

# Fast and Efficient 2-bit LLM Inference on GPU: 2/4/16-bit in a Weight Matrix with Asynchronous Dequantization

Jinhao Li<sup>1\*</sup>, Jiaming Xu<sup>13\*</sup>, Shiyao Li<sup>23</sup>, Shan Huang<sup>1</sup>, Jun Liu<sup>1</sup>, Yaoxiu Lian<sup>1</sup>, Guohao Dai<sup>13†</sup>

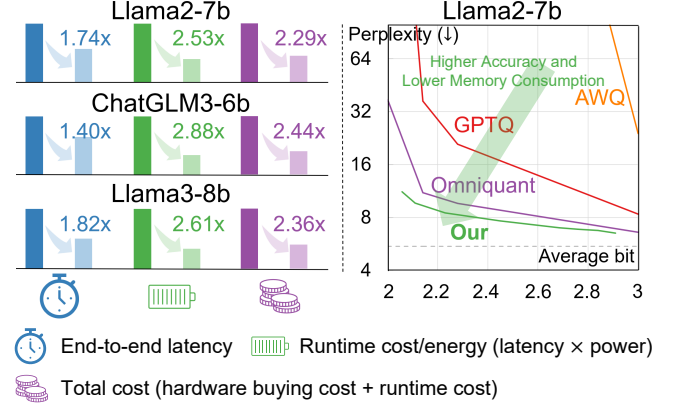
<sup>1</sup>Shanghai Jiao Tong University, <sup>2</sup>Tsinghua University, <sup>3</sup>Infinigence-AI, \*Equal Contributions

<sup>†</sup>Corresponding author: daiguohao@sjtu.edu.cn

## ABSTRACT

Large language models (LLMs) have demonstrated impressive abilities in various domains while the inference cost is expensive. Many previous studies exploit quantization methods to reduce LLM inference cost by reducing latency and memory consumption. Applying 2-bit single-precision weight quantization brings >3% accuracy loss, so the state-of-the-art methods use mixed-precision methods for LLMs (e.g. Llama2-7b, etc.) to improve the accuracy. However, challenges still exist: (1) **Uneven distribution in weight matrix**. Weights are quantized by groups, while some groups contain weights with large range. Previous methods apply inter-weight mixed-precision quantization and neglect the range difference inside each weight matrix, resulting in >2.7% accuracy loss (e.g. LLM-MQ and APTQ). (2) **Large speed degradation by adding sparse outliers**. Reserving sparse outliers improves accuracy but slows down the speed affected by the outlier ratio (e.g. 1.5% outliers resulting in >30% speed degradation in SpQR). (3) **Time-consuming dequantization operations on GPUs**. Mainstream methods require a dequantization operation to perform computation on the quantized weights, and the 2-order dequantization operation is applied because scales of groups are also quantized. These dequantization operations lead to >50% execution time.

To tackle these challenges and enable fast and efficient LLM inference on GPUs, we propose the following techniques in this paper. (1) **Intra-weight mixed-precision quantization**. We only quantize a small fraction of groups with higher sensitivity (larger Hessian value and range variation) using 4-bit. Meanwhile, we also take the memory alignment into consideration on GPUs. (2) **Exclusive 2-bit sparse outlier with minimum speed degradation**. We only reserve a small fraction of large weights in 2-bit groups as sparse outliers using 16-bit, which leads to a lower average bit increment and speed degradation. (3) **Asynchronous dequantization**. We point out that calculating the scales of each group in 2-order dequantization is independent of the loading weights of each group in 1-order dequantization. Thus, we design the asynchronous dequantization on GPUs. We conduct extensive experiments on different model families (e.g. Llama3, etc.) and model sizes. We achieve 2.91-bit for each weight considering all scales/zeros for different models with negligible loss. As a result, with our 2/4/16 mixed-precision



**Figure 1: Summary of this paper compared with state-of-the-art quantization designs on mainstream large language models. Detailed data can be found in Table II and Section 6.1.**

quantization for each weight matrix and asynchronous dequantization during inference, our design achieves an end-to-end speedup for Llama2-7b is 1.74× over the original model, and we reduce both runtime cost and total cost by up to 2.53× and 2.29× with less GPU requirements.

## ACM Reference Format:

Jinhao Li<sup>1\*</sup>, Jiaming Xu<sup>13\*</sup>, Shiyao Li<sup>23</sup>, Shan Huang<sup>1</sup>, Jun Liu<sup>1</sup>, Yaoxiu Lian<sup>1</sup>, Guohao Dai<sup>13†</sup>. 2024. Fast and Efficient 2-bit LLM Inference on GPU: 2/4/16-bit in a Weight Matrix with Asynchronous Dequantization. In *IEEE/ACM International Conference on Computer-Aided Design (ICCAD '24)*, October 27–31, 2024, New York, NY, USA. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3676536.3676796>

## 1 INTRODUCTION

Large language models (LLMs) demonstrate remarkable capabilities in various domains, excelling in tasks like natural language understanding and generation [7, 20, 22, 26, 29]. However, their computational expense during inference is noteworthy. For instance, OpenAI reveals that the GPT-4 Turbo incurs a cost of millions of dollars each day [32]. This substantial cost underscores the economic considerations associated with deploying and utilizing such advanced language models [33]. Many previous studies exploit quantization methods to reduce LLM inference cost by reducing storage and accelerating computation [5, 10, 17, 19, 42]. Quantization with the lower bit-width can further reduce the storage and memory access overheads, leading to a more economic inference. However, for 2-bit weight quantization, these methods including Greenbit [13] still fail to prevent the accuracy loss (>3%). Compared with single-precision methods, the state-of-the-art methods use mixed-precision methods [12, 18] for LLMs to achieve better accuracy exemplified by the Llama-2 family [38].

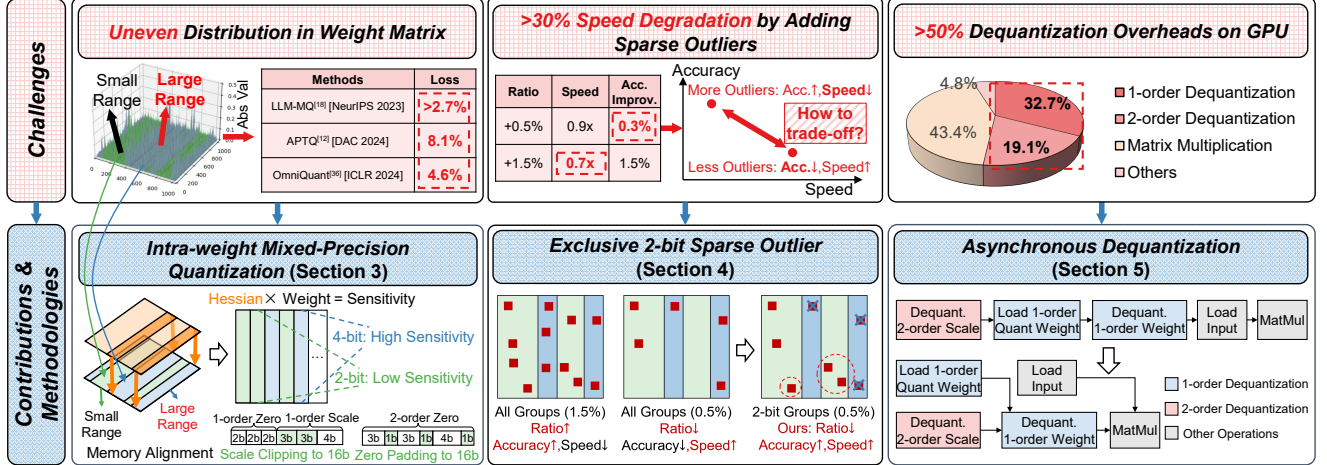
Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components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

