
FineQuant: Unlocking Efficiency with Fine-Grained Weight-Only Quantization for LLMs

Young Jin Kim*

Microsoft
youki@microsoft.com

Rawn Henry*

NVIDIA
rhenry@nvidia.com

Raffy Fahim

Microsoft
raffybekheit@microsoft.com

Hany Hassan Awadalla

Microsoft
hanyh@microsoft.com

Abstract

Large Language Models (LLMs) have achieved state-of-the-art performance across various language tasks but pose challenges for practical deployment due to their substantial memory requirements. Furthermore, the latest generative models suffer from high inference costs caused by the memory bandwidth bottleneck in the auto-regressive decoding process. To address these issues, we propose an efficient `weight-only` quantization method that reduces memory consumption and accelerates inference for LLMs. To ensure minimal quality degradation, we introduce a simple and effective heuristic approach that quantizes only the model weights of a pre-trained model with finer granularity. This approach is applicable to both Mixture-of-Experts (MoE) and dense models without requiring additional fine-tuning. Furthermore, we implement highly efficient GPU GEMMs that perform on-the-fly matrix multiplication and dequantization, supporting the multiplication of `fp16` or `bf16` activations with `int8` or `int4` weights. We evaluate our approach on large-scale open source models such as OPT-175B and internal MoE models, showcasing minimal accuracy loss while achieving up to 3.65 times higher throughput on the same number of GPUs.

1 Introduction

Large Language Models (LLMs) have proven their efficacy in various language tasks by increasing the number of trainable parameters and pre-training models on large-scale data to be used in different downstream tasks (Devlin et al., 2018; Radford et al., 2018; Brown et al., 2020; Zhang et al., 2022; Chowdhery et al., 2022). However, deploying such large models comes with a significant cost which increases proportionally with the model size. Model size growth has increased several orders of magnitude over the last few years (1,588 times larger from BERT `large` - 340 million to PaLM 540 billion)(Devlin et al., 2018; Chowdhery et al., 2022).

Quantization is a well-known compression technique and numerous studies have demonstrated the effectiveness of quantization in accelerating neural network model inference, particularly in natural language processing domain (Kim et al., 2019; Aji and Heafield, 2020; Kim and Awadalla, 2020; Fan et al., 2021; Park et al., 2022; Kim et al., 2022). However, it is still under-explored how weight-only quantization can be effectively utilized in the context of large language models. Also, the existing methods introduce complex and costly procedures such as additional Quantization Aware Training (QAT) and/or calibration on additional data. To more effectively solve the challenge, we focus on

*Equal contribution.

simple weight-only quantization method that requires no additional training in this study because it has multiple advantages - (i) the accuracy could be maintained well because its underlying numerical computation is done in floating-point precision which is more accurate. As a result, we can effectively push the precision to very low bit-ranges. (ii) it could be used for various hardware and GPU architectures without needing specific hardware instructions dealing with low-bit multiplications. (iii) it can avoid expensive additional training steps.

In this paper, we make the following contributions:

1. **Empirical Analyses of Quantization Behaviors and Fine-Grained Quantization Algorithm:** We provide comprehensive analyses of the quantization behaviors on Language Model Models (LLMs). We investigate the impact of applying low-bit quantization (down to 3-bits) on LLM accuracy. Then, we propose a fine-grained quantization algorithm that incorporates group-wise quantization and adaptive selection of granularity. This approach helps preserve the original floating-point precision accuracy even when there is loss due to quantization.
2. **Accelerated Inference with Large-Scale Models with Highly Efficient GPU Kernels:** We demonstrate the effectiveness of the proposed method by applying it to a large-scale open-source dense transformer model called OPT. With its 175 billion parameters and internal MoE models utilizing optimized GPU kernels, our method enables deployment of the 175 billion parameter model on only 2 GPUs, resulting in a significant reduction of overhead and cost by 64%. Moreover, our method achieves 3.65 times higher throughput on the same number of GPUs.

2 Designing Quantization Methods for LLMs - Adaptive Fine-grained Quantization


Considering the design choices described in Appendix C, the granularity of quantization emerges as the most crucial component of the quantization algorithm. For efficient computation and reduced memory consumption, it is typical to have 1 quantization scale per tensor or 1 quantization scale for each column in the tensor. However, to maintain a close approximation of the original numerical values with the quantized values, it is desirable to have smaller groups of parameters sharing scales. This is necessary because outliers in the distribution have the potential to significantly skew the data, leading to decreased quantization precision, especially for smaller numerical values.

