



Tilus: A Tile-Level GPGPU Programming Language for Low-Precision Computation

Yaoyao Ding^{*†}

University of Toronto
Toronto, ON, Canada
yaoyao@cs.toronto.edu

Bohan Hou

Carnegie Mellon University
Pittsburgh, PA, USA
bohanhou@andrew.cmu.edu

Xiao Zhang^{*†}

University of Toronto
Toronto, ON, Canada
bohanhou@andrew.cmu.edu

Allan Lin

University of Waterloo
Waterloo, ON, Canada
a8lin@uwaterloo.ca

Tianqi Chen^{*}

Carnegie Mellon University
Pittsburgh, PA, USA
tqchen@cmu.edu

Cody Hao Yu

Independent Researcher
Santa Clara, CA, USA
hao.yu.cody@gmail.com

Yida Wang

Amazon Web Services
Santa Clara, CA, USA
wangyida@amazon.com

Gennady Pekhimenko^{*†}

University of Toronto
Toronto, ON, Canada
pekhimenko@cs.toronto.edu

Abstract

Serving Large Language Models (LLMs) is critical for AI-powered applications, yet it demands substantial computational resources, particularly in memory bandwidth and computational throughput. Low-precision computation has emerged as a key technique to improve efficiency while reducing resource consumption. Existing approaches for generating low-precision kernels are limited to weight bit widths that are powers of two and suffer from suboptimal performance because of high-level GPU programming abstractions. These abstractions restrict critical optimizations, such as fine-grained register management and optimized memory access patterns, that are essential for efficient low-precision computations. In this paper, we introduce **Tilus**, a domain-specific language designed for General-Purpose GPU (GPGPU) computing that supports low-precision data types with arbitrary bit widths from 1 to 8 while maintaining GPU programmability. Tilus features a thread-block-level programming model, a hierarchical memory space, a novel algebraic layout system, and extensive support for diverse low-precision data types. Tilus programs are compiled into

highly efficient GPU programs through automatic vectorization and instruction selection. Extensive experiments demonstrate that Tilus efficiently supports a full spectrum of low-precision data types, and outperforms state-of-the-art low-precision kernels. Compared to existing compilers such as Triton and Ladder, as well as hand-optimized kernels such as QuantLLM and Marlin, Tilus achieves performance improvements of: 1.75 \times , 2.61 \times , 1.29 \times and 1.03 \times , respectively. We open-source Tilus at <https://github.com/NVIDIA/tilus>.

CCS Concepts: • Computing methodologies → Parallel programming languages; Machine learning.

Keywords: GPU; programming language; parallel computation; low-precision computation; quantization

ACM Reference Format:

Yaoyao Ding, Bohan Hou, Xiao Zhang, Allan Lin, Tianqi Chen, Cody Hao Yu, Yida Wang, and Gennady Pekhimenko. 2026. Tilus: A Tile-Level GPGPU Programming Language for Low-Precision Computation. In *Proceedings of the 31st ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1 (ASPLOS '26)*, March 22–26, 2026, Pittsburgh, PA, USA. ACM, New York, NY, USA, 17 pages. <https://doi.org/10.1145/3760250.3762219>

^{*}Also with NVIDIA.

[†]Also with Vector Institute.



This work is licensed under a Creative Commons Attribution 4.0 International License.

ASPLOS '26, Pittsburgh, PA, USA

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2165-6/26/03

<https://doi.org/10.1145/3760250.3762219>

1 Introduction

The development of Large Language Models (LLMs) has revolutionized natural language processing tasks, enabling advanced capabilities in areas such as text generation [9], summarization [31], translation [59], and conversational AI [43]. However, serving LLMs poses substantial computational challenges due to the large model sizes and high computational demands. Efficient LLM serving demands innovative computational strategies to manage latency and

power consumption constraints. As such, optimizing LLM inference has become a priority in both industry and research to reduce latency and increase throughput of LLM serving.

Quantization [8, 11, 20, 30, 32, 54] has emerged as a leading method for enhancing the efficiency of LLM serving. By reducing the bit-width of model parameters and activations, quantization reduces weight storage, DRAM bandwidth usage, and achieves faster computation. For instance, A16W4 quantization (16-bit activation and 4-bit weight) reduces DRAM consumption and throughput by 4× compared to the A16W16 scheme, thereby reducing the time to generate a token by about 4× [21]. However, the state-of-the-art 4-bit quantization methods [8, 11, 32] still suffer from non-negligible accuracy degradation. While using 5- to 7-bit quantization can mitigate this accuracy loss [3, 60], the lack of efficient GPU kernels for these bit widths has hindered their adoption. Generating optimized kernels for these hardware-unfriendly bit widths (e.g., 3, 5, 6, and 7) remains an important open problem.

Existing methods for generating computation kernels fall into two main categories: manually written kernels [21, 60] and compiler-generated kernels [12, 53, 58]. While manually written kernels are highly optimized for specific hardware, they are time-consuming and error-prone to develop, and difficult to generalize to new architectures and evolving quantization methods. For example, QuantLLM [60] only supports floating-point quantization for 5- and 6-bit data types but lacks support for sub-channel quantization granularity. Similarly, Marlin [21] is limited to 4-bit signed integer quantization and does not support Hopper GPUs [33].

To address these limitations, compiler-based approaches [12, 53, 58] have been proposed to automate kernel generation. Among them, Triton [53] simplifies the GPGPU programming through a tile-based model. A Triton program defines the computation of a *thread block* for tensor *tiles*. However, Triton lacks built-in support for low-precision data types, and thus requires users to implement low-level bitwise operations manually. Additionally, it does not expose the full GPU memory hierarchy, limiting optimization opportunities for low-precision LLM inference. Ladder [58], on the other hand, extends TVM's scheduling system [12] to support low-precision computation but is restricted to data types with bit widths that are powers of two. Moreover, its primitives cannot express crucial optimizations such as software pipelining [26], leading to suboptimal performance, particularly for batch sizes greater than one during LLM decoding with *continuous batching* [63] enabled.

In response to these challenges, we propose **Tilus**, a tile-level GPGPU programming language backed by a virtual machine with dedicated support for low-precision computation. Tilus abstracts GPU program execution into thread-block-level instructions, simplifying GPGPU programming while exposing hierarchical memory spaces for fine-grained manipulation of sub-tensors in on-chip memory. This dual

approach enables efficient handling of arbitrary-precision data types while reducing the complexity of GPU programming. To achieve these goals, Tilus introduces: (1) **an algebraic layout system** that specifies how tensor elements within a tile are distributed across GPU threads. This layout system provides a unified way to represent the storage of tile elements among different threads, making it possible to expose the registers to the kernel developers and simplify the code generation during compilation. (2) **a thread-block-level programming model with fine-grained memory management**, providing explicit control over data movement, placement, and computation across different levels of the GPU memory hierarchy; and comprehensive support for (3) **arbitrary low-precision data types** including signed integers, unsigned integers, and floating-point numbers with bit widths ranging from 1 to 8.

Extensive experiments show that Tilus extends the spectrum of efficient low-precision kernels to support arbitrary bit widths (from 1 to 8) and data type kinds (e.g., integer and floating-point numbers). At the same time, Tilus outperforms the state-of-the-art compilers, such as Triton [53] and Ladder [58], as well as hand-crafted kernels from QuantLLM [60] and Marlin [21], achieving performance improvements of: 1.75×, 2.61×, 1.29× and 1.03×, respectively.

We summarize our key contributions as follows:

- We propose a GPGPU programming language Tilus with dedicated support for low-precision computation, addressing the critical bit-width coverage and performance gap in existing approaches.
- Within Tilus, we introduce a novel layout system, a thread-block-level programming model with hierarchical memory space, and support low-precision data types with arbitrary bit-width from 1 to 8.
- Through extensive evaluation, we demonstrate that Tilus generates a full spectrum of highly efficient low-precision kernels, achieving up to 2.6× performance improvement over state-of-the-art approaches on their supported kernels.

2 Background

2.1 LLM Serving and Quantization

LLM serving consists of two inference stages: prefill and decode. The prefill stage processes the input prompt to establish context, while the decode stage iteratively generates output tokens based on prior tokens. Key LLM layers include multi-head attention, feed-forward networks, and layer normalization [56]. Among these, matrix multiplications dominate computation time and memory consumption, making their optimization crucial for efficient LLM serving. Quantization [13, 20] improves their efficiency by reducing model weights and activations to lower-precision formats, such as 8-bit or 4-bit integers. It reduces memory usage, bandwidth requirements, and inference latency while aiming to preserve

model accuracy. While 4-bit quantization provides significant computational savings, state-of-the-art methods [8, 11, 32] still suffer from accuracy degradation. Increasing precision to 5-bit, 6-bit, or 7-bit quantization [3, 60] can help preserve accuracy while maintaining efficiency, but these bit widths lack optimized GPU support, limiting their adoption. Current GPU architectures and software stacks primarily optimize for power-of-two bit widths (e.g., 4-bit and 8-bit), making arbitrary bit widths computationally inefficient. However, demand for flexible quantization is growing, as 4-bit can be too aggressive for some models while 8-bit wastes resources. Supporting a broader spectrum of bit widths enables better accuracy-efficiency trade-offs in LLM serving, driving the need for new kernel generation techniques that can efficiently handle non-standard low-precision formats (e.g., those with 3, 5, 6, 7 bit widths) on modern GPUs.

2.2 GPGPU Programming

General-Purpose GPU (GPGPU) programming enables parallel computation by organizing tasks within a structured execution and memory hierarchy [37]. The execution hierarchy begins with the *thread*, the smallest unit of execution, which performs instructions independently, using its own registers and local memory. Threads are grouped into *thread blocks*, which enable data sharing through shared memory and support synchronized execution. A *grid* consists of multiple independent thread blocks, enabling large-scale parallelism by organizing thousands or millions of threads. The GPU memory hierarchy comprises registers, shared memory, and global memory. *Registers* provide the fastest and thread-private storage. *Shared memory* is accessible by all threads within a thread block and faster than global memory. *Global memory* is accessible across the entire grid with high latency. This structure allows for highly efficient parallel execution by leveraging both the execution and memory hierarchies.

2.3 The GPGPU Languages and Compilers

2.3.1 GPGPU Programming Languages. GPGPU programming involves various languages and compilers that balance hardware abstraction with control. Low-level languages like SASS [42] and CDNA3 [6] offer direct hardware access for fine-grained optimizations but require deep architectural knowledge. Slightly higher in abstraction, NVIDIA’s PTX [41] serves as an intermediate representation that links high-level languages like CUDA [40] to GPU-specific instructions while preserving optimization flexibility. High-level languages like CUDA [40] and HIP [7] simplify programming by extending the C programming language. Despite these languages, GPGPU programming remains complex. It is constrained by hardware-specific memory and computation hierarchies and requires workload-specific optimizations. To address these challenges, researchers have introduced higher-level languages and compilers, classified into two categories: *tile-oriented compilers*, which simplify programming through

abstractions beyond CUDA [40], and *schedule-oriented compilers*, which optimize computation-hardware mappings via declarative scheduling primitives.

2.3.2 Tile-Oriented Compilers. This type of compilers, such as Graphene [23], Hidet [14], and Triton [53], enables programmers to write kernels directly, offering abstractions like tile types or tile-level task/element distribution to simplify the process. Triton [53], for instance, introduces the tile programming model, where thread block behavior is defined programmatically, and tiles replace scalars as the basic data type. This approach combines programming simplicity with high-performance kernel generation, making Triton widely adopted. However, Triton lacks native support for low precision data types like `uint4`. Handling these types requires manually unpacking sub-byte data from larger storage types (e.g., `uint32`) [25]. Additionally, Triton does not expose the GPU memory hierarchy, limiting programmers’ control over data loading and memory scope usage, which complicates performance optimization for low-precision kernels. These limitations result in inefficient low-precision kernel execution. Figure 1(a) illustrates the inefficiencies in Triton-generated low-precision kernels, using a `uint4` weight loading pipeline as an example. The process includes four steps: ❶ weights are asynchronously copied from global memory to shared memory using pipelined `cp.async` instructions [26]; ❷ shared memory data is loaded into registers; ❸ unpacking and casting operations are performed; and ❹ the register tensor layout is converted to meet the requirements of tensor core instructions. Among these, step ❹ is a major bottleneck due to the reliance on shared memory for layout conversion, which incurs significant overhead.

