The Anatomy of Efficient FFT and Winograd Convolutions on Modern CPUs

Aleksandar Zlateski*
Massachusetts Institute of Technology

Kai Li
Princeton University

Zhen Jia*
Princeton University

Fredo Durand
Massachusetts Institute of Technology

Abstract

Winograd-based convolution has quickly gained traction as a preferred approach to implement convolutional neural networks (ConvNets) on various hardware platforms because it could require fewer floating point operations than FFT-based or direct convolutions.

In this paper, we analyze the theoretical performances of three methods (regular FFT-, Gauss-FFT-, and Winograd-based convolutions), as well as compare their highly optimized implementations on modern multi- and many-core CPUs. With all three implementations employing the same optimizations on modern CPUs, our experimental results with modern ConvNets show that the FFT-based implementations generally outperform the Winograd-based approach, which is contrary to the popular belief.

To understand the results, we use a Roofline performance model to analyze the three implementations in detail, by looking at each of their computation phases and by considering not only the number of floating point operations, but also the memory bandwidth and the cache sizes. The performance analysis explains why, and under what conditions, the FFT-based implementations outperform the Winograd-based one, on modern CPUs.

CCS Concepts

• **General and reference** → **Performance**; • **Computing methodologies** → **Neural networks**; • **Software and its engineering** → *Just-in-time compilers*.

1 Introduction

Convolutional neural networks (ConvNets) have emerged as a widespread machine learning method for many application domains, soon after the demonstration of its superior performance for classification and localization tasks in the ImageNet[12] competition in 2012 [21]. Since convolutional layers are computationally intensive and dominate the total execution time of modern deep ConvNets [19, 21, 26, 29, 30], many efforts have been made to improve the performance of the convolutional primitives for CPUs [1, 10, 31, 38, 40], GPUs [6, 11, 25, 32] or both [39].

*Both authors contributed equally to the paper

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

ICS '19, June 26–28, 2019, Phoenix, AZ, USA

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-6079-1/19/06...\$15.00

<https://doi.org/10.1145/3330345.3330382>

An important class of optimization is to reduce the number of calculations required for a convolution. Initially, several efforts used FFT-based convolutions to reduce the required computations for GPUs [25, 32] and CPUs [37–39]. Recently Lavin et al. [22] proposed a Winograd-based method for performing convolutions, and demonstrated that it can save more floating point operations than FFT, especially for small 2D kernels (e.g. 3×3). This prompted a large number of implementations to employ Winograd-based convolutions. For example, Nervana [6] and Nvidia's cuDNN [11] implemented Winograd-based convolution for GPUs. FALCON [1], LIBXSMM[7] and Intel MKL-DNN [2] provided CPU implementations of Winograd-based convolutions. Budden et al. [10] extended the 2D Winograd-based convolution to arbitrary dimensions and kernel sizes. A highly optimized implementation for modern many-core CPUs was presented by [18]. However, the community has not put in similar optimizing efforts into FFT-based implementations.

This paper presents results of comparing Winograd-based with FFT-based convolutions on modern CPUs. We have extended the highly optimized Winograd-based implementation for Intel Xeon Phi processors [18], to support arbitrary modern multi- and many-core CPUs (that support the AVX2 or the AVX512 instruction set). Using the same optimization techniques, we additionally implemented two FFT-based algorithms for modern CPUs. One is based on the regular FFT algorithm (Regular-FFT). Another is also FFT-based but uses Gauss' multiplication method (Gauss-FFT). All implementations are open-sourced at [4].

We have compared all three implementations with two popular ConvNets (VGG and AlexNet) on 10 different systems with modern CPUs. Our results show that, contrary to the popular belief, the regular-FFT or Gauss-FFT implementations are generally faster than the Winograd implementation. And in some cases, they are substantially faster.

To understand the experimental results, we have used the Roofline performance model [34] to analyze each algorithm and its implementation in detail by carefully analyzing its computational phases. Such analysis provided two key explanations. First, Winograd-based approach, in most cases, requires fewer floating point operations than FFT-based approach because it works with real numbers instead of complex numbers. However, due to its numerical instability, the Winograd method can use only very small transform sizes [10, 22, 33], which limits the speedups achieved in practice. The FFT-based convolutions do not suffer from such instabilities, allowing for arbitrary large tile sizes. Large tile sizes allow the FFT-based approach to reduce a large number of redundant

or unnecessary computations. Thus, in certain scenarios, the FFT-based method requires fewer operations than the Winograd-based one.

Second, our analysis considers not only the number of floating point operations, but also the total amount of data movements (to and from memory) and their costs, the arithmetic intensity (operations per moved byte), the processor’s speed, as well as the memory bandwidth. We also analyze the effects of cache sizes, which determine the arithmetic intensity of both methods, and indirectly affect the running times.

Large arithmetic intensity allows for better utilization of hardware systems whose compute-to-memory ratios are increasing over time, because of the imbalance of the evolution between processor speeds and memory bandwidths. The speeds for computations typically improve much faster than memory bandwidths [36]¹. This benefits the FFT-based method more, as its arithmetic intensity is generally higher due to computing with complex numbers.

Our analysis suggests that whether the Winograd-based or a FFT-based approach is faster depend on the specific convolutional layer and the particular system it is executed on. However, on average, the FFT-based approaches outperform the Winograd-based one with commonly used neural networks, with the margin increasing as the system’s compute-to-memory ratio increases.

The findings in this paper challenge the current belief that Winograd-based convolutions are better in nearly all practical cases.

2 Background

2.1 Winograd- and FFT- Based Convolutions

Winograd – Based Convolution As recently illustrated by Lavin et al. [22], “valid” convolution of discrete signals, the main operation used in modern convolutional neural networks, can be performed using Winograd’s minimal filtering algorithm [35] via

$$f \underset{\text{Valid}}{*} g = \mathbb{A}^T [(\mathbb{G}g) \odot (\mathbb{B}^T f)] \quad (1)$$

Where f and g are 1D signals; $\odot$ represents element-wise multiplication; and $\mathbb{A}^T$, $\mathbb{B}^T$ and $\mathbb{G}$ are special matrices, derived from Vandermonde matrices for Homogeneous Coordinate polynomials [33]. By convention, Winograd convolutions have the matrices $\mathbb{A}^T$, $\mathbb{B}^T$ and $\mathbb{G}$ in real-space. In “valid” convolution the filter slides across every “valid” location in the filtered images – such that the filter is fully contained inside the image. When the size of the filter is $|g| = r$, $f \underset{\text{Valid}}{*} g$ will have a length of $|f| - |g| + 1 = m$, and we refer to the method above as Winograd convolution $F(m, r)$.

FFT – Based Convolution The convolution theorem states that a convolution can be performed using Fourier transforms via

$$f \underset{\text{Circ}}{*} g = \mathcal{F}^{-1}(\mathcal{F}(f) \cdot \mathcal{F}(g)) \quad (2)$$

Here, $\mathcal{F}$ and $\mathcal{F}^{-1}$ are Fourier and inverse Fourier transforms. In the discrete case, f and g need to have the same number of elements, which can be accomplished by padding zeros to the shorter signal.