2.1 Catastrophic collapse of model performance


We have noted a significant decline in performance when employing matrix-wise quantization compared to column-wise quantization across various layers, as demonstrated in Appendix D. However, even with column-wise quantization, we have encountered instances of catastrophic collapse in LLM performance, particularly when certain outliers exist in the model weights. Figure 1a depicts the relationship between the Mean Squared Error (MSE) of quantized values and the translation BLEU scores as we modify the group size of quantization (quantization granularity) in the OPT 30B model. While increasing granularity leads to a gradual rise in MSE values, the model quickly loses its capability in terms of task BLEU score beyond a certain point. Consequently, it is crucial to determine the optimal granularity for each matrix to preserve the task performance while maximizing the size of the parameter groups which share scales.

2.2 Adaptive fine-grained quantization

Upon further investigation into the catastrophic failure of a quantized model, we have discovered that the failure could be rectified by adjusting the granularity of four specific matrices out of the 288 quantized matrices in the entire model. Merely increasing the granularity of these four matrices by a factor of two allowed for the recovery of over 94% of the lost accuracy. Based on this observation, we have developed a simple heuristic-based method to assign varying granularity to different model weight matrices. In the process of quantizing a matrix, we start from the column-wise quantization and compute the range of the values that must be quantized. We then halve the quantization group size and compute the range of each group. If for any group, $\frac{new_range}{old_range} > \alpha$ we halve the quantization group size again. We repeat this process until the quantization range differences between two granularities becomes smaller than α . Figure 1b illustrates the impact of adaptive group size on BLEU scores and



(a) MSE and BLEU changes with quantization granularity.



(b) BLEU score and model size comparison with adaptive group quantization with reference lines of fp16 and fixed group size (64). X-axis represents threshold value α of adaptive fine-grained quantization.

Figure 1: Impact analyses of quantization granularity on translation accuracy of OPT-30B.

model sizes in gigabytes (GB). With the adaptive fine-grained quantization approach, there is only a marginal 0.1% difference in BLEU score, while the model size is reduced to a mere 26% of the original FP16 model size.

3 Experiments

Our latency and throughput experiments are conducted using NVIDIA A100 SXM4 GPUs. We benchmark various open-source dense models including GPT-2-XL (1.5B) (Radford et al., 2019), OPT (13B and 30B) (Zhang et al., 2022), and OPT-IML (Max 30B and Max 175B) (Iyer et al., 2022) as well as internal pre-trained MoE models (5.3B) on various natural language tasks. More details of the models and tasks are available in Appendix B .

3.1 Accuracy

The impact of quantization on various natural language tasks using different models is presented in Appendix E (Table 2 and Table 4). In general, 8-bit weight-only quantization does not significantly affect the accuracy compared to fp16. This is observed across different language tasks, indicating that the models produce similar outputs. However, 4-bit quantization with column-wise granularity leads to some degradation in accuracy due to outliers in the weight distribution, as discussed in Section 2. To recover the accuracy, we adopt a more fine-grained group-wise quantization strategy, which shows similar accuracy to the original fp16.

3.2 Dense model performance results

3.2.1 Microbenchmarks

To understand how our weight-only quantization accelerates the matrix multiplies, we collect micro-benchmarks from OPT-13B and OPT-30b and present the results in Figure 2. We find that the matrix multiplies can be accelerated by up to 2.5X for those models when the number of tokens in the activation is small. This is typically the case for the auto-regressive part of LLMs which tends to dominate the overall run-time.

3.2.2 End to End Benchmarks

Table 1 shows end to end times (constructed from Table 6 in Appendix F) and associated throughput increases. To calculate the throughput increase, we assume the original FP16 model was sharded across 8-GPUs within a single node and that same node is used to serve INT8 or INT4 models. We measure the throughput per node by assuming that the model is replicated twice on the node for INT8 and 4 times for INT4 (64) and that requests are served to the independent model instances concurrently. We highlight that our compression technique allows serving 4 instances of OPT-175B on a single A100 node with 8 GPUs which can result in up to 3.65X throughput increase.