2.3.3 Schedule-Oriented Compilers. Schedule-oriented compilers decouple computation from scheduling to optimize the computation-to-hardware mappings. Halide [46] pioneered this approach, which was later extended by TVM [12] and subsequent works [19, 26, 49, 57, 58, 67, 68] in the domain of deep learning. Among them, Ladder [58] is the first one to support low-precision computation by introducing dedicated primitives to pack low-precision data (e.g., 4-bit integers) into larger types (e.g., 8-bit integers). However, Ladder [58] has two limitations. First, it cannot handle non-power-of-two bit widths efficiently due to *type-level packing*, packing low-precision types into storage types. Second, its primitive-style scheduling prevents optimizations like *software pipelining* [26], resulting in suboptimal performance. Figure 1 (b) illustrates the weight loading process in Ladder’s low-precision kernels. This process includes ❶ loading weights from global memory to registers without pipelining; ❷ vectorized casting; ❸ storing the cast results in shared memory; and finally ❹ using the `ldmatrix` instruction to load weights from shared memory to registers for subsequent tensor core operations. This lack of pipelining between weight loading and computation significantly hinders performance.

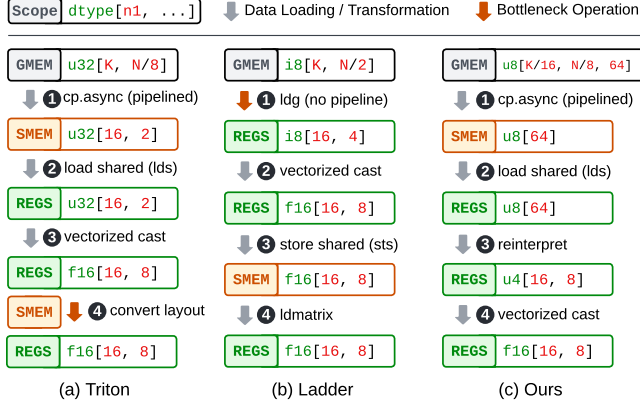


Figure 1. The weight loading pipeline of Triton, Ladder, and our approach. The tensors could be in global memory (GMEM), shared memory (SMEM), or registers (REGS).

3 System Overview

3.1 Key Ideas

Our work introduces a domain-specific language, Tilus, that provides fine-grained control over shared memory and registers, making it possible to program efficient low-precision deep learning kernels. Tilus supports low-precision data types with arbitrary bit widths ranging from 1 to 8, enabling efficient weight loading and computation. Figure 1 (c) shows the weight loading pipeline of a Tilus program, using `uint4` as an example. It begins with ① a pipelined asynchronous memory copy from global memory to shared memory, followed by ② loading the register tensor from shared memory. Next, it ③ *reinterprets* the register tensor into a different data type and layout at no cost, before finally ④ performing vectorized casting. This pipeline achieves superior efficiency compared to the other methods in Figure 1, as it eliminates layout conversion (unlike Triton [53]) and incorporates pipelining (unlike Ladder [58]). More importantly, our pipeline is generic, making our work the first to seamlessly support arbitrary low-precision data types with bit widths ranging from 1 to 8 bits.

To achieve this efficiency, our design is built on several key ideas. **A GPGPU Virtual Machine:** every Tilus program is a program for an abstract GPGPU virtual machine (VM) that contains an instruction set. This decision stems from the need for greater flexibility in GPU programming. By abstracting GPU functionalities, such as memory loading and computation, into instructions, it becomes easier to add support for new architectural features while keeping support for older ones. **A Thread-Block-Level Programming Model with Hierarchical Memory Spaces:** The underlying VM explicitly exposes the GPU memory hierarchy — including registers, shared memory, and global memory — that existing solutions like Triton [53] abstract away. By granting programmers fine-grained control over data placement

and movement, our approach enables memory pipelining and eliminates unnecessary layout conversions, as shown in Figure 1. **An Algebraic Layout System:** we introduce an algebraic layout system that precisely defines how elements within a register tensor are distributed among threads. This structured representation simplifies the construction, analysis, and interpretation of tensor layouts. Notably, it enables seamless reinterpretation of low-precision register tensors into standard data types, as demonstrated in Step ③ of Figure 1(c). **Native Support for Arbitrary Low-Precision Data Types:** Tilus provides built-in support for a wide range of low-precision data types, including both signed and unsigned integers and floating-point numbers with bit widths from 1 to 8. Supported types include `int2` to `int8`, `uint1` to `uint8`, and `float3` to `float8`, with arbitrary *exponent* and *mantissa* distribution for floating-point types. These innovations collectively enhance the programmability, efficiency, and flexibility of low-precision kernel development on modern GPUs. We chose not to extend Triton [53] because its programming model inherently abstracts away tensor layouts, making it incompatible with our approach of explicit layout control. Similarly, Ladder [58] relies on type-level packing, whereas Tilus employs tile-level reinterpretation, making the two fundamentally incompatible. The next section presents a Tilus-programmed example of low-precision matrix multiplication.

3.2 An Example of Tilus Program

Figure 2 illustrates a low-precision matrix multiplication in Tilus. Matrix multiplication is defined as $C_{M,N} = A_{M,K} \times B_{K,N}$, where A and B are `float16` (a 16-bit floating-point number [27]) and `int6` (a 6-bit signed integer), respectively. The kernel performs matrix multiplication with dimensions M , N , and K , where each thread block computes a $BM \times BN$ tile of the C matrix (Line 1). Therefore, a grid of $(M / BM, N / BN)$ thread blocks must be launched (Line 2). Inside the kernel, the `BlockIndices` instruction retrieves the thread block indices bi and bj (Line 3), which determine the offset $(bi * BM, bj * BN)$ for computing the corresponding C tile. Three tensor views are created for the input and output tensors in global memory by specifying their addresses and shapes (Line 4-6). Then, a register tensor of type `f16[16, 8]` is created with the following layout:

```
local(2, 1).spatial(8, 4).local(1, 2).
```

It distributes $16 \times 8 = 128$ elements across 32 threads. Each thread stores 4 elements (Line 7). This layout is *composed* of three *primitive layouts* (Section 4) and aligns with the C matrix layout used by the `mma.m16n8k16` tensor core instruction in PTX [41]. The reduction loop over the k dimension (Line 8-13) repeatedly loads tiles of A and B from global memory into registers and accumulates their product. For each iteration, we first load a `f16[16, 16]` tile from global memory to register with a `LoadGlobal` instruction (Line 9). The layout of

```

1 const M, N, K, BM, BN, BK = 1024, 1024, 1024, 16, 8, 16 # M, N, K are matmul sizes while BM, BN, BK are tile sizes
2 def matmul<M / BM, N / BN>(f16* a_ptr, u8* transformed_b_ptr, f16* c_ptr):
3     bi, bj = BlockIndices() # block index
4     ga = ViewGlobal(a_ptr, dtype=f16, shape=[M, K]) # create global tensor views
5     gb = ViewGlobal(transformed_b_ptr, dtype=u8, shape=[K/BK, N/BN, BK*BN*6/8])
6     gc = ViewGlobal(c_ptr, dtype=f16, shape=[M, N])
7     acc = AllocateRegister(f32, layout=local(2, 1).spatial(8, 4).local(1, 2), init=0.0)
8     for bk in range(K / BK):
9         a = LoadGlobal(ga, layout=column_local(2, 2).spatial(8, 4).local(1, 2), offset=[bi * BM:, bk * BK:])
10        b = LoadGlobal(gb, layout=local(3).spatial(32), offset=[bk, bj, 0:]) # load sub-tensor from global memory to registers
11        b1 = View(b, dtype=i6, layout=local(2, 1).column_spatial(4, 8).local(2, 1)) # reinterpret register tensor
12        b2 = Cast(b1, dtype=f16)
13        Dot(a, b2, acc, out=acc) # acc = dot(a, b2) + acc
14        acc = Cast(acc, f16)
15        StoreGlobal(acc, gc, offset=[bi * BM:, bj * BN:])
16 )

```

Algebraic Layout System (Section 4)

Thread-Block-Level Programming Model with Hierarchical Memory Space (Section 5)

(Type, Variables, Constant, Layout, Instructions, Language Constructs)

(a) The VM Program of Low-Precision Matmul (FP16 x INT6)

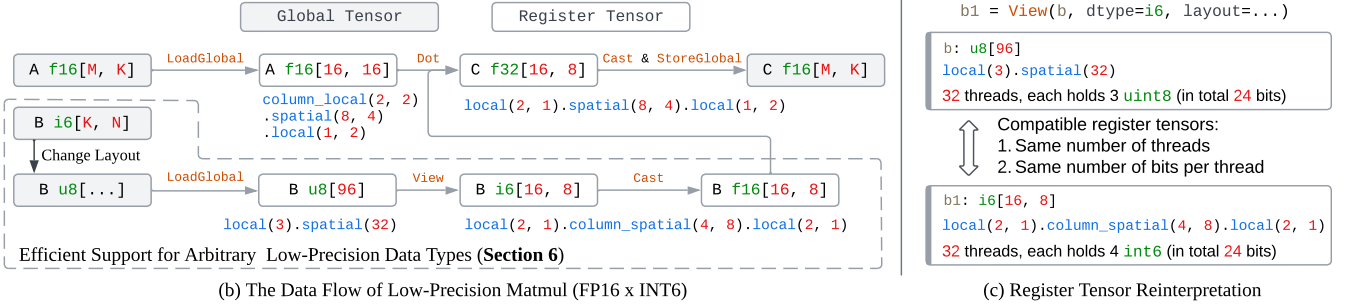


Figure 2. This figure provides a concrete example of how the Tilus is used to implement low-precision matrix multiplication (FP16 × INT6). Figure (a) illustrates the virtual machine program, highlighting key features such as the algebraic layout system (Section 4), thread-block-level instructions (Section 6), and efficient low-precision data support. Figure (b) illustrates the kernel’s data flow, emphasizing tensor movement across the memory hierarchy and intermediate operations such as tensor reinterpretation and type casting. This similar weight-loading strategy can be applied to arbitrary type widths (Section 7). Finally, Figure (c) demonstrates register tensor reinterpretation, showing how tensors with compatible bit distributions across threads (e.g., 24 bits per thread) can be efficiently reinterpreted into different data types and layouts.

the loaded register tile is specified and required by the tensor core instruction. The offset parameter specifies the position of the loaded tile within the global tensor. Loading tensor B, with data type `int6`, involves a more complex process, detailed in Section 7. We summarize the high-level ideas here. As a pre-processing step before launching the kernel, the weight tensor’s layout in global memory is transformed from `i6[K, N]` to `u8[K / BK, N / BN, BK * BN * 6 / 8]`, enabling efficient loading via the `LoadGlobal` instruction (the ‘Change Layout’ step in Figure 2 (b)). Next, in the kernel, the transformed tile is loaded into a register tensor (Line 10) and then *reinterpreted* to a tensor with a different data type and layout (Line 11). This reinterpretation is valid because both tensors are stored across the same number of threads (32), with each thread holding exactly 24 bits: $3 \times \text{u8}$ or $4 \times \text{i6}$, as shown in Figure 2 (c). Following this, the `i6` tensor is cast to an `f16` tensor (Line 12), which is then fed to the tensor core to perform matrix-multiply accumulate (mma) (Line 13). Finally, the accumulation tensor is cast from `f32` to `f16` and stored in global memory (Line 14-15). For simplicity, this program does not use shared memory and omits optimizations like software pipelining [26]. Additionally, each *k*-iteration performs only a single tensor core instruction [40].

