compute a parallel reduction of elements within a warp. CUB also provides `cub::BlockReduce` to perform reduction within a block. These two primitives require users to partition the data and construct the kernel. Since CUB's device-wide segmented reduction does not perform well for segment size smaller then 2^{13}, we evaluate our TCU implementations against `cub::WarpReduce` and `cub::BlockReduce` implementations, shown in Figure 11. The `cub::WarpReduce` implementation is tunable on block size, wheras the `cub::BlockReduce` implementation is tunable on block size, thread coarsening factor, and reduction algorithms. We compare our implementation against the best CUB implementation. We find that our TCU implementations is still faster for segment size smaller than 1024, and is comparable to `cub::BlockReduce` for the other cases.

For segmented scan, we evaluate the TCU algorithms against Thrust's implementation (`inclusive_scan_by_key`), since CUB has no user visible API for segmented scan. The Thrust implementation utilizes CUB's internal warp- and block-level scan to implement the scan-by-key operation. We evaluate different segment sizes with a fixed number of input elements — the results are shown in Figure 12. Thrust, consistent with previous work [17], has constant performance irrespective of the segment size. Whereas, our scan TCU implementations achieve more than 89% of the peak throughput and is $3\times$ faster than thrust for small segment sizes. We observe lower power consumption compared to Thrust — observing it to be either equivalent in power usage or up to 17% less. Our segmented scan is not ideal for large segment sizes since, as explained in Section 4, only a small number of blocks get launched and thus the GPU is underutilized. This inefficiency can be remedied using the same strategy described for reduction.

CUB provides `cub::WarpScan` to compute a parallel scan of data partitioned within a warp, and `cub::BlockScan` within a block.
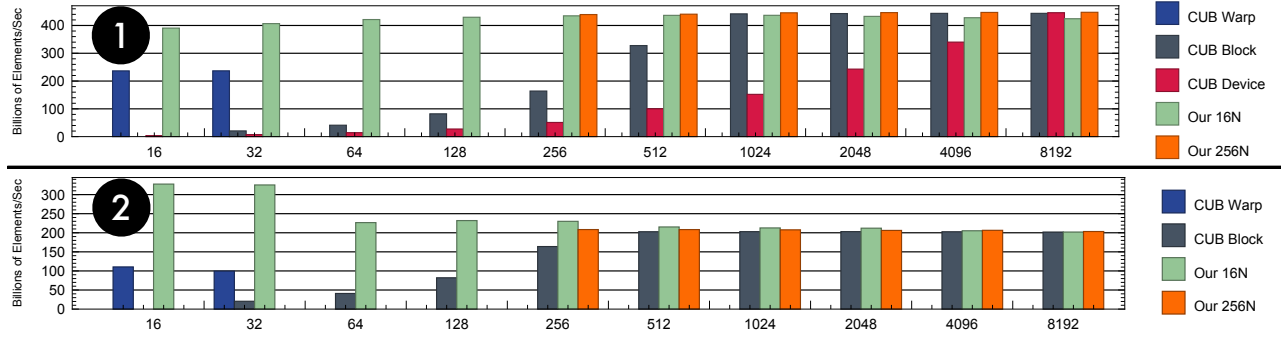
Figure 11: Segmented ❶ reduction and ❷ scan are evaluated in terms of billions of half-precision elements per second (*y*-axis) for segment sizes between 2^4 and 2^{13} (*x*-axis). The best configurations for our implementation as well as CUB are selected.
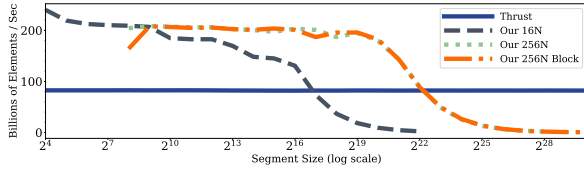

Figure 12: We evaluate the segmented scan for the algorithms presented on different segment sizes for a fixed 2^{31} element list. Through a combination of the algorithms presented, for the range between 16 and 2^{19} we are able to achieve throughput within 89% and 97% of ideal throughput (the theoretical peak is 225 billion half precision elements per second).
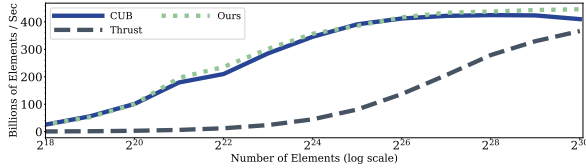

Figure 13: A full reduction implementation based on the description in Section 4 achieves performance on par to CUB.

