


# High Performance Im2win and Direct Convolutions using Three Tensor Layouts on SIMD Architectures

Xiang Fu\* Xinpeng Zhang\* Jixiang Ma\* Peng Zhao† Shuai Lu\* Xu T. Liu‡§ 

\*Nanchang Hangkong University, China †Microsoft, China ‡University of Washington, USA §Amazon Web Services, USA

**Abstract**—Convolution is the core component within deep neural networks and it is computationally intensive and time consuming. Tensor data layouts significantly impact convolution operations in terms of memory access and computational efficiency. Yet, there is still a lack of comprehensive performance characterization on data layouts on SIMD architectures concerning convolution methods. This paper proposes three novel data layouts for im2win convolution: NHWC, CHWN, and CHWN8, and introduces a set of general optimization techniques for both direct and im2win convolutions. We compare the optimized im2win convolution with the direct convolution and PyTorch’s im2col-based convolution across the aforementioned layouts on SIMD machines. The experiments demonstrated that the im2win convolution with the new NHWC layout achieved up to 355% performance speedup over NCHW layout. Our optimizations also significantly improve the performance of both im2win and direct convolutions. Our optimized im2win and direct convolutions achieved up to 95% and 94% of machine’s theoretical peak performance, respectively.

**Index Terms**—direct convolution, im2win convolution, NHWC layout, CHWN layout, CHWN8 layout

## I. INTRODUCTION

Convolution is the essential component of deep neural networks for computer vision tasks such as feature extraction from large-scale image data [1]. It not only comprises 50%-90% of computational operations including convolutional, pooling, ReLU, and fully-connected layers [2], but also consumes more than 90% of the total execution time of many popular neural networks [3]–[5]. Hence, optimizing convolution operations is crucial for enhancing the performance of neural networks.

Convolution methods can be classified into three main categories based on how they transform the input tensor: direct, im2col-based, and im2win. Direct convolution performs convolution operations directly on the tensor without changing its format [6]. This approach avoids extra memory consumption compared to im2col-based and im2win convolutions, but it suffers from nonconsecutive memory access. Im2col-based convolution transforms the convolution into general matrix-matrix multiplications (GEMM) [7], [8], leveraging optimized Basic Linear Algebra Subprograms (BLAS) [9] for excellent performance. It has regular memory access but a significant extra memory footprint, which greatly limits its applicability on memory-constrained devices. Previously, we propose a memory-efficient convolution called image-to-window (im2win), which reorganizes the input tensor into a row of dot product windows and flattens the unique elements of these windows into a row that correspond to the convolution operation’s receptive fields [10], [11]. It provides sequential

memory access and data reuse, and thus greatly reduces memory overhead.

A tensor memory/data layout (referred as layout henceforth) refers to how the data of a tensor is physically arranged in memory. There are commonly three layouts for tensors: NCHW, NHWC and CHWN, where N is the batch size, C is the number of channels, H is the image height, and W is the image width. It significantly impacts convolution operations in terms of memory access, computational efficiency and compatibility with deep learning frameworks [12]–[14].




Fig. 1. The original input tensor ( $N = 1, H_{original} = W_{original} = C_{original} = 3$ ) and its corresponding im2win tensor ( $N = 1, C_{im2win} = 3, H_{im2win} = 2, W_{im2win} = 6$ ) in the NCHW layout

Tensors are stored as one-dimensional arrays in memory, although logically represented as four-dimensional arrays. An input tensor and an im2win tensor in a NCHW layout are illustrated in Figure 1. Different colors represent different channels. In the NCHW layout, the elements in the width dimension are contiguous in memory by prioritizing the width dimension first, followed by height, channel and batch. Assuming a stride of 1 and a filter tensor in the NCHW layout with dimensions of  $1 \times 3 \times 2 \times 2$ , the elements in the solid-lined boxes are used to compute the first output element in the output tensor, while the elements in the dashed-lined boxes are for the next. Further details will be provided in Section III.

In general, the memory efficiency and performance implications of various tensor layouts with different convolution algorithms on single-input, multiple-data (SIMD) architectures have received limited attention. The previous im2win works optimize the NCHW layout on CPU and GPU but have not tried the NHWC and CHWN layouts [10], [11]. Li et al. reveal the performance impact of the NCHW layout with the im2col-based convolution and the CHWN layout with the direct convolution in different CNN layers, and propose a fast multi-dimension layout transformation algorithm on GPU [15]. Li et al. optimize the cache utilization of the NCHW and NHWC layouts in DNN training along with DNN pruning

on CPU [16]. To date, no performance characterization have been conducted on NCHW, NHWC and CHWN layouts on SIMD architectures with the three convolution methods.