ICCAD '24, October 27–31, 2024, New York, NY, USA

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-1077-3/24/10

<https://doi.org/10.1145/3676536.3676796>



**Figure 2: Challenges in fast and efficient quantized LLM inference on GPU: (1) uneven distribution in weight matrix, (2) large speed degradation by adding sparse outliers, and (3) time-consuming dequantization overheads. We propose three novel contributions: (1) intra-matrix mixed-precision quantization, (2) exclusive 2-bit sparse outlier with minimum speed degradation, and (3) asynchronous dequantization, to solve these challenges.**

However, there still exist challenges of reducing LLM inference cost with the mixed-precision weight quantization method. **(1) Uneven distribution in weight matrix.** Weights are quantized by groups, while some groups contain weights with large range. Previous works like GPTQ [10], AWQ [19], LLM-QAT [21] and Greenbit [13] quantize weight matrix with single-precision and incur the 3.2% to 5.6% accuracy loss for Llama2-7b with 2-bit quantization. And the recent works apply mixed-precision methods to improve accuracy such as APTQ [12] and LLM-MQ [18]. They apply inter-matrix mixed-precision quantization and neglect the difference among the groups in each weight, still resulting in >2.7% accuracy loss, as shown in Figure 2 left top. **(2) Large speed degradation by adding sparse outliers.** Reserving sparse outliers improves accuracy but introduces speed degradation affected by the outlier ratio. Figure 2 middle top illustrates that the relationship between the speed and accuracy with different sparse outlier ratios (0.5% and 1.5%). SpQR reserves 1.5% sparse outliers to improve the accuracy, resulting in >30% speed degradation and hindering the inference speedup [5] while reserving less brings limited accuracy improvement. Consequently, it is necessary to improve accuracy and speed simultaneously by adding a small fraction of sparse outliers in specific regions. **(3) Time-consuming dequantization operations on GPUs.** Mainstream methods require a dequantization operation to perform computation on 2-bit weights. Furthermore, a 2-order dequantization operation is applied because the scales of groups are also quantized. Figure 2 right top shows that, the dequantization operation takes over 50% of the total execution time. These challenges become the crucial factors impeding fast and efficient LLM inference.

In response to these challenges, we enable fast and efficient LLM inference on GPUs with the following contributions in this paper:

**(1) Intra-matrix mixed-precision quantization.** We point out that the range of weights by groups varies and these groups always exhibit high sensitivity (large Hessian value and range variation). Because 4-bit quantization can prevent accuracy loss [10, 19], we only quantize 25% of sensitive groups with large Hessian value and range variation using 4-bit, and we also apply scale clipping and

zero padding techniques to achieve the memory alignment. Such a method reduces the accuracy loss for 2-bit quantization from 4.6% to 2.2% for Llama2-7b and from 8.6% to 2.5% for Llama2-7b.

**(2) Exclusive 2-bit sparse outlier with minimum speed degradation.** We point out that the speed degradation by reserving sparse outliers consists of launching GPU kernel and calculation. The time of launching kernel is fixed while the calculation time is proportional with sparse outlier ratio. Thus, we only reserve a small fraction (<0.5% with only launching kernel time) of large weights in 2-bit groups as sparse outliers, which leads to a lower average bit increment and higher speed. Such design improves the accuracy by >0.5% with <0.15 increased average weight bit and only 20% speed degradation.

**(3) Asynchronous dequantization.** We point out that calculating the scales of each group in 2-order dequantization is independent of the loading group weights in 1-order dequantization. Thus, we design the asynchronous dequantization and accelerate the GPU kernel by up to 3.92x.

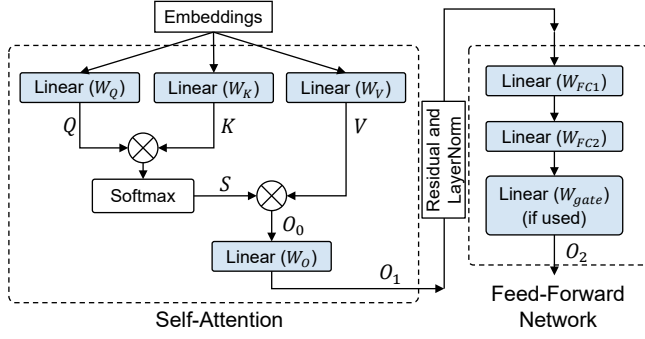
We conduct extensive experiments on different model families (e.g., Llama1 [25], Llama2 [38], Llama3 [24], ChatGLM3 [45] and BERT [6]) and model sizes (i.e., Llama2-7b to Llama2-70b). We achieve 2.91-bit for each weight considering all scales/zeros with 25% 4-bit for different models with negligible accuracy loss (e.g. 2.2%/1.1% for Llama2-7b/13b, respectively). The end-to-end speedup for Llama2-7b is 1.74x, and we reduce both runtime cost and total cost by up to 2.53x and 2.29x with less GPU requirements.

## 2 BACKGROUND

### 2.1 Transformer

The transformer model [39] is a typical backbone architecture primarily used for LLM tasks such as language translation, text summarization, and question answering [2, 16, 28]. The transformer model consists of several transformer blocks, and each block contains two components, self-attention and feed-forward network.

**Self-attention:** The self-attention mechanism in the Transformer model involves four computations. First, the input embedding  $X$  is multiplied by four weights matrix  $W_Q$ ,  $W_K$ ,  $W_V$  to get



**Figure 3: Overview of each transformer block. The self-attention contains four linear projections and feed-forward network contains at least two linear projections.**

Query, Key, and Value vectors ( $Q$ ,  $K$ , and  $V$ ) separately. Second, the attention scores are computed by taking the dot product of  $Q$  with  $K$  and scaled down by  $\sqrt{d_k}$ , the square root of the dimension of  $K$ . Then, the scaled attention scores are normalized by a softmax function to obtain a relative attention scores, which allows the model to focus on relevant parts of the input sequence. And the softmax-normalized attention scores are multiplied by  $V$  to obtain the weighted output  $O_0$ . In addition,  $O_0$  is multiplied by  $W_O$  to obtain the final output  $O_1$  of self-attention.

$$O_0 = \text{Softmax}\left(\frac{QK^T}{\sqrt{d_k}}\right)V \quad (1)$$

**Feed-forward Network:** After the self-attention, the output  $O_1$  is passed through a feed-forward network within each layer. This network applies two or three linear transformations to  $O_1$  and obtain the final output  $O_2$ .

## 2.2 Quantization

A group of floating weights with range of  $(w_{min}, w_{max})$  are quantized to integers ( $w_{int}$ ) with the range of  $(0, 2^N - 1)$ , a zero-point ( $z$ ) and a scaling factor ( $s$ ) [27]. Scaling factor is half and zero-point is an  $N$  bits integer. The quantization and dequantization operations are:

$$s = \frac{w_{max} - w_{min}}{2^N - 1} \quad (2)$$

$$z = \text{round}\left(\frac{-w_{min}}{s}\right) \quad (3)$$

$$w_{int} = \text{clamp}(0, 2^N - 1, \text{round}\left(\frac{w}{s}\right) + z) \quad (4)$$

$$w_{float} = (w_{int} - z) \times s \quad (5)$$

Weights along the direction of input dimensions are quantized by group, while the scaling factors of groups along the direction of the output dimensions are also quantized to further reduce the average bit. Here, the average bit, denoted as  $\overline{bit}$  with 1-order and 2-order quantization is as follows:

$$\overline{bit} = \begin{cases} N + \frac{N+16}{g_1}, & \text{1-order} \\ N + \frac{N+N_s}{g_1} + \frac{N_s+16}{g_1 \times g_2}, & \text{2-order} \end{cases} \quad (6)$$

where  $g_i$  represents the size of group and  $N_s$  represents the bit width of the quantized scaling factors with 2-order.