Similar to reduction, these two primitives require more programming effort from the users to partition the data and construct the kernel. The CUB scan implementations have the same tunable parameters as CUB's reduction. We evaluate our TCU segmented scan against the best `cub::WarpScan` and `cub::BlockScan` parameters, shown in Figure 11. We can see that our TCU implementations are still faster for small segment size, and are at least comparable to `cub::BlockScan` for other cases.

6.4 Grid-level Reduction and Scan

Unlike the warp- and block-level operations, this paper does not attempt to optimize grid-level operations — opting to use a naïve

implementation for the grid-level collectives. The naïve implementation involves multiple kernel launches. We include the evaluation results to show that even our naïve grid-level implementation achieves performance that is better or comparable to that of CUB and Thrust.
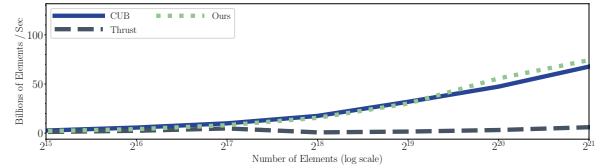

Figure 14: A full scan implementation based on the description in Section 5 achieves performance comparable to CUB.

We compare against both CUB and Thrust for full reduction (Figure 13), and scan (Figure 14). For both cases, our implementation uses the 256*N* block-level algorithm. Even with our naïve grid-level implementation, we are able to mach the performance of CUB and are considerably faster than the Thrust implementation. For reduction and scan, the TCU implementation is slightly faster than CUB with large input sizes being bounded by memory bandwidth and is within 98% (for reduction) of peek memory copy bandwidth. For scan, our current unoptimized implementation uses CUB to reduce the partial values (kernel 2 described in Section 5.3). Future implementations would not use CUB, since it fails for inputs larger than 2^{21} which causes our implementation to fail as well.

7 RELATED WORK

The mapping of algorithms onto matrix multiplication has been well studied [22, 30, 60, 68]. Similarly, both reduction and scan are well studied from a performance and application [6, 7, 21, 29, 65] aspect on a wide range of architectures and have been heavily evaluated across GPU architectures [14, 16, 37, 64, 70]. To our knowledge however, there has been no attempt at mapping either reduction or scan in terms of matrix multiplication.

Considerable research has been done on the development of performance portable compilers for reduction and scan kernels [2, 10, 13, 31, 66]. These compilers express the algorithms as systems of

alternative building blocks that are then composed and auto-tuned at compile time for both the target architecture and the input characteristics. These tools are synergistic with our technique, since we are able to add our algorithm as another building block to implement reduction or scan.

Previous work [11, 38, 39, 61, 70] has also shown that optimizations can be made to either avoid or hide the overhead of multi-kernel launches. These optimizations would enable our grid-level operations to be competitive for large sizes when compared to state-of-the-art methods. Other research looked at specific cases of scan, in [34] the authors look at performing scan on tuples while minimizing global reads and facilitating latency hiding.

Work describing NVIDIA's Tensor Cores is scarce. In [26], the authors use microbenchmarks to discern micro-architectural details of the V100 architecture. This work was extended in [62] where authors' study expand on the micro-architectural study and show a proposed NVIDIA TCU architecture. In [35, 20] the authors use half precision and TCUs to implement iterative solvers. They use half precision along with low quality solvers to compute the initial conditions and then switch to both higher precision solvers for subsequent iterations. The authors also examine the numerical error incurred when using TCUs and half-precision for HPC workloads.