To address aforementioned gaps, we propose three new layouts for the im2win convolution: NHWC, CHWN, and CHWN8, and propose a set of optimization techniques based on the roofline model for both direct and im2win convolutions. We compare the optimized im2win convolution with the optimized direct and PyTorch’s im2col-based convolutions with the above layouts on SIMD machines. Our experiments demonstrate that the new NHWC layout on the im2win convolution achieves 11% to 355% performance speedup compared with the NCHW layout. Our proposed optimizations have been empirically validated to enhance the performance of both im2win and direct convolutions. Our optimized im2win convolution and our optimized direct convolution achieve up to 95% and 94% of the theoretical peak performance of the machine, respectively. As the main contributions, this work

- 1) proposes three novel layouts for im2win convolution.
- 2) comes up with a set of general optimization techniques that not only can be applied to im2win convolution and but also to direct convolution on different layouts.
- 3) compares the optimized im2win convolution with the im2col-based and the optimized direct convolutions using four tensor layouts on SIMD machines.

## II. PRELIMINARY AND RELATED WORKS

### A. Notation

The input of a convolution operation includes an input tensor ( $\mathcal{I}$ ), a filter tensor ( $\mathcal{F}$ ), and an output tensor ( $\mathcal{O}$ ). In these tensors,  $N_i$  is the batch size,  $s$  is the stride size,  $C_i$  and  $C_o$  are the number of input and output channels (also known as feature maps),  $H_{i/f/o}$  and  $W_{i/f/o}$  denote the height and width of a feature map.

### B. Tensor Layouts: NCHW, NHWC, CHWN

In the NCHW layout, the input ( $\mathcal{I}$ ), the filter ( $\mathcal{F}$ ), and the output tensors ( $\mathcal{O}$ ) are expressed as  $\mathcal{I}[N_i][C_i][H_i][W_i]$ ,  $\mathcal{F}[C_o][H_f][W_f][C_i]$ , and  $\mathcal{O}[N_i][C_o][H_o][W_o]$ , respectively. The convolution is defined as:

$$\mathcal{O}_{(i,j,m,n)} = \sum_{j=1}^{C_i} \sum_{m=1}^{H_f} \sum_{n=1}^{W_f} (\mathcal{I}_{(i,j,m \times s+u, m \times s+v)}) \times \mathcal{F}_{(j,r,u,v)}, \quad (1)$$

In the NHWC layout, the input ( $\mathcal{I}$ ), the filter ( $\mathcal{F}$ ), and the output tensors ( $\mathcal{O}$ ) are expressed as  $\mathcal{I}[N_i][H_i][W_i][C_i]$ ,  $\mathcal{F}[C_o][H_f][W_f][C_i]$ , and  $\mathcal{O}[N_i][H_o][W_o][C_o]$ , respectively. The convolution is defined as:

$$\mathcal{O}_{(i,m,n,j)} = \sum_{j=1}^{C_i} \sum_{m=1}^{H_f} \sum_{n=1}^{W_f} (\mathcal{I}_{(i,m \times s+u, m \times s+v, j)}) \times \mathcal{F}_{(j,u,v,r)}, \quad (2)$$

In the CHWN layout, the input ( $\mathcal{I}$ ), the filter ( $\mathcal{F}$ ) and the output tensors ( $\mathcal{O}$ ) are expressed as  $\mathcal{I}[C_i][H_i][W_i][N_i]$ ,

$\mathcal{F}[C_i][H_f][W_f][C_o]$ , and  $\mathcal{O}[C_o][H_o][W_o][N_i]$  respectively. The convolution is defined as:

$$\mathcal{O}_{(j,m,n,i)} = \sum_{j=1}^{C_i} \sum_{m=1}^{H_f} \sum_{n=1}^{W_f} (\mathcal{I}_{(j,m \times s+u, m \times s+v, i)}) \times \mathcal{F}_{(r,u,v,j)}, \quad (3)$$

The definitions (1), (2) and (3) above are all subject to

$$\begin{aligned} j &= 1, 2, \dots, C_o, m = 1, 2, \dots, H_o, n = 1, 2, \dots, W_o, \\ i &= 1, 2, \dots, N_i, u = 1, 2, \dots, H_f, v = 1, 2, \dots, W_f, \\ r &= 1, 2, \dots, C_i. \end{aligned}$$

### C. Convolution Algorithms and Related Works

Direct convolution performs on the original  $\mathcal{I}$  and  $\mathcal{F}$  without any tensor transformation. It has seven nested for loops and an AXPY operation in the innermost loop. Based on the tensor layouts that direct convolution works with, the AXPY operation needs to read at different indices of  $\mathcal{F}$  and  $\mathcal{I}$ , and writes at different indices of  $\mathcal{O}$ .

Direct convolution can compute with the original input tensors, so they usually adopt the NCHW layout of raw images. A study shows that, for convolution instances with large  $C$ , NHWC layout outperforms NCHW layout [15]. Grouping a certain dimension of the input tensors by a fixed size can also enhance the performance of direct convolution, such as NC32HW32 layout [17]. Several works [6], [18] have shown that the performance of direct convolution can be greatly improved by designing specific layouts based on the loop ordering of the algorithm on SIMD architecture.

