



HACK: Homomorphic Acceleration via Compression of the Key-Value Cache for Disaggregated LLM Inference

Zeyu Zhang
University of Virginia
Charlottesville, Virginia, USA
qxc4fh@virginia.edu

Haiying Shen
University of Virginia
Charlottesville, Virginia, USA
hs6ms@virginia.edu

Shay Vargaftik
VMware Research
Palo Alto, California, USA
vargaftik@gmail.com

Ran Ben Basat
University College London
London, England, UK
r.benbasat@cs.ucl.ac.uk

Michael Mitzenmacher
Harvard University
Boston, Massachusetts, USA
michaelm@eecs.harvard.edu

Minlan Yu
Harvard University
Boston, Massachusetts, USA
minlanyu@g.harvard.edu

Abstract

Disaggregated Large Language Model (LLM) inference decouples the compute-intensive prefill stage from the memory-intensive decode stage, allowing low-end, compute-focused GPUs for prefill and high-end, memory-rich GPUs for decode, which reduces cost while maintaining high throughput. However, transmitting Key-Value (KV) data between the two stages can be a bottleneck, especially for long prompts. Additionally, the computational overhead in the two stages is key for optimizing Job Completion Time (JCT), and KV data size can become prohibitive for long prompts and sequences. Existing KV quantization methods can alleviate transmission and memory bottlenecks, but they introduce significant dequantization overhead, exacerbating the computation time.

We propose Homomorphic Acceleration via Compression of the KV cache (HACK) for disaggregated LLM inference. HACK eliminates the heavy KV dequantization and directly computes on quantized KV data to approximate and reduce the cost of expensive matrix multiplication. Extensive trace-driven experiments show that HACK reduces JCT by up to 70.9% compared to disaggregated LLM inference baseline and by up to 52.3% compared to state-of-the-art KV quantization methods.

CCS Concepts

• **Computing methodologies** → **Natural language generation**;
• **Networks** → **Application layer protocols**; • **Information systems** → **Information systems applications**.

Keywords

Large Language Models, Disaggregation, KV Cache, Compression

ACM Reference Format:

Zeyu Zhang, Haiying Shen, Shay Vargaftik, Ran Ben Basat, Michael Mitzenmacher, and Minlan Yu. 2025. HACK: Homomorphic Acceleration via Compression of the Key-Value Cache for Disaggregated LLM Inference. In *ACM SIGCOMM 2025 Conference (SIGCOMM '25)*, September 8–11, 2025, Coimbra, Portugal. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3718958.3750481>



This work is licensed under a Creative Commons Attribution 4.0 International License. *SIGCOMM '25, Coimbra, Portugal*

© 2025 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-1524-2/2025/09

<https://doi.org/10.1145/3718958.3750481>

1 Introduction

Disaggregated LLM inference improves cost-efficiency by assigning low-end GPUs (e.g., NVIDIA A10G, V100, T4, and L4) to the prefill stage and high-end GPUs (e.g., A100 and H100) to the decode stage [8, 16, 17, 19, 23, 24, 26, 27]. However, Key-Value (KV) transmission between the two stages can be a bottleneck, since low-end GPU instances often lack high-speed networking for cost savings. For example, AWS's A10G, V100, T4, and L4 instances cost roughly 10–20 times less than A100 instances—which typically offer 400 Gbps bandwidth—but their networks are limited to 10–50 Gbps or lower [10]. Similarly, Tencent Cloud's A100 instances are configured with only 5–50 Gbps bandwidth to cut costs [12]. Computation can also become a bottleneck due to the attention mechanism. Moreover, during the decode stage, GPU memory is constrained by the large volume of cached KV data [23, 27].