Table 1: Shows the throughput improvement for batch 1 on a 8-GPU node for different input and output lengths. We assume that the model is replicated twice on the node for INT8 weight-only quantization and 4 times for INT4 (64) weight-only quantization. We show the throughput increase relative to FP16 in parentheses next to through-puts for INT8 and INT4. The table is constructed using data from Table 6.

Input Length	Output Length	FP16 throughput per 8 GPU node (generated tokens per sec)	INT8 throughput per 8 GPU node (generated tokens per sec)	INT4 (64) throughput per 8 GPU node (generated tokens per sec)
128	32	24	49 (2.04×)	85 (3.54×)
128	128	25	52 (2.08×)	91 (3.64×)
512	32	21	41 (1.95×)	69 (3.29×)
512	128	23	47 (2.04×)	84 (3.65×)
1024	32	20	37 (1.85×)	57 (2.85×)
1024	128	23	47 (2.04×)	79 (3.43×)




Figure 2: Demonstrates the speed up over FP16 on only the matrix multiplies for OPT-13B (left) and OPT-30B (right) on a single GPU. We measure the performance of the QKV Projection, Attention Output, FFN1 and FFN2 matrix multiplies and compare our CUTLASS FP16 x INT GEMM against cuBLAS performing FP16 x FP16 GEMM. We report the geometric mean of the speedups across those 4 GEMMs while varying the number of rows in the activation (which represents batch_size × sequence_length). We highlight that when the number of rows is small (such as the decoding phase of GPT), we achieve up to 2.5X GEMM speedup when doing int4 quantization with block size 64.

3.3 MoE model performance results

We evaluate the performance of our weight-only quantization method on an MoE model in Appendix E (Table 5). Figure 3 shows the end-to-end speed improvements with various batch size with 8-bit and 4-bit quantization.




Figure 3: MoE model speed-up with quantization methods.

4 Conclusions and Limitations

This paper presents a method for accelerating large language models through the use of low-bit quantization. The proposed weight-only quantization technique demonstrates promising results in compressing very large models with up to 175 billion parameters, while still maintaining accuracy. To address the issue of outliers affecting the quantized weight distribution, fine-grained quantization is employed.

Despite its strengths, the study does have a few limitations. Firstly, optimized GPU kernels are only implemented for group size 64. However, we plan to expand support for any power of 2 group size greater than 16. Secondly, the performance benchmarking is conducted solely on A100 GPUs, so the

speed improvements may vary on different GPU architectures. Lastly, the proposed method does not leverage integer instructions even when they are available. These limitations suggest potential directions for future research.

References

- Alham Fikri Aji and Kenneth Heafield. 2020. Compressing neural machine translation models with 4-bit precision. In *NGT*.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. Palm: Scaling language modeling with pathways. *arXiv preprint arXiv:2204.02311*.
- 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*.
- Angela Fan, Pierre Stock, Benjamin Graham, Edouard Grave, Rémi Gribonval, Hervé Jégou, and Armand Joulin. 2021. Training with quantization noise for extreme model compression. *ArXiv*, abs/2004.07320.
- William Fedus, Barret Zoph, and Noam Shazeer. 2021. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *arXiv preprint arXiv:2101.03961*.
- Leo Gao, Jonathan Tow, Stella Biderman, Sid Black, Anthony DiPofi, Charles Foster, Laurence Golding, Jeffrey Hsu, Kyle McDonell, Niklas Muennighoff, Jason Phang, Laria Reynolds, Eric Tang, Anish Thite, Ben Wang, Kevin Wang, and Andy Zou. 2021. A framework for few-shot language model evaluation.
- Srinivasan Iyer, Xi Victoria Lin, Ramakanth Pasunuru, Todor Mihaylov, Dániel Simig, Ping Yu, Kurt Shuster, Tianlu Wang, Qing Liu, Punit Singh Koura, et al. 2022. Opt-1ml: Scaling language model instruction meta learning through the lens of generalization. *arXiv preprint arXiv:2212.12017*.
- Guolin Ke, Di He, and Tie-Yan Liu. 2021. Rethinking positional encoding in language pre-training. *ArXiv*, abs/2006.15595.
- Young Jin Kim and Hany Hassan Awadalla. 2020. Fastformers: Highly efficient transformer models for natural language understanding. *arXiv preprint arXiv:2010.13382*.
- Young Jin Kim, Ammar Ahmad Awan, Alexandre Muzio, Andres Felipe Cruz Salinas, Liyang Lu, Amr Hendy, Samyam Rajbhandari, Yuxiong He, and Hany Hassan Awadalla. 2021. Scalable and efficient moe training for multitask multilingual models. *arXiv preprint arXiv:2109.10465*.
- Young Jin Kim, Rawn Henry, Raffy Fahim, and Hany Hassan Awadalla. 2022. Who says elephants can’t run: Bringing large scale moe models into cloud scale production. *arXiv preprint arXiv:2211.10017*.
- Young Jin Kim, Marcin Junczys-Dowmunt, Hany Hassan, Alham Fikri Aji, Kenneth Heafield, Roman Grundkiewicz, and Nikolay Bogoychev. 2019. From research to production and back: Ludicrously fast neural machine translation. In *Proceedings of the 3rd Workshop on Neural Generation and Translation*, pages 280–288.
- Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. 2020. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*.
- Rui Liu, Young Jin Kim, Alexandre Muzio, and Hany Hassan. 2022. Gating dropout: Communication-efficient regularization for sparsely activated transformers. In *International Conference on Machine Learning*, pages 13782–13792. PMLR.