The im2col-based convolution transforms a convolution operation into a GEMM operation.  $\mathcal{I}[N_i][C_i][H_i][W_i]$  is processed in  $N_i$  batches, each batch contains data  $\mathcal{I}'[C_i][H_i][W_i]$  (that is, a single image). The im2col algorithm flattens the elements of each dot product window of  $\mathcal{I}'$  and copies them into a single row of a 2D matrix [7]. In addition to the conventional im2col data transformation algorithm, the MEC algorithm compresses the matrix layout, while still enabling the utilization of high-performance BLAS algorithms to perform convolution operations [19].

## III. HIGH-PERFORMANCE IM2WIN AND DIRECT CONVOLUTION USING THREE TENSOR LAYOUTS

In this section, we first review three layouts in the context of the direct convolution. Then we present three new layouts for the im2win convolution, following by how to determine the loop ordering based on the layouts and a set of optimizations for the im2win and direct convolutions on SIMD architectures.

### A. Motivations for Different Tensor Layouts

An example of direct convolution on an original input tensor in the NHWC layout is illustrated in Figure 2. The NHWC layout prioritizes the storage of elements in the last logical dimension— $C_i$ —followed by that of  $W_i$ ,  $H_i$ , and  $N_i$ . In the NHWC layout, the elements with the same  $N_i$ ,  $H_i$ , and  $W_i$  but different  $C_i$  are contiguous in memory, which have unit stride access. Assuming a stride of 1, the red




Fig. 2. The original input tensor ( $N_i = 1, H_i = W_i = C_i = 3$ ) and its corresponding im2win tensor ( $N_i = 1, C_i = 3, H_i = 2, W_i = 6$ ) in the NHC layout, the filter tensor ( $N_f = 1, C_f = 3, H_f = W_f = 2$ ),  $s = 1$ , the output tensor ( $N_o = 1, C_o = 1, H_o = W_o = 2$ )

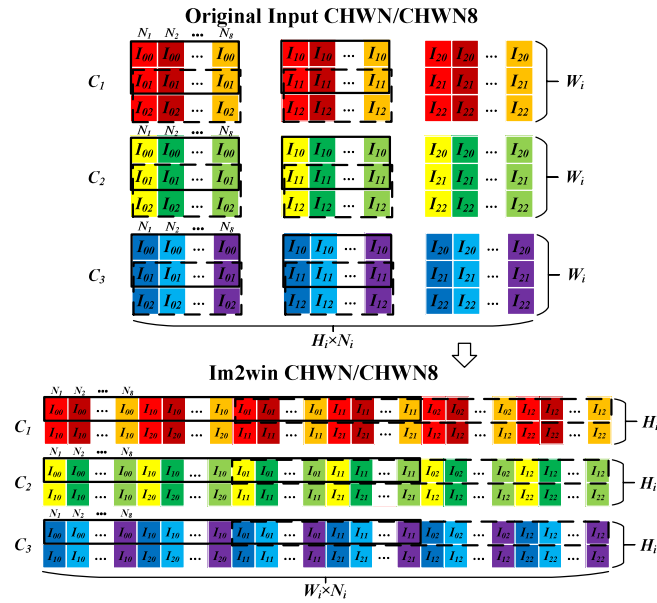


Fig. 3. The original tensor ( $N_i = 8, H_i = W_i = C_i = 3$ ) and its corresponding im2win tensor ( $N_i = 8, C_i = 3, H_i = 2, W_i = 6$ ) in the CHWN/CHWN8 layout

and green elements (outlined by solid lines, representing the convolutional windows to compute a single output element) in the input tensor are multiplied with the corresponding elements in the filter tensor, and the results are summed up (i.e., an AXPY) to obtain  $O_{00}$ . Next, the green and yellow elements (outlined by dashed lines) are used to calculate  $O_{01}$ . This process continues, with the convolutional window moving by the stride length in  $H_i$  or  $W_i$ , until all elements of the output tensor are computed.

Recall in Figure 1, the NCHW layout prioritizes the storage of elements in the last logical dimension— $W_i$ —followed by that of  $H_i$ ,  $C_i$ , and  $N_i$ . Note that all  $C_i$  but not all  $W_i$  are used during the AXPY operation to obtain one output element. NCHW has non-unit stride access during the tensor convolution. However, the non-unit stride access may not be

### Algorithm 1: Im2win Tensor Transformation Algorithm

**Input:** input tensor  $\mathcal{I}$  in the NHC layout, filter tensor  $\mathcal{F}$ , Stride  $s$

**Output:** Im2win tensor  $\hat{\mathcal{I}}$  in the NHC layout

- 1  $H_o = (H_i - H_f)/s + 1$
- 2 **for**  $i = 1$  to  $N_i$  **do**
- 3     **for**  $m = 1$  to  $H_o$  **do**
- 4         **for**  $k = 1$  to  $W_i$  **do**
- 5             **for**  $u = 1$  to  $H_f$  **do**
- 6                 **for**  $r = 1$  to  $C_i$  **do**
- 7                      $\hat{\mathcal{I}}[i][m][u + k \times H_f][r] = \mathcal{I}[i][m \times s + u][k][r]$

harmful, depending on the effects of caching and the access patterns used (determined by the loop ordering).