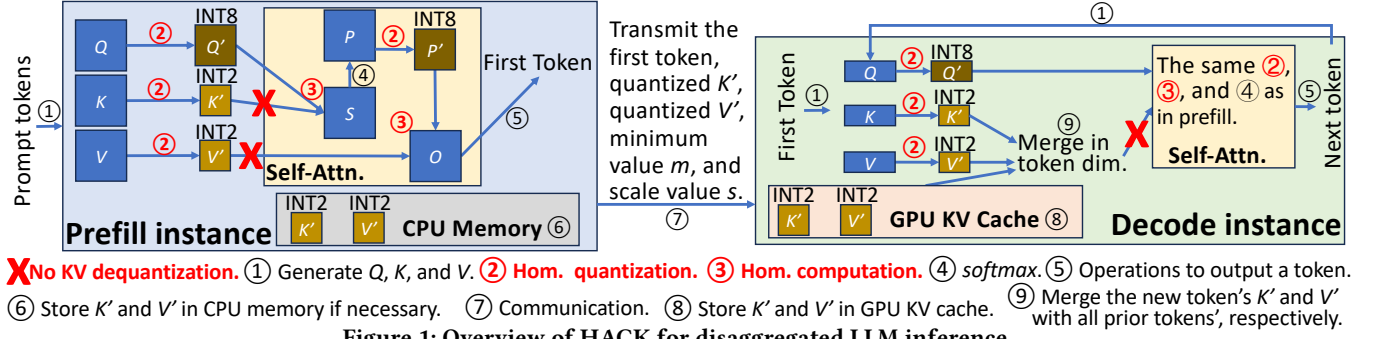
KV quantization (e.g., CacheGen [22] and KVQuant [15]) can alleviate transmission and memory bottlenecks. They quantize KV data after each iteration before storing it in the cache and then retrieve and dequantize all tokens' KV data in the next decode iteration. However, they introduce significant KV dequantization overhead and cannot reduce computation time.

Ideally, arithmetic operations should be executable directly on quantized KV data, eliminating dequantization and accelerating computation through smaller data elements. A similar idea has been explored for gradient aggregation [21], but this was limited to addition operations and unsuitable for the matrix multiplications required in the attention mechanism. To this end, we propose Homomorphic Acceleration via Compression of the KV cache (HACK) for disaggregated LLM inference. HACK addresses the KV transmission bottleneck by enabling computation on quantized data (INT2/8) while maintaining comparable inference accuracy and reducing computation and memory constraints. HACK is compatible with any quantization method that dequantizes data by linear transformation (e.g., MXFP4/8 [25]). We open-sourced the code of HACK [2]. This work does not raise any ethical issues.

2 Motivation

We show the networking, computation, and memory bottlenecks and demonstrate the limitations of current KV quantization methods in addressing these issues. The default experiment settings are detailed in §4.

In our measurements with the default settings, KV transmission can contribute up to 42.2% of JCT. Prefill and decode times can



reach up to 45.6% and 84.3% of JCT. GPU memory usage can reach up to 93.7%. KV memory access can reach up to 33.1% of JCT.

Although KV quantization (e.g., CacheGen and KVQuant) can reduce KV transmission overhead, memory usage, and KV memory access time, they introduce substantial dequantization costs per decode iteration. In our measurements, CacheGen and KVQuant introduce additional KV dequantization overhead up to 37.9% of JCT, which can be even higher for long sequences. In addition, they cannot reduce computation time. This highlights the need for a quantization method that simultaneously lowers communication and memory overhead, avoids the cost of dequantization, and reduces computation time.

3 Design

HACK avoids KV dequantization and reduces attention computation time via homomorphic matrix multiplication on quantized data. Fig. 1 illustrates the workflow. The most critical step is Step ②, which quantizes KV to INT2 and query Q to INT8 using homomorphic quantization, followed by Step ③, which performs homomorphic computation on the quantized data. HACK consists of the following components.

Homomorphic quantization for matrix multiplication. Attention primarily involves matrix multiplications. For any matrix multiplication $C = AB$, HACK quantizes A and B to obtain A' and B' and then finds $C' = A'B'$ to obtain a quantized output C' . C' is subsequently turned into an approximation of C with a minimal overhead. We use an asymmetric 2-bit/8-bit stochastic quantization [20] when performing homomorphic quantization to reduce quantization error. It identifies the minimum (\min_i) and maximum (\max_i) values of the matrix elements and computes the $scale = \frac{\max_i - \min_i}{2^2 - 1}$. Each original value x is quantized to an integer $x' = \text{round}(\frac{x - \min_i}{scale})$. The stochastic rounding $\text{round}(\cdot)$ rounds \cdot to $\lfloor \cdot \rfloor$ with probability $(\lceil \cdot \rceil - \cdot) / (\lceil \cdot \rceil - \lfloor \cdot \rfloor)$ and to $\lceil \cdot \rceil$ otherwise. We explain how to estimate C given C' . Let a_{iz} represent the element in the i -th row and z -th column of A , and b_{zj} represent the element in the z -th row and j -th column of B . The matrix multiplication $C = AB$ can then be expressed as $c_{ij} = \sum_z a_{iz} b_{zj}$, $\forall i, j$. Let m_{a_i} and s_{a_i} denote the minimum and scale values of a_{iz} . Since $a'_{iz} = \text{round}(\frac{a_{iz} - m_{a_i}}{s_{a_i}})$