8 CONCLUSION

This paper leveraged the Tensor Core Units (TCUs) (a specialized accelerator developed to optimize matrix multiplication for deep learning) to implement both reduction and scan. We showed a novel, simple, and efficient mapping of the reduction and scan primitives onto TCUs. We believe we are the first to formulate these algorithms to exercise the TCU. Unlike existing work which designs ASICs to map reduction and scan onto hardware, we develop an algorithmic solution to map both reduction and scan on existing TCUs. An algorithmic solution is relevant when using preexisting TCU designs (as is the case for the NVIDIA TCU). We also pointed out directions for future API and architectural changes to relax some of the TCU constraints such as loading fragments from constant, extracting single row or column, etc. — resulting in a simplified implementation.

We implemented the proposed algorithms onto V100 TCUs, achieved up to $100\times$ speedup for reduction and up to $3\times$ for scan, and showed performance that rivals state of the art implementation in the worst cases. We observed up to 22% less power consumption for reduction and 16% for scan using NVPROF. As a result of the algorithms, we were able to make use of the otherwise idle TCUs— enabling better GPU utilization for kernels that exercise the general purpose ALUs.

Future work would leverage the techniques described in this paper to map more algorithms and functions onto TCUs. We are specifically interested in transcendental and special functions, since the NVIDIA special function units have been observed to be the bottleneck in HPC applications. We also want to express neural network layers in terms of TCUs, where some layer implementations and layer fusion opportunities would be enabled by our work: such as the computation of variance in batch norm [23, 24] or the evaluation of special functions in activation layers.

REFERENCES

[1] Emmanuel Agullo, Jim Demmel, Jack Dongarra, Bilel Hadri, Jakub Kurzak, Julien Langou, Hatem Ltaief, Piotr Luszczek, and Stanimire Tomov. 2009. Numerical linear algebra on emerging architectures: the plasma and magma projects. In *Journal of Physics: Conference Series* number 1. Vol. 180. IOP Publishing, 012037.

[2] Jason Ansel, Cy Chan, Yee Lok Wong, Marek Olszewski, Qin Zhao, Alan Edelman, and Saman Amarasinghe. 2009. *PetaBricks: a language and compiler for algorithmic choice*. Number 6. Vol. 44. ACM.

[3] Apple. 2019 (accessed January 14, 2019). *A11 Bionic*. https://www.apple.com/iphone-x.

[4] Arm. 2019 (accessed January 14, 2019). *Arm Machine Learning Processor*. https://developer.arm.com/products/processors/machine-learning/arm-ml-processor.

[5] Nathan Bell and Jared Hoberock. 2011. Thrust: a productivity-oriented library for cuda. In *GPU computing gems Jade edition*. Elsevier, 359–371.

[6] Guy E Blelloch. 1989. Scans as primitive parallel operations. *IEEE Transactions on computers*, 38, 11, 1526–1538.

[7] Guy E Blelloch, Michael A Heroux, and Marco Zagha. 1993. Segmented operations for sparse matrix computation on vector multiprocessors. Tech. rep. Carnegie-Mellon Univ Pittsburgh PA School of Computer Science.

[8] Caffe2. 2019 (accessed January 14, 2019). *Caffe2*. https://caffe2.ai.

[9] Timothy M Chan. 2010. More algorithms for all-pairs shortest paths in weighted graphs. *SIAM Journal on Computing*, 39, 5, 2075–2089.

[10] Li-Wen Chang, Izzat El Hajj, Christopher Rodrigues, Juan Gómez-Luna, and Wen-mei Hwu. 2016. Efficient kernel synthesis for performance portable programming. In *The 49th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE Press, 12.

[11] Gaurav Chaurasia, Jonathan Ragan-Kelley, Sylvain Paris, George Drettakis, and Fredo Durand. 2015. Compiling high performance recursive filters. In *Proceedings of the 7th Conference on High-Performance Graphics*. ACM, 85–94.