As shown in Figure 3, the CHWN layout stores elements by prioritizing  $N_i$ , followed by  $W_i$ ,  $H_i$ , and  $C_i$  in memory. Previous research on GPU recommends to use  $N$  as the lowest dimension for coalesced memory access and data reuse in registers [15]. It has been observed that the performance is sensitive to the value of  $N$ . The elements within the solid-lined boxes are used to compute the first eight output elements, while the elements within the dashed-lined box are used for the next eight output elements. This facilitates the use of vector registers for vectorization. Assuming  $s = 1$ , the convolutional window moves by one element in both the  $H_i$  and  $W_i$  dimensions until all output elements are computed.

### B. Im2win Tensor Transformation on Three Tensor Layouts

In this subsection, we present the im2win tensor transformation on three tensor layouts. The im2win tensor transformation process for the NHC layout is shown as Algorithm 1. For other layouts, slight modifications need to be made on it. The im2win transformation flattens the elements of a convolutional window, storing them contiguously in memory and prioritizing them in the  $C_i$  dimension, as outlined in Algorithm 1 from Line 4 to Line 7. An example of transforming the original input tensor into an im2win tensor in the NHC layout is shown in Figure 2. It is important to note that the green elements can be reused between two adjacent windows, which

---

**Algorithm 2:** Naive Im2win Convolution

---

**Input:** input tensor  $\mathcal{I}$  in the NHWC layout, filter tensor  $\mathcal{F}$  in the NHWC layout, Stride  $s$   
**Output:** output tensor  $\mathcal{O}$  in the NHWC layout

```
1  $\hat{\mathcal{I}} = \text{Function IM2WIN}(\mathcal{I}, \mathcal{F}, s)$ 
2 transform  $\mathcal{F}$  in NHWC to  $\hat{\mathcal{F}}$  in NWHC
3 for  $i = 1$  to  $N_i$  do
4   for  $m = 1$  to  $H_o$  do
5     for  $j = 1$  to  $C_o$  do
6       for  $n = 1$  to  $W_o$  do
7         for  $v = 1$  to  $W_f$  do
8           for  $u = 1$  to  $H_f$  do
9             for  $r = 1$  to  $C_i$  do
10               $\mathcal{O}[i][m][n][j] += \hat{\mathcal{I}}[i][m][n \times$   

 $s \times W_f + v \times W_f + u][r] \times$   

 $\hat{\mathcal{F}}[j][v][u][r]$ 
```

---

eliminates the need for repetitive storage in memory like the im2col transformation [7]. In the output im2win tensor of Algorithm 1, the elements involved in each convolution operation are contiguous and more compact in memory, improving spatial locality, cache and SIMD efficiency.

An example of the im2win convolution using the NHWC layout is shown in Figure 2. A naive im2win convolution method is shown as Algorithm 2. It is similar to the direct convolution, which involves seven nested loops and an AXPY in the innermost loop, except prior to that, the im2win convolution performs a tensor transformation upon  $\mathcal{I}$  to get  $\hat{\mathcal{I}}$ , and  $\mathcal{F}$  in the NHWC layout is transformed into a NWHC layout to match  $\hat{\mathcal{I}}$ . All output elements can be computed through the outer four loops of Algorithm 2 from Line 3 to Line 6.

An im2win tensor in CHWN/CHWN8 layout is shown in Figure 3. The characteristics of  $C_i$ ,  $H_i$ , and  $W_i$  in this layout are similar to NCHW, but the CHWN layout prioritizes storage in  $N_i$ . The CHWN layout’s efficiency is constrained by 256-bit vector registers, which process only 8 outputs at once. When  $N_i > 8$ , this leads to low cache utilization, as the cache holds unnecessary data for these 8 output calculations. Motivated by the approach of dividing  $C_i$  into blocks in the direct convolution [18], we propose a new CHWN8 layout. CHWN8 lays 8  $N_i$  in the innermost layer and remaining  $N_i$  in the outermost layer, which takes full advantage of the vector registers without sacrificing the cache utilization.  $N_i$  can be set to a multiple of 8 (with padding if necessary), CHWN8 layout may provide better data continuity than the NHWC layout when  $C_i$  is relatively small.