## 2.3 Sparse Outlier

In LLMs, weights exhibit large magnitudes compared to the others. By applying the normal uniform quantization, the quantization range is expanded by these outliers and thus, the majority of weights are quantized with only a few representations, leading to large quantization error. Mixed-precision quantization methods like SpQR, SqueezeLLM and LLM-MQ filter and retain outliers that negatively impact quantization error and store outliers with 16-bit sparse matrix [5, 17, 18]. These methods quantize the weights matrix into a dense matrix and a 16-bit sparse matrix. And the Compressed Sparse Row (CSR) is used to represent sparse matrices, optimizing memory usage.


In CSR, three arrays encapsulate the matrix data. The values array (values) efficiently arranges non-zero elements in row-major order, while the column indices (col\_ind) array meticulously records the corresponding column index for each non-zero element. The row pointer (row\_ptr) array serves as a guide, storing the starting index for each row in the values array. Outlier retention of weights is common for low-bit quantization in LLMs. Each outlier is stored using one 32-bit combination: a 16-bit weight value and a corresponding 16-bit column index. Additionally, for every row, a 32-bit number is allocated to store the total count of outliers up to that row. This results in an overall 32-bit storage overheads for each weight outlier.

## 3 INTRA-WEIGHT MIXED-PRECISION QUANTIZATION

The quantization method is widely used to reduce LLM inference cost [5, 10, 19] by reducing storage and accelerating computation. It uses discrete values to represent continuous weights, leading to quantization errors. For 2-bit quantization, weights are quantized by groups, while some groups contain weights with large range.

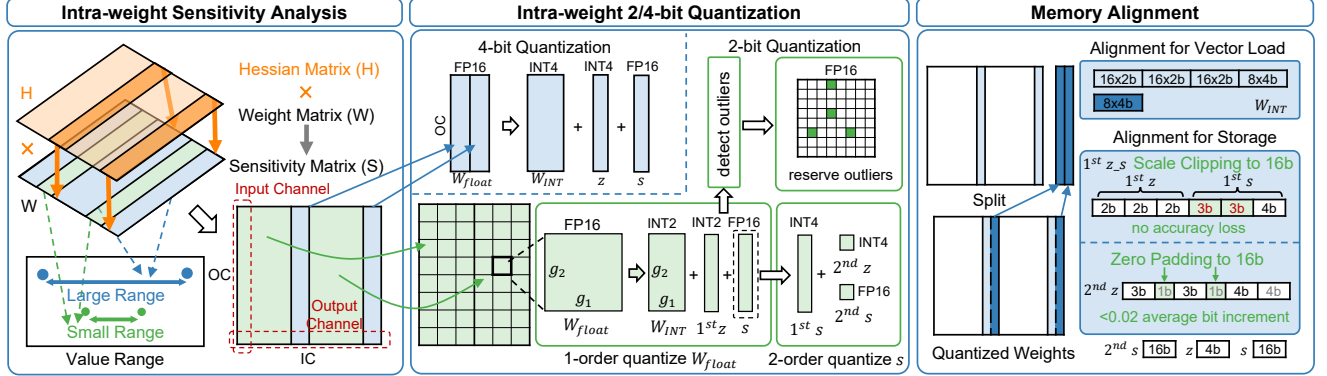
### 3.1 Challenge

Previous works [10, 13, 19, 21] only quantize weight matrix with 2-bit single-precision and incur the 3.2% to 5.6% accuracy loss for Llama2-7b. The state-of-the-art methods use inter-weight mixed-precision methods [12, 18] for LLMs to improve the accuracy with <3 average bit considering all scaling factors and zero-points. However, these methods neglect the range variance of quantization groups (uneven distribution) inside each weight matrix, which leads to >2.7% accuracy loss for Llama2-7b. Thus, the key challenge is that even the state-of-the-art mixed-precision methods suffer from handling the uneven distribution inside weight matrices.



**Figure 4: The range of weight matrix by groups: the blue groups have larger range while the green groups are smaller.**





**Figure 5: Overview of range-aware quantization method.** We first analyze the sensitivity  $S$  of different groups by the multiplication of the original weight matrix  $W$  with different range distribution and Hessian matrix  $H$ , and then apply intra-weight 2/4-bit mixed-precision quantization. Last, we propose three techniques for memory alignment.

### 3.2 Motivation and Insight

Some previous methods like SpQR [5], SparseGPT [9] and AWQ [19] have analyzed that the quantization errors are large in specific groups. So we depict the range of weights by groups, as shown in Figure 4. The range of weights exhibits different distributions for different groups, and only a small fraction of groups show a large range while the others show a small range. Applying 4-bit quantization to these groups with a large range can reduce the quantization error. Thus, our key insight is, **the range by groups varies inside weight matrix and some groups require 4-bit quantization to reduce large quantization error.**

### 3.3 Approach

**Intra-weight Quantization.** In Figure 5, we depict the quantization flow of weights. Notably, weights are quantized by group to reduce accuracy loss while the groups with a large range still require larger bit-width to improve accuracy. To determine which values are more sensitive to affect the loss caused by quantization, we perform Taylor series [8] expansion to analyze how the model output changes in response to perturbations in the parameters  $W$ :

$$\begin{aligned} L(W_Q) &\approx L(W) - g^T (W - W_Q) + \frac{1}{2} (W - W_Q)^T H (W - W_Q) \\ &\approx L(W) + \frac{1}{2} (W - W_Q)^T H (W - W_Q) \end{aligned} \quad (8)$$

where  $g$  is the gradient and  $H$  is the second derivative (i.e., Hessian matrix) of the loss at  $W$ . Assuming that the model has converged to a local minimum, the gradient  $g$  can be approximated as zero, which is mentioned in previous works [10, 14, 17]. The quantization error  $W - W_Q$  is weighted by the second-order derivative  $H$ . This highlights the importance of minimizing perturbations for weights that have large Hessian values, as they have a greater impact on the overall perturbation of the final output. In other words, the second-order derivative serves as a measure of importance for each weight value. **To reduce the accuracy loss, we should consider both the range variation and the second derivative variation by groups.**

Therefore, we first analyze the range variation by group. The range of groups in the same input channels is merged to streamline the complexity of hardware computation. Concretely, weights are divided by the Hessian inverse matrix to get the sensitivity

of weights, leveraging the Hessian inverse matrix's ability to discern weight importance. Subsequently, the calibrated weights are squared to magnify the sensitivity variance. The sensitivity of each input channel is computed as follows:

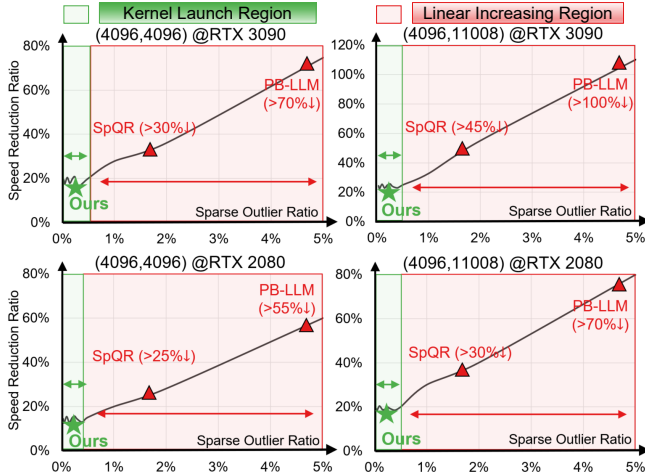
$$S_i = \sum_{m=1}^M \sum_{j=1}^{OC} \frac{w_{Mi+m,j}^2}{[H^{-1}]_{Mi+m}^2} \quad (9)$$

where  $H^{-1}$  represents the Hessian inverse matrix of each weight matrix,  $M$  represents the group size,  $OC$  represents the dimension of output channels, and  $w_{Mi+m,j}$  represents the original value of each element in the group  $m$ . The Hessian inverse matrix  $H^{-1}$  can be calculated offline [10] to reduce the quantization overheads. And  $S_i$  represents the sensitivity of group  $i$ , which serves as a criterion and an indicator to determine which group should be quantized with higher bitwidth.

Thus, the variance of weights by groups is equivalent to the sensitivity variance. Then, we analyze the variance and obtain the information for  $S_i$  that the groups with large sensitivity require 4-bit quantization and others require 2-bit quantization. According to the analysis results, we apply the finetuning approach [21, 36] like quantization aware training (QAT) to further reduce the quantization error. In the forward process of QAT, the weights with a large range are quantized and then dequantized with 4-bit by group while others are with 2-bit. After finetuning, we use the same analysis results to quantize different grouped input channels with different bit-width. For small range input channels, we first do 2-bit 1-order quantization for the weights  $W_{float}$  by group  $g_1$  in the direction of input channels to get the quantized weight  $W_{INT}$ , 1-order zeros  $1^{st}z$  and scales  $s$ . To further reduce the average bit, the scales  $s$  are also quantized with 4-bit by group  $g_2$  in the direction of output channels (2-order quantization) to get 1-order quantized scales  $1^{st}s$ , 2-order zeros  $2^{nd}z$  and scales  $2^{nd}s$ . During 2-bit quantization, outliers are also detected and reserved to reduce accuracy loss. Then, the weights in large range input channels are quantized with 4-bit. Therefore, the average bit-width of each weight element considering all scaling factors and zero-points is:

$$\overline{bit} = \alpha \times \left( 2 + \frac{2 + N_s}{g_1} + \frac{N_s + 16}{g_1 g_2} \right) + (1 - \alpha) \times 4 \quad (10)$$

where  $\alpha$  is the ratio of 2-bit and other parameters are as the same definition in equation (6) and (7).



**Figure 6: The GPU kernel speed degradation ratio of SpMV+GEMV with dequantization compared to GEMV with dequantization with different shape on different GPUs. Kernel launch region is irrelevant with sparse outlier ratio. Linear increasing region is increased linearly with the ratio.**

**Memory Alignment.** As illustrated in Figure 5 right, the quantized weights are split into two quantized weight matrixes. The first matrix stores 3 groups of 2-bit quantized weights (the size of group is 16) and 8 4-bit quantized weights into continuous memory. The other matrix stores 8 remaining 4-bit quantized weights into independent memory. Thus, the memory access for each thread in GPU is memory-aligned. Then, to maximize memory utilization without compromising accuracy, a nuanced approach is taken with scale factors and zero values. 1-order zeros  $1^{st}z$  and scales  $1^{st}s$  are coherently packed together and two of the three scales are compressed into 3-bit representations to occupy a 16-bit memory space efficiently (Scale Clipping). In contrast, 2-order zeros are stored with paddings (Zero Padding) to maintain alignment with memory boundaries. These meticulous memory alignment techniques effectively preserve accuracy loss while incurring a negligible  $<0.02$  average bit overhead.

## 4 EXCLUSIVE 2-BIT SPARSE OUTLIER

The intra-weight mixed-precision quantization method can efficiently reduce quantization errors of groups with large variations. However, there are still some outliers sparsely distributed in the weight matrices. Reserving these sparse outliers can further improve accuracy but it also introduces speed degradation affected by the sparse outlier ratio.

### 4.1 Challenge

Reserving these sparse outliers can further improve accuracy but it also introduces speed degradation affected by the outlier ratio. In order to reduce the quantization error caused by these random outliers, the straightforward approach is to reserve a large amount of outliers from all groups in weight matrix, which introduces speed degradation and large average bit increment. Previous works reserve 1.5% sparse outliers to improve the accuracy, and apply a

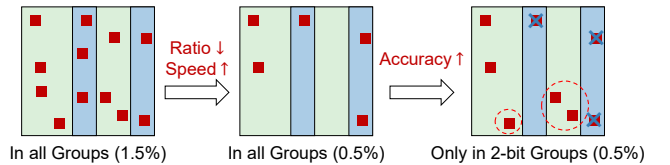
sparse matrix-vector multiplication (SpMV) to compute these outliers. Due to the differences in computational processes between the SpMV and the general matrix-vector multiplication (GEMV) with dequantization, these two parts need to be calculated separately and then combined at the end, resulting in  $>30\%$  speed degradation and hindering the inference speedup [5]. And reserving less outliers maintains the speed but brings limited accuracy improvement. Thus, the key challenge is to maximize the accuracy improvement and maintain speed simultaneously by reserving sparse outliers.

### 4.2 Motivation and Insight

We depict the GPU kernel speed degradation ratio caused by sparse outliers in Figure 6. In kernel launch region, the speed degradation ratio is fixed while in linear increasing region it is increased linearly with sparse outlier ratio. Besides straightforwardly reserving a large ratio of outliers from all groups in the weight matrix, another way is to reserve a small ratio of outliers from some specific groups. Previous designs like SpQR [5] and SqueezeLLM [17] apply the corresponding sparse format (e.g., CSR) to store and compute the sparse outliers. Each sparse outlier requires at least one 16-bit for the weight representation and one 16-bit for the position. Therefore, besides the speed degradation, introducing 1.5% 16-bit sparse outliers leads to  $(16 + 16 - 2) \times 1.5\% = 0.45$  extra average bit. Thus, our key insight is, **reserving a small fraction (e.g.,  $\leq 0.5\%$ ) of sparse outliers from 2-bit groups can maintain inference speed and the algorithm accuracy simultaneously.**

### 4.3 Approach

The detailed method of the exclusive 2-bit sparse outlier is illustrated in Figure 7. The green regions represent weights quantized to 2-bit precision, while the blue regions represent weights quantized to 4-bit precision. The red squares represent sparse outliers. The original ratio of outliers is 1.5% and they are reserved from all groups including 2-bit groups and 4-bit groups. First, we constrain the ratio of sparse outliers to 0.5%, which decreases the speed degradation ratio from  $>30\%$  to the minimum (about 20%, for GEMV with the size of (4096,4096) on NVIDIA RTX 3090). Then, we constrain the reservation regions from all groups to only 2-bit groups. Thus, with the minimum inference overheads, our method reserves more outliers in 2-bit groups. In practice, we only reserve 0.2-0.5% outliers and increase 0.06-0.15 average bit. We utilize the CSR format used in previous designs [5, 17] for sparse outliers representation, and apply GPU kernels in the NVIDIA cuSPARSE library [30] to perform matrix-vector multiplication on these outliers.