[12] Leonardo Dagum and Ramesh Menon. 1998. Openmp: an industry standard api for shared-memory programming. *IEEE computational science and engineering*, 5, 1, 46–55.

[13] Simon Garcia De Gonzalo, Sitao Huang, Juan Gómez-Luna, Simon Hammond, Onur Mutlu, and Wen-mei Hwu. 2019. Automatic generation of warp-level primitives and atomic instructions for fast and portable parallel reduction on gpus. In *Proceedings of the 2019 IEEE/ACM International Symposium on Code Generation and Optimization*. IEEE Press, 73–84.

[14] Yuri Dotsenko, Naga K Govindaraju, Peter-Pike Sloan, Charles Boyd, and John Manferdelli. 2008. Fast scan algorithms on graphics processors. In *Proceedings of the 22nd annual international conference on Supercomputing*. ACM, 205–213.

[15] Zidong Du et al. 2017. An accelerator for high efficient vision processing. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 36, 2, 227–240.

[16] Martin Dybdal, Martin Elsman, Bo Joel Svensson, and Mary Sheeran. 2016. Low-level functional gpu programming for parallel algorithms. In *Proceedings of the 5th International Workshop on Functional High-Performance Computing*. ACM, 31–37.

[17] Marco Eilers. 2014. Multireduce and multiscan on modern gpus. *Department of Computer Science, University of Copenhagen. Master's thesis*.

[18] Google. 2019 (accessed January 14, 2019). *Edge TPU*. https://cloud.google.com/edge-tpu.

[19] Google. 2019 (accessed January 14, 2019). *Google Cloud TPU*. https://cloud.google.com/tpu.

[20] Azzam Haidar, Panruo Wu, Stanimire Tomov, and Jack Dongarra. 2017. Investigating half precision arithmetic to accelerate dense linear system solvers. In *Proceedings of the 8th Workshop on Latest Advances in Scalable Algorithms for Large-Scale Systems*. ACM, 10.

[21] Mark Harris, Shubhabrata Sengupta, and John D Owens. 2007. Parallel prefix sum (scan) with cuda. *GPU gems*, 3, 39, 851–876.

[22] Xiaohan Huang and Victor Y Pan. 1998. Fast rectangular matrix multiplication and applications. *Journal of complexity*, 14, 2, 257–299.

[23] Sergey Ioffe and Christian Szegedy. 2015. Batch normalization: accelerating deep network training by reducing internal covariate shift. *arXiv preprint arXiv:1502.03167*.

[24] X. Jia et al. 2018. Highly Scalable Deep Learning Training System with Mixed-Precision: Training ImageNet in Four Minutes. *ArXiv e-prints*, (July 2018). arXiv: 1807.11205.

[25] Yangqing Jia, Evan Shelhamer, Jeff Donahue, Sergey Karayev, Jonathan Long, Ross Girshick, Sergio Guadarrama, and Trevor Darrell. 2014. Caffe: convolutional architecture for fast feature embedding. *arXiv preprint arXiv:1408.5093*.

[26] Zhe Jia, Marco Maggioni, Benjamin Staiger, and Daniele P Scarpazza. 2018. Dissecting the nvidia volta gpu architecture via microbenchmarking. *arXiv preprint arXiv:1804.06826*.

[27] Norman P Jouppi et al. 2017. In-datacenter performance analysis of a tensor processing unit. In *Computer Architecture (ISCA), 2017 ACM/IEEE 44th Annual International Symposium on*. IEEE, 1–12.

[28] Julien Demouth. 2019 (accessed January 14, 2019). *Kepler Shuffle: Tips and Tricks*. http://on-demand.gputechconf.com/gtc/2013/presentations/S3174-Kepler-Shuffle-Tips-Tricks.pdf.

[29] Hee-Seok Kim, Shengzhao Wu, Li-wen Chang, and W Hwu Wen-mei. 2011. A scalable tridiagonal solver for gpus. In *Parallel Processing (ICPP), 2011 International Conference on*. IEEE, 444–453.