The following sections introduce the three core components of Tilus. Section 4 introduces an algebraic layout formulation to systematically define how the elements of a tile are stored in the registers among the block threads. Section 6 introduces the thread-block-level programming model with a hierarchical memory space exposed explicitly. Section 7 introduces the native support for arbitrary low-precision data types to address the growing demand for low-precision computation in deep learning workloads.

4 Algebraic Layout System

Tilus exposes a hierarchical memory space to programmers, comprising global memory, shared memory, and registers. We need a way to model the mapping between the logical index of a tensor element and the location of the corresponding element in memory for all three memory scopes. Such a mapping is usually called the *layout* of the tensor. Of the memory scopes, the layout for register tensors is the most complicated. Figure 3 illustrates an example of the layout used by a tensor core instruction: `mma.m16n8k8.f32.f16.f16.f32 D, A, B, C`. It performs the following computation: $D_{16,8} = A_{16,8}B_{8,8} + C_{16,8}$ where *A, B, C, D* are tensors stored in thread registers and distributed across the 32 threads in a warp. Since the elements are spread across different threads, we

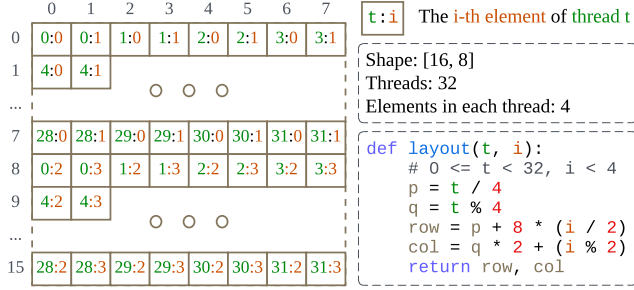


Figure 3. Layout of operand A in a Tensor Core instruction. The operand, with 16×8 elements, is distributed across 32 threads, with each thread storing four elements. The logical index of each element is determined by a layout function given the thread index t and the local element index i .

refer to this layout as a *distributed layout* [53]. Such a layout can be defined as a function f that maps a thread index t and a local index i to the logical index $f(t, i)$ of the tensor element. For example, the layout in Figure 3 can be represented as:

$$f(t, i) = (t/4 + i/2 \times 8, t \% 4 \times 2 + i \% 2)^1$$

Here, t ranges from 0 to 31, and i from 0 to 3. The function $f(t, i)$ represents the logical index of the element stored in the local element i in thread t .

4.1 Parameterized Primitive Layouts

0:0	0:1	0:2	local(2, 3)	Local Layout
0:3	0:4	0:5	$f(t, i) = (i / 3, i \% 3)$	
0:0	1:0	2:0	spatial(2, 3)	Spatial Layout
3:0	4:0	5:0	$f(t, i) = (t / 3, t \% 3)$	

Figure 4. Two types of primitive layouts: local and spatial. A local layout stores all tile elements within a single thread, whereas a spatial layout distributes them across multiple threads, with each thread holding only a single element.

With the formal definition of layout, we introduce parameterized primitive layouts that serve as the fundamental building blocks of our layout algebra. Given a tile² with shape (n_1, n_2) , there are two primary ways to store it: (1) to store all $n_1 n_2$ elements in a single thread, or (2) to distribute all elements across $n_1 n_2$ threads, with each thread holding a single element. We refer to the first type as *local layouts*, denoted as `local(n_1, n_2)`, and the second as *spatial layouts*, denoted as `spatial(n_1, n_2)`. This concept naturally extends to tiles with arbitrary dimensions. Figure 4 illustrates these two primitive layouts. The `local(2, 3)` layout maps the i -th local element of thread t to the logical index $(i/3, i \% 3)$, while the `spatial(2, 3)` layout maps it to $(t/3, t \% 3)$. We

²We use the terms *tile* and *tensor* interchangeably.

observe that the layouts for all common operators in LLMs can be constructed using these two primitive layouts. In the next section, we introduce the Kronecker product of layouts to construct more complex layouts.

4.2 Kronecker Product

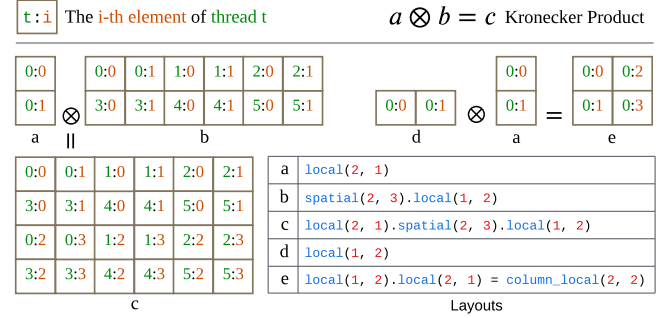


Figure 5. Examples of Kronecker products over layouts. In the figure, layout (c) is the product of layouts (a) and (b), while layout (e) is the product of layouts (d) and (a).

The layouts used in modern deep learning workloads, as well as those defined by hardware instructions, typically exhibit a hierarchical structure. Consider layout (c) in Figure 5 as an example. This layout has shape $(4, 6)$, storing 24 elements across 6 threads. Each thread holds four elements. We denote the four elements stored in each thread as a_0, a_1, a_2, a_3 . Comparing its first two rows with the last two, we observe a similar structure, except that the last two rows store elements in a_2 and a_3 instead of a_0 and a_1 . To model this structural invariance, layout (c) can be viewed as a *Kronecker product* of layouts (a) and (b). Indeed, layouts (a) and (b) can be multiplied to represent layout (c):

$$c(t, i) = a(t/6, i/2) \odot (2, 6) + b(t \% 6, i \% 2),$$

where $0 \leq t < 6, 0 \leq i < 4$, \odot denotes the element-wise product, and $(2, 6)$ represents the shape of layout (c). The Kronecker product can be generalized. Given two layouts f and g with the same number of dimensions, we define their Kronecker product, $h = f \otimes g$ as

$$h(t, i) = f(t/T_g, i/N_g) \odot S_g + g(t \% T_g, i \% N_g),$$

where T_g, N_g, S_g represent the number of threads, the number of local elements per thread, and the shape of layout g , respectively. We can prove that this operation is *associative*, meaning that for any three layouts f, g , and h , the equality $f \otimes (g \otimes h) = (f \otimes g) \otimes h$ holds. However, the operation is not commutative, meaning that, in general, $f \otimes g \neq g \otimes f$. Spatial and local layouts follow a row-major ordering for threads and local elements, respectively. Using this operation, we can construct their column-major counterparts, `column_spatial(...)` and `column_local(...)`,

as demonstrated by layout (e) in Figure 5. Returning to the tensor core instruction layout in Figure 3, it can be expressed as a Kronecker product $\text{local}(2, 1) \cdot \text{spatial}(8, 4) \cdot \text{local}(1, 2)$. We can also define its inverse operation. If $h = f \otimes g$, we define $f = h/g$ as the result of dividing layout h by layout g . For example, dividing $\text{local}(2, 4)$ by $\text{local}(1, 2)$ results in $\text{local}(2, 2)$. We also refer to the Kronecker product as layout composition in some contexts.

5 Unified Layout Representation

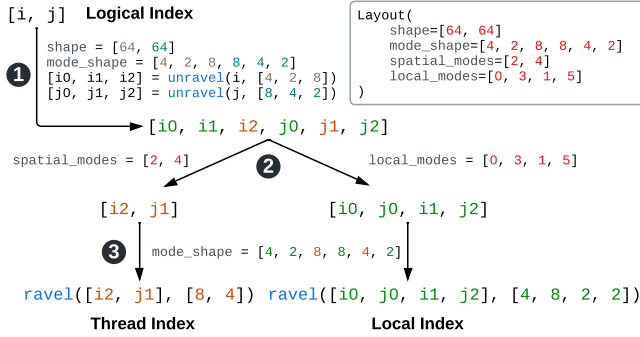


Figure 6. Example of the unified layout representation.

We use a unified representation for all layouts of register tensors in Tilus. This representation gives each layout four attributes: shape, mode_shape, spatial_modes, and local_modes. The shape is a sequence of integers that defines the shape of the register tensor. We can split each dimension of the register tensor into some sub-dimensions (we call them mode following prior work [23]) and concatenate these sub-dimensions to get the mode_shape. Then we use spatial_modes and local_modes to specify the sub-dimensions assigned to spatial threads and to the local storage of each thread. The dimension split-distribute-merge method uniquely defines a register layout.

Figure 6 shows an example of a layout and how we map the logical index of a register tensor element to the pair of thread_index and local_index. There are three steps: 1) split dimensions, 2) distribute sub-dimensions; and 3) merge sub-dimensions. Given a logical index $[i, j]$, we first ① split each index into the indices of its sub-dimensions (i.e., $[i0, i1, i2]$ for i and $[j0, j1, j2]$ for j) with the `unravel` operation. After that, we ② distribute the sub-dimension indices to get the indices for spatial threads (i.e., $[i2, j1]$) and local storage ($[i0, j0, i1, j2]$). Finally, we ③ convert the multi-dimensional indices for threads and local storage into linear index to get the thread index and local index. The `ravel` and `unravel` functions are used to convert between multi-dimensional index in a grid with given shape and its row-major linear index. For example: `unravel(i, [4, 2, 8]) = [i / 16, i / 8 % 2, i % 8]`, and `ravel([i2, j1], [8, 4]) = i2 * 4 + j1`.

The layouts represented in this form are closed under the Kronecker product, meaning that the product result of two layouts in this form can also be represented in this form.

6 Thread-Block-Level Programming Model

Modern GPU programming models, such as PTX [41] and CUDA [40], define operations at the thread level, following the Single-Instruction-Multiple-Thread (SIMT) paradigm [1]. To simplify GPU programming, we adopt the thread-block-level programming model, defining operations at the granularity of thread blocks rather than individual threads. Additionally, building on the layout system introduced previously, we propose explicitly exposing the hierarchical memory structure in modern GPUs, enabling fine-grained memory control while reducing programming complexity. We refer to this model as Single-Instruction-Multiple-Block (SIMB). In this section, we will introduce the type system, program structure, and instruction set of Tilus.

6.1 State Space and Type System

Tilus supports three variable types. *Scalar variables* store individual values, such as integers (e.g., `int32`) or floating-point numbers (e.g., `float16`). *Pointer variables* store memory addresses rather than direct data values. *Tensor variables* represent multi-dimensional arrays, with types that specify their shape, element type, memory scope, and layout. Tensors reside in different memory scopes, including global memory, shared memory, and registers. The tensor layout determines how high-dimensional tensor elements are mapped to linear memory. All variables in Tilus operate at the thread-block level, meaning that all threads within a block collaboratively allocate and maintain their state.

6.2 Program Structure and Control Flow

Figure 7 illustrates the structure of a Tilus program. Each program consists of a program name, a grid shape, a list of parameters, and a program body. The grid shape is specified as a list of expressions enclosed in `<...>`, where each expression is either a positive integer or an integer expression based on the program parameters. If the grid shape contains parameter-based expressions, its dimensions are determined at runtime based on the program's launch arguments. The program body consists of a sequence of statements, including if-else statements, range-based for-loops, and while-loops. Unlike other low-level virtual machines [41] or instruction set architectures (ISAs) [42], our virtual machine does not abstract control-flow statements into jump instructions. Instead, it retains high-level control structures to improve readability and ease of programming for human developers. In addition to control-flow statements, individual instructions can also serve as statements. Most of Tilus's functionality is implemented as instructions within its instruction set.

Table 1. The thread-block-level instruction set of Tilus’s virtual machine. Each instruction specifies an operation applied to the entire thread block. Parameters enclosed in [...] are optional. Instructions that return a new register tensor also have an in-place variant, which writes the result to an existing register tensor using the *out* parameter instead of creating a new tensor.