**Figure 7: Exclusive 2-bit sparse outlier reservation method. The green regions represent 2-bit groups while the blue regions represent 4-bit groups. The red squares represent sparse outliers. We constrain the ratio of sparse outliers from 1.5% to 0.5% and from all groups to only 2-bit groups.**

**Algorithm 1** GEMV with Asynchronous Dequantization GPU Kernel Algorithm

**Input:** Quantized weights  $W_{int}$ , 1-order zero\_points  $1^{st}Z$ , 1-order scaling\_factors  $1^{st}S_{int}$ , 2-order zero\_points  $2^{nd}Z$ , 2-order scaling\_factors  $2^{nd}S$ , input  $X$ , input channel  $IC$ .

```

1: /*Hyperparameters for 2-bit dequantization*/
2:  $N_{worker} = 2048$ ,  $N_{pack} = 16$ 
3:  $L_o = \text{Ceil}(IC/N_{worker})$ 
4:  $L_i = 4$  /*Equal with vector load double4*/
5: for _ from 1 to  $L_o$  do
6:   /*load  $1^{st}S_{int}$ ,  $2^{nd}Z$ ,  $2^{nd}S$ ,  $X$  to Shared Memory*/
7:    $sm\_s\_1st \leftarrow 1^{st}S_{int}$ ,  $sm\_x \leftarrow X$ 
8:    $sm\_z\_2nd \leftarrow 2^{nd}Z$ ,  $sm\_s\_2nd \leftarrow 2^{nd}S$ 
9:    $\_syncthreads()$ 
10:  for _ from 1 to  $L_i$  do
11:    /*load  $W_{int}$ ,  $1^{st}Z$  to Shared Memory*/
12:     $sm\_w\_int \leftarrow W_{int}$ ,  $sm\_z\_1st \leftarrow 1^{st}Z$ 
13:    /*load to registers*/
14:     $1^{st}s_{int} \leftarrow sm\_s\_1st$ ,  $2^{nd}z \leftarrow sm\_z\_2nd$ ,  $2^{nd}s \leftarrow sm\_s\_2nd$ 
15:     $1^{st}s \leftarrow \text{dequantize}(1^{st}s_{int}, 2^{nd}z, 2^{nd}s)$ 
16:     $\_syncthreads()$ 
17:    for _ from 1 to  $N_{pack}$  do
18:      /*load input to register*/
19:       $w_{int} \leftarrow sm\_w\_int$ ,  $1^{st}z \leftarrow sm\_z\_1st$ ,  $x \leftarrow sm\_x$ 
20:       $w_{float} \leftarrow \text{dequantize}(w_{int}, 1^{st}z, 1^{st}s)$ 
21:       $psum \leftarrow psum + w_{float} \times x$ 
22:    end for
23:  end for
24: end for
25:  $\text{parallel\_reduce}(O, psum)$ 
Output: Output vector  $O$ .

```

## 5 ASYNCHRONOUS DEQUANTIZATION

### 5.1 Challenge

Because the weights are quantized by 2/4-bit intra-weight mixed-precision, it requires the dequantization operation to restore the weights to half data type (16-bit) before performing the multiplication between the input and weights. We apply 1-order and 2-order quantization to quantize the weights of each group and the scales of adjacent groups, respectively. Previous designs (*e.g.*, SpQR[5], Greenbit[13]) use the synchronous dataflow (*i.e.*, performing dequantization after loading all weights), resulting in >50% overheads of end-to-end execution time. Thus, the key challenge is the time-consuming synchronous dequantization operation becomes the bottleneck in accelerating LLM inference after quantization.

### 5.2 Motivation and Insight

The 1-order and 2-order quantization requires two synchronous dequantizations in dataflow. However, the 2-order dequantization for calculating the scales of each group is independent of the weights of each group for the 1-order dequantization. Moreover, the 1-order dequantization for restoring the weights to 16-bit is independent of the input data of multiplication. Thus, our key insight is **calculating the scales of each group in 2-order dequantization can be overlapped with loading weights of each group in 1-order dequantization in GPU kernel**.

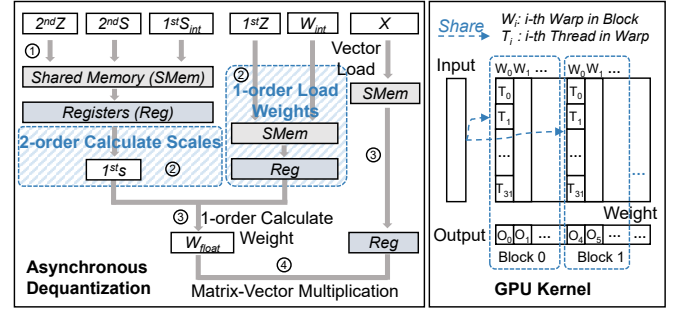


Figure 8: Asynchronous dequantization on GPUs.

### 5.3 Approach

Based on the insights above, we design the asynchronous dequantization dataflow as illustrated in Figure 8 on GPUs. With the help of the shared memory, we can overlap the calculating scales of each group and the loading weights of each group. Further, we use CUDA primitive `__shfl_down_sync()` to reduce the partial result inside a warp in parallel. From the perspective of memory access, we apply vector load (*e.g.* double4) technique to load quantized weights, input, scales and zeros to minimize the numbers of memory access. Algorithm 1 shows the four main parts including outer loop, inner loop, packed loop and parallel reduction. The outer loop (Line 5-24) loads the 2-order zero-points  $2^{nd}Z$ , quantized 1-order scaling factors  $1^{st}S_{int}$  and 2-order scaling factors  $2^{nd}S$  from global memory to shared memory. Then, it prepares for the 2-order dequantization, corresponding to ① in Figure 8. The inner loop (Line 10-23) loads the quantized weight  $W_{int}$  and the 1-order zero-points  $1^{st}Z$  from global memory to shared memory. Simultaneously, it also performs 2-order dequantization for scaling factors  $1^{st}s$ , corresponding to ②. The packed loop (Line 17-22) uses the dequantized scaling factors  $1^{st}s$ , 1-order zeros-point  $2^{nd}Z$  and quantized weights  $W_{int}$  to perform 1-order dequantization, and then loads input vector  $X$  for multiplication, corresponding to operation ③ and ④. Last, parallel reduction is performed to add partial results from all threads.

## 6 EXPERIMENTAL RESULT

### 6.1 Experimental Setup

**Benchmarks.** We conduct comprehensive experiments on the Llama1 [25], Llama2 [38], Llama3 [24] model families and the ChatGLM3-6b model [45], which are owing to critical and efficient influence in recent model advancements. We focus on two primary metrics: *perplexity* and *zero-shot performance*. The perplexity (PPL) is evaluated by the WikiText-2 [23] benchmark. The zero-shot performance is assessed across four zero-shot benchmarks, namely Piqa [1], HellaSwag [44], WinoGrande [34], and Arc-e [3]. We also conduct experiments on the BERT-base [6] model on two datasets (MNLI and STS-B) of the GLUE benchmark [40].

**Baselines.** We compare our method with state-of-the-art single-precision and mixed-precision quantization methods, including AWQ [19], GPTQ [10], OmniQuant [36], APTQ [12], LLM-MQ [18] and Greenbit [13]. For kernel and end-to-end performance, the original PyTorch implementation on HuggingFace [15] is also used as the baseline.