[30] Tamara G Kolda and Brett W Bader. 2009. Tensor decompositions and applications. *SIAM review*, 51, 3, 455–500.

[31] Rasmus Wriedt Larsen and Troels Henriksen. 2017. Strategies for regular segmented reductions on gpu. In *Proceedings of the 6th ACM SIGPLAN International Workshop on Functional High-Performance Computing*. ACM, 42–52.

[32] Lawrence Livermore National Laboratory. 2019 (accessed January 14, 2019). *Sierra Supercomputer*. https://computation.llnl.gov/computers/sierra.

[33] Chris Leary and Todd Wang. 2017. Xla: tensorflow, compiled. *TensorFlow Dev Summit*.

[34] Sepideh Maleki, Annie Yang, and Martin Burtscher. 2016. *Higher-order and tuple-based massively-parallel prefix sums*. Number 6. Vol. 51. ACM.

[35] Stefano Markidis, Steven Wei Der Chien, Erwin Laure, Ivy Bo Peng, and Jeffrey S Vetter. 2018. Nvidia tensor core programmability, performance & precision. *arXiv preprint arXiv:1803.04014*.

[36] Michael D McCool, Arch D Robison, and James Reinders. 2012. *Structured parallel programming: patterns for efficient computation*. Elsevier.

[37] Trevor L McDonell, Manuel MT Chakravarty, Gabriele Keller, and Ben Lippmeier. 2013. Optimising purely functional gpu programs. *ACM SIGPLAN Notices*, 48, 9, 49–60.

[38] D Merrill. 2018. CUB v1.8.0: CUDA Unbound, a library of warp-wide, block-wide, and device-wide GPU parallel primitives. *NVIDIA Research*.

[39] Duane G Merrill and Andrew S Grimshaw. 2010. Revisiting sorting for gpgpu stream architectures. In *Proceedings of the 19th international conference on Parallel architectures and compilation techniques*. ACM, 545–546.

[40] Duane Merrill and Michael Garland. 2016. Single-pass parallel prefix scan with decoupled look-back. Tech. rep. NVIDIA Technical Report NVR-2016-002.

[41] Bruce Merry. 2015. A performance comparison of sort and scan libraries for gpus. *Parallel Processing Letters*, 25, 04, 1550007.

[42] Rory Mitchell and Eibe Frank. 2017. Accelerating the xgboost algorithm using gpu computing. *PeerJ Computer Science*, 3, e127.

[43] MXNet. 2019 (accessed January 14, 2019). *MXNet*. https://mxnet.apache.org.

[44] Wilt Nicholas. 2013. The cuda handbook: a comprehensive guide to gpu programming. (2013).

[45] NVIDIA. 2019 (accessed January 14, 2019). *Cooperative Groups: Flexible CUDA Thread Programming*. https://devblogs.nvidia.com/cooperative-groups/.

[46] NVIDIA. 2019 (accessed January 14, 2019). *cuBLAS*. https://developer.nvidia.com/cublas.

[47] NVIDIA. 2019 (accessed January 14, 2019). *CUDA C Programming Guide*. https://docs.nvidia.com/cuda/cuda-c-programming-guide/index.html.

[48] NVIDIA. 2019 (accessed January 14, 2019). *cuDNN*. https://developer.nvidia.com/cudnn.

[49] NVIDIA. 2019 (accessed January 14, 2019). *CUTLASS*. https://devblogs.nvidia.com/cutlass-linear-algebra-cuda.

[50] NVIDIA. 2019 (accessed January 14, 2019). *Mixed Precision Training*. https://docs.nvidia.com/deeplearning/sdk/mixed-precision-training.

[51] NVIDIA. 2019 (accessed January 14, 2019). *NVPROF*. https://docs.nvidia.com/cuda/profiler-users-guide/index.html#system-profiling.