- Gunho Park, Baeseong Park, Se Jung Kwon, Byeongwook Kim, Youngjoo Lee, and Dongsoo Lee. 2022. nuqmm: Quantized matmul for efficient inference of large-scale generative language models. *arXiv preprint arXiv:2206.09557*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. 2018. Improving language understanding by generative pre-training.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Ashish Vaswani, Noam M. Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *NIPS*.
- Ruibin Xiong, Yunchang Yang, Di He, Kai Zheng, Shuxin Zheng, Chen Xing, Huishuai Zhang, Yanyan Lan, Liwei Wang, and Tie-Yan Liu. 2020. On layer normalization in the transformer architecture. In *ICML*.
- Susan Zhang, Stephen Roller, Naman Goyal, Mikel Artetxe, Moya Chen, Shuohui Chen, Christopher Dewan, Mona Diab, Xian Li, Xi Victoria Lin, et al. 2022. Opt: Open pre-trained transformer language models. *arXiv preprint arXiv:2205.01068*.

A Quantization method formulation

Linear quantization with absolute maximum. We used linear quantization with absolute maximum as the main method. Given a matrix A and b bits, this method encodes A as follows:

$$s_j = \frac{2 \times \max(|A_{:,j}|)}{2^b - 1}$$
$$Q_{:,j} = \text{int}\left(\frac{A_{:,j}}{s_j}\right)$$

Here, s is the scaling factor, which can be chosen per channel, as shown, or per the whole tensor. At inference time, the quantized Q is dequantized back to A' with the scaling factor s as follows:

$$A'_{:,j} = Q_{:,j} \times s_j$$

Log-scale quantization. Another quantization method we experimented is log-scale quantization where 1 bit is kept for the sign and $(b - 1)$ bits are used to encode the log-scaled values. Given a matrix A , the quantization formula is as follows:

$$P = \text{sign}(A)$$
$$T = \text{clip}\left(\frac{|A|}{s}, 1, 2^{1-2^{b-1}}\right)$$
$$Q = \lceil \log_2\left(\frac{2}{3}T\right) \rceil$$

where s can be chosen in two ways, either (i) the absolute maximum or (ii) the optimal value to minimize the mean squared error (MSE) between the quantized and original values which is described in Aji and Heafield (2020). We use the second algorithm which we observe a better accuracy with the quantization. At inference time, the quantized weight values are dequantized based on the formula as follows:

$$A' = P \times s \times 2^Q$$

Figure 4 shows the performance comparison of two quantization methods.

B Experimental setup

Our latency and throughput experiments are conducted using NVIDIA A100 SXM4 GPUs inside a Docker container running Ubuntu 20.04 and CUDA 11.8. All code is compiled using nvcc 11.8.89 and gcc/g++ 9.3. To carry out the experiments, we use a modified version of FasterTransformer² v5.3. The weight-only quantization kernels for per-column quantization are already open source.

Task and datasets. For the dense models, We utilize various open-source language tasks, including LAMBADA, HellaSwag, PiQA, WinoGrande, OpenBookQA, RTE, COPA from the lm-evaluation harness (Gao et al., 2021), as well as WMT machine translation task (WMT16 German and English)³.