### C. Loop Reordering

In this subsection, we reorder the loops for the im2win and direct convolutions based on different tensor layouts. Ideally, we arrange the inner loops to access data closer in memory to enjoy the unit stride access as long as possible. In both direct and im2win convolutions, for NCHW, CHWN, and CHWN8 layouts, as they have the CHW memory access pattern, we use the width of the convolution window as the innermost loop, followed by the height, and the channel ( $W_f$ ,  $H_f$ , and  $C_i$ ). Conversely, for the NHWC layout, the innermost three loops

---

**Algorithm 3:** High Performance Im2win Convolution

---

**Input:** input tensor  $\mathcal{I}$  in NHWC layout, filter tensor  $\mathcal{F}$  in NHWC layout, Stride  $s$   
**Output:** output tensor  $\mathcal{O}$  in NHWC layout

```
1  $\hat{\mathcal{I}} = \text{Function IM2WIN}(\mathcal{I}, \mathcal{F}, s)$ 
2 transform  $\mathcal{F}$  in NHWC to  $\hat{\mathcal{F}}$  in NWHC
3  $H_o = (H_i - H_f)/s + 1$ 
4 for  $im = 1$  to  $N_i \times H_o$  in parallel do
5    $i = im/H_o$ ,  $m = im \% H_o$ 
6   for  $j = 1$  to  $C_o$  do
7     for  $n = 1$  to  $W_o/W_{o,b}$  do
8       DOT_PRODUCT( $i, j, m, n, W_{o,b}, s$ )
9 Function DOT_PRODUCT( $i, j, m, n, W_{o,b}, s$ ):
10   $ymm_1 = \dots = ymm_{W_{o,b}} = 0$ 
11  for  $r = 1$  to  $C_i$  do
12    for  $v = 1$  to  $W_f/N_{vec}$  do
13      for  $u = 1$  to  $H_f$  do
14        FMA( $\hat{\mathcal{I}}[i][r][u + m][v \times N_{vec}]$ ,
15           $\hat{\mathcal{F}}[j][r][u][v]$ ,  $ymm_1$ )
16        ...
17        FMA( $\hat{\mathcal{I}}[i][r][u + m][v \times N_{vec} + s \times$   

 $(W_{o,b} - 1)]$ ,  $\hat{\mathcal{F}}[j][r][u][v]$ ,  $ymm_{W_{o,b}}$ )
```

---

iterate over the channel, the width, and the height ( $C_i$ ,  $W_f$ , and  $H_f$ ) of the convolution window. For the im2win convolution, since the im2win transformation flattens the elements of a convolutional window, we usually coalesce the height and the width layers ( $W_f$ ,  $H_f$ ) into a  $W_f \times H_f$ .

Next we determine the order of the outer four loops. It not only determines the order in which the elements of the output tensor are produced but also influences the order in which the element in the input tensor are accessed. The layout mainly affects the inner three levels of the loop order, for the outer four levels the loop order applies to all data layouts. Recall in Figure 2, the green elements of the input tensor are shared between the convolution windows of adjacent output tensor elements. Therefore, we consider to position the width ( $W_o$ ) of the traversed output tensor in the fourth layer of loops at Line 6 in Algorithm 2. Because the memory access of the input tensor is expensive in a convolution operation, we set the channel ( $C_o$ ) as the third layer to reduce its access. Since the batch of the output tensor corresponds to the batch of the input tensor, we place the batch ( $N_i$ ) in the first layer of the loop. Finally we have the order of the outer four loops as NHCW in Algorithm 2.

### D. Optimizations for the Im2win and Direct Convolutions

In this subsection, we use the Roofline Model [20] to determine how to optimize the im2win and direct convolutions on SIMD systems. We propose a set of optimizations for both the im2win and direct convolutions and apply the optimizations to them (with a slight modification on the direct convolution). The optimizations are classified into two categories: reducing the memory bottleneck and increasing the arithmetic intensity of the convolution kernel. The former includes hoist, memory alignment, register and cache blocking. The latter includes loop unrolling, vectorization and FMA

instructions, loop coalescing, and parallelization strategies. We have three level parallelization: Non-Uniform Memory Access (NUMA) level, thread level using OpenMP and instruction level using SIMD.

We present the optimized im2win convolution as Algorithm 3. Since each element of the output tensor in each batch ( $N_i$ ) can be computed independently, there is abundant parallelism available [21] and this is NUMA-friendly. Within each batch, the shared elements of adjacent windows can be maximized across the threads. When the dimension we choose to parallelize is small and the number of CPU cores is large, this leads to a workload imbalance in many core architecture. The solution is to coalesce multiple dimensions/loops into one parallel loop to achieve better load balance. In practice, we find that coalescing two dimensions of the output tensor yields the best load balance. Hence, we apply parallel strategies and coalesce  $N_i$  and  $H_o$  in a parallel loop at Line 4 in Algorithm 3.

In Algorithm 3, we hoist three things: the indices of the elements in the 1D array of the tensor, the elements of the input tensor, and the entire filter tensor at Line 12. We use register and cache blocking [22], [23] at Line 7 to reduce cache misses because the neighboring convolutional windows share duplicate elements of the input tensor, which  $W_{o,b}$  is the blocking size. Next, we unroll the loop at Line 12 to take advantage of the spatial locality, because the innermost loop accesses consecutive elements. Since we are using FP32 and AVX2, we set the loop unrolling size  $N_{vec}=8$ . Finally, we vectorize the input tensor and filter tensor in units of eight to use the FMA instruction in AVX2 [10] at Line 14.