and $b'_{zj} = \text{round}(\frac{b_{zj} - m_{b_j}}{s_{b_j}})$, we have $a_{iz} \approx s_{a_i} q_{a_{iz}} + m_{a_i}$ and $b_{zj} \approx s_{b_j} q_{b_{zj}} + m_{b_j}$. Thus, $(AB)_{ij}$ can be extended to:

$$\begin{aligned} \sum_z a_{iz} b_{zj} &\approx s_{a_i} s_{b_j} \sum_z a'_{iz} b'_{zj} + m_{b_j} s_{a_i} \sum_z a'_{iz} + \\ &\quad m_{a_i} s_{b_j} \sum_z b'_{zj} + Z m_{a_i} m_{b_j}, \end{aligned} \quad (1)$$

where $\{\sum_z a'_{iz} b'_{zj}, \forall i, j\}$ is the quantized matrix multiplication that can be accelerated by INT8 computation. The other terms in Eq. (1) allow computing an approximation of $\sum_z a_{iz} b_{zj}$ (C) from $\sum_z a'_{iz} b'_{zj}$ (C'). Eq. (1) provides homomorphic computation for multiplication. **Summation elimination.** We store the sum $\sum_z b'_{zj}$ in Eq. (1) for K and V during decode and reuse them every iteration to avoid the recomputation cost for the decode stage. This only needs a little extra memory, up to $\sim 2.7\%$ of the GPU memory capacity.

Requantization elimination for the last block of V . Quantization is applied to groups of elements. For the value matrix V , each group spans the sequence dimension. During decode, if the last block of V has fewer tokens than the group size, quantization metadata (e.g., the minimum) is undefined, requiring recomputation and requantization at each iteration until the group is full. To avoid this, we store the original FP16 values of the last group V in a cache, which consumes at most 0.51% of the GPU memory capacity.

4 Evaluation

Amazon EC2 provides a wide selection of GPU instances [9]. We use two AWS p4de.24xlarge (8 A100 and 400 Gbps for each) for decode [23, 27]; ten g5.12xlarge (4 A10G and 40 Gbps for each), sixteen p3.8xlarge (4 V100 and 10 Gbps for each), sixteen g4dn.12xlarge (4 T4 and 50 Gbps for each), ten g6.12xlarge (4 L4 and 40 Gbps for each), or two p4de.24xlarge for prefill to avoid underutilizing decode instances [23]. We evaluate HACK using Mistral-v0.3 7B [7], Phi-3 14B [6], Yi 34B [1], Llama-3.1 70B [5], and Falcon 180B [3] with their recommended and empirically validated Tensor Parallelism (TP) and Pipeline Parallelism (PP) sizes [23, 27] across various datasets (IMDb [18], HumanEval [11], GSM8K [4], arXiv [13], and Cocktail [14]).

When achieving 99% of the accuracy of the disaggregated baseline without quantization, HACK provides up to 86% KV size reduction, up to 70.9% JCT reduction over the disaggregated baseline, and up to 52.3% JCT reduction over quantization methods (CacheGen and KVQuant).

Acknowledgments

This work was supported in part by U.S. NSF grants NSF-2421782, NSF-2350425, NSF-2319988, NSF-2206522, Microsoft Research Faculty Fellowship 8300751, Amazon research award, AWS Cloud Credit for Research, and the Commonwealth Cyber Initiative (CCI), an investment in the advancement of cyber research, innovation and workforce development. For more information about CCI, visit cyberinitiative.org. This work was also supported in part by ACE, one of the seven centers in JUMP 2.0, a Semiconductor Research Corporation (SRC) program sponsored by DARPA.