Category	Instruction	Semantics
Indexing	<code>indices = BlockIndices()</code>	Get the indices of the current thread block in the grid.
Tensor Creation	<code>global_tensor = AllocateGlobal(dtype, shape, [layout])</code>	Allocate a tensor in global memory with the given data type, shape, and optional layout.
	<code>shared_tensor = AllocateShared(dtype, shape, [layout])</code>	Allocate a tensor in shared memory with the given data type, shape, and optional layout.
	<code>register_tensor = AllocateRegister(dtype, shape, [layout])</code>	Allocate a tensor in registers with the given data type, shape, optional layout and init value.
	<code>global_tensor = ViewGlobal(ptr, [dtype], shape, [layout])</code>	Create a tensor view in global memory given the pointer, data type, shape, and layout.
Tensor Transferring	<code>register_tensor = LoadGlobal(global_tensor, layout, offset)</code>	Load a tensor from global memory to register, given the layout and offset.
	<code>register_tensor = LoadShared(shared_tensor, layout, offset)</code>	Load a tensor from shared memory to register, given the layout and offset.
	<code>StoreGlobal(register_tensor, global_tensor, offset)</code>	Store a register tensor in the global tensor at the given offset.
	<code>StoreShared(register_tensor, shared_tensor, offset)</code>	Store a register tensor in the shared tensor at the given offset.
	<code>CopyAsync(shared_tensor, global_tensor)</code>	Issue an asynchronous copy task from global tensor to shared tensor .
	<code>CopyAsyncCommitGroup(), CopyAsyncWaitGroup(n)</code>	Commit CopyAsync instructions as a group, or wait until there are only n ongoing groups.
Register Tensor Computation	<code>c = {Add, Sub, Mul, Div, Mod}(a, b); b = Neg(a)</code>	Arithmetic operations.
	<code>b = Cast(a, dtype)</code>	Cast a register tensor from one data type to another without changing the layout.
	<code>b = View(a, [dtype], [layout])</code>	Reinterpret a register tensor with another data type and layout at no cost.
	<code>d = Dot(a, b, c)</code>	Compute the dot product: $d = \text{dot}(a, b) + c$.
Control	<code>Synchronize(), Exit()</code>	Synchronize or exit the thread block execution.
Debug	<code>Print(tensor)</code>	Print a tensor to standard output.

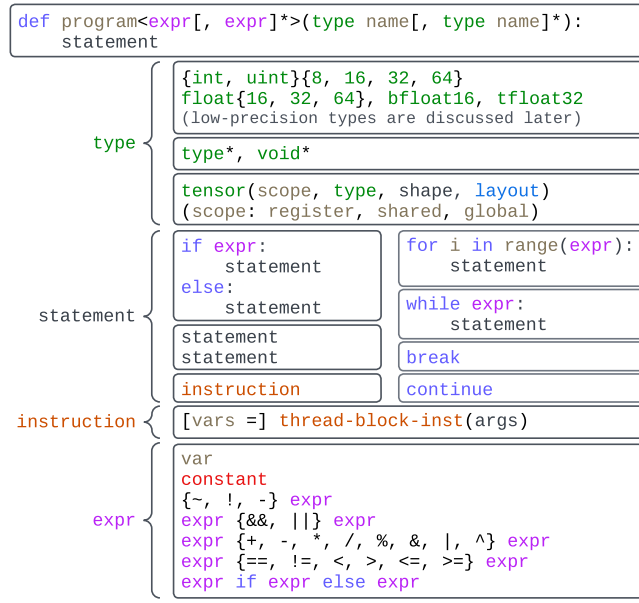


Figure 7. A Tilus program contains parameters and a body. The body is a list of control-flow statements or block-level instructions. The majority of functionality, such as tensor allocation and computation, is provided by instructions.

6.3 Thread-Block-Level Instruction Set

Each instruction in the Tilus’s instruction set operates at the thread-block level rather than the thread level. Table 1 shows a list of the instructions in the instruction set, with the

signature of each instruction and a brief description of the instruction semantics. These instructions allocate tensors with specific data types, shapes, and layouts in designated memory spaces (e.g., global memory, shared memory, registers), transfer tensors between memory spaces, and perform computations or transformations on register tensors. The execution model of modern GPUs allows different warps to execute different instructions at the same time. Similarly, the execution of instructions in Tilus exhibits this behavior: certain subsequent instructions may begin execution before the current instruction completes, resulting in multiple block-level instructions being in-flight simultaneously. Generally, this behavior does not pose significant issues. However, an exception occurs when two instructions access the same region of shared or global memory, and the second instruction depends on the completion of the first. In such cases, a Synchronize instruction must be inserted to ensure all preceding instructions complete before subsequent ones execute. Instructions like Print are used for debugging.

7 Arbitrary Low-Precision Data Types

Modern processors use bytes (8 bits each) as the smallest processing unit. As a result, standard data types in modern programming languages typically have bit widths that are multiples of 8. However, the high computational and memory demands of LLMs make low-precision data types with less than 8 bits essential for reducing resource consumption. This section describes how Tilus efficiently supports low-precision data types with bit width from 1 to 8.

7.1 Storage of Low-Precision Data

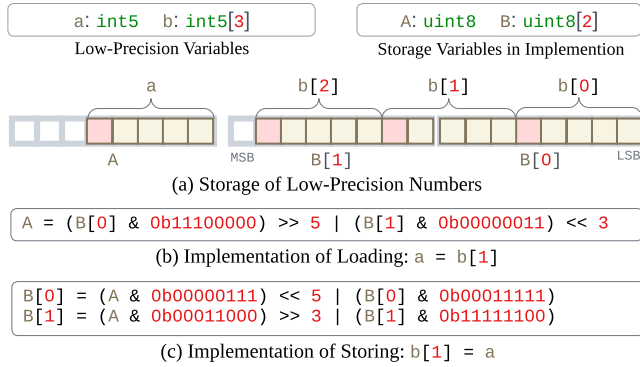


Figure 8. Compact storage and access of low-precision data. Figure (a) illustrates the use of the uint8 type to store low-precision data, where some elements may span two consecutive bytes. Figure (b–c) illustrate the implementation of loading and storing low-precision elements.

Since modern processors, including CPUs and GPUs, use bytes as the smallest unit for memory access and computation, we store low-precision data (fewer than 8 bits per element) *compactly* within bytes, as shown in Figure 8. Compact storage eliminates bit gaps between consecutive low-precision values, which may result in a single value spanning two uint8 entries (e.g., $b[1]$ in Figure 8). Bitwise operations are employed to extract, manipulate, and store low-precision values within packed byte arrays. To load a low-precision value, we first extract relevant bits using bitwise AND, adjust their position with bitwise SHIFT operations, and finally combine separated parts using bitwise OR if the value spans multiple bytes. Similarly, to store a low-precision value, we first clear the target bit positions using a bitwise mask, then insert the new value using bitwise OR while preserving the other bits. Low-precision data is cast to standard data types before arithmetic computations and is cast back afterward. While these methods enable support for arbitrary bit-width data types, they are often inefficient. They serve only as a fall-back mechanism. More efficient handling of low-precision data is necessary for LLM serving.

7.2 Efficient Low-Precision Support in LLMs

Low-precision kernels in LLMs typically follow two steps before computation: (1) loading weights into on-chip memory (registers or shared memory) from global memory, and (2) casting and de-quantizing low-precision weights to high-precision (e.g., float16). Efficient memory loading and casting are thus critical for performance. **Efficient Low-Precision Weight Loading.** With the low-precision support as discussed in the previous subsection, we can use the LoadGlobal instruction to load low-precision tensors. However, directly loading in this way is inefficient due to multiple bitwise

```

1 const M, N, K, BM, BN, BK = 1024, 1024, 1024, 16, 8, 16
2 def transform_b<K / BK, N / BN>(
3     i6* b_ptr, u8* transformed_b_ptr
4 ):
5     bk, bj = BlockIndices()
6     b_in = ViewGlobal(b_ptr, dtype=i6, shape=[K, N])
7     b_out = ViewGlobal(transformed_b_ptr, dtype=u8,
8         shape=[K / BK, N / BN, BK * BN * 6 / 8])
9 )
10 # the register tensor b is stored in 32 threads, each
11 # thread holds 4 int6 elements (in total 24 bits)
12 b = LoadGlobal(b_in,
13     layout=column(2, 2).spatial(8, 4).local(1, 2),
14     offset=[bk*BK:, bj*BM:])
15 )
16 # reinterpret the tensor with uint8 data type and a
17 # new layout. The 24 bits held by each thread will
18 # be reinterpreted into 3 uint8 elements.
19 b = View(b, dtype=u8, layout=local(3).spatial(32))
20 StoreGlobal(b, b_out, offset=[bk, bj, 0:])
21 )

```

Figure 9. Program to rearrange tensor B with data type int6, used in the "Change Layout" step of Figure 2 (b).

operations and non-coalesced memory accesses [40]. To address this, we transform the weight tensor layout in global memory to facilitate more efficient loading. Without transformation, loading a register tensor with dtype i6 and layout `local(2, 1).column_spatial(4, 8).local(2, 1)` results in non-contiguous memory accesses, causing multiple *memory access transactions* [40]. Moreover, extracting low-precision bits requires additional bitwise operations. To optimize this, we identify a *compatible* tensor type with dtype uint8 and layout `local(3).spatial(32)`, which retains the number of threads and thread local elements, while enabling efficient memory loading. As illustrated in Figure 9, we partition the weight tensor $[K, N]$ into tiles of shape $[BK, BN]$. Each tile is reinterpreted from $i6[BK, BN]$ to $u8[BK * BN * 6 / 8]$ (Line 19) and stored contiguously (Line 20). This allows us to load tiles efficiently using the hardware-friendly instructions in Figure 2 (Line 10, 11), while also enabling pipelined asynchronous memory transfers like standard data types and avoiding any layout conversion that relies on shared memory. This method generalizes to loading any low-precision tensor with arbitrary layout. More formally, given a tensor with n bytes per thread and T threads, we reinterpret it using dtype uint8 and layout `local(n2).spatial(T).local(n1)`, where $n_1 = \gcd(n, 16)$ and $n_2 = n / \gcd(n_1, 16)^3$.

Efficient Casting. After loading, weights must be cast from low-precision to high-precision (e.g., float16) for computation, especially if hardware lacks native support for the given low-precision format. We leverage target-specific instructions for efficient vectorized casting. On CUDA, we use the PRMT (permute bytes in a 32-bit register), LOP3 (arbitrary logical operation on three inputs), and bitwise instructions to execute casting with minimal overhead, as all operations are performed within registers and do not require any communication between threads.

³ $\gcd(a, b)$ represents the greatest common divisor of a and b

8 Implementation

Tilus comprises five main components: a domain specific language (DSL) in Python, an intermediate representation (IR), optimization passes, a code generator, and a runtime system. The DSL enables developers to write Tilus programs in Python, which are then translated into the VM's IR for further processing. Optimization passes refine the IR by eliminating redundancies and simplifying arithmetic expressions. The code generator translates the optimized IR into Hidet IR [14], a CUDA C-like intermediate representation. Subsequently, we apply the transformations from Section 7 to implement low-precision operations using standard precision types while preserving original semantics. The final CUDA C code is generated from Hidet IR and compiled into a hardware binary using the `nvcc` compiler [40]. The runtime system manages dynamically loaded binaries and provides the execution environment. The entire system consists of approximately 35K lines of Python and C++ code.

8.1 Program Compilation and Runtime

Tilus provides a user-friendly interface that allows programmers to write Tilus programs directly in Python, simplifying the integration of produced kernels with the rich deep learning ecosystem. Given a program, we take several steps to compile it to GPU executable code.

Step 1: Global and Shared Memory Planning Each GPU kernel can use a shared memory space of a size determined at launch-time. To simplify GPU programming, we allow the users to allocate shared memory multiple times in the program on demand. Thus, we need a shared memory planner to calculate the size of shared memory the Tilus program needs, and map the shared tensor to one region of the kernel's shared memory space. Similar to shared memory planning, we also require a global memory planner to manage the allocation of global memory shared by all thread blocks. We request Tilus's runtime system to allocate a workspace in global memory, enabling the kernel to use this workspace via the `AllocateGlobal` instruction during its execution.