A cache-line is 512 bits on contemporary x86\_64 architectures, which is the minimum data quantity that can be fetched from memory to cache. Without memory alignment of the tensor data structure, the CPU needs to issue two memory requests to access an element, because the element may be in the middle of the cache-line. This hurts the performance greatly because not only the CPU has to wait for memory access but also more caches are used. This leads to higher cache miss and lower cache utilization. We store the elements of the tensor in a memory-aligned way using `posix_memalign` in C during the memory allocation.

## IV. EXPERIMENTS

### A. Experimental Setup

**Architectures.** We use a server with two Intel<sup>®</sup> Xeon<sup>®</sup> Gold 6330 CPUs and 251 GB RAM. Each CPU has 28 physical cores, running at 2.0 GHz, with 48 KB L1d cache, 32 KB L1i cache, 1.28MB L2 cache and 43 MB L3 cache.

**Benchmarks** We aim to cover the majority of the convolutional layers in commonly used DNNs. Hence for our experimental evaluation, we select an state-of-the-art DNN benchmark [19] shown in Table I, which includes twelve unique convolution layers, conv1-conv12.

**Software** We compare the direct, im2win, and im2col-based convolutions with four tensor layouts: NHWC, NCHW, CHWN, CHWN8. We intend to compare with the state-of-the-art implementation of direct convolution and the layout

TABLE I  
TWELVE CONVOLUTION LAYERS OF THE DNN BENCHMARKS.

NAME	INPUT	FILTER, STRIDE	OUTPUT
	$C_i \times H_i \times W_i$	$C_o \times H_f \times W_f, s_h/s_w$	$C_o \times H_o \times W_o$
conv1	$3 \times 227 \times 227$	$96 \times 11 \times 11, 4$	$96 \times 55 \times 55$
conv2	$3 \times 231 \times 231$	$96 \times 11 \times 11, 4$	$96 \times 56 \times 56$
conv3	$3 \times 227 \times 227$	$64 \times 7 \times 7, 2$	$64 \times 111 \times 111$
conv4	$64 \times 224 \times 224$	$64 \times 7 \times 7, 2$	$64 \times 109 \times 109$
conv5	$96 \times 24 \times 24$	$256 \times 5 \times 5, 1$	$256 \times 20 \times 20$
conv6	$256 \times 12 \times 12$	$512 \times 3 \times 3, 1$	$512 \times 10 \times 10$
conv7	$3 \times 224 \times 224$	$64 \times 3 \times 3, 1$	$64 \times 222 \times 222$
conv8	$64 \times 112 \times 112$	$128 \times 3 \times 3, 1$	$128 \times 110 \times 110$
conv9	$64 \times 56 \times 56$	$64 \times 3 \times 3, 1$	$64 \times 54 \times 54$
conv10	$128 \times 28 \times 28$	$128 \times 3 \times 3, 1$	$128 \times 26 \times 26$
conv11	$256 \times 14 \times 14$	$256 \times 3 \times 3, 1$	$256 \times 12 \times 12$
conv12	$512 \times 7 \times 7$	$512 \times 3 \times 3, 1$	$512 \times 5 \times 5$

proposed in [18], but their implementation is not open-sourced. We use the im2col-based convolution in PyTorch 2.1 [24] with MKL [25]. Note that PyTorch only supports the NHWC and NCHW layouts. Our code is compiled with GCC 9.5.0 compiler and `-O3 -mavx2 -mfma -fopenmp -march=native` compilation flags. We use OpenMP 4.0 for parallelization with *guided* scheduling.

### B. Performance of Different Convolution Algorithms

We run each algorithm 50 times on each benchmark with  $N_i=128$  and report the best runtime. We exclude the im2win data transformation time from FLOPS calculations because this transformation occurs only once at the first convolutional layer in a DNN. Figure 4 shows the performance results in TFLOPS and Figure 5 shows the memory usage of our high-performance direct convolution and high-performance im2win-based convolution, and the im2col-based convolution using MKL in PyTorch. In Figure 4, the left y-axis shows the performance in TFLOPS, and the right y-axis shows the performance of the machine peak. Overall, our optimized im2win convolution achieves eight best TFLOPS out of twelve benchmarks; our optimized direct convolution achieves three best TFLOPS out of twelve benchmarks and the im2win convolution achieves close performance on these three; the im2col-based convolution in PyTorch achieves one out of twelve best TFLOPS on conv12. All twelve best TFLOPS are all yielded from the NHWC layout across these three methods. Our proposed optimization techniques are proven effective on both im2win and direct convolutions. Our im2win convolution achieves 95% and 91% of the theoretical peak performance of the architecture on conv5 and conv6 respectively. Our direct convolution achieves 91% and 94% of the theoretical peak performance on conv5 and conv6 respectively.

With the NHWC layout and performance normalization, excluding conv6 and conv12, our im2win convolution achieves between  $1.1\times$  and  $4.6\times$  performance speedup, and our direct convolution achieves between  $1.1\times$  and  $3.8\times$  performance speedup against the im2col-based convolution. All the best performance of the im2win convolution on twelve benchmarks




Fig. 4. Performance results in TFLOPS of the direct convolution, the im2win convolution and the im2col-based convolution using different layouts. Note that the theoretical peak performance of the server is 3584 GFLOPS.




Fig. 5. Memory usage of the direct, the im2win and the im2col-based convolutions using different tensor layouts. Note that in conv4, the im2col-based convolutions with the NWHC and NCHW layouts use 21GB of memory.

is achieved using the NHWC layout on CPU. Our im2win convolution using the NHWC layout outperforms the NCHW layout by at least 11% and up to 355% across all benchmarks, showcasing significant efficiency gains in various scenarios. For the direct convolution, nine out of twelve best TFLOPS are achieved using the NHWC layout on CPU, the rest three are achieved using the CHWN8 layout.

With the NCHW layout, our im2win-based convolution achieves between 1.4 $\times$  and 2.4 $\times$  performance speedup on all benchmarks, and PyTorch’s im2col-based convolution achieves between 1.1 $\times$  and 7.5 $\times$  performance speedup (excluding conv7) against the direct convolution. The direct convolution performs poorly on the NCHW layout.

With the CHWN/CHWN8 layout, im2win\_CHWN8 outperforms im2win\_CHWN on all benchmarks from 3.7 $\times$  to 16 $\times$ ; direct\_CHWN8 outperforms 2.3 $\times$  and 8 $\times$  over direct\_CHWN except in conv7. This shows our proposed CHWN8 layout overwhelmingly beats the CHWN layout on these two methods. The direct convolution with the CHWN8 layout performs better than the CHNW, NCHW and NHWC layouts when  $C_i$  is small ( $C_i = 3$  for conv1, conv2, conv3). Similar results have also been observed in the previous GPU work [15].

In Figure 5, using the same tensor layout, the memory usage of each convolution is the same hence we only annotate one number per method. In all benchmarks, the direct convolution

uses the least memory, while the im2col-based convolution consumes the most memory. On average, the im2col-based convolution has 3.9 $\times$  more memory usage than direct convolution, and im2win-based convolution has 1.5 $\times$  more memory usage than direct convolution. The im2win-based convolution uses on average 39% of the memory of the im2col-based method, and in conv5, it uses only 24% of the memory of the im2col-based convolution.

## V. CONCLUSION

We proposed three new layouts for the im2win convolution: NHWC, CHWN and CHWN8, and a set of general optimization techniques for both direct and im2win convolutions on SIMD architectures. We applied these optimizations on the im2win and the direct convolutions, and compared with PyTorch’s im2col-based convolution on the above layouts along with the NCHW layout. Our experiments demonstrated that im2win convolution using the new NHWC layout achieved 11% to 355% performance speedup compared to NCHW layout. The proposed optimizations were proven to boost the performance of the im2win convolution and direct convolution. Our optimized im2win and direct convolutions achieved up to 95% and 94% of the theoretical peak performance of the machine, respectively.

## ACKNOWLEDGMENT

This research was partially supported by the NSF OAC-2402542 and the National Natural Science Foundation of China (Grant No. 42261070). This work is not related to Xu T. Liu's position at Amazon.

## REFERENCES

- [1] Y. Cui, Y. Sun, J. Hu, and G. Sheng, "A Convolutional Auto-Encoder Method for Anomaly Detection on System Logs," in *2018 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, Oct. 2018, pp. 3057–3062, iSSN: 2577-1655.
- [2] N. Ma, X. Zhang, H.-T. Zheng, and J. Sun, "Shufflenet v2: Practical guidelines for efficient cnn architecture design," in *Proceedings of the European conference on computer vision (ECCV)*, 2018, pp. 116–131.
- [3] V. Sze, Y.-H. Chen, T.-J. Yang, and J. S. Emer, "Efficient processing of deep neural networks: A tutorial and survey," *Proceedings of the IEEE*, vol. 105, no. 12, pp. 2295–2329, 2017.
- [4] Y. Yan, F. Guo, H. Mo, and X. Huang, "Hierarchical tracking control for a composite mobile robot considering system uncertainties," in *2024 16th International Conference on Computer and Automation Engineering (ICCAE)*. IEEE, 2024, pp. 512–517.
- [5] J. Chen, W. Xu, and J. Wang, "Prediction of car purchase amount based on genetic algorithm optimised bp neural network regression algorithm," *Preprints*, June 2024. [Online]. Available: <https://doi.org/10.20944/preprints202406.1304.v1>
- [6] E. Georganas, S. Avancha, K. Banerjee, D. Kalamkar, G. Henry, H. Pabst, and A. Heinecke, "Anatomy of high-performance deep learning convolutions on simd architectures," in *SC18: International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE, 2018, pp. 830–841.
- [7] K. Chellapilla, S. Puri, and P. Simard, "High performance convolutional neural networks for document processing," in *10th Int'l Workshop on Frontiers in Handwriting Recog.*, 2006.
- [8] A. V. Trusov, E. E. Limonova, D. P. Nikolaev, and V. V. Arlazarov, "p-im2col: Simple yet efficient convolution algorithm with flexibly controlled memory overhead," *IEEE Access*, vol. 9, pp. 168 162–168 184, 2021.
- [9] J. J. Dongarra, J. Du Croz, S. Hammarling, and I. S. Duff, "A set of level 3 basic linear algebra subprograms," *ACM Transactions on Mathematical Software (TOMS)*, vol. 16, no. 1, pp. 1–17, 1990.
- [10] S. Lu, J. Chu, and X. T. Liu, "Im2win: Memory efficient convolution on SIMD architectures," in *2022 IEEE High Performance Extreme Computing Conference (HPEC)*, 2022, pp. 1–7.
- [11] S. Lu, J. Chu, L. Guo, and X. T. Liu, "Im2win: An Efficient Convolution Paradigm on GPU," in *Euro-Par 2023: Parallel Processing*, ser. Lecture Notes in Computer Science, J. Cano, M. D. Dikaiakos, G. A. Papadopoulos, M. Pericàs, and R. Sakellariou, Eds. Cham: Springer Nature Switzerland, 2023, pp. 592–607.
- [12] A. Paszke, S. Gross, F. Massa *et al.*, "Pytorch: An imperative style, high-performance deep learning library," in *Proceedings of the 33rd International Conference on Neural Information Processing Systems (NeurIPS)*, 2019, pp. 8024–8035.
- [13] M. Abadi, P. Barham, J. Chen *et al.*, "Tensorflow: A system for large-scale machine learning," in *Proceedings of the 12th USENIX Conference on Operating Systems Design and Implementation*, ser. OSDI'16. USA: USENIX Association, 2016, p. 265–283.
- [14] S. Chetlur, C. Woolley, P. Vandermersch, J. Cohen, J. Tran, B. Catanzaro, and E. Shelhamer, "cuDNN: Efficient primitives for deep learning," *CoRR*, vol. abs/1410.0759, 2014.
- [15] C. Li, Y. Yang, M. Feng, S. Chakradhar, and H. Zhou, "Optimizing Memory Efficiency for Deep Convolutional Neural Networks on GPUs," in *SC16: International Conference for High Performance Computing, Networking, Storage and Analysis*. Salt Lake City, UT, USA: IEEE, Nov. 2016, pp. 633–644.
- [16] B. Li, Q. Xue, G. Yuan, S. Li, X. Ma, Y. Wang, and X. Tang, "Optimizing Data Layout for Training Deep Neural Networks," in *Companion Proceedings of the Web Conference 2022*, ser. WWW '22. New York, NY, USA: ACM, Aug. 2022, pp. 548–554.
- [17] V. Ferrari, R. Sousa, M. Pereira, J. P. L. Carvalho, J. N. Amaral, J. Moreira, and G. Araujo, "Advancing direct convolution using convolution slicing optimization and isa extensions," in *Advancing Direct Convolution using Convolution Slicing Optimization and ISA Extensions*, March 2023.
- [18] J. Zhang, F. Franchetti, and T. M. Low, "High performance zero-memory overhead direct convolutions," in *International Conference on Machine Learning*. PMLR, 2018, pp. 5776–5785.
- [19] M. Cho and D. Brand, "MEC: memory-efficient convolution for deep neural network," in *International Conference on Machine Learning (ICML)*. PMLR, 2017, pp. 815–824.
- [20] S. Williams, A. Waterman, and D. Patterson, "Roofline: an insightful visual performance model for multicore architectures," *Commun. ACM*, vol. 52, no. 4, p. 65–76, apr 2009. [Online]. Available: <https://doi.org/10.1145/1498765.1498785>
- [21] E. Georganas, S. Avancha, K. Banerjee, D. Kalamkar, G. Henry, H. Pabst, and A. Heinecke, "Anatomy of high-performance deep learning convolutions on simd architectures," *Proceedings - International Conference for High Performance Computing, Networking, Storage, and Analysis, SC 2018*, pp. 830–841, 2019.
- [22] M. Wolfe, "Iteration space tiling for memory hierarchies," in *Proceedings of the Third SIAM Conference on Parallel Processing for Scientific Computing*. USA: SIAM, 1987, p. 357–361.
- [23] —, "More iteration space tiling," in *Supercomputing '89: Proceedings of the 1989 ACM/IEEE Conference on Supercomputing*, 1989, pp. 655–664.
- [24] "PyTorch," 2024. [Online]. Available: <https://github.com/pytorch/pytorch>
- [25] "MKL," 2024. [Online]. Available: <https://www.intel.com/content/www/us/en/developer/tools/oneapi/onemkl.html>