References

- [1] 2025. 01-ai Model Yi. <https://huggingface.co/01-ai/Yi-34B-200K>.
- [2] 2025. The Code of HACK. <https://github.com/pcl-projects/HACK>.
- [3] 2025. Falcon-180B. <https://huggingface.co/tiiuae/falcon-180B>.
- [4] 2025. GSM8K. <https://huggingface.co/datasets/openai/gsm8k>.
- [5] 2025. Meta Llama-3.1. <https://llama.meta.com/>.
- [6] 2025. Microsoft Phi-3. <https://huggingface.co/microsoft/Phi-3-medium-128k-instruct>.
- [7] 2025. Mistral-v0.3. <https://huggingface.co/mistralai/Mistral-7B-Instruct-v0.3>.
- [8] 2025. NVIDIA Dynamo. <https://developer.nvidia.com/blog/introducing-nvidia-dynamo-a-low-latency-distributed-inference-framework-for-scaling-reasoning-ai-models/>.
- [9] Inc. Amazon Web Services. 2025. Amazon EC2 Instance Types. <https://aws.amazon.com/ec2/instance-types/>.
- [10] Inc. Amazon Web Services. 2025. Recommended AWS GPU Instances. <https://docs.aws.amazon.com/dlami/latest/devguide/gpu.html>.
- [11] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. 2021. Evaluating Large Language Models Trained on Code. [arXiv:2107.03374](https://arxiv.org/abs/2107.03374) [cs.LG] <https://arxiv.org/abs/2107.03374>
- [12] Tencent Cloud. 2025. Tencent Cloud - A100 Instances. <https://www.tencentcloud.com/document/product/560/19701#GT4>.
- [13] Arman Cohan, Franck Dernoncourt, Doo Soon Kim, Trung Bui, Seokhwan Kim, Walter Chang, and Nazli Goharian. 2018. A Discourse-Aware Attention Model for Abstractive Summarization of Long Documents. [arXiv:1804.05685](https://arxiv.org/abs/1804.05685) [cs.CL] <https://arxiv.org/abs/1804.05685>
- [14] Sunhao Dai, Weihao Liu, Yuqi Zhou, Liang Pang, Rongju Ruan, Gang Wang, Zhenhua Dong, Jun Xu, and Ji-Rong Wen. 2024. Cocktail: A Comprehensive Information Retrieval Benchmark with LLM-Generated Documents Integration. [arXiv:2405.16546](https://arxiv.org/abs/2405.16546) [cs.IR] <https://arxiv.org/abs/2405.16546>
- [15] Coleman Richard Charles Hooper, Sehoon Kim, Hiva Mohammadzadeh, Michael W. Mahoney, Sophia Shao, Kurt Keutzer, and Amir Gholami. 2024. KVQuant: Towards 10 Million Context Length LLM Inference with KV Cache Quantization. In *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. <https://openreview.net/forum?id=0LXotew9Du>
- [16] Cunchen Hu, Heyang Huang, Junhao Hu, Jiang Xu, Xusheng Chen, Tao Xie, Chenxi Wang, Sa Wang, Yungang Bao, Ninghui Sun, and Yizhou Shan. 2024. MemServe: Context Caching for Disaggregated LLM Serving with Elastic Memory Pool. [arXiv:2406.17565](https://arxiv.org/abs/2406.17565) [cs.DC] <https://arxiv.org/abs/2406.17565>
- [17] Cunchen Hu, Heyang Huang, Liangliang Xu, Xusheng Chen, Jiang Xu, Shuang Chen, Hao Feng, Chenxi Wang, Sa Wang, Yungang Bao, Ninghui Sun, and Yizhou Shan. 2024. Inference without Interference: Disaggregate LLM Inference for Mixed Downstream Workloads. [arXiv:2401.11181](https://arxiv.org/abs/2401.11181) [cs.DC] <https://arxiv.org/abs/2401.11181>
- [18] IMDb. 2020. Genre Classification Dataset IMDb. <https://www.kaggle.com/datasets/hijest/genre-classification-dataset-imdb>.
- [19] Philip Kiely. 2024. NVIDIA A10 vs A10G for ML model inference. <https://www.baseten.co/blog/nvidia-a10-vs-a10g-for-ml-model-inference/>.