Discrete Fourier transform (DFT) results in a circular (also known as cyclic) convolution. The result of the “valid” convolution can be extracted from the last $|f| - |g| + 1 = m$ elements of the circular convolution.

“Valid” convolution using discrete FFTs can also be regarded as a special case of the Eq 1, where the matrices $\mathbb{A}^T$, $\mathbb{B}^T$ and $\mathbb{G}$ are in complex space and are derived from Vandermonde matrices with polynomial points being the roots of unity [33]. $\mathbb{B}^T$ and $\mathbb{G}$ perform, implicitly zero-padded (to size of $m + r - 1 = |f|$), DFT transforms of f and g , and $\mathbb{A}^T$ computes the last m elements of the inverse DFT transform. Using the FFT algorithm allows for efficient computation of matrix-vector products with matrices $\mathbb{A}^T$, $\mathbb{B}^T$ and $\mathbb{G}$. We refer to this special case as Regular-FFT $\mathcal{F}(m, r)$.

Multi-Dimensional Convolution Both Winograd and FFT convolutions can easily be extended to an arbitrary number of dimensions [10]. N -dimensional convolution is performed via

$$f \underset{\text{Valid}}{*} g = \left[(g \times_{n=1}^N \mathbb{G}) \odot (f \times_{n=1}^N \mathbb{B}) \right] \times_{n=1}^N \mathbb{A}^T \quad (3)$$

Here, the operation $x \times_n^N \mathbb{Y}$ is short for $x \times_1 \times_2 \cdots \times_n \mathbb{Y}$, where $\times_n$ represents tensor-matrix mode- n multiplication as defined in [10, 20]. For the 2D case, $x \times_1 \mathbb{M} = \mathbb{M}x$, and $x \times_2 \mathbb{M} = x\mathbb{M}^T$. The formula above reduces to

$$f \underset{\text{Valid}}{*} g = \mathbb{A}^T \left[(\mathbb{G}g\mathbb{G}^T) \odot (\mathbb{B}f\mathbb{B}^T) \right] \mathbb{A} \quad (4)$$

Which is consistent with Lavin et al. [22].

2.2 Winograd- and FFT- Based Convolution Layers

A convolutional layer transforms an input tuple of C images into an output tuple of C' images. A batch of B inputs yielding a batch of B outputs is processed at the time via

$$I'_{b,c'} = \sum_{c=1}^C I_{b,c} * W_{c',c} \quad (5)$$

Where C and C' denote the number of input/output images (also called channels or feature-maps). $I_{b,c}$ and $I'_{b,c'}$ are the (arbitrary dimensional) input and output images of the b -th batch. In total, $B \cdot C \cdot C'$ convolutions are performed.

Using Winograd or Regular FFT convolution, the output images are computed via

$$\begin{aligned} I'_{b,c'} &= \sum_{c=1}^C \left[(W_{c,c'} \times_{n=1}^N \mathbb{G}_n) \odot (I_{b,c} \times_{n=1}^N \mathbb{B}_n) \right] \times_{n=1}^N \mathbb{A}_n^T \\ &= \left[\sum_{c=1}^C (W_{c,c'} \times_{n=1}^N \mathbb{G}_n) \odot (I_{b,c} \times_{n=1}^N \mathbb{B}_n) \right] \times_{n=1}^N \mathbb{A}_n^T \end{aligned} \quad (6)$$

Note that we can have different sizes for matrices $\mathbb{A}_n$, $\mathbb{B}_n$ and $\mathbb{G}_n$ for each dimension.

$F_n(m_n, r_n)$ and $\mathcal{F}_n(m_n, r_n)$ assume a particular size of I ($m_n + r_n - 1$) and I' (m_n) along the n -th dimension. For larger image sizes, the convolution is performed using the overlap-add method (OLA) [28]. With OLA, the input images are divided into tiles with sizes of $m_n + r_n - 1$, and an overlap of $r_n - 1$ along the n -th dimension. Considering tiles at the same location from all the input images, tiles of size m_n of the output images are computed using the formula above.

¹For instance, the 4.5 TFLOPS Intel Knights Landing processor [17] has a compute-to-memory ratio of 11, whereas the latest Skylake Xeon and SkylakeX family of processors have reached the ratio in the range between 20 to almost 40.

The main savings in computation in both the Winograd and FFT methods comes from the fact that both the kernel transforms ($W_{c,c'} \times_{n=1}^N \mathbb{G}_n$), and image (tiles) transforms ($I_{b,c} \times_{n=1}^N \mathbb{B}_n$) can be precomputed and reused many times. The computation is dominated by computing the dot products – accumulation of element-wise product inside the square brackets in Eqn. 6. Computing all the dot products in Eqn. 6 is an equivalent problem to matrix multiplications, with real matrices for the case of Winograd and complex matrices for the case of Regular FFT convolution.

2.3 Gauss' Multiplication of Complex Numbers

In the Regular FFT convolution, the computation is dominated by complex matrix multiplications, where each complex number pair multiplication requires 4 real multiplications, when computed directly, and 3 real multiplications when using Gauss' multiplication algorithm [22, 24].

Using Gauss' multiplication algorithm, the product of two complex numbers $u_r + u_i i$ and $v_r + v_i i$ is computed by first computing $tmp_1 = v_r(u_r + u_i)$, $tmp_2 = u_r(v_i - v_r)$ and $tmp_3 = u_i(v_r + v_i)$. The real part of the result is then equal to $tmp_1 - tmp_3$ and the imaginary part to $tmp_1 + tmp_2$. Similarly, an element-wise product of two complex tensors $U \odot V$ ($U = U_r + U_i i$ and $V = V_r + V_i i$) can be performed using three element-wise products of real-valued tensors.

For the Regular-FFT convolutional layer, element-wise product of complex tensors representing the image (tile) transforms and kernel transforms are performed, and each tensor is reused many times (Eqn. 6). After performing a transform of an image tile, which yields a complex tensor $U = U_r + U_i i$, a real-valued tensor $U_r + U_i$ can be computed and stored alongside U_r and U_i . Similarly, tensors $V_i - V_r$ and $V_r + V_i$ can be computed during kernel transforms and stored alongside V_r (V_i does not have to be stored). Each element-wise products of complex tensors can then be replaced with three independent element-wise products of real-valued tensors.

The resulting three real tensors are implicitly converted back to a single complex tensor during the computation of inverse transform ($\times_{n=1}^N A_n^T$).

Computing all the dot products in Eqn. 6 is then performed using three real-valued matrix multiplications instead of a single complex matrix multiplication, reducing the number of required operations by 25%.

We refer to the FFT method using Gauss' multiplication as Gauss-FFT ($\mathbb{G}(m, r)$).

3 Implementations

Both Winograd and FFT approaches perform computations in four distinct stages: input transform, kernel transform, element-wise computation, and inverse transform. The first two stages convert the images/kernels from a spatial domain into Winograd or FFT domain. The third stage is equivalent to matrix multiplications. The inverse transform stage converts the results back to the spatial domain.