Step 2: Code Emitting for Each Instruction. We emit low-level GPU code for each Tilus instruction sequentially. In our implementation, we use the Hidet IR [14] to represent the low-level GPU code. During this process, we perform instruction selection to choose the most efficient low-level instructions where feasible. For example, we use `lds` PTX instruction [41] to load the data from shared memory to register. However, a more efficient PTX instruction `ldmatrix` could also be used if the layout of the loaded register tensor is compatible with the layout `spatial(8, 4).repeat(1, 4)`, which means it can be divided properly by the latter. Besides, we also perform automatic vectorization for memory loading and storing instructions. For example, we use vectorized instructions, such as `cp.async.v4`, `lds128`, and `ldg128`, to maximize memory access efficiency.

Step 3: Lowering Low-Precision Data Types. After we emit thread-block-level instructions to the low-level IR, we will apply the passes that implement the rules discussed in Section 7.1 to transform all low-precision operations in the low-level IR to corresponding operations on hardware-friendly types. In most cases, only the vectorized type casting from low-precision type to standard type (e.g., `float16`) will be applied since the memory loading of low-precision data will be replaced by standard types thanks to our layout formalization and register tensor reinterpretation. Subsequently, we generate CUDA code (for NVIDIA GPUs) from the low-level IR and then use the `nvcc` compiler to produce the hardware binary that can be dynamically loaded.

Step 4: Loading by Runtime System. The compiled binary is dynamically loaded by the Tilus's runtime system. The runtime system also maintains internal states to serve the kernel execution: 1) a workspace memory allocated on-demand, which compiled kernels can request via `AllocateGlobal` instruction; 2) an execution context that stores the CUDA stream on which the kernel are launched; and 3) the kernels cached in memory to avoid unnecessary recompilation.

9 Evaluation

9.1 Experimental Setup

Workloads. We benchmark three representative LLMs with varying model sizes: Gemma-2-9B [52], QWen2.5-32B [62], and Llama-3.3-70B-Instruct [22]. Both the prefill and decode stages are evaluated. For operator-level analysis, we focus on matrix multiplication kernels extracted from these models. Tilus supports all kernels supported by Triton in principle, but we focus on quantized matmul in this work.

Baselines. We compare our approach, Tilus, against the vendor library cuBLAS [39], state-of-the-art DL compilers Triton [53] and Ladder [58], and hand-crafted kernels QuantLLM [60] and Marlin [21]. Auto-tuning for Triton [53] and Ladder [58] was enabled, while QuantLLM [60] used its heuristic policy to select kernel hyperparameters. For end-to-end evaluations, we integrate our quantized kernels into the state-of-the-art LLM serving framework vLLM [29] and compare them against vLLM [29] and Ladder [58] in end to end execution. The specific versions of the tools are: vLLM v0.5.3, Triton v3.1.0, bitblas v0.0.1.dev15 (Ladder), QuantLLM with commit 9802c5a, and Marlin v0.1.1.

Hardware Configuration. Experiments were primarily conducted on a server equipped with an NVIDIA L40S GPU (48 GiB), with GPU driver 565.57.01 and CUDA Toolkit 12.6.3. Benchmarks were also performed on NVIDIA A100 and H100 GPUs to demonstrate the general applicability of our approach across different hardware platforms.

Experimental Protocol. For operator experiments, each kernel was executed 50 times, while for model experiments, each model was executed 10 times. In both cases, latency was measured using CUDA Events [40], and the median latency

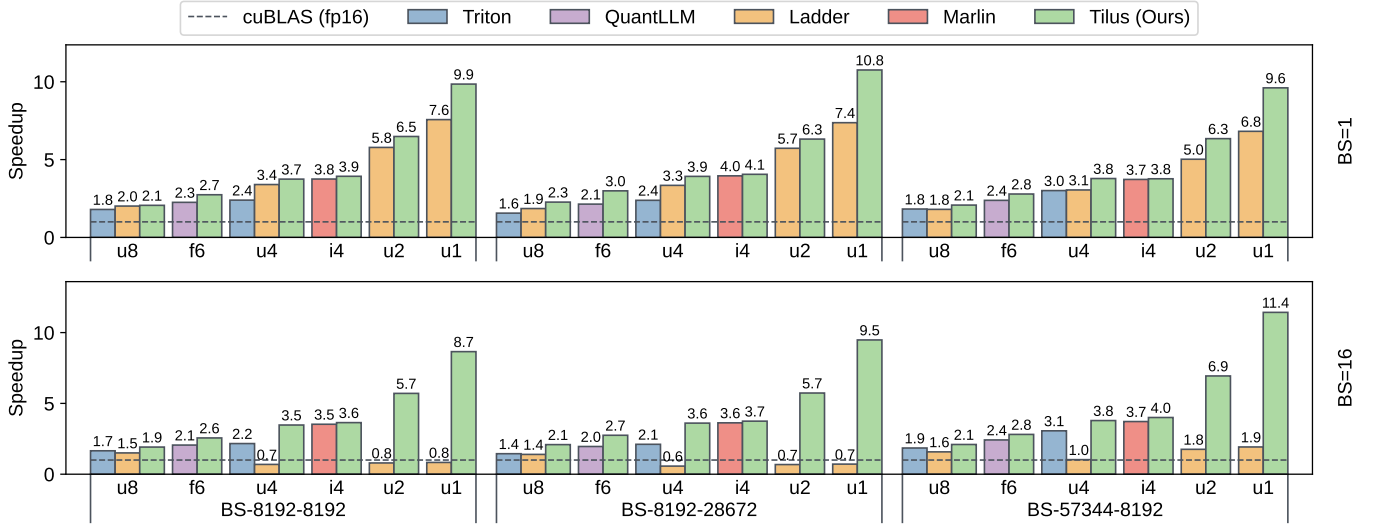


Figure 10. Speedup of low-precision kernels in Triton, QuantLLM, Ladder, and Tilus (Ours) compared against the standard half-precision kernel from cuBLAS. Benchmarked data types include uint8 (u8), float6_e3m2 (f6), int4 (i4), uint4 (u4), uint2 (u2), and uint1 (u1). Each workload (BS-N-K) corresponds to a matrix multiplication in Llama-3.3-70B, with batch sizes 1 and 16.

was reported. To eliminate artifacts from consecutive runs, the L2 cache was cleared before each execution.

9.2 Performance of Low-Precision Kernels

A single virtual machine program template is implemented to support matrix multiplication with all quantized types, taking tile sizes as tunable hyperparameters. We denote the performance of this auto-tuned program as Tilus in the evaluation. Figure 10 compares the speedup of Triton [53], Ladder [58], QuantLLM [60], Marlin [21], and Tilus (ours) against cuBLAS [39] for various low-precision matrix multiplications: uint8 (u8), float6_e3m2 (f6), uint4 (u4), int4 (i4), uint2 (u2), and uint1 (u1). While each baseline supports a limited set of quantized data types, Tilus consistently achieves speedups across all cases. For small batch sizes, the primary bottleneck is loading weights from global memory to registers for computation on SIMT or Tensor Cores. Triton struggles here due to costly layout conversions after weights are loaded into registers. Although changing layout in global memory could mitigate this, Triton’s programming model lacks explicit layout control, making such optimizations infeasible. Ladder improves upon Triton by modifying data layouts in global memory, avoiding redundant conversions. However, it lacks critical optimizations such as software pipelining [26, 38], and its type-level packing limits efficient support for arbitrary bit widths, leading to underutilized memory bandwidth. Expert-crafted kernels from QuantLLM [60] and Marlin [21] are optimized for specific quantization schemes but lack flexibility and maintainability. In contrast, Tilus consistently outperforms all baselines using

a single parameterized Tilus program template, which efficiently supports a full range of quantization types through a well-abstracted programming model.

9.3 Arbitrary Data Type Support

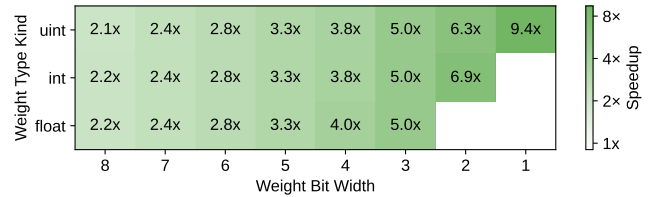


Figure 11. Speedup of quantized matrix multiplication compared against the cuBLAS FP16 kernel. A full spectrum of quantized data types is evaluated.

Tilus supports low-precision matrix multiplications of the form $\text{matmul}(A, B)$, where operand A can have data types with 32, 16, or 8 bits, and weight B supports a wide range of bit widths, from 32 bits down to 1 bit. Standard data types such as float32, float16, and int8 are supported, along with customized low-precision types with fewer than 8 bits, which include signed integers, unsigned integers, and floating-point formats with arbitrary exponent and mantissa distributions. Leveraging the algebraic layout system (Section 4 and 7.2), Tilus enables efficient memory access for low-precision data. Figure 11 illustrates the speedup achieved for the full spectrum of quantized weight data types: uint1 to uint8, int2 to int8, and float3 to float8. Representative exponent-mantissa distribution of floating-point data types such as e4m3, e3m3, e3m2, e2m2, e2m1, and e1m1 are chosen.

Each row represents the type kind (e.g., unsigned integer, signed integer or floating data type) while each column represents the bit width. Using matrix multiplication dimensions of $BS=16$, $K=8192$, and $N=57344$ the results demonstrate substantial speedups. These findings validate Tilus’s effectiveness in supporting arbitrary low-precision types with high efficiency, making it a robust solution for low-precision computations in modern GPUs. Notably, all kernels are generated from the same program template by parameterizing tile sizes, which limits the required programming effort. There are around 200 configurations per operator, and it takes around one minute to compile. We used `float16` as the activation data type in the experiment and we also support `bfloat16` and `int8`.

9.4 End-to-End Performance

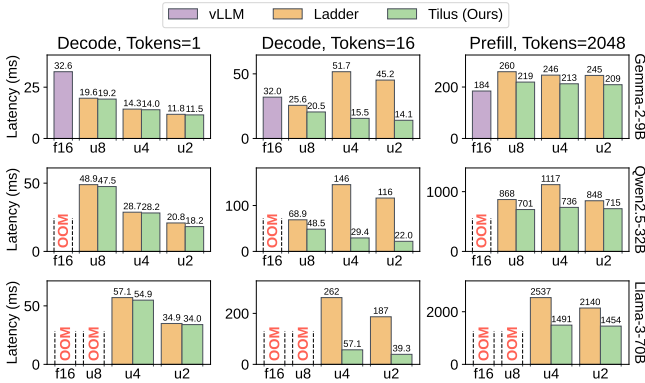


Figure 12. End-to-end performance across representative LLMs. The first two columns correspond to the decode stage with 1 and 16 tokens, respectively, while the third column shows latency for the prefill stage with 2048 prompt tokens.

We evaluated the end-to-end performance of representative LLMs: Gemma-2-9B [52], Qwen-2.5-32B [62], and Llama-3.3-70B [22], across both prefill and decode stages. The prefill stage processes all prompt tokens at once, generating the kv-cache for subsequent token generation. The decode stage then iteratively generates one token at a time. Prefill latency determines the time-to-first-token (TTFT), while decode latency impacts the speed of subsequent token generation. Both stages are critical for optimizing user experience and system utilization. Contiguous batching [29, 63] was used to efficiently batch multiple decode requests. Figure 12 shows the latency of both stages across these models. Our method consistently outperforms Ladder [58], particularly in the decode stage for batch sizes greater than one (middle column of Figure 12). Analysis of Ladder’s generated kernels revealed suboptimal use of CUDA Cores for 1–15 tokens and Tensor Cores for 16 or more tokens, as key optimizations like software pipelining [26] and k-dimension parallelization [44] were not implemented, leading to poor performance. For the

prefill stage, quantized weights are decoded to `float16`, and computations are performed using standard `f16xf16` matrix multiplication kernels, as computation becomes the bottleneck at this stage. Our efficient handling of quantized weight layouts ensures minimal overhead for decoding, contributing to the superior performance observed.