**Hardware Platforms.** We implement our design and compare with other baselines on NVIDIA RTX 2080, NVIDIA RTX 3090, and

**Table 1: LLM Algorithm Perplexity and Zero-shot Accuracy**

Method	$\bar{bit}$	PPL(↓)	Piqa/Hella./Wino./Arc-e	Avg.(↑)
Llama1-7b	16	5.68	79.2/76.2/69.9/72.8	75.1
APTQ[12]	3.13	6.76	74.5/68.3/65.3/57.9	66.5
<b>Our</b>	<b>2.84</b>	<b>6.61</b>	76.6/72.8/68.5/70.4	<b>72.1</b>
<b>Our (0.2%)</b>	<b>2.91</b>	<b>6.56</b>	76.8/73.8/68.7/71.0	<b>72.6</b>
Llama1-13b	16	5.09	80.3/79.0/72.7/74.8	76.7
APTQ[12]	3.13	/	74.4/71.2/68.0/64.1	69.4
<b>Our</b>	<b>2.84</b>	<b>5.92</b>	79.2/76.9/71.6/72.7	<b>75.1</b>
<b>Our (0.2%)</b>	<b>2.91</b>	<b>5.89</b>	79.3/77.8/71.9/73.0	<b>75.5</b>
Llama2-7b	16	5.47	78.5/56.7/67.2/69.3	67.9
GPTQ[10]	3	8.37	71.7/48.2/61.2/56.1	59.3
OmniQuant[36]	3	6.65	74.1/51.9/63.5/63.8	63.3
LLM-MQ[18]	2.91	/	76.5/53.2/65.0/65.8	65.1
Greenbit[13]	2.91	<b>6.09</b>	77.2/53.8/65.8/62.5	64.8
<b>Our</b>	<b>2.84</b>	6.62	75.9/51.3/66.4/66.9	<b>65.1</b>
<b>Our (0.2%)</b>	<b>2.91</b>	6.59	76.1/52.0/66.8/67.2	<b>65.7</b>
Llama2-13b	16	4.88	78.3/59.7/69.6/73.3	70.2
GPTQ[10]	3	6.44	75.2/56.1/64.3/64.2	64.0
OmniQuant[36]	3	5.59	77.2/56.8/66.5/66.7	66.8
LLM-MQ[18]	2.91	/	77.6/56.9/68.0/72.6	68.8
<b>Our</b>	<b>2.84</b>	<b>5.48</b>	77.2/57.6/68.5/71.2	<b>68.7</b>
<b>Our (0.2%)</b>	<b>2.91</b>	<b>5.37</b>	77.8/58.8/68.6/71.3	<b>69.1</b>
Llama3-8b	16	6.10	79.9/60.2/72.8/80.1	73.3
GPTQ[10]	3	13.03	60.8/41.8/60.9/38.8	50.6
<b>Our</b>	<b>2.84</b>	<b>8.54</b>	75.3/54.9/70.3/69.4	<b>67.5</b>
<b>Our (0.2%)</b>	<b>2.91</b>	<b>8.46</b>	76.8/55.4/70.5/70.3	<b>68.3</b>
ChatGLM3-6b	16	/	70.8/49.4/61.3/51.1	58.2
<b>Our</b>	<b>2.91</b>	/	66.7/44.7/59.4/48.5	<b>54.8</b>

NVIDIA A100 (80G) GPUs with CUDA version 12.2. As mentioned in Figure 1 and Table 3, the runtime cost is evaluated considering the power consumption using the NVML library [31], and the hardware buying cost is evaluated considering the price (\$12,500 for A100, \$1499 for RTX 3090, and \$699 for RTX 2080). With the 5-year service life period, the buying cost for each GPU is: \$0.285/h for A100, \$0.034/h for RTX 3090, and \$0.016/h for RTX 2080. We assume the electricity cost is 0.09\$/kWh [11], and the total cost for each GPU can be calculated by:

$$\text{buying cost} = \text{price} \div \text{years}$$

$$\text{runtime cost} = \text{electricity cost} \times \text{measured runtime power} \quad (11)$$

$$\text{total cost} = \text{buying cost} + \text{runtime cost}$$

## 6.2 Perplexity and Accuracy Evaluation

Table 1 shows algorithm perplexity and accuracy with zero-shot performance on Llama1, Llama2, Llama3 families and ChatGLM3-6b. We set the 1-order group size to 16 and 2-order group size to 16 for Llama1-7b/13b/70b, Llama2-7b/13b and Llama3-8b with 0.2% sparse outliers, and the 1-order group size to 8 and 2-order group size to 16 for and ChatGLM3-6b. For Llama1 family, we only compare with APTQ because its code is not open-source and it only provides the results for these LLMs. And with the emergent abilities [41] of LLMs, we can achieve lower accuracy loss for Llama2-70b model. Table 2 shows the algorithm accuracy on BERT-base.

**Llama1.** Compared with APTQ [12], we achieve 5.6% and 5.7% higher accuracy with lower 2.84 average bit and 13b models,

**Table 2: BERT Algorithm Accuracy and Compression Ratio**

Method	$\bar{bit}$	CR	MNLI	STS-B	Avg.(↑)
BERT <sub>base</sub>	32	1×	84.9	89.7	87.3
GOBO[43]	3.26	9.8×	83.8	88.3	86.1
Q-BERT[37]	3.52	9.1×	81.8	/	/
<b>Our</b>	<b>2.85</b>	<b>11.2×</b>	<b>84.2</b>	<b>88.7</b>	<b>86.5</b>

respectively. By reserving 0.2% sparse outliers, we achieve 5.56% and 5.7% higher accuracy with lower 2.91 average bit, respectively.

**Llama2.** Compared with GPTQ [10], our design significantly reduces the accuracy loss from 8.7% to 2.2% for Llama2-7b, and from 6.4% to 1.1% Llama2-13b with lower average bit. Compared with OmniQuant [36], our design significantly reduces the accuracy loss from 4.6% to 2.2% for Llama2-7b, and from 3.4% to 1.1% Llama2-13b. Compared with LLM-MQ [18], we achieve 0.6% and 0.3% higher accuracy for Llama2-7b and Llama2-13b, respectively. Compared with Greenbit [13] on Llama2-7b, we achieve 0.9% higher accuracy with the same 2.91 average bit. And due to the emergent abilities [35], Llama2-70b can prevent more accuracy loss than Llama2-13b.

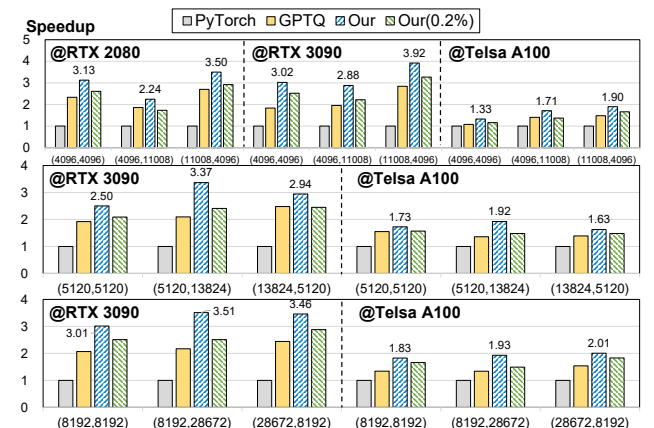
**Llama3.** On Llama3-8b, compared with GPTQ, our design reduces the accuracy loss from 22.7% to 5.0% with lower average bit.

**ChatGLM3.** On ChatGLM3-6b, the accuracy loss 3.4% is also controllable and it can be lower if the bilingual dataset is used for calibration.