For the MoE models, we use a multilingual machine translation task that covers 10 language translation directions from and into English covering German (de), French (fr), Italian (it), Spanish (es), Dutch (nl), and English (en). We use a 128,000 sub-word vocabulary, built with the sentencepiece library⁴. The number of training sentences is included in Appendix E. To measure the accuracy of the models, we utilized sacrebleu⁵ on the detokenized output.

Dense model architecture. For the dense model experiments, we utilize various open-source large language models that share a similar architecture, which consists of decoder-only with multiple transformer layers. To evaluate the accuracy of these models, we include GPT-2-XL (1.5B) (Radford et al., 2019), OPT (13B and 30B) (Zhang et al., 2022), and OPT-IML (Max 30B and Max 175B) (Iyer

²<https://github.com/NVIDIA/FasterTransformer>

³<https://statmt.org/wmt16/>

⁴<https://github.com/google/sentencepiece>

⁵<https://github.com/mjpost/sacrebleu>

et al., 2022). The number of model parameters ranges from 1.5 billion to 175 billion. The detailed number of layers and hidden dimensions can be found in the original papers.

MoE model architecture. For our MoE model experiments, we utilize internal pre-trained MoE models (5.3B) with a few modifications to the transformer model architecture (Vaswani et al., 2017). These modifications encompass the following: (i) a deep encoder consisting of 24 transformer layers and a shallow decoder comprising 12 transformer layers, (ii) adoption of Transformer with Untied Positional Encoding (TUPE) proposed in Ke et al. (2021) instead of the conventional sinusoidal positional embedding, and (iii) implementation of pre-layer normalization from Xiong et al. (2020). For the MoE models, we employ top-1 learned gating from Fedus et al. (2021) and an MoE layer with 32 experts at every other layer, specifically the even-numbered layers, as utilized in Lepikhin et al. (2020); Fedus et al. (2021); Kim et al. (2021). Additionally, we apply jittering noise, balancing loss (ratio of 0.01) (Lepikhin et al., 2020; Fedus et al., 2021) to more uniformly distribute expert utilization and gating dropout (0.2) (Liu et al., 2022) to prevent overfitting and improve regularization.

GPU kernel implementations. We utilized the kernel implementations developed by Kim et al. (2022), which rely on CUTLASS to create efficient kernels for fused dequantization and matrix multiplication. These kernels can process either FP16 or BF16 activations, a vector of scales of the same data type as the activation, and int8 or int4 weights. The kernels dequantize the weights to match the data type of the activation and perform floating-point tensor core math. The final output of the kernel is also of the same data type as the input activation. These kernels are available as open source code in FasterTransformer. To support multiple scaling factors for each column, we extended these kernels to process a matrix of scales, enabling us to implement int4 block quantization kernels. We set the block size to 64 for all performance analyses below, since it matches the K tile size of our fused gemm + dequantize kernels.

In compute-bound cases such as an encoder or the context creation phase of GPT, the conversions from integer to float bottlenecks our kernels, rather than tensor core math. As a result, our weight-only quantization GEMMs slower than equivalent FP16xFP16 GEMMs in compute bound cases but offer significant speedup in memory bound cases as seen in Figure 2. We argue that this kernel is useful because:

1. Large language models (LLMs) usually spend a lot more time in the memory-bound decoding phase than in the compute-bound context creation phase, especially when the output sequence length is long.
2. LLMs are typically served with small batch sizes in most practical cases, which puts significant pressure on the memory system during matrix multiplication as the weights need to be read from the GPU’s HBM. However, our kernel utilizes int4 compression, which reduces the number of bytes needed to load the weights by up to 4X. The overhead of loading the scales is small, even for block quantization with block size 64 as shown in Figure 2.

Quantization Method. All quantization experiments have one scaling factor for each column of the weight matrix, unless a block size B is specified. In that case, each contiguous block of B elements in a given column has its own scaling factor. This means we have multiple scaling factors per column.

C Quantization methodology: basic settings

Uniformity of quantization

We conducted experiments involving two quantization techniques that focus on the uniformity of the quantized range. Firstly, we employed linear quantization, which uniformly maps quantized integer values to their corresponding original float values. Secondly, we explored log-based quantization, inspired by Aji and Heafield (2020), where both integer and float ranges are mapped in a logarithmic scale. In both cases, we applied column-wise quantization to assess the impact of quantization uniformity on model accuracy. Detailed formulations for those two techniques are described in Appendix A.