9.5 Case Studies

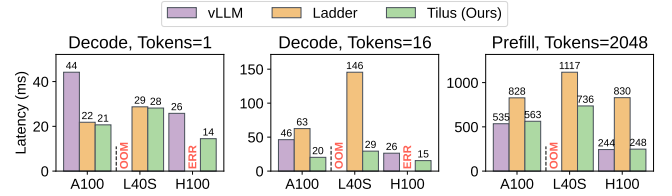


Figure 13. End-to-end performance of the Qwen2.5-30B model across NVIDIA A100, L40S, and H100 GPUs. The weight data types for vLLM, Ladder, and Tilus are `float16`, `uint4`, and `uint4`, respectively. OOM indicates out-of-memory error, and ERR indicates a runtime error.

9.5.1 Speedup over Different Hardware. We evaluate the end-to-end performance of the Qwen2.5-30B model on NVIDIA A100, L40S, and H100 GPUs, which correspond to the Ampere, Ada Lovelace, and Hopper architectures, respectively. Figure 13 presents a performance comparison of vLLM [29] (`float16`), Ladder [58] (`uint4`), and Tilus (`uint4`, ours) across the decode and prefill stages. On the Hopper architecture (H100), Ladder is unable to generate valid kernels, leading to a CUDA error (‘an illegal instruction was encountered’), which we denote as ERR in the figure. On the L40S GPU, vLLM [29] exceeds the available 48 GiB DRAM capacity, leading to out-of-memory (OOM) errors. In all other configurations, Tilus consistently outperforms Ladder across all GPUs and both processing stages, highlighting its robust performance and adaptability across architectures.

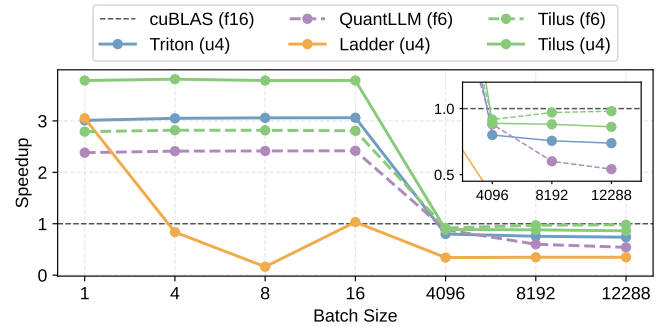


Figure 14. Speedup of quantized matmuls across different batch sizes from both prefill and decode stages.

9.5.2 Speedup over Different Batch Sizes. We analyze the relationship between speedup and batch size by benchmarking matrix multiplication performance under different batch sizes. For the decode stage, we evaluate batch sizes of 1, 4, 8, and 16, while for the prefill stage, we use batch sizes of 4096, 8192, and 12,288. The batch size corresponds to the number of tokens processed in one step. In the prefill stage, it equals the sum of sequence length of all requests, while in decode stage, it equals to the number of requests since each request only generates one token each time. Experiments are conducted on Llama-3.3-70B-Instruct [22] model with quantized data types `float6_e3m2` (f6) and `uint4` (u4), using $k = 8192$ and $n = 57344$. As shown in Figure 14, Tilus consistently outperforms baselines across all batch sizes that are used in both decode and prefill stages of LLM serving.

10 Related Work

Many deep learning compilers adopt loop-oriented scheduling [12, 46] and build auto-tuning frameworks on top of it [2, 4, 19, 49, 55, 57, 58, 61, 66–69]. In contrast, Tilus employs a procedure-oriented approach that better models GPU hardware, improving programmability and flexibility. Beyond loop-oriented scheduling, tensor programs are often optimized using vendor libraries (e.g., cuBLAS [39]), predefined templates for efficient matrix multiplication [38], hardware-aware tiling strategies [71], and domain-specific compilers for linear algebra [48]. While these methods prioritize performance, they lack extensibility for arbitrary low-precision data types. Other research focuses on optimizing irregular or ragged tensor programs [17, 51], operator fusion [65, 70], dynamic shape handling [16, 50, 66, 72], and scheduling independent operators [15, 28, 34]. Prior works [18, 35, 45, 64] endeavor to support quantization with arbitrary bit-width by independently processing different bits. However, these methods are primarily limited to integer quantization and often exhibit suboptimal execution performance. PartIR [5] also introduces a layout system for tiles; however, its abstraction level is higher than Tilus’s, and it does not detail the assignment of tile elements to threads. ExTensor [24] and FuseMax [36] leverage einsum to define computations, implicitly determining memory access patterns. Nevertheless, they do not specify how tile elements are distributed among GPU threads. Tilus can also accommodate codebook quantization, such as LCQ [10], to facilitate low-precision computation through the addition of a new lookup instruction. Microscaling data types [47] can be thought as a more fine-grained quantization thus we could also support it. These techniques are complementary to our focus on efficient low-precision computation. Triton [53] introduces a tile-based programming model. However, it lacks explicit support for low-precision data types and does not expose the GPU memory hierarchy, limiting optimization opportunities. Similarly, Hidet [14], which serves as our backend,

does not provide built-in support for low-precision types. Graphene [23] presents an intermediate representation (IR) with a layout representation. Unlike Graphene’s focus on strides and computation, our algebraic layout system emphasizes hierarchical organization. In fact, we can express Graphene’s layout representation as a primitive component within our system.

11 Conclusion

We introduced Tilus, a tile-level GPGPU programming language designed to expose shared memory and registers to developers, which enables the creation of efficient low-precision kernels for LLM serving. Tilus features an algebraic layout system for managing tensor distribution across threads, a thread-block-level programming model with fine-grained memory management, and comprehensive support for sub-byte data types, enabling arbitrary precision from 1 to 8 bits. Our experimental results demonstrate substantial performance gains over state-of-the-art approaches, showcasing the flexibility and scalability of our method. This work establishes a foundation for efficient and extensible LLM inference, paving the way for future optimizations in emerging hardware, advanced quantization techniques, and diverse low-precision formats.

Acknowledgement

We would like to thank the members of the EcoSystem research laboratory at the University of Toronto for their feedback on the early manuscript. We also thank the anonymous ASPLOS reviewers and our shepherd, Jacques Pienaar, for their valuable feedback and suggestions, as well as the artifact evaluation reviewers for reproducing our experiments. The authors affiliated with the University of Toronto were supported by the Canada Foundation for Innovation JELF grant, the NSERC Discovery grant RGPIN-2024-06017, the Vector CIFAR AI Chair Award, the Google Scholar Research Award, the VMware Early Career Faculty Grant, and the Ontario Early Researcher Award ER22-17-191.

Appendix

A Artifact Appendix

A.1 Abstract

We provide artifacts to reproduce all experimental results discussed in the evaluation section. These artifacts include the compiler implementation, kernels in our domain-specific languages, and scripts for running experiments and plotting figures. To simplify reproduction, all necessary dependencies are packaged within a Docker image. A Dockerfile is also included to demonstrate the image building process. We do not provide the models directly; instead, they and their metadata will be automatically downloaded from Hugging Face Hub when the experiment scripts are launched.

A.2 Artifact check-list (meta-information)

- **Model:** Gemma, Llama, Qwen
- **Run-time environment:** Linux, CUDA
- **Hardware:** NVIDIA L40s
- **Metrics:** Latency (ms), Speedup
- **Output:** Efficient kernels, Numerical results
- **Experiments:** Automated scripts in docker
- **How much disk space required (approximately)?:** 25 GiB
- **How much time is needed to prepare workflow (approximately)?:** 10 minutes
- **How much time is needed to complete experiments (approximately)?:** 3 hours
- **Publicly available?:** Yes
- **Code licenses (if publicly available)?:** Apache 2.0
- **Workflow automation framework used?:** Docker
- **Archived (provide DOI)?:**
<https://doi.org/10.5281/zenodo.16756859>

A.3 Description

A.3.1 How to access. We have open-sourced our artifacts at <https://github.com/yaoyaoding/tilus-artifacts>. To pull the Docker image and perform experiments, follow the guide in the README.md file. The code itself is several megabytes, while the Docker image is approximately 21 GiB. We only fetch model meta-information (e.g., number of layers, layer size) from Hugging Face Hub, and dummy weights are used; therefore, the models do not consume significant disk space.

A.3.2 Hardware dependencies. Our experiments were primarily conducted on NVIDIA L40s. To perform a hardware ablation study, we also ran some experiments on NVIDIA A100 and NVIDIA H100. Any NVIDIA GPU with compute capability ≥ 8.0 should be able to run our artifacts and observe speedup, though the specific numbers might vary slightly.

A.3.3 Software dependencies. We provide a Docker image with all software dependencies pre-installed. Therefore, only the following software is required:

- NVIDIA GPU driver $\geq 565.57.01$
- NVIDIA container toolkit
- Docker

We have the following packages pre-installed:

- PyTorch v2.5.1
- Triton v3.1.0
- BitBLAS v0.0.1.dev15
- Marlin v0.1.1
- vLLM 0.7.3

A.3.4 Data sets. We use dummy inputs and weights because we are solely focused on system performance, which is independent of the input and weight content.

A.3.5 Models. The artifact uses three models for end-to-end evaluation: Gemma-2-9B, QWen-2.5-32B, and Llama-3.3-70B. The meta-information of these models will be automatically fetched from Hugging Face Hub (some may require a Hugging Face token).

A.4 Installation

First, clone the artifact Git repository:

```
git clone https://github.com/yaoyaoding/tilus-artifacts.git tilus
```

Then, install Docker and the NVIDIA Container Toolkit by following the README.md in the artifact.

A.5 Experiment workflow

All experiments can be executed with:

```
bash run.sh
```

This command will create a Docker container and sequentially run all experiments within it.

A.6 Evaluation and expected results

Upon completion of the experiments, a folder named results or precompiled-results will be created under the artifact directory. This folder will contain four figures of the evaluation results, corresponding to those presented in the evaluation section.