[52] NVIDIA. 2019 (accessed January 14, 2019). *NVVM IR Specification 1.5*. https://docs.nvidia.com/cuda/nvvm-ir-spec/index.html.

[53] NVIDIA. 2019 (accessed January 14, 2019). *Parallel Thread Execution ISA Version 6.2*. https://docs.nvidia.com/cuda/parallel-thread-execution/index.html.

[54] NVIDIA. 2019 (accessed January 14, 2019). *Programming Tensor Cores in CUDA 9*. https://devblogs.nvidia.com/programming-tensor-cores-cuda-9.

[55] NVIDIA. 2019 (accessed January 14, 2019). *Tensor Cores*. https://www.nvidia.com/en-us/data-center/tensorcore.

[56] NVIDIA. 2019 (accessed January 14, 2019). *TensorRT*. https://developer.nvidia.com/tensorrt.

[57] Oak Ridge National Laboratory. 2019 (accessed January 14, 2019). *Summit Supercomputer*. https://www.olcf.ornl.gov/summit.

[58] Paavo Pärssinen. [n. d.] Modern mobile graphics processors. *Science: Internet, Data and Things (CS-E4000), Spring 2018*, 211.

[59] PyTorch. 2019 (accessed January 14, 2019). *PyTorch*. https://pytorch.org.

[60] Stephan Rabanser, Oleksandr Shchur, and Stephan Günnemann. 2017. Introduction to tensor decompositions and their applications in machine learning. *arXiv preprint arXiv:1711.10781*.

[61] Jonathan Ragan-Kelley, Andrew Adams, Dillon Sharlet, Connelly Barnes, Sylvain Paris, Marc Levoy, Saman Amarasinghe, and Frédo Durand. 2017. Halide: decoupling algorithms from schedules for high-performance image processing. *Communications of the ACM*, 61, 1, 106–115.

[62] Md Aamir Raihan, Negar Goli, and Tor M. Aamodt. 2018. Modeling deep learning accelerator enabled gpus. *CoRR*, abs/1811.08309. arXiv: 1811.08309. http://arxiv.org/abs/1811.08309.

[63] Brandon Reagen, Robert Adolf, Paul Whatmough, Gu-Yeon Wei, and David Brooks. 2017. Deep learning for computer architects. *Synthesis Lectures on Computer Architecture*, 12, 4, 1–123.

[64] Shubhabrata Sengupta, Mark Harris, and Michael Garland. 2008. Efficient parallel scan algorithms for gpus. *NVIDIA, Santa Clara, CA, Tech. Rep. NVR-2008-003*, 1, 1, 1–17.

[65] Shubhabrata Sengupta, Aaron E Lefohn, and John D Owens. 2006. A work-efficient step-efficient prefix sum algorithm. In *Workshop on edge computing using new commodity architectures*, 26–27.

[66] Michel Steuwer, Toomas Remmelg, and Christophe Dubach. 2017. Lift: a functional data-parallel ir for high-performance gpu code generation. In *Code Generation and Optimization (CGO), 2017 IEEE/ACM International Symposium on*. IEEE, 74–85.

[67] TensorFlow. 2019 (accessed January 14, 2019). *TensorFlow*. https://www.tensorflow.org.

[68] Charles Van Loan. 1992. A survey of matrix computations. *Handbooks in operations research and management science*, 3, 247–321.

[69] WikiChip. 2019 (accessed January 14, 2019). *Cascade Lake - Microarchitectures - Intel*. https://en.wikichip.org/wiki/intel/microarchitectures/cascade_lake.

[70] Shengen Yan, Guoping Long, and Yunquan Zhang. 2013. Streamscan: fast scan algorithms for gpus without global barrier synchronization. In *ACM SIGPLAN Notices* number 8. Vol. 48. ACM, 229–238.

[71] Yuhao Zhu, Matthew Mattina, and Paul Whatmough. 2018. Mobile machine learning hardware at arm: a systems-on-chip (soc) perspective. *arXiv preprint arXiv:1801.06274*.