**BERT.** On BERT-base model, compared with GOBO [43] and Q-BERT [37] methods, we achieve 2.4% and 0.4% higher accuracy with lower 2.85 average bit. Compared with the original full precision model, the accuracy loss is only 0.8%.

## 6.3 Kernel Evaluation

The detailed speedup of the GEMV kernel with dequantization is shown in Figure 9. We compare our kernel performance with some mainstream quantized kernels (e.g., GPTQ [10]) and the original FP16 models in common LLM GEMV cases. Because the large accuracy loss of AWQ method with 3-bit quantization, we exclude the comparison with its kernel. Without computing sparse outliers, our kernel achieves  $1.33\times \sim 3.92\times$  over FP16 on various NVIDIA



**Figure 9: GEMV kernel with dequantization speedup.** The tuple in the x-label represents the matrix shape of GEMV kernel.



**Table 3: Comparison of end-to-end speedup, runtime cost, and total cost**

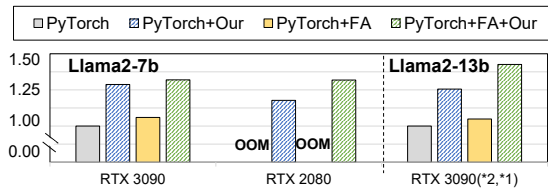
Model	Method	GPU	Power (W)	Token/s	End-to-end speedup	Runtime cost reduction	Total cost reduction
Llama2-7b	PyTorch	3090×1	290/350	25.9	1×	1×	1×
	GPTQ	3090×1	320/350	37.5	1.45×	1.31×	1.12×
	Our	3090×1	350/350	45.2	<b>1.74×</b>	1.44×	1.16×
	Our	2080×1	150/215	34.0	1.31×	<b>2.53×</b>	<b>2.29×</b>
Llama2-13b	PyTorch	3090×2	240/350	22.0	1×	1×	1×
	GPTQ	3090×1	270/350	23.1	1.06×	1.89×	1.95×
	Our	3090×1	250/350	25.0	<b>1.14×</b>	<b>2.19×</b>	<b>2.07×</b>
Llama2-70b	PyTorch	A100×2	400/400	38.5	1×	1×	1×
	GPTQ	A100×1	400/400	46.5	1.21×	2.42×	2.04×
	Our	A100×1	400/400	50.5	<b>1.31×</b>	<b>2.62×</b>	<b>2.06×</b>
Llama3-8b	PyTorch	3090×1	290/350	26.5	1×	1×	1×
	Our	3090×1	350/350	48.2	<b>1.82×</b>	1.51×	1.22×
	Our	2080×1	150/215	34.0	1.30×	<b>2.61×</b>	<b>2.36×</b>
ChatGLM3-6b	PyTorch	3090×1	350/350	32.5	1×	1×	1×
	Our	3090×1	350/350	40.0	1.23×	1.23×	1.09×
	Our	2080×1	170/215	45.5	<b>1.40×</b>	<b>2.88×</b>	<b>2.44×</b>

GPUs. With 0.2% sparse outliers, our kernel achieves  $1.16\times \sim 3.23\times$  over FP16 on various NVIDIA GPUs.

#### 6.4 Performance Evaluation

We compare the inference cost (end-to-end speedup, runtime cost reduction, and total cost reduction) including 0.2% sparse outliers among our method, PyTorch, and GPTQ in various model families and model sizes in Table 3. All results are normalized to PyTorch. In the Llama2 model family, our design achieves  $1.14\times \sim 1.74\times$  end-to-end speedup and far outperforms other previous works (e.g., GPTQ [10]) in terms of running cost reduction, reaching  $2.19\times \sim 2.53\times$  over the original model. The performance of Llama1 models is the same as Llama2 because they have the same model parameter. Moreover, the hardware requirement is dramatically lowered in our design for all models in Table 3, bringing up to  $2.88\times$  runtime cost reduction and  $2.44\times$  total cost reduction. The performance underscores the effectiveness of our design’s optimizations and the practical utility of our approach tailored to optimize the inference and deployment of LLMs.

We also show that our method is also compatible with FlashAttention [4], the most widely used LLM inference engine. Figure 10 shows that our method achieves higher speedup compared with adopting FlashAttention. Integrating our method with FlashAttention further accelerates LLM inference by up to  $1.31\times$ , and even enables LLM inference on GPUs with limited memory (e.g., NVIDIA RTX 2080).



**Figure 10: FlashAttention (FA) [4] integration and comparison.**

#### 6.5 Ablation Study

We also present an ablation study to validate the superiority of our method over manual inter-layer quantization schemes. The most intuitive mixed-precision quantization strategy is to uniformly quantize all layers within each block with single-precision. The results in Table 4 reveal our method’s efficacy over manual inter-layer quantization for Llama2-7B on WikiText-2, reflected in its consistently lower perplexity across various quantization ratios.

**Table 4: Ablation study: Perplexity comparison on Llama2-7b**

Method	$\overline{bit}$	Ratio of 4-bit	Perplexity ( $\downarrow$ )
Manual inter-layer	2.85	25%	7.13
<b>Our</b>	2.85	25%	6.62
Manual inter-layer	2.60	10%	9.32
<b>Our</b>	2.60	10%	7.13

## 7 CONCLUSIONS

We enable fast and efficient 2-bit LLM inference on GPUs in this paper with three novel techniques. We apply intra-weight mixed-precision quantization for weight matrices with 2-bit and 4-bit groups. We also introduce exclusive 16-bit sparse outliers in the 2-bit group with the minimum GPU kernel launching overhead. We further design the asynchronous GPU kernel to accelerate LLM dequantization. As a result, with our 2/4/16 mixed-precision quantization for each weight matrix and asynchronous dequantization during inference, our design achieves an end-to-end speedup for Llama2-7b is  $1.74\times$  over the original model, and we reduce both runtime cost and total cost by up to  $2.53\times$  and  $2.29\times$  with less GPU requirements.

## 8 ACKNOWLEDGMENTS

This work was supported by the National Natural Science Foundation of China (No. 62104128, U21B2031), Beijing Douyin Information Service Co., Ltd.



