

Can Low-Rank Knowledge Distillation in LLMs be Useful for Microelectronic Reasoning?

Nirjhor Rouf*

ECE Department,
North Carolina State University.
nrouf2@ncsu.edu

Fin Amin*

ECE Department,
North Carolina State University.
samin2@ncsu.edu

Paul D. Franzon

ECE Department,
North Carolina State University.
paulf@ncsu.edu

Abstract- In this work, we present empirical results regarding the feasibility of using offline large language models (LLMs) in the context of electronic design automation (EDA). The goal is to investigate and evaluate a contemporary language model's (Llama-2-7B) ability to function as a microelectronic Q&A expert as well as its reasoning, and generation capabilities in solving microelectronic-related problems. Llama-2-7B was tested across a variety of adaptation methods, including introducing a novel low-rank knowledge distillation (LoRA-KD) scheme. Our experiments produce both qualitative and quantitative results. Furthermore, we release our evaluation benchmark along with the code necessary to replicate our experiments at github.com/FinAminToastCrunch.

Index Terms- LLMs for EDA education, LLM fine-tuning, knowledge-distillation, RAG, Low-Rank adaptation

I. INTRODUCTION AND MOTIVATION

The emergence of Large Language Models (LLM) has revolutionized the field of natural language processing. At present, LLMs are garnering significant research interests for domain-specific tasks. In the field of electronic design automation (EDA) in particular, applications of LLMs are still at the nascent stage. However, it is very apparent that the effective use of LLMs in EDA can improve manufacturing yields by streamlining the design flow when it comes to IC design. Recently published works showed the successful use of LLMs in chip design [2], [3], [13]. Additionally, LLMs have also shown significant proficiency in the analysis of designed systems [8] and even in reviewing and analysis of design specifications of VLSI systems [11]. Development of open-source benchmarks such as VerilogEval [14] is also facilitating future research in this field. Similarly, LLMs can be useful in enhancing productivity. Internal studies carried out at Nvidia have shown that checklist related tasks can take up to 60% of an engineer's time and thus bottleneck productivity [13]. An LLM-based engineering assistant can certainly reduce this bottleneck by helping with engineering knowledge dissemination.

However, several key challenges must be addressed for more effective and efficient application of LLMs in EDA. One big concern is the unintentional data retention of DNNs from training sets [9]. There are two aspects of this issue. Firstly, classified IP designs can be leaked if the API stores user input.

*These authors contributed equally to this work.

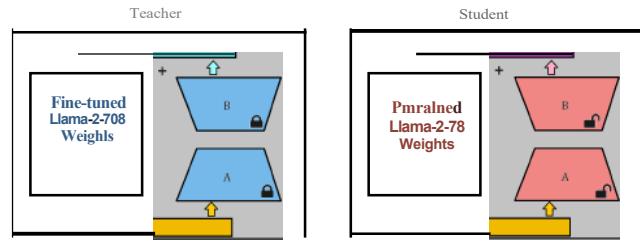


Fig. I. LoRA-KD works by first fine-tuning the teacher model using LoRA. Afterward, the teacher is frozen and its outputs are used for equation 4. Note that only the low-rank A and B parameters of the student are updated.

Secondly, when trying to complete a user request, the LLM can inadvertently use copy-righted IP designs without attributing references to them—potentially causing downstream legal trouble. Another major challenge is the heavy computational resource requirement of LLMs. For example, Meta's Llama2-70B requires 130 GB memory to load [17].

Choosing the appropriate LLM for EDA applications is also a big challenge and here, the proprietary vs open-source debate must be addressed. While proprietary models, such as ChatGPT-4 [1] are powerful, they have limited accessibility, store user data/designs, and are pay-to-use. Additionally, the inability to fine-tune them hinders their capabilities in domain-specific EDA tasks. On the other hand, open-source LLMs offering better accessibility are restricted by limited scale and resources compared to their proprietary counterparts resulting in lower performance [21]. In this work, we explore the feasibility of adapting the open-source Llama-2-7B for use in EDA education. We focus on this model in particular because it can be used on consumer hardware. Our contributions are as follows:

- 1) A quantitative and qualitative analysis of Llama-2-7B adapted in various ways for EDA usage. This investigation allows us to understand the impact of fine-tuning, distillation, and retrieval augmentation on the model's performance in the context of EDA knowledge.
- 2) We introduce and evaluate a novel fine-tuning method, Low-Rank Knowledge Distillation (LoRA-KD).
- 3) The release of a benchmark, *RAQ*, designed for evaluating LLMs on EDA knowledge, aimed at facilitating future research and development in the field.

Histogram of Survey Rankings For Various Model/Adaptation Combinations