Figure 4 illustrates the performance comparison between two quantization techniques applied to FFN layers using low bits. For 3 and 4 bits, both techniques exhibit similar performance. However, with 2-bit quantization, log-scale quantization shows a significant decrease in accuracy. Considering these

observations and the computational simplicity, we opt to use uniform quantization for all subsequent experiments.




Figure 4: A comparison of how the quality of the model, as measured by BLEU, changes when quantizing with different precisions using different quantization methods.

Symmetry - numerical distribution of model weights

In order to determine the most appropriate quantization approach, we have conducted further analysis on the weight parameter distribution across various layers. Figure 5 presents example distributions of model weights, which generally exhibit a normal distribution centered around zero. However, in some cases, outliers can distort the weight distribution, potentially leading to an inaccurate quantization range. Based on our observations and considering implementation efficiency, we choose to employ symmetric quantization around zero.




Figure 5: A comparison of example weight distributions from MoE and dense FFN layers.

D Channel-wise vs matrix-wise quantization

Scaling factors are calculated by the quantization algorithm and stored in half precision floating-point (fp16) numbers to dequantize the matrices with. These factors can be chosen on the channel scale or the whole matrix scale. As shown in figure 6, channel-wise quantization gives quite higher scores than tensor-wise especially for low precision. Additional parameters to store channel-wise scaling factors is small, because only one value is needed for a channel and less than 1% of total parameters in a matrix. Therefore, we use channel-wise quantization for all the quantization experiments.




Figure 6: Linear quantization of expert FFNs with channel-wise and matrix-wise scaling factors.

E Dense and MoE model quantized accuracy

Table 2: Accuracy of various models with low-bit weight only quantization on different natural language tasks. We also include the perplexity on the wikitext dataset for each model. We note that with int4 per-column for OPT-30B actually performs worse than FP16 for OPT-13B. Using block quantization (with block size 64) improves the accuracy by 2.3 % over just using per-column

Model type	Precision	LAMBADA	HellaSwag	PiQA	WinoGrande	OBQA	RTE	COPA	Average \uparrow	Wikitext \downarrow
GPT2-XL	fp16	51.1%	40.0%	70.7%	58.2%	22.4%	52.3%	73.0%	52.5%	20.4
	int8	51.1%	40.0%	70.7%	58.3%	22.6%	52.7%	73.0%	52.6%	20.4
	int4 (64)	49.3%	39.6%	70.7%	58.4%	20.6%	50.9%	74.0%	51.9%	20.9
	int4	47.5%	37.4%	69.4%	57.1%	19.4%	51.9%	73.0%	50.8%	21.7
OPT-13B	fp16	68.6%	52.5%	75.9%	65.0%	26.6%	58.1%	86.0%	61.8%	11.5
	int8	68.5%	52.4%	76.0%	65.4%	27.2%	57.0%	86.0%	61.8%	11.5
	int4 (64)	67.4%	50.7%	75.6%	65.4%	25.8%	59.2%	84.0%	61.2%	12.0
	int4	65.5%	50.2%	75.5%	64.8%	26.4%	56.0%	85.0%	60.5%	12.8
OPT-30B	fp16	71.5%	54.3%	77.6%	68.2%	30.2%	57.4%	82.0%	63.0%	10.7
	int8	71.4%	54.3%	77.6%	67.9%	30.2%	58.1%	82.0%	63.0%	10.7
	int4 (64)	69.9%	53.4%	77.5%	67.3%	30.0%	56.0%	83.0%	62.4%	11.1
	int4	69.5%	51.9%	75.8%	66.3%	26.8%	54.9%	79.0%	60.1%	11.6

Table 3: Perplexity using LM Eval Harness and FasterTransformer. OPT 66B suffers from catastrophic collapse with INT4 per column quantization, but recovers with block quantization with a size of 64.