## REFERENCES

- [1] Yonatan Bisk, Rowan Zellers, et al. 2020. Piqa: Reasoning about physical commonsense in natural language. In *AAAI*.
- [2] Necati Cihan Camgoz, Oscar Koller, Simon Hadfield, and Richard Bowden. 2020. Sign language transformers: Joint end-to-end sign language recognition and translation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 10023–10033.
- [3] Peter Clark, Isaac Cowhey, et al. 2018. Think you have solved question answering? try arc, the ai2 reasoning challenge. *arXiv preprint arXiv:1803.05457* (2018).
- [4] Tri Dao, Dan Fu, et al. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *NeurIPS* (2022).
- [5] Tim Dettmers et al. 2023. SpQR: A Sparse-Quantized Representation for Near-Lossless LLM Weight Compression. *arXiv preprint arXiv:2306.03078* (2023).
- [6] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805* (2018).
- [7] Mengnan Du, Fengxiang He, Na Zou, Dacheng Tao, and Xia Hu. 2023. Shortcut learning of large language models in natural language understanding. *Commun. ACM* 67, 1 (2023), 110–120.
- [8] Wade H Foy. 1976. Position-location solutions by Taylor-series estimation. *IEEE transactions on aerospace and electronic systems* 2 (1976), 187–194.
- [9] Elias Frantar and Dan Alistarh. 2023. SparseGPT: Massive Language Models Can Be Accurately Pruned in One-Shot. (2023).
- [10] 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).
- [11] FRED. 2024. Average Price: Electricity per Kilowatt-Hour in U.S. City Average. <https://fred.stlouisfed.org/series/APU000072610>.
- [12] Ziyi Guan, Hantao Huang, Yupeng Su, Hong Huang, Ngai Wong, and Hao Yu. 2024. APTQ: Attention-aware Post-Training Mixed-Precision Quantization for Large Language Models. *arXiv preprint arXiv:2402.14866* (2024).
- [13] Nianhui Guo et al. 2023. Advanced Ultra-Low Bitrate Compression Techniques for the LLaMA Family of LLMs. [https://github.com/GreenBitAI/low\\_bit\\_llama](https://github.com/GreenBitAI/low_bit_llama) (2023).
- [14] Wei Huang, Yangdong Liu, Haotong Qin, Ying Li, Shiming Zhang, Xianglong Liu, Michele Magno, and Xiaojuan Qi. 2024. BiLLM: Pushing the Limit of Post-Training Quantization for LLMs. *arXiv preprint arXiv:2402.04291* (2024).
- [15] HuggingFace. 2024. <https://huggingface.co/>.
- [16] Gunjan Keswani, Wani Bisen, Hirkani Padwad, Yash Wankhedkar, Sudhanshu Pandey, and Ayushi Soni. 2024. Abstractive Long Text Summarization Using Large Language Models. *International Journal of Intelligent Systems and Applications in Engineering* 12, 12s (2024), 160–168.
- [17] Sehoon Kim, Coleman Hooper, et al. 2023. SqueezeLLM: Dense-and-Sparse Quantization. *arXiv preprint arXiv:2306.07629* (2023).
- [18] Shiyao Li, Xuefei Ning, Ke Hong, Tengxuan Liu, Luning Wang, Xiuhong Li, Kai Zhong, Guohao Dai, Huazhong Yang, and Yu Wang. 2023. Llm-mq: Mixed-precision quantization for efficient llm deployment. In *The Efficient Natural Language and Speech Processing Workshop with NeurIPS*, Vol. 9.
- [19] 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 On-Device LLM Compression and Acceleration. *Proceedings of Machine Learning and Systems* 6 (2024), 87–100.
- [20] Jiawei Liu, Chunqiu Steven Xia, Yuyao Wang, and Lingming Zhang. 2024. Is your code generated by chatgpt really correct? rigorous evaluation of large language models for code generation. *Advances in Neural Information Processing Systems* 36 (2024).
- [21] Zechun Liu, Barlas Oguz, et al. 2023. LLM-QAT: Data-Free Quantization Aware Training for Large Language Models. *arXiv preprint arXiv:2305.17888* (2023).
- [22] Brady D Lund, Ting Wang, Nishith Reddy Mannuru, Bing Nie, Somipam Shimray, and Ziang Wang. 2023. ChatGPT and a new academic reality: Artificial Intelligence-written research papers and the ethics of the large language models in scholarly publishing. *Journal of the Association for Information Science and Technology* 74, 5 (2023), 570–581.
- [23] Stephen Merity, Caiming Xiong, et al. 2016. Pointer sentinel mixture models. *arXiv preprint arXiv:1609.07843* (2016).
- [24] Meta. 2024. Build the future of AI with Meta Llama 3. <https://llama.meta.com/llama3/>.
- [25] AI Meta. 2023. Introducing LLaMA: A foundational, 65-billion-parameter large language model. *Meta AI* (2023).
- [26] Bonan Min, Hayley Ross, Elior Sulem, Amir Pouran Ben Veyseh, Thien Huu Nguyen, Oscar Sainz, Eneko Agirre, Ilana Heintz, and Dan Roth. 2023. Recent advances in natural language processing via large pre-trained language models: A survey. *Comput. Surveys* 56, 2 (2023), 1–40.
- [27] Markus Nagel, Marios Fournarakis, et al. 2021. A white paper on neural network quantization. *arXiv preprint arXiv:2106.08295* (2021).
- [28] Khalid Nassiri and Moulay Akhloufi. 2023. Transformer models used for text-based question answering systems. *Applied Intelligence* 53, 9 (2023), 10602–10635.
- [29] Ansong Ni, Srini Iyer, et al. 2023. Lever: Learning to verify language-to-code generation with execution. In *International Conference on Machine Learning*.
- [30] NVIDIA. 2024. NVIDIA CUDA Sparse Matrix Library. <https://docs.nvidia.com/cuda/cusparse/>.
- [31] NVIDIA. 2024. NVIDIA Management Library (NVML) | NVIDIA Developer. <https://developer.nvidia.com/nvidia-management-library-nvml>.
- [32] Dylan Patel and Afzal Ahmad. 2024. The Inference Cost Of Search Disruption - Large Language Model Cost Analysis. <https://www.semianalysis.com/p/the-inference-cost-of-search-disruption>.
- [33] Baolin Peng, Chunyuan Li, et al. 2023. Instruction tuning with gpt-4. *arXiv preprint arXiv:2304.03277* (2023).
- [34] Keisuke Sakaguchi, Ronan Le Bras, et al. 2021. Winogrande: An adversarial winograd schema challenge at scale. *Commun. ACM* (2021).
- [35] Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. 2024. Are emergent abilities of large language models a mirage? *Advances in Neural Information Processing Systems* 36 (2024).
- [36] Wenqi Shao, Mengzhao Chen, Zhaoyang Zhang, Peng Xu, Lirui Zhao, Zhiqian Li, Kaipeng Zhang, Peng Gao, Yu Qiao, and Ping Luo. 2023. Omniquant: Omnidirectionally calibrated quantization for large language models. *arXiv preprint arXiv:2308.13137* (2023).
- [37] Sheng Shen, Zhen Dong, Jiayu Ye, Linjian Ma, Zhewei Yao, Amir Gholami, Michael W Mahoney, and Kurt Keutzer. 2020. Q-bert: Hessian based ultra low precision quantization of bert. In *Proceedings of the AAAI Conference on Artificial Intelligence*, Vol. 34. 8815–8821.
- [38] Hugo Touvron, Louis Martin, et al. 2023. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288* (2023).
- [39] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in neural information processing systems* 30 (2017).
- [40] Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R Bowman. 2018. GLUE: A multi-task benchmark and analysis platform for natural language understanding. *arXiv preprint arXiv:1804.07461* (2018).
- [41] Jason Wei, Yi Tay, Rishi Bommasani, Colin Raffel, Barret Zoph, Sebastian Borgeaud, Dani Yogatama, Maarten Bosma, Denny Zhou, Donald Metzler, et al. 2022. Emergent abilities of large language models. *arXiv preprint arXiv:2206.07682* (2022).
- [42] Guangxuan Xiao, Ji Lin, et al. 2023. Smoothquant: Accurate and efficient post-training quantization for large language models. In *ICML*.
- [43] Ali Hadi Zadeh, Isak Edo, Omar Mohamed Awad, and Andreas Moshovos. 2020. Gobo: Quantizing attention-based nlp models for low latency and energy efficient inference. In *2020 53rd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO)*. IEEE, 811–824.
- [44] Rowan Zellers, Ari Holtzman, et al. 2019. Hellaswag: Can a machine really finish your sentence? *arXiv preprint arXiv:1905.07830* (2019).
- [45] Aohan Zeng, Xiao Liu, et al. 2022. Glm-130b: An open bilingual pre-trained model. *arXiv preprint arXiv:2210.02414* (2022).