- [20] John R. Klauder. 1983. Stochastic Quantization. In *Recent Developments in High-Energy Physics*, H. Mitter and C. B. Lang (Eds.). Springer Vienna, Vienna, 251–281.
- [21] Minghao Li, Ran Ben Basat, Shay Vargaftik, ChonLam Lao, Kevin Xu, Michael Mitzenmacher, and Minlan Yu. 2024. THC: Accelerating Distributed Deep Learning Using Tensor Homomorphic Compression. In *21st USENIX Symposium on Networked Systems Design and Implementation (NSDI 24)*. USENIX Association, Santa Clara, CA, 1191–1211. <https://www.usenix.org/conference/nsdi24/presentation/li-minghao>
- [22] Yuhua Liu, Hanchen Li, Yihua Cheng, Siddhant Ray, Yuyang Huang, Qizheng Zhang, Kuntai Du, Jiayi Yao, Shan Lu, Ganesh Ananthanarayanan, Michael Maire, Henry Hoffmann, Ari Holtzman, and Junchen Jiang. 2024. CacheGen: KV Cache Compression and Streaming for Fast Large Language Model Serving. In *Proceedings of the ACM SIGCOMM 2024 Conference* (Sydney, NSW, Australia) (ACM SIGCOMM '24). Association for Computing Machinery, New York, NY, USA, 38–56. doi:10.1145/3651890.3672274
- [23] Pratyush Patel, Esha Choukse, Chaojie Zhang, Aashaka Shah, Íñigo Goiri, Saeed Maleki, and Ricardo Bianchini. 2024. Splitwise: Efficient Generative LLM Inference Using Phase Splitting. In *2024 ACM/IEEE 51st Annual International Symposium on Computer Architecture (ISCA)*. 118–132. doi:10.1109/ISCA59077.2024.00019
- [24] Ruoyu Qin, Zheming Li, Weiran He, Mingxing Zhang, Yongwei Wu, Weimin Zheng, and Xinran Xu. 2024. Mooncake: A KVCache-centric Disaggregated Architecture for LLM Serving. [arXiv:2407.00079](https://arxiv.org/abs/2407.00079) [cs.DC] <https://arxiv.org/abs/2407.00079>
- [25] Bitu Darvish Rouhani, Ritchie Zhao, Ankit More, Mathew Hall, Alireza Khodamoradi, Summer Deng, Dhruv Choudhary, Marius Cornea, Eric Dellinger, Kristof Denolf, Stosic Dusan, Venmugil Elango, Maximilian Golub, Alexander Heinecke, Phil James-Roxby, Dharmesh Jani, Gaurav Kolhe, Martin Langhammer, Ada Li, Levi Melnick, Maral Mesmakhosroshahi, Andres Rodriguez, Michael Schulte, Rasoul Shafipour, Lei Shao, Michael Siu, Pradeep Dubey, Paulius Micikevicius, Maxim Naumov, Colin Verrilli, Ralph Wittig, Doug Burger, and Eric Chung. 2023. Microscaling Data Formats for Deep Learning. [arXiv:2310.10537](https://arxiv.org/abs/2310.10537) [cs.LG] <https://arxiv.org/abs/2310.10537>
- [26] Foteini Strati, Sara Mcallister, Amar Phanishayee, Jakub Tarnawski, and Ana Klimovic. 2024. DéjàVu: KV-cache Streaming for Fast, Fault-tolerant Generative LLM Serving. In *Proceedings of the 41st International Conference on Machine Learning (Proceedings of Machine Learning Research, Vol. 235)*, Ruslan Salakhutdinov, Zico Kolter, Katherine Heller, Adrian Weller, Nuria Oliver, Jonathan Scarlett, and Felix Berkenkamp (Eds.). PMLR, 46745–46771. <https://proceedings.mlr.press/v235/strati24a.html>
- [27] Yinmin Zhong, Shengyu Liu, Junda Chen, Jianbo Hu, Yibo Zhu, Xuanzhe Liu, Xin Jin, and Hao Zhang. 2024. DistServe: Disaggregating Prefill and Decoding for Goodput-optimized Large Language Model Serving. In *18th USENIX Symposium on Operating Systems Design and Implementation (OSDI 24)*. USENIX Association, Santa Clara, CA, 193–210. <https://www.usenix.org/conference/osdi24/presentation/zhong-yinmin>