We extended the publicly available Winograd implementation from [3, 18], which was highly optimized for many-core CPUs, to support arbitrary AVX2 and AVX512 multi-core CPUs. We used it as a base for our FFT implementations. We reused most of the code (90%) in order to leverage the same optimization methods.

Optimizations In order to achieve high utilization of the hardware, both Winograd and FFT methods use the identical optimizations as proposed in [18], including software prefetching, memory and cache blocking, using aligned vector data access and interleaving memory access with computation.

We adopted the data layout proposed in [17, 18, 40] for input images, kernels and output, where 16 images are interleaved in memory for easy vectorization. In [18] 16 was used due to the size of the AVX512 vector register, we keep it to 16 regardless of the vector register size, as 16 is the cache-line width (16 32-bit floats), to facilitate efficient utilization of the memory subsystem.

For data hand-off between the four stages of the algorithm, streaming stores to main memory were used, since the data will not be used in near future. This saves memory bandwidth and avoids cache pollution.

Overall, all three implementations achieved high utilization of the available hardware, as discussed in the following sections.

Transforms To perform transforms of the input images and kernels, as well as the output images, the implementation of [18] provides C++ codelets that perform transforms of 16 tiles at the same time. The codelets are created by generating Winograd transforms using Wincnn [8], after which a computation graph is created and optimized, yielding codelets that utilize AVX512 instructions to transform 16 tiles at the time.

For the FFT-based implementations, the codelets were replaced by C++ codelets generated using "genfft" supplied with FFTW [13]. "Genfft" was modified so that it can generate codelets that perform implicitly zero-padded FFT transforms, as well as computing only a subset of elements of the inverse transform. Multidimensional (forward) transforms were implemented by combining codelets performing implicitly zero-padded real-to-complex transforms along one dimension, and ones performing complex-to-complex transforms along other dimensions. Backward transforms combined complex-to-complex transform codelets and complex-to-real ones. Different from existing FFT-based convolutions, which limit transform size to small prime factors [39] or only numbers that are powers of two [9, 25], our implementations support arbitrary sizes.

Element-Wise Multiplications For the element-wise stage, where matrix-matrix multiplications are performed, the implementation of [18] provides JIT routines for real-matrix multiplications optimized for AVX512 instruction set. Following the same principles of [18] we implemented JIT real-matrix multiplication routines optimized for the AVX2 instruction set, as well as complex-number multiplication routines for both AVX512 and AVX2 instruction sets, which are required for the Regular-FFT method.

Parallelization Through Static Scheduling Each stage of our algorithm is parallelized using static scheduling originally proposed in [40], using the generalized implementation provided by [18]. To achieve optimal performance, each core is assigned roughly the same amount of computation. The work is then executed using a single fork-join routine.

4 Performance Comparisons

To compare running times of the FFT and Winograd methods, we benchmarked our FFT-based implementations and the improved Winograd implementation of Jia et al. [18] by using the two most

Figure 1: Convolution layers' running time on Xeon Gold CPU.

Table 1: Machine configurations used for benchmarking. The MB represents the theoretical peak memory bandwidth, and CMR represents the ratio between the available FLOPS and the memory bandwidth.

CPU	Cores	GFLOPS	AVX	Cache	MB(GB/s)	CMR
Xeon Phi 7210	64	4506	512	512 KB	409.6	11
i7-6950X	10	960	2	1 MB	68.3	14.06
i9-7900X	10	2122	512	1 MB	96	22.1
Xeon Gold 6148	20	3072	512	1 MB	128	24
E7-8890v3	18	1440	2	256 KB	51.2	28.13
Xeon Platinum 8124M	18	3456	512	1MB	115.2	30
i9-7900X	10	2122	512	1 MB	68.3	31.07
i9-7900X	10	2122	512	1 MB	51.2	41.45
Xeon Phi 7210	64	4506	512	512 KB	102.4	44
Xeon Phi 7210	48	3379	512	512 KB	102.4	33

popular ConvNets, AlexNet [21] and VGG [29], which are frequently used for benchmarking [5]. We benchmarked the time required for the forward propagation of each distinct layer of the two networks.

Additional, publicly available, implementations were included for reference – The Winograd implementations provided by LIBXSMM [7] and MKL-DNN [2], and the direct convolution provided by MKL-DNN [2]. To the best of our knowledge, no other implementation of FFT-based methods for CPUs, beside our own, is currently available.

Both Winograd and FFT methods can work with arbitrary transform sizes. Generally, larger transform sizes decrease the number of required operations. However, the numerical inaccuracy of Winograd convolutions increases exponentially with transform (tile) sizes [18, 22, 27]. In practice, there is an upper bound on the transform size for which the Winograd produces satisfactory results. All major vendors, such as FALCON, MKL-DNN, LIBXSMM, and cuDNN [1, 2, 7, 9] implement the Winograd algorithm with transform sizes less than or equal to 6×6 ². For these tile sizes, the

numerical error of computation is comparable to the one that implements convolution directly.

We follow the same convention, and consider Winograd convolutions only with tile sizes of at most 6×6 . However, we allow the FFT methods to have an arbitrary large tile size, as they don't suffer from such numerical instability.

Running Time Experiments The benchmarks were performed on a total of 10 systems showed in Tbl. 1.

In Fig. 1 we show detailed results on one of the systems. Note that both LIBXSMM's and MKL-DNN's Winograd implementation support only kernel sizes of 3×3 , and thus can not use the Winograd method for the second layer of AlexNet.

The Winograd-based method outperformed in only 3 out of 12 layers, whereas the a FFT-based method outperformed on 6 layers; and on 3 layers, they had roughly the same performances. More importantly, in the cases when the FFT methods outperformed, they did it with a larger margin, and on the layers that require more computation time. This suggests that the FFT methods can allow for significant savings in the overall computation time, when compared to the Winograd. For instance, the time spent on all convolutional layers in AlexNet using the Winograd method would consume 58.79 milliseconds, whereas the Regular-FFT method requires only 31.96 milliseconds; that is a speedup of 1.84x.

Additionally, in Fig. 2 we show the normalized running time on all other AVX512 systems (neither LIBXSMM, nor MKL-DNN support the AVX2 instruction set). The results for each individual layer are scaled based on the slowest implementation, as we are only interested in the relative performances. Detailed results are available in our open-sourced repository [4]. In all cases, our two FFT-based implementations as well as the modified Winograd implementation of [18] outperformed other publicly available implementations. Thus, for the rest of the paper, we will only focus on these three implementations.

FFT transform sizes An important observation was that the optimal transform sizes for the FFT method were not always powers of two; they were 27 for VGG1.2, 25 for VGG2.1 and VGG2.2, 21 for VGG 3.1 and VGG3.2, 16 for VGG4.1 and VGG4.2, and 9 for VGG5.1/5.2. For AlexNet, the optimal sizes were 31 for the second

²With the transform sizes of 6×6 , the average numerical error of the Winograd method on benchmarked layers was $7.03 \cdot 10^{-6}$, which is similar to the error of direct-convolution $1.11 \cdot 10^{-6}$. When the transform size is increased to 8×8 , the error is increased to $1.24 \cdot 10^{-3}$, which was expected [18]. The numerical error of the FFT-based method was no larger than $2.88 \cdot 10^{-7}$, regardless of the tile size.