Dataset	OPT 66B				OPT 175B			
	FP16	INT8 per col	INT4 per col	INT4 (64)	FP16	INT8 per col	INT4 per col	INT4 (64)
Wikitext	10.15	10.15	143.16	10.66	9.08	9.08	11.08	9.84

Table 4: Accuracy of OPT-IML models (30B and 175B) with various weight only quantization settings on machine translation tasks. The BLEU score is used as a metric and higher number represent a better result. The group size for the group-wise quantization is specified together with quantization bits.

Model type	Attention (group)	Others (group)	WMT 2016 German to English	Model Footprint (GB)
OPT-IML Max 30B	fp16	fp16	38.20	55.21
	int8 (7,168)	int8 (7,168)	38.20	27.66
	int4 (16)	int4 (16)	38.10	17.32
	int4 (64)	int4 (64)	37.96	14.73
	int4 (7,168)	int4 (7,168)	16.86	13.88
	int4 (adaptive)	int4 (adaptive)	38.12	14.62
	int3 (64)	int4 (64)	37.75	13.87
	int4 (64)	int3 (64)	37.06	12.15
	int3 (16)	int3 (16)	37.57	13.87
	int3 (64)	int3 (64)	36.95	11.29
	int3 (7,168)	int3 (7,168)	0.00 (degenerate)	10.43
OPT-IML Max 175B	fp16	fp16	41.14	324.16
	int8 (12,288)	int8 (12,288)	41.18	162.18
	int4 (64)	int4 (64)	40.86	86.23
	int4 (12,288)	int4 (12,288)	0.00 (degenerate)	81.19
	int3 (64)	int4 (64)	40.93	81.16
	int4 (64)	int3 (64)	37.02	71.04

Table 5: Accuracy of MoE models with quantization. Speed-up comparison is presented in Figure 3. The optimized kernels are implemented for 8-bit and 4-bit precisions.

Model type	Precision	BLEU (Δ BLEU compared to fp16)	Size (X times compared to fp16)
MoE 5.3B	fp16	46.35 (0.0)	1.00X
	int8	46.34 (-0.01)	0.55X
	int4	46.18 (-0.17)	0.32X
	int3	46.01 (-0.34)	0.26X

Table 6: We show the time taken to construct the context and the time per decoder step for OPT-175B on 8, 4 and 2 GPUs using our different weight-only quantization schemes. The numbers for int4 per-column quantization are similar to int4 (64) so they are omitted. The compute bound context creation phase is up to 3.5X slower when using INT4 block quantization, but running on 4x fewer GPUs. For INT8, it is up to 1.9X slower but runs on 2X fewer GPUs. In addition, the time per decoder step is typically within 20% of FP16 despite using 2X for 4X fewer GPUs with weight-only quantization. End to end times for different numbers of generated tokens can be estimated from this table by identifying the batch size and input length of interest and computing: $\text{context_time} + \text{num_generated_tokens} * \text{time_per_decoder_step}$. The batch sizes and sequence lengths shown are the maximum sizes that could fit in GPU memory.

Batch Size	Input length	FP16 (8 GPUs)		INT8 (4 GPUs)		INT4 (64) (2 GPUs)	
		Context time (ms)	Avg time per decoder step (ms)	Context time (ms)	Avg time per decoder step (ms)	Context time (ms)	Avg time per decoder step (ms)
1	128	60	40	76	38	121	43
2	128	82	41	134	38	226	42
4	128	148	41	283	38	431	43
8	128	272	41	468	40	835	45
12	128	372	42	743	41	1173	48
16	128	491	42	890	42	1627	49
32	128	935	44	1776	47	3261	58
1	512	148	42	280	40	427	44
2	512	273	43	470	40	838	45
4	512	493	43	892	41	1637	46
8	512	939	43	1784	43	3291	49
1	1024	271	41	465	39	829	44
2	1024	498	42	899	39	1648	46
4	1024	945	42	1795	42	3307	50

F OPT-175B Speed measurements with various input matrix sizes

We construct Table 6 as a reference to compute end to end times for different input and output lengths for OPT-175B on 8, 4 and 2 GPUs. Our table shows that the context phase slows down which is primarily due to running on fewer GPUs. Additionally, our weight-only quantization kernels have some slowdown for compute bound cases. However, we show that the time per decoder step is typically within 20 % of FP16 despite using 2X or 4x fewer GPUs. The per-token latency does not scale with the number of GPUs since fewer GPUs need to communicate and our kernels provide significant acceleration (as shown in Table 2) in the decoder phase.