A.7 Methodology

Submission, reviewing and badging methodology:

- <https://www.acm.org/publications/policies/artifact-review-and-badging-current>
- <https://cTuning.org/ae>

References

- [1] Tor M Aamodt, Wilson Wai Lun Fung, and Timothy G Rogers. 2018. The SIMT Core: Instruction and Register Data Flow. In *General-Purpose Graphics Processor Architectures*. Springer, 21–66.
- [2] Andrew Adams, Karima Ma, Luke Anderson, Riyadh Baghdadi, Tzu-Mao Li, Michaël Gharbi, Benoit Steiner, Steven Johnson, Kayvon Fatahalian, Frédo Durand, and Jonathan Ragan-Kelley. 2019. Learning to Optimize Halide with Tree Search and Random Programs. *ACM Trans. Graph.* 38, 4, Article 121 (jul 2019), 12 pages. doi:10.1145/3306346.3322967
- [3] Aditya Agrawal, Matthew Hedlund, and Blake Hechtman. 2024. eXmY: A Data Type and Technique for Arbitrary Bit Precision Quantization. arXiv:2405.13938 [cs.LG] <https://arxiv.org/abs/2405.13938>
- [4] Byung Hoon Ahn, Prannoy Pilligundla, Amir Yazdanbakhsh, and Hadi Esmaeilzadeh. 2020. Chameleon: Adaptive Code Optimization for Expedited Deep Neural Network Compilation. In *International Conference on Learning Representations*. <https://openreview.net/forum?id=rygG4AVFvH>
- [5] Sami Alabed, Daniel Belov, Bart Chrzaszcz, Juliana Franco, Dominik Grewe, Dougal Maclaurin, James Molloy, Tom Natan, Tamara Norman, Xiaoyue Pan, Adam Paszke, Norman A. Rink, Michael Schaarschmidt, Timur Sitdikov, Agnieszka Swietlik, Dimitrios Vytiniotis, and Joel Wee. 2025. PartIR: Composing SPMD Partitioning Strategies for Machine Learning. In *Proceedings of the 30th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 1* (Rotterdam, Netherlands) (ASPLOS '25). Association

- for Computing Machinery, New York, NY, USA, 794–810. doi:10.1145/3669940.3707284
- [6] AMD Corporation. 2024. *CDNA 3 Architecture for Accelerated Computing*. Available at <https://www.amd.com/en/technologies/cdna.html>.
 - [7] AMD Corporation. 2024. *HIP: Heterogeneous-Compute Interface for Portability*. Available at <https://rocm.docs.amd.com/projects/HIP/en/latest/>.
 - [8] Saleh Ashkboos, Amirkeivan Mohtashami, Maximilian L. Croci, Bo Li, Pashmina Cameron, Martin Jaggi, Dan Alistarh, Torsten Hoefer, and James Hensman. 2024. QuaRot: Outlier-Free 4-Bit Inference in Rotated LLMs. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=dfqsW38v1X>
 - [9] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. Language Models are Few-Shot Learners. In *Advances in Neural Information Processing Systems*, H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (Eds.), Vol. 33. Curran Associates, Inc., 1877–1901. https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf
 - [10] Wen-Pu Cai, Ming-Yang Li, and Wu-Jun Li. 2024. Lcq: Low-rank codebook based quantization for large language models. *arXiv preprint arXiv:2405.20973* (2024).
 - [11] Jerry Chee, Yaohui Cai, Volodymyr Kuleshov, and Christopher De Sa. 2023. QuIP: 2-bit quantization of large language models with guarantees. In *Proceedings of the 37th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (*NIPS '23*). Curran Associates Inc., Red Hook, NY, USA, Article 196, 34 pages.
 - [12] Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Q. Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, Carlos Guestrin, and Arvind Krishnamurthy. 2018. TVM: An Automated End-to-End Optimizing Compiler for Deep Learning. In *OSDI*.
 - [13] Tim Dettmers, Mike Lewis, Younes Belkada, and Luke Zettlemoyer. 2024. LLM.int8(): 8-bit matrix multiplication for transformers at scale. In *Proceedings of the 36th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (*NIPS '22*). Curran Associates Inc., Red Hook, NY, USA, Article 2198, 15 pages.
 - [14] Yaoyao Ding, Cody Hao Yu, Bojian Zheng, Yizhi Liu, Yida Wang, and Gennady Pekhimenko. 2023. Hidet: Task-Mapping Programming Paradigm for Deep Learning Tensor Programs. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2* (Vancouver, BC, Canada) (*ASPLOS 2023*). Association for Computing Machinery, New York, NY, USA, 370–384. doi:10.1145/3575693.3575702
 - [15] Yaoyao Ding, Ligeng Zhu, Zhihao Jia, Gennady Pekhimenko, and Song Han. 2021. Ios: Inter-operator scheduler for cnn acceleration. *Proceedings of Machine Learning and Systems* 3 (2021), 167–180.
 - [16] Pratik Fegade, Tianqi Chen, Phillip Gibbons, and Todd Mowry. 2021. Cortex: A compiler for recursive deep learning models. *Proceedings of Machine Learning and Systems* 3 (2021), 38–54.
 - [17] Pratik Fegade, Tianqi Chen, Phillip Gibbons, and Todd Mowry. 2022. The CoRa Tensor Compiler: Compilation for Ragged Tensors with Minimal Padding. In *Proceedings of Machine Learning and Systems*, D. Marculescu, Y. Chi, and C. Wu (Eds.), Vol. 4. 721–747. <https://proceedings.mlsys.org/paper/2022/file/d3d9446802a44259755d38e6d163e820-Paper.pdf>
 - [18] Boyuan Feng, Yuke Wang, Tong Geng, Ang Li, and Yufei Ding. 2021. APNN-TC: accelerating arbitrary precision neural networks on ampere GPU tensor cores. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis* (St. Louis, Missouri) (*SC '21*). Association for Computing Machinery, New York, NY, USA, Article 37, 13 pages. doi:10.1145/3458817.3476157
 - [19] Siyuan Feng, Bohan Hou, Hongyi Jin, Wuwei Lin, Junru Shao, Ruihang Lai, Zihao Ye, Lianmin Zheng, Cody Hao Yu, Yong Yu, and Tianqi Chen. 2023. TensorIR: An Abstraction for Automatic Tensorized Program Optimization. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 2* (Vancouver, BC, Canada) (*ASPLOS 2023*). Association for Computing Machinery, New York, NY, USA, 804–817. doi:10.1145/3575693.3576933
 - [20] Elias Frantar, Saleh Ashkboos, Torsten Hoefer, and Dan Alistarh. 2022. Gptq: Accurate post-training quantization for generative pre-trained transformers. *arXiv preprint arXiv:2210.17323* (2022).
 - [21] Elias Frantar, Roberto L. Castro, Jiale Chen, Torsten Hoefer, and Dan Alistarh. 2024. MARLIN: Mixed-Precision Auto-Regressive Parallel Inference on Large Language Models. *arXiv:2408.11743* [cs.LG] <https://arxiv.org/abs/2408.11743>
 - [22] Aaron Grattafiori, Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Alex Vaughan, et al. 2024. The llama 3 herd of models. *arXiv preprint arXiv:2407.21783* (2024).
 - [23] Bastian Hagedorn, Bin Fan, Hanfeng Chen, Cris Cecka, Michael Garland, and Vinod Grover. 2023. Graphene: An IR for Optimized Tensor Computations on GPUs. In *Proceedings of the 28th ACM International Conference on Architectural Support for Programming Languages and Operating Systems, Volume 3* (Vancouver, BC, Canada) (*ASPLOS 2023*). Association for Computing Machinery, New York, NY, USA, 302–313. doi:10.1145/3582016.3582018
 - [24] Kartik Hegde, Hadi Asghari-Moghaddam, Michael Pellauer, Neal Crago, Amer Jaleel, Edgar Solomonik, Joel Emer, and Christopher W. Fletcher. 2019. ExTensor: An Accelerator for Sparse Tensor Algebra. In *Proceedings of the 52nd Annual IEEE/ACM International Symposium on Microarchitecture* (Columbus, OH, USA) (*MICRO '52*). Association for Computing Machinery, New York, NY, USA, 319–333. doi:10.1145/3352460.3358275
 - [25] Adnan Hoque, Less Wright, Chih-Chieh Yang, Mudhakar Srivatsa, and Raghu Ganti. 2024. Accelerating a Triton Fused Kernel for W4A16 Quantized Inference with SplitK work decomposition. *arXiv preprint arXiv:2402.00025* (2024).
 - [26] Guyue Huang, Yang Bai, Liu Liu, Yuke Wang, Bei Yu, Yufei Ding, and Yuan Xie. 2023. Alcop: Automatic load-compute pipelining in deep learning compiler for ai-gpus. *Proceedings of Machine Learning and Systems* 5 (2023), 680–694.
 - [27] IEEE Standards Association. 2019. *IEEE Standard for Floating-Point Arithmetic (IEEE 754-2019)*. IEEE, New York, NY. doi:10.1109/IEEESTD.2019.8766229 Available at <https://standards.ieee.org/standard/754-2019.html>.
 - [28] Zhihao Jia, Oded Padon, James Thomas, Todd Warszawski, Matei Zaharia, and Alex Aiken. 2019. TASO: optimizing deep learning computation with automatic generation of graph substitutions. In *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. 47–62.
 - [29] Woosuk Kwon, Zhuohan Li, Siyuan Zhuang, Ying Sheng, Lianmin Zheng, Cody Hao Yu, Joseph E. Gonzalez, Hao Zhang, and Ion Stoica. 2023. Efficient Memory Management for Large Language Model Serving with PagedAttention. In *Proceedings of the ACM SIGOPS 29th Symposium on Operating Systems Principles*.
 - [30] Ji Lin, Jiaming Tang, Haotian Tang, Shang Yang, Wei-Ming Chen, Wei-Chen Wang, Guangxuan Xiao, Xingyu Dang, Chuang Gan, and Song Han. 2024. AWQ: Activation-aware Weight Quantization for LLM Compression and Acceleration. In *MLSys*.
 - [31] Yang Liu and Mirella Lapata. 2019. Text Summarization with Pre-trained Encoders. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*,