Figure 2: Normalized running time of different convolution implementations. (The systems with the same CPUs are distinguished by their CMR).

layer, and 15 for all other layers. This is counter intuitive, as the common belief is that optimal FFT transforms should use sizes that only contain factors that are small primes [13, 39] or transforms with sizes equal to the powers of two, which is suggested by the two GPU FFT-based implementations, i.e., fbfft [25, 32] and CuDNN [11].

Relative performances While an FFT-method outperformed the Winograd more often than not, the relative performances varied among different layers and systems. In the rest of the paper we focus on the theoretical analysis of all the methods. We would like to understand why our findings are not aligned with the popular belief that the Winograd method should be strictly faster.

5 Performance Analysis

Our experimental observations suggested that some of the stages in both Winograd- and FFT-based approach have relatively low utilizations of the system’s available FLOPS. In most, but not all, cases, these were the transform stages, which have relatively small amount of compute, but access relatively large amount of data, suggesting that their running time is bound by the memory bandwidth, and not the computational capabilities of the CPU.

For this reason, we used the Roofline [34] performance model to analyze Winograd- and FFT- based convolutions.

Roofline Performance Model is an ISA oblivious, throughput oriented, performance model used to estimate performance of an application running on processing units, and is often used to predict performances on CPUs, GPUs, TPUs (Tensor Processing Units), etc.

It is suitable for analyzing particular methods on systems where the memory bandwidth often becomes the constraining resource. It estimates the performance bound of an algorithm as a function of its arithmetic intensity (AI), which is defined as the ratio of total floating-point operations (FPO) and the total data movement (DM) in bytes (AI= FPO/DM) between different levels of the memory

hierarchy. For each level of memory hierarchy, the performance ceiling (attainable FLOPS) is determined by:

$$\text{Attainable FLOPS} = \min(\text{Peak FLOPS}, \text{MB} \times \text{AI}) \quad (7)$$

Where Peak FLOPS is defined as the maximal number of floating operations per second of a given system, and MB as system’s peak memory bandwidth. When plotted, the performance ceiling line resembles a roofline.

Here, we are interested in the DM between the highest level of on-chip, core-exclusive cache (typically L2 for CPUs) and off-chip, shared memory³. In systems where the L2 is shared among a small number of cores (e.g. Intel Xeon Phi series share L2 cache between two cores.), the L2 is assumed to be equally divided for exclusive usage among the cores.

DM then accounts for all of regular and streaming stores to main memory, regular loads from main memory, and pre-fetches from main memory to either L1 or L2.

Our analysis does not consider the presence of higher level, shared caches, as on modern processors, the sizes of such caches are very limited, and can not fit the working set of a typical convolutional layers.

Additional performance ceilings for lower levels of cache, are not necessary, since in both Winograd-based and FFT-based convolutional algorithms the computation is either bounded by transfers to and from main memory, or the processor peak FLOPS, as shown in the following sections.

As the performance ceiling is set by algorithm’s arithmetic intensity (Eqn. 7), its running time can be estimated by:

$$\begin{aligned} \text{running time} &= \frac{\text{FPO}}{\text{Attainable FLOPS}} \\ &= \frac{\text{FPO}}{\min(\text{Peak FLOPS}, \text{MB} \times \text{AI})} = \frac{\text{FPO}/\text{MB}}{\min(\text{CMR}, \text{AI})} \quad (8) \\ &= \begin{cases} \frac{\text{FPO}}{\text{Peak FLOPS}}, & \text{CMR} \leq \text{AI} \\ \frac{\text{DM}}{\text{MB}}, & \text{CMR} > \text{AI} \end{cases} \end{aligned}$$

Here, CMR is the system’s compute-to-memory ratio, defined as the ratio of it’s Peak FLOPS and MB – the memory bandwidth. FPO represents the total number of floating point operations required, and DM (the total amount of data movements in bytes), as defined above. The running time is *compute bound* when $\text{CMR} \leq \text{AI}$, in which case it can be computed as $\frac{\text{FPO}}{\text{Peak FLOPS}}$; otherwise, the running time is *memory bound*, and can be computed as $\frac{\text{DM}}{\text{MB}}$.

Estimating running time For the Winograd- and FFT-based convolutional layers, where the computation is composed of several sequential stages (S), each with a unique arithmetic intensity (AI), the running time is calculated by accumulating the running times of each stage $s \in S$:

$$\text{running time}_{\text{Total}} = \sum \text{running time}_s = \sum \frac{\text{FPO}_s/\text{MB}}{\min(\text{CMR}, \text{AI}_s)} \quad (9)$$

5.1 Relative Performances

We are interested in the relative performance between Winograd and FFT methods. We define $\text{Speedup}(\mathcal{A}, \mathcal{B})$ as the ratio of the

³The term off-chip refers to the main memory of the system, typically large, but much lower throughput and higher latency than the on-chip caches. The HBM MCDRAM of the Knights Landing processor would be considered off-chip.

running times of an algorithm $\mathcal{B}$ and an algorithm $\mathcal{A}$.

$$\text{Speedup}(\mathcal{A}, \mathcal{B}) = \frac{\text{running time}_{\text{Total}}^{\mathcal{B}}}{\text{running time}_{\text{Total}}^{\mathcal{A}}} \quad (10)$$

A speedup greater than one indicates that the algorithm $\mathcal{A}$ is faster, and smaller than one means that the algorithm $\mathcal{B}$ is faster. While Eqn. 9 estimates the running time of an algorithm assuming perfect utilization of the hardware, which is rarely possible in practice. However, the Eqn. 10 is valid when the compared algorithms are equally optimized, meaning that they utilize the same percentage of the hardware capabilities.

In the Eqn. 10 the value of AI will differ between Winograd and FFT based approaches, and will also depend on the amount of available cache size [18]. Therefore, when comparing performance of two algorithms, the relative performance on a given system will depend on the CMR ratio and the amount of available cache, but not on the absolute values of the system’s compute speed or memory throughput.

Detailed analysis on obtaining the values of AI, DM, and FPO are presented in the Appendix A.

Using Eqn. 10 and the values of AI, DM, and FPO calculated in the Appendix A (Tbl. 2), we estimate the speedup between the FFT-based methods and the Winograd-based one. For all three methods, the tile sizes are chosen, in a way that minimizes the total running time (Eqn. 9).

5.2 Verification of Theoretical Analysis

In Fig. 3, we plot the estimated theoretical speedup of the two FFT-based methods over the Winograd-based one. The solid lines represent the theoretical estimates of the relative performances as a function of the system’s CMR. The color of the line represents the amount of available L2 cache. The lines are drawn in a semi-translucent fashion, as they overlap in many places.

The empirical results from the benchmarks described in Section 4 are overlaid; each cross-hair represents the measurement of the relative performances on a single system. The x coordinate of the cross-hairs represents the system’s compute-to-memory (CMR) ratio (see Tbl. 1), and the color represents the L2 cache sizes.

Our empirical results are consistent with theoretical estimates. The overall relative root mean square error (**rRMSE**) was 0.079 for Regular-FFT vs Winograd and 0.1 for Gauss-FFT vs Winograd. The fitness⁴ was 92.68% for Regular-FFT vs Winograd and 90% for Gauss-FFT vs Winograd.

5.3 Discussions

Hardware utilization We have also measured the system utilization of each stage of the three methods. While the utilization varied across benchmarked layers and systems, on average, during the compute bound stages, 75% of theoretical peak FLOPS were attained; in memory bound stages slightly more than 85% of the theoretical memory bandwidth was achieved. This resulted in a somehow lower *effective* CMR, which is consistent with the results shown in Fig. 3. The empirical results are slightly shifted to the left (towards lower values of CMR), when compared to the theoretical predictions, which assume equal utilization of both the FLOPS and the memory bandwidth.

⁴fitness = $\frac{100}{1+rRMSE}$

Optimality of Tile Transforms Both “wincnn” [8] and “genfft” [13] use heuristics to minimize the number of operations for performing transforms, and might not be optimal. However, as their AIs are much smaller than the CMRs of modern systems, the computation is memory bound, and the running time of the transform stages will depend solely on the amount of data movement. The largest AI of the FFT transforms is 5.55, and for Winograd 2.38, much lower than CMRs of the modern CPUs. The Xeon Phi Knights Landing processor has the CMR of 11 (due to on-chip fast MCDRAM), and the Xeon Server processor family has CMRs in the range of 20–40. Hence, our theoretical analysis yield identical estimates as if the optimal number of operations were provided.

This is consistent with our experimental findings, where in some cases, tile sizes that were large primes, such as 31, were optimal. Using such sizes, the images could be divided into overlapping tiles with minimal overhead (minimal padding of the original image), which reduces both the number of operations, and the amount of data movements.

6 Conclusions

This paper presents theoretical analysis and experimental findings of FFT- and Winograd-based convolutions on modern CPUs which are extensively used for fast evaluation of modern ConvNets.

In contrast to the popular belief that the Winograd-based method is more efficient, our experimental results of a highly optimized Winograd-based implementation and two similarly optimized FFT-based implementations on modern CPUs show that the FFT convolutions are, more often than not, faster than Winograd ones for most commonly used ConvNet layers.

Our analysis using the Roofline model shows that whether the Winograd- or FFT-based approach is faster depends on both the layer and target hardware. However, with the tendency of increasing compute-to-memory ratio of systems with modern CPUs, the FFT-based methods tend to be faster than Winograd.

While, we primarily focused on modern multi- and many-core CPUs, which generally have larger caches, but smaller memory bandwidths than modern GPUs, we believe that our performance analysis can be applied to GPUs. Future work might include implementing and benchmarking efficient GPU-based implementations and validating performance analysis based on the Roofline model.

A Detailed Analysis

For simplicity, we assume 2D and isotropic images and kernels, as well as computation being performed with 32-bit floating point numbers (4 bytes per number). Extension to non-isotropic and N -dimensional images and kernels, as well as lower precision floats, is straightforward. The benchmarked implementations, described in Section 3 support an arbitrary number of dimensions, as well as non-isotropic kernels.

We follow the same notation as in Section 2, and for an arbitrary layer, use Winograd convolution $F(m^2, r^2)$, Regular-FFT convolution $\mathcal{F}(m^2, r^2)$ and Gauss-FFT convolution $\mathcal{G}(m^2, r^2)$. Here, m can take an arbitrary value, whereas r^2 has to equal to the sizes of the kernels in the layer.

We proceed to present the details of all three methods and estimate the total number of operations and data movement required

Figure 3: Theoretical estimates and empirical measurements for the speedup of Regular- and Gauss-FFT methods over the Winograd method on VGG and AlexNet. The x axis represent the system’s CMR.

in each stage, from which we also calculate the arithmetic intensity of each of the three method.

A.1 Input Transform Stage

In Winograd and both FFT approaches, each of the $B \cdot C$ images (B batches, each with C input channels) is divided into overlapping tiles, which are then transformed. Each tile has the size of t^2 with $t = (m + r - 1)$, and there is an overlap of $r - 1$ pixels between adjacent tiles along both dimensions.

Each image of size $x \times x$ will thus be divided into $N = \lceil (x - r + 1)/m \rceil^2$ tiles. If necessary, the image will be implicitly zero-padded. This yields a total of $BC \lceil (x - r + 1)/m \rceil^2 = BCN$ image tiles.

Number of required operations Each tile ($I_{b,c,n}$) is transformed via $\tilde{I}_{b,c,n} = \mathbb{B}I_{b,c,n}\mathbb{B}^T$, which is a composition of two matrix-matrix multiplications. Here $I_{b,c,n}$ is the n -th tile of the c -th channel in batch b , and $\tilde{I}_{b,c,n}$ is its transform. Both $\mathbb{B}$ and $I_{b,c,n}$ have the size t^2 , requiring a total of $4t^3$ operations ($2t^3$ per matrix-matrix multiplication). Operations are on real numbers

for the Winograd approach, and on complex numbers for the FFT approaches.

The transforms can, however, be performed with much fewer operations. In the Winograd method, the matrices $\mathbb{B}$ are sparse, and contain pairs of columns with similar numbers, allowing for reuse of many intermediate results. In the FFT methods, instead of matrix multiplications, 2D real-to-complex DFT transforms are performed, which require much fewer operations ($\mu n^2 \log n$ instead of $O(n^3)$). Here, the constant μ can vary significantly based on the size of the transform performed [13]. It takes small values, when the size of the transform is a power of two, and larger values when the transform size contains large prime factors.

Instead of using a closed-form expression that gives bounds on the number of operations required, we counted the number of operations in real, optimized, implementations and stored them in a lookup tables. For Winograd, we used the Winograd matrix generator [8], as proposed in [10, 18] and further reduced the number of operations using the simple optimizer provided by [18]. For the FFT methods, we used the FFT codelet generator “genfft” from the FFTW library [13]. We denote the number of operations required for transforming an image tile as $F^I(m^2, r^2)$ for the Winograd, $\mathcal{F}^I(m^2, r^2)$ for Regular-FFT, and $\mathcal{G}^I(m^2, r^2)$ for Gauss-FFT. The pre-computed lookup tables are included in our open-sourced repository.

The total number of operations required for performing all the transforms for all three methods are given in Tbl. 2.

Data movement The tile sizes are relatively small, much smaller than available cache; thus, once a tile is fetched from main memory, all the computation is done in-cache. Additionally, the overlapping regions between tiles need to be fetched from memory only once, as they can be stored in cache and reused for adjacent tiles.

The total data that has to be moved from main memory to cache is thus $BCx^2 \cdot 4$ bytes – reading each of the $B \cdot C$ images of size x^2 only once. Each of the BCN transformed tiles is stored back to main memory. The size for each transformed tile will depend on the method.

In Winograd convolution, transformed tiles are real tensors, and require a single 32-bit float per element, for a total of $4t^2$ bytes.

The FFT methods perform 2D FFT transforms of each tile. As the FFTs are performed on real tensors, the resulting transforms will be complex, conjugate-symmetric (along one of the dimensions). Only $t[(t+1)/2]$ complex numbers need to be stored. The Regular-FFT requires two real numbers per complex number, for a total of $8t[(t+1)/2]$ bytes per tile, and the Gauss-FFT requires three reals per complex number, a total of $12t[(t+1)/2]$ bytes per tile.

The total amount of data movement for the three methods, as well as their arithmetic intensities (AIs), is shown in Tbl. 2

A.2 Kernel Transform Stage

In all methods, each of the $C \times C'$ kernels (with size $r \times r$) is transformed via $\tilde{W}_{cc'} = \mathbb{G}W_{cc'}\mathbb{G}^T$. The matrix $\mathbb{G}$ has size $t \times r$.

Computing the transform of a kernel is an equivalent procedure to performing the input transform of a zero-padded kernel, to match the size of the input tiles. In practice, the transforms are computed with implicit zero-padding. As in the input stage, we

have pre-computed the number of operations required for transforming a kernel in all three methods: $F^K(m^2, r^2)$ for the Winograd, $\mathcal{F}^K(m^2, r^2)$ for Regular-FFT. For the Gauss-FFT method, $\mathcal{G}^K(m^2, r^2) = \mathcal{F}^K(m^2, r^2) + 2t[(t+1)/2]$, as we need two extra operations per complex number as described in Sec. 2.3.

All three methods need to fetch all kernel data from memory; read a total of $4CC'r^2$ bytes, and store CC' transformed kernels back to main memory. The transformed kernels have size $t \times t$, and are stored in the same fashion as the transformed input tiles, requiring total of $4t^2$, $8t[(t+1)/2]$, and $12t[(t+1)/2]$ bytes for the Winograd, Regular-FFT and Gauss-FFT methods respectively.

The total number of operations, data movements, and AIs for transforming all kernels of a particular layer, using any of the three methods, are given in Tbl. 2.

A.3 Element-wise Products

Having all transformed input and kernel tiles, the pre-transformed output tiles $\tilde{I}'_{b,c',n}$ are computed via

$$\tilde{I}'_{b,c',n} = \sum_{c=1}^C \tilde{I}_{b,c,n} \odot \tilde{W}_{c,c'} \quad (11)$$

Here, all of the pre-transformed output tiles ($\tilde{I}'_{b,c',n}$), transformed input tiles $\tilde{I}_{b,c,n}$, and transformed kernels $\tilde{W}_{c,c'}$ have the same size $t \times t$. Computing an element of the pre-transformed output tile at a location $\vec{e}$ depends only on elements of the transformed input tiles and transformed kernels at the same location, and is computed via:

$$\tilde{I}'_{b,c',n}(\vec{e}) = \sum_{c=1}^C \tilde{I}(\vec{e})_{b,c,n} \cdot \tilde{W}(\vec{e})_{c,c'} \quad (12)$$

Note that the equation above can be interpreted as multiplication of a matrix $U_{\vec{e}}$ of size $BN \times C$ with a matrix $V_{\vec{e}}$ of size $C \times C'$ resulting in a matrix $X_{\vec{e}}$ of size $BN \times C'$. For layers of modern ConvNets, values of BN are much larger than C , which results in multiplication of tall and skinny matrices [18, 22]. Such matrix multiplications have to be performed for each element location of $\tilde{I}'_{b,c',n}$.

A.3.1 Operations Required In the Winograd method, t^2 real matrices are multiplied, requiring a total of $2t^2BNCC'$ operations for the whole layer, where N is the number of tiles per image.

For the FFT based methods, complex matrices are multiplied. However, only $t[(t+1)/2]$ locations of the pre-transformed output tiles need to be computed. As both the transformed input tiles and transformed kernels are conjugate symmetric, the pre-transformed output tiles will also be conjugate symmetric.

In the Regular-FFT method, 4 real multiply-adds are required for performing a complex multiply-add. A total of $8t[(t+1)/2]BNCC'$ FLOPs are required for the whole layer. Gauss-FFT reduces complex matrix multiplications to 3 real matrix multiplications of matrices of the same size, totaling $6t[(t+1)/2]BNCC'$ FLOPs required for the whole layer.

A.3.2 Data Movement In the Winograd method, t^2 real matrix multiplications $X = U \times V$ are performed (for the rest of the section we will omit the subscript $\vec{e}$, for clarity). In the Regular-FFT method, $t[(t+1)/2]$ complex matrix multiplications are performed, with the same sizes. The Gauss-FFT method replaces one complex matrix

Table 2: FLOPS, DM and AI for different stages of the Winograd ($F(m^2, r^2)$), Regular-FFT ($\mathfrak{F}(m^2, r^2)$) and Gauss-FFT ($\mathfrak{G}(m^2, r^2)$) methods. N is the number of tiles per single image, the tiles have size of t^2 ($t = m + r - 1$). The values of c and c' depend on the amount of available cache, and are obtained by solving the optimization problem given in Eqn. 13.

Stage		Winograd	Regular-FFT	Gauss-FFT
FLOPS	Input image transform	$BCNF^I(m^2, r^2)$	$BCN\mathfrak{F}^I(m^2, r^2)$	$BCN\mathfrak{G}^I(m^2, r^2)$
	Kernel transform	$CC'F^K(m^2, r^2)$	$CC'\mathfrak{F}^K(m^2, r^2)$	$CC'\mathfrak{G}^K(m^2, r^2)$
	Element-wise Product	$2t^2BNCC'$	$8t[(t+1)/2]BNCC'$	$6t[(t+1)/2]BNCC'$
	Output transform	$BC'NF^O(m^2, r^2)$	$BC'N\mathfrak{F}^O(m^2, r^2)$	$BC'N\mathfrak{G}^O(m^2, r^2)$
DM	Input image transform	$4BCx^2 + 4BCNt^2$	$4BCx^2 + 8BCNt[(t+1)/2]$	$4BCx^2 + 12BCNt[(t+1)/2]$
	Kernel transform	$4CC'(r^2 + t^2)$	$4CC'(r^2 + 2t[(t+1)/2])$	$4CC'(r^2 + 3t[(t+1)/2])$
	Element-wise Product	$4t^2BN(c + \alpha c') \frac{CC'}{cc'}$	$8t[(t+1)/2]BN(c + \alpha c') \frac{CC'}{cc'}$	$12t[(t+1)/2]BN(c + \alpha c') \frac{CC'}{cc'}$
	Output transform	$4BC'N(t^2 + m^2)$	$4BC'N(2t[(t+1)/2] + m^2)$	$4BC'N(3t[(t+1)/2] + m^2)$
AI	Input image transform	$\frac{NF^I(m^2, r^2)}{4x^2 + 4Nt^2}$	$\frac{N\mathfrak{F}^I(m^2, r^2)}{4x^2 + 8Nt[(t+1)/2]}$	$\frac{N\mathfrak{G}^I(m^2, r^2)}{4x^2 + 12Nt[(t+1)/2]}$
	Kernel transform	$\frac{F^K(m^2, r^2)}{4r^2 + 4t^2}$	$\frac{\mathfrak{F}^K(m^2, r^2)}{4r^2 + 8t[(t+1)/2]}$	$\frac{\mathfrak{G}^K(m^2, r^2)}{4r^2 + 12t[(t+1)/2]}$
	Element-wise Product	$\frac{2(c + \alpha c')}{cc'}$	$\frac{c + \alpha c'}{cc'}$	$\frac{2(c + \alpha c')}{cc'}$
	Output transform	$\frac{F^O(m^2, r^2)}{4t^2 + 4m^2}$	$\frac{\mathfrak{F}^O(m^2, r^2)}{8t[(t+1)/2] + 4m^2}$	$\frac{\mathfrak{G}^O(m^2, r^2)}{12t[(t+1)/2] + 4m^2}$

multiplication with three real matrix multiplications; thus, a total of $3t[(t+1)/2]$ real matrix multiplications are performed.

For modern ConvNets, the matrices X , U and V may not fit in cache, in which case standard cache-blocking approaches are employed [14–16, 18, 23], and might require some data to be moved to and/or from the main memory multiple times.

In the standard cache-blocking technique, V (of size $C \times C'$) is subdivided into equally sized matrices of size $c \times c'$,

To minimize transfers to and from main memory, a sub-matrix of V is kept in cache. A small sub-matrix of U , with size of $\rho \times c$, is fetched from memory, multiplied with the sub-matrix of V stored in-cache, and accumulated to the appropriate sub-matrix of X . Here ρ is a small number, required for efficient in-cache computation [14–16, 18, 23].

This requires transferring a total of ρc numbers from main memory, and $\rho c'$ numbers bytes of X from and then back to the main memory. A total of $\rho(c + 2c')$ numbers. In the special case, when $c = C$, only $\rho(c + c')$ numbers need to be moved, as each sub-matrix multiplication produces the final result (no accumulation is necessary).

Total of $BNCC'/\rho cc'$ such sub-matrix multiplications are performed, and require $BNCC'/cc'(c + \alpha c')$ numbers to be moved. With α being 1 when $c = C$ and 2 when $c < C$.

In the Winograd method t^2 real matrix multiplications ($X = U \times V$) are performed, thus transferring a total of $4t^2BN(c + \alpha c')CC'/cc'$ bytes to be moved.

In the Regular-FFT method $t[(t+1)/2]$ complex matrix multiplications are performed, requiring a total of $8t[(t+1)/2]BN(c + \alpha c')CC'/cc'$ bytes to be moved.

The Gauss-FFT replaces each of the $t[(t+1)/2]$ complex matrix multiplications with 3 real matrix multiplications. Total of $12t[(t+1)/2]BN(c + \alpha c')CC'/cc'$ bytes need to be transferred.

Figure 4: Arithmetic intensities for the element-wise stage given various amounts of cache. In the Winograd and Gauss-FFT method, real matrix multiplications are performed, whereas, in the Regular-FFT complex matrix multiplications are performed.

The total number of operations is fixed. To minimize the amount of data movements we need to choose optimal values for c and c' , while allowing for in-cache computation.

As the values of t , B , C , C' and N are constant, the optimal values of c and c' can be chosen by solving the following optimization problem:

$$\begin{aligned}
 &\underset{c, c'}{\text{minimize}} && \frac{(c + \alpha c')}{cc'} \\
 &\text{subject to} && c \mid C && (C \text{ is divisible by } c) \\
 &&& c' \mid C' && (C' \text{ is divisible by } c') \\
 &&& 4\beta cc' \leq \frac{\text{Cache Size}}{2} && (\text{fits in half cache})
 \end{aligned} \tag{13}$$

Where α equals 1 when $c = C$ and 2 when $c < C$; β is 1 for real matrices, and 2 for complex ones. Half the cache is allowed for sub-matrices of V . This is typical practice, to ensure enough space for sub-matrices of U and X .

The arithmetic intensities can be then computed by dividing the number of required operations with the amount of required data movements for optimal values of c and c' . This results in the AIs of $cc'/2(c + ac')$ for the Winograd and Gauss-FFT methods, and $cc'/(c+ac')$ for the Regular-FFT method, which are equal to the AIs of real matrix, and complex matrix multiplications, respectively [14–16, 18].

In Fig. 4, we show the arithmetic intensities for layers with different numbers of channels as a function of cache size. The AIs of both complex (used in Regular-FFT method) and real matrix multiplication (used in Winograd and Gauss-FFT) increase with the cache size and the number of channels (C and C'). For a fixed cache size, the complex matrix multiplication allows for a higher arithmetic intensity. This indicates that the element-wise stage of the Regular-FFT method may achieve better hardware utilization than the ones of Winograd or Gauss-FFT convolution, when the cache size is a limiting resource.

A.4 Output Transform

In the output transform stage, each of the BNC' pre-transformed output tiles $\tilde{I}'_{b,c',n}$ is transformed from the Winograd/FFT domain back to the spatial domain via $I'_{b,c,n} = \mathbb{A}^T \tilde{I}'_{b,c',n} \mathbb{A}$. The matrix $\mathbb{A}$ has size of $t \times m$ ($t = m + r - 1$), resulting in $I'_{b,c',n}$ having a size of $m \times m$. Here $I'_{b,c',n}$ is the n -th (non-overlapping) tile of the c' -th image in batch b of the final result of the convolutional layer.

Performing the inverse transforms is similar to performing the input transforms, except that: 1) Only a subset of $m \times m$ elements need to be computed, and 2) Different, (inverse transform) matrices (or inverse FFT transforms) are used.

Given the analysis of the input and kernel transforms, analyzing the output transforms is straight-forward. The total number of operations, data movements, and AIs for transforming output tiles of a particular layer, using any of the three methods are given in Tbl. 2.

Acknowledgement

We thank Sebastian Seung for helpful discussions. This research was supported by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DoI/IBC) contract number D16PC0005, NIH/NIMH (U01MH114824, U01MH117072, RF1MH117815), NIH/NINDS (U19NS104648, R01NS104926), NIH/NEI (R01EY027036), and ARO (W911NF-12-1-0594). The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/IBC, or the U.S. Government. We are grateful to Toyota Research Institute for supporting the Toyota - CSAIL Joint Research Center at MIT. We are grateful for assistance from Google, Amazon, and Intel.

References

- [1] 2016. FALCON Library: Fast Image Convolution in Neural Networks on Intel Architecture. <https://colfaxresearch.com/falcon-library/>. (2016).
- [2] 2016. Intel(R) Math Kernel Library for Deep Neural Networks. <https://github.com/01org/mkl-dnn>. (2016).

- [3] 2018. N-Dimensional Winograd-based convolution framework. <https://bitbucket.org/poozh/ond-winograd>. (2018).
- [4] 2018. Winograd and FFT Convolution. <https://bitbucket.org/jiazhentim/winograd-fft-conv/src/master>.
- [5] Accessed: 01-31-2019. Imagenet Winners Benchmarking. <https://github.com/soumith/convnet-benchmarks>. (Accessed: 01-31-2019).
- [6] Accessed: 01-31-2019. Intel® Nervana reference deep learning framework. <https://github.com/NervanaSystems/neon>. (Accessed: 01-31-2019).
- [7] Accessed: 01-31-2019. LIBXSMM. <https://github.com/hfp/libxsmm>. (Accessed: 01-31-2019).
- [8] Accessed: 01-31-2019. Wincnn. <https://github.com/andravin/wincnn>. (Accessed: 01-31-2019).
- [9] Jeremy Appleyard. Accessed: 01-31-2019. Optimizing Recurrent Neural Networks in cuDNN 5. <https://devblogs.nvidia.com/parallelforall/optimizing-recurrent-neural-networks-cudnn-5/>. (Accessed: 01-31-2019).
- [10] David Budden, Alexander Matveev, Shibani Santurkar, Shraman Ray Chaudhuri, and Nir Shavit. 2016. Deep Tensor Convolution on Multicores. *arXiv preprint arXiv:1611.06565* (2016).
- [11] Sharan Chetlur, Cliff Woolley, Philippe Vandermersch, Jonathan Cohen, John Tran, Bryan Catanzaro, and Evan Shelhamer. 2014. cudnn: Efficient primitives for deep learning. *arXiv preprint arXiv:1410.0759* (2014).
- [12] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. 2009. Imagenet: A large-scale hierarchical image database. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. IEEE, 248–255.
- [13] Matteo Frigo and Steven G Johnson. 1998. FFTW: An adaptive software architecture for the FFT. In *Acoustics, Speech and Signal Processing, 1998. Proceedings of the 1998 IEEE International Conference on*, Vol. 3. IEEE, 1381–1384.
- [14] Dennis Gannon, William Jalby, and Kyle Gallivan. 1988. Strategies for Cache and Local Memory Management by Global Program Transformation. In *Proceedings of the 1st International Conference on Supercomputing*. Springer-Verlag, London, UK, 229–254. <http://dl.acm.org/citation.cfm?id=647970.761024>
- [15] Alexander Heinecke, Greg Henry, Maxwell Hutchinson, and Hans Pabst. 2016. LIBXSMM: accelerating small matrix multiplications by runtime code generation. In *High Performance Computing, Networking, Storage and Analysis, SC16: International Conference for*. IEEE, 981–991.
- [16] Alexander Heinecke, Hans Pabst, and Greg Henry. 2015. Libxsmm: A high performance library for small matrix multiplications. *Poster and Extended Abstract Presented at SC* (2015).
- [17] James Jeffers, James Reinders, and Avinash Sodani. 2016. *Intel Xeon Phi Processor High Performance Programming: Knights Landing Edition*. Morgan Kaufmann.
- [18] Zhen Jia, Aleksandar Zlateski, Fredo Durand, and Kai Li. 2018. Optimizing N-Dimensional, Winograd-Based Convolution for Manycore CPUs. In *Proceedings of the 23rd ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*.
- [19] Zhen Jia, Aleksandar Zlateski, Fredo Durand, and Kai Li. 2018. Towards Optimal Winograd Convolution on Manycores. *SysML* 2018.
- [20] Tamara G Kolda and Brett W Bader. 2009. Tensor decompositions and applications. *SIAM review* 51, 3 (2009), 455–500.
- [21] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*. 1097–1105.
- [22] Andrew Lavin and Scott Gray. 2016. Fast algorithms for convolutional neural networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4013–4021.
- [23] Jiajia Li, Casey Battaglini, Ioakeim Perros, Jimeng Sun, and Richard Vuduc. 2015. An input-adaptive and in-place approach to dense tensor-times-matrix multiply. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. ACM, 76.
- [24] M Donald MacLaren. 1970. The art of computer programming. Volume 2: Seminumerical algorithms (Donald E. Knuth). *SIAM Rev.* 12, 2 (1970), 306–308.
- [25] Michael Mathieu, Mikael Henaff, and Yann LeCun. 2013. Fast training of convolutional networks through ffts. *arXiv preprint arXiv:1312.5851* (2013).
- [26] Guido F Montufar, Razvan Pascanu, Kyunghyun Cho, and Yoshua Bengio. 2014. On the number of linear regions of deep neural networks. In *Advances in neural information processing systems*. 2924–2932.
- [27] Victor Y Pan. 2016. How bad are Vandermonde matrices? *SIAM J. Matrix Anal. Appl.* 37, 2 (2016), 676–694.
- [28] Lawrence R Rabiner and Bernard Gold. 1975. Theory and application of digital signal processing. *Englewood Cliffs, NJ, Prentice-Hall, Inc.*, 1975. 777 p. (1975).
- [29] Karen Simonyan and Andrew Zisserman. 2014. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).
- [30] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. 2015. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*. 1–9.
- [31] Vincent Vanhoucke, Andrew Senior, and Mark Z Mao. 2011. Improving the speed of neural networks on CPUs. In *Proc. Deep Learning and Unsupervised Feature Learning NIPS Workshop*, Vol. 1. 4.

- [32] Nicolas Vasilache, Jeff Johnson, Michael Mathieu, Soumith Chintala, Serkan Piantino, and Yann LeCun. 2014. Fast convolutional nets with fbfft: A GPU performance evaluation. *arXiv preprint arXiv:1412.7580* (2014).
- [33] Kevin Vincent, Kevin Stephano, Michael Frumkin, Boris Ginsburg, and Julien Demouth. 2017. On Improving the Numerical Stability of Winograd Convolutions.
- [34] Samuel Williams, David Patterson, Leonid Oliker, John Shalf, and Katherine Yelick. 2008. The roofline model: A pedagogical tool for auto-tuning kernels on multicore architectures. In *Hot Chips*, Vol. 20. 24–26.
- [35] Shmuel Winograd. 1980. *Arithmetic complexity of computations*. Vol. 33. Siam.
- [36] Wm A Wulf and Sally A McKee. 1995. Hitting the memory wall: implications of the obvious. *ACM SIGARCH computer architecture news* 23, 1 (1995), 20–24.
- [37] Aleksandar Zlateski, Zhen Jia, Kai Li, and Fredo Durand. 2018. A Deeper Look at FFT and Winograd Convolutions. SysML 2018.
- [38] Aleksandar Zlateski, Kisuk Lee, and H Sebastian Seung. 2016. ZNN–A Fast and Scalable Algorithm for Training 3D Convolutional Networks on Multi-core and Many-Core Shared Memory Machines. In *Parallel and Distributed Processing Symposium, 2016 IEEE International*. IEEE, 801–811.
- [39] Aleksandar Zlateski, Kisuk Lee, and H Sebastian Seung. 2016. ZNNi: maximizing the inference throughput of 3D convolutional networks on CPUs and GPUs. In *High Performance Computing, Networking, Storage and Analysis, SC16: International Conference for*. IEEE, 854–865.
- [40] Aleksandar Zlateski and H Sebastian Seung. 2017. Compile-time optimized and statically scheduled ND convnet primitives for multi-core and many-core (Xeon Phi) CPUs. In *Proceedings of the International Conference on Supercomputing*. ACM, 8.