- Kentaro Inui, Jing Jiang, Vincent Ng, and Xiaojun Wan (Eds.). Association for Computational Linguistics, Hong Kong, China, 3730–3740. doi:10.18653/v1/D19-1387
- [32] Zechun Liu, Changsheng Zhao, Igor Fedorov, Bilge Soran, Dhruv Choudhary, Raghuraman Krishnamoorthi, Vikas Chandra, Yuandong Tian, and Tijmen Blankevoort. 2025. SpinQuant: LLM Quantization with Learned Rotations. In *The Thirteenth International Conference on Learning Representations*. <https://openreview.net/forum?id=ogO6DGE6FZ>
- [33] Weile Luo, Ruibo Fan, Zeyu Li, Dayou Du, Hongyuan Liu, Qiang Wang, and Xiaowen Chu. 2025. Dissecting the NVIDIA Hopper Architecture through Microbenchmarking and Multiple Level Analysis. *arXiv preprint arXiv:2501.12084* (2025).
- [34] Lingxiao Ma, Zhiqiang Xie, Zhi Yang, Jilong Xue, Youshan Miao, Wei Cui, Wenxiang Hu, Fan Yang, Lintao Zhang, and Lidong Zhou. 2020. RAMMER: Enabling Holistic Deep Learning Compiler Optimizations with Rtasks. USENIX Association, USA, 17.
- [35] Shaobo Ma, Chao Fang, Haikuo Shao, and Zhongfeng Wang. 2025. Efficient Arbitrary Precision Acceleration for Large Language Models on GPU Tensor Cores. In *Proceedings of the 30th Asia and South Pacific Design Automation Conference* (Tokyo, Japan) (ASPDAC '25). Association for Computing Machinery, New York, NY, USA, 1181–1187. doi:10.1145/3658617.3697668
- [36] Nandeeeka Nayak, Xinrui Wu, Toluwanimi O. Odemuyiwa, Michael Pel-lauer, Joel S. Emer, and Christopher W. Fletcher. 2024. FuseMax: Leveraging Extended Einsums to Optimize Attention Accelerator Design. In *2024 57th IEEE/ACM International Symposium on Microarchitecture (MICRO)*. 1458–1473. doi:10.1109/MICRO61859.2024.00107
- [37] John Nickolls and William J. Dally. 2010. The GPU Computing Era. *IEEE Micro* 30, 2 (2010), 56–69. doi:10.1109/MM.2010.41
- [38] NVIDIA Corporation. 2021. CUTLASS: CUDA Templates for Linear Algebra Subroutines and Solvers. <https://github.com/NVIDIA/cutlass>.
- [39] NVIDIA Corporation. 2023. *NVIDIA cuBLAS Library*. <https://developer.nvidia.com/cublas> Version 12.2..
- [40] NVIDIA Corporation. 2024. *CUDA C++ Programming Guide*. Version 12.0. Available at <https://docs.nvidia.com/cuda/cuda-c-programming-guide/>.
- [41] NVIDIA Corporation. 2024. *Parallel Thread Execution ISA Version 12.0*. Available at <https://docs.nvidia.com/cuda/parallel-thread-execution/index.html>.
- [42] NVIDIA Corporation. 2024. *SASS: Streaming Assembler for NVIDIA GPUs*. Available at <https://docs.nvidia.com/cuda/cuda-binary-utilities/index.html>.
- [43] OpenAI. 2024. ChatGPT. <https://chat.openai.com/>. Accessed: 2024-11-12; Generative AI language model.
- [44] Muhammad Osama, Duane Merrill, Cris Cecka, Michael Garland, and John D. Owens. 2023. Stream-K: Work-Centric Parallel Decomposition for Dense Matrix-Matrix Multiplication on the GPU. In *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming* (Montreal, QC, Canada) (PPoPP '23). Association for Computing Machinery, New York, NY, USA, 429–431. doi:10.1145/3572848.3577479
- [45] Yeonhong Park, Jake Hyun, SangLyu Cho, Bonggeun Sim, and Jae W. Lee. 2024. Any-Precision LLM: Low-Cost Deployment of Multiple, Different-Sized LLMs. In *Proceedings of the 41st International Conference on Machine Learning*.
- [46] Jonathan Ragan-Kelley, Connelly Barnes, Andrew Adams, Sylvain Paris, Frédo Durand, and Saman Amarasinghe. 2013. Halide: a language and compiler for optimizing parallelism, locality, and recomputation in image processing pipelines. In *Acm Sigplan Notices*, Vol. 48. ACM, 519–530.
- [47] Bitu Darvish Rouhani, Ritchie Zhao, Ankit More, Mathew Hall, Alireza Khodamoradi, Summer Deng, Dhruv Choudhary, Marius Cornea, Eric Dellinger, Kristof Denolf, et al. 2023. Microscaling data formats for deep learning. *arXiv preprint arXiv:2310.10537* (2023).
- [48] Amit Sabne. 2020. XLA : Compiling Machine Learning for Peak Performance.
- [49] Junru Shao, Xiyu Zhou, Siyuan Feng, Bohan Hou, Ruihang Lai, Hongyi Jin, Wuwei Lin, Masahiro Masuda, Cody Hao Yu, and Tianqi Chen. 2024. Tensor program optimization with probabilistic programs. In *Proceedings of the 36th International Conference on Neural Information Processing Systems* (New Orleans, LA, USA) (NIPS '22). Curran Associates Inc., Red Hook, NY, USA, Article 2593, 14 pages.
- [50] Haichen Shen, Jared Roesch, Zhi Chen, Wei Chen, Yong Wu, Mu Li, Vin Sharma, Zachary Tatlock, and Yida Wang. 2021. Nimble: Efficiently Compiling Dynamic Neural Networks for Model Inference. In *Proceedings of Machine Learning and Systems*, A. Smola, A. Dimakis, and I. Stoica (Eds.), Vol. 3. 208–222. <https://proceedings.mlsys.org/paper/2021/file/4e732ced3463d06de0ca9a15b6153677-Paper.pdf>
- [51] Shizhi Tang, Jidong Zhai, Haojie Wang, Lin Jiang, Liyan Zheng, Zhenhao Yuan, and Chen Zhang. 2022. FreeTensor: A Free-Form DSL with Holistic Optimizations for Irregular Tensor Programs. In *Proceedings of the 43rd ACM SIGPLAN International Conference on Programming Language Design and Implementation (PLDI '22)* (San Diego, CA, USA) (PLDI '22). New York, NY, USA, 16 pages. doi:10.1145/3519939.3523448
- [52] Gemma Team, Morgane Riviere, Shreya Pathak, Pier Giuseppe Sessa, Cassidy Hardin, Surya Bhupatiraju, Léonard Hussenot, Thomas Mesnard, Bobak Shahriari, Alexandre Ramé, et al. 2024. Gemma 2: Improving open language models at a practical size. *arXiv preprint arXiv:2408.00118* (2024).
- [53] Philippe Tillet, H. T. Kung, and David Cox. 2019. Triton: an intermediate language and compiler for tiled neural network computations. In *Proceedings of the 3rd ACM SIGPLAN International Workshop on Machine Learning and Programming Languages* (Phoenix, AZ, USA) (MAPL 2019). Association for Computing Machinery, New York, NY, USA, 10–19. doi:10.1145/3315508.3329973
- [54] Albert Tseng, Jerry Chee, Qingyao Sun, Volodymyr Kuleshov, and Christopher De Sa. 2024. QuIP#: even better LLM quantization with hadamard incoherence and lattice codebooks. In *Proceedings of the 41st International Conference on Machine Learning* (Vienna, Austria) (ICML '24). JMLR.org, Article 1987, 27 pages.
- [55] Nicolas Vasilache, Oleksandr Zinenko, Theodoros Theodoridis, Priya Goyal, Zach DeVito, William S. Moses, Sven Verdoolaege, Andrew Adams, and Albert Cohen. 2018. Tensor Comprehensions: Framework-Agnostic High-Performance Machine Learning Abstractions. *ArXiv abs/1802.04730* (2018).
- [56] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- [57] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All you Need. In *Advances in Neural Information Processing Systems*, I. Guyon, U. Von Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (Eds.), Vol. 30. Curran Associates, Inc. https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf
- [58] Lei Wang, Lingxiao Ma, Shijie Cao, Quanlu Zhang, Jilong Xue, Yining Shi, Ningxin Zheng, Ziming Miao, Fan Yang, Ting Cao, Yuqing Yang, and Mao Yang. 2024. Ladder: Enabling Efficient Low-Precision Deep Learning Computing through Hardware-aware Tensor Transformation. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*. USENIX Association, Santa Clara, CA, 307–323. <https://www.usenix.org/conference/osdi24/presentation/wang-lei>

- [59] BigScience Workshop, Teven Le Scao, Angela Fan, Christopher Akiki, Ellie Pavlick, Suzana Ilić, Daniel Hesslow, Roman Castagné, Alexandra Sasha Luccioni, François Yvon, et al. 2022. Bloom: A 176b-parameter open-access multilingual language model. *arXiv preprint arXiv:2211.05100* (2022).
- [60] Haojun Xia, Zhen Zheng, Xiaoxia Wu, Shiyang Chen, Zhewei Yao, Stephen Youn, Arash Bakhtiari, Michael Wyatt, Donglin Zhuang, Zhongzhu Zhou, Olatunji Ruwase, Yuxiong He, and Shuaiwen Leon Song. 2024. Quant-LLM: Accelerating the Serving of Large Language Models via FP6-Centric Algorithm-System Co-Design on Modern GPUs. In *2024 USENIX Annual Technical Conference (USENIX ATC 24)*. USENIX Association, Santa Clara, CA, 699–713. <https://www.usenix.org/conference/atc24/presentation/xia>
- [61] Jiarong Xing, Leyuan Wang, Shang Zhang, Jack Chen, Ang Chen, and Yibo Zhu. 2022. Bolt: Bridging the Gap between Auto-tuners and Hardware-native Performance. In *Proceedings of Machine Learning and Systems*, Vol. 4.
- [62] An Yang, Baosong Yang, Binyuan Hui, Bo Zheng, Bowen Yu, Chang Zhou, Chengpeng Li, Chengyuan Li, Dayiheng Liu, Fei Huang, et al. 2024. Qwen2 technical report. *arXiv preprint arXiv:2407.10671* (2024).
- [63] Gyeong-In Yu, Joo Seong Jeong, Geon-Woo Kim, Soojeong Kim, and Byung-Gon Chun. 2022. Orca: A Distributed Serving System for Transformer-Based Generative Models. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. USENIX Association, Carlsbad, CA, 521–538. <https://www.usenix.org/conference/osdi22/presentation/yu>
- [64] Chao Zeng, Songwei Liu, Yusheng Xie, Hong Liu, Xiaojian Wang, Miao Wei, Shu Yang, Fangmin Chen, and Xing Mei. 2025. ABQ-LLM: Arbitrary-Bit Quantized Inference Acceleration for Large Language Models. *Proceedings of the AAAI Conference on Artificial Intelligence* 39, 21 (Apr. 2025), 22299–22307. doi:10.1609/aaai.v39i21.34385
- [65] Jie Zhao, Xiong Gao, Ruijie Xia, Zhaochuang Zhang, Deshi Chen, Lei Chen, Renwei Zhang, Zhen Geng, Bin Cheng, and Xuefeng Jin. 2022. Apollo: Automatic Partition-based Operator Fusion through Layer by Layer Optimization. In *Proceedings of Machine Learning and Systems*, D. Marculescu, Y. Chi, and C. Wu (Eds.), Vol. 4. 1–19. <https://proceedings.mlsys.org/paper/2022/file/069059b7ef840f0c74a814ec9237b6ec-Paper.pdf>
- [66] Bojian Zheng, Ziheng Jiang, Cody Hao Yu, Haichen Shen, Joshua Fromm, Yizhi Liu, Yida Wang, Luis Ceze, Tianqi Chen, and Gennady Pekhimenko. 2022. DietCode: Automatic Optimization for Dynamic Tensor Programs. In *Proceedings of Machine Learning and Systems*, D. Marculescu, Y. Chi, and C. Wu (Eds.), Vol. 4. 848–863. <https://proceedings.mlsys.org/paper/2022/file/fa7cdfad1a5aaf8370ebeda47a1ff1c3-Paper.pdf>
- [67] Lianmin Zheng, Chengfan Jia, Minmin Sun, Zhao Wu, Cody Hao Yu, Ameer Haj-Ali, Yida Wang, Jun Yang, Danyang Zhuo, Koushik Sen, Joseph E. Gonzalez, and Ion Stoica. 2020. Ansor: Generating High-Performance Tensor Programs for Deep Learning. In *14th USENIX Symposium on Operating Systems Design and Implementation (OSDI 20)*. 863–879.
- [68] Size Zheng, Renze Chen, Anjiang Wei, Yicheng Jin, Qin Han, Liqiang Lu, Bingyang Wu, Xiuhong Li, Shengen Yan, and Yun Liang. 2022. AMOS: enabling automatic mapping for tensor computations on partial accelerators with hardware abstraction. In *Proceedings of the 49th Annual International Symposium on Computer Architecture (New York, New York) (ISCA '22)*. Association for Computing Machinery, New York, NY, USA, 874–887. doi:10.1145/3470496.3527440
- [69] Size Zheng, Yun Liang, Shuo Wang, Renze Chen, and Kaiwen Sheng. 2020. FlexTensor: An Automatic Schedule Exploration and Optimization Framework for Tensor Computation on Heterogeneous System. *Proceedings of the Twenty-Fifth International Conference on Architectural Support for Programming Languages and Operating Systems* (2020).
- [70] Zhen Zheng, Xuanda Yang, Pengzhan Zhao, Guoping Long, Kai Zhu, Feiwen Zhu, Wenyi Zhao, Xiaoyong Liu, Jun Yang, Jidong Zhai, Shuaiwen Leon Song, and Wei Lin. 2022. AStitch: Enabling a New Multi-Dimensional Optimization Space for Memory-Intensive ML Training and Inference on Modern SIMT Architectures. In *Proceedings of the 27th ACM International Conference on Architectural Support for Programming Languages and Operating Systems (Lausanne, Switzerland) (ASPLOS '22)*. Association for Computing Machinery, New York, NY, USA, 359–373. doi:10.1145/3503222.3507723
- [71] Hongyu Zhu, Ruofan Wu, Yijia Diao, Shanbin Ke, Haoyu Li, Chen Zhang, Jilong Xue, Lingxiao Ma, Yuqing Xia, Wei Cui, Fan Yang, Mao Yang, Lidong Zhou, Asaf Cidon, and Gennady Pekhimenko. 2022. ROLLER: Fast and Efficient Tensor Compilation for Deep Learning. In *16th USENIX Symposium on Operating Systems Design and Implementation (OSDI 22)*. USENIX Association, Carlsbad, CA, 233–248. <https://www.usenix.org/conference/osdi22/presentation/zhu>
- [72] K. Zhu, W.Y. Zhao, Z. Zheng, T.Y. Guo, P.Z. Zhao, J.J. Bai, J. Yang, X.Y. Liu, L.S. Diao, and W. Lin. 2021. DISC: A Dynamic Shape Compiler for Machine Learning Workloads. In *Proceedings of the 1st Workshop on Machine Learning and Systems (Online, United Kingdom) (EuroMLSys '21)*. Association for Computing Machinery, New York, NY, USA, 89–95. doi:10.1145/3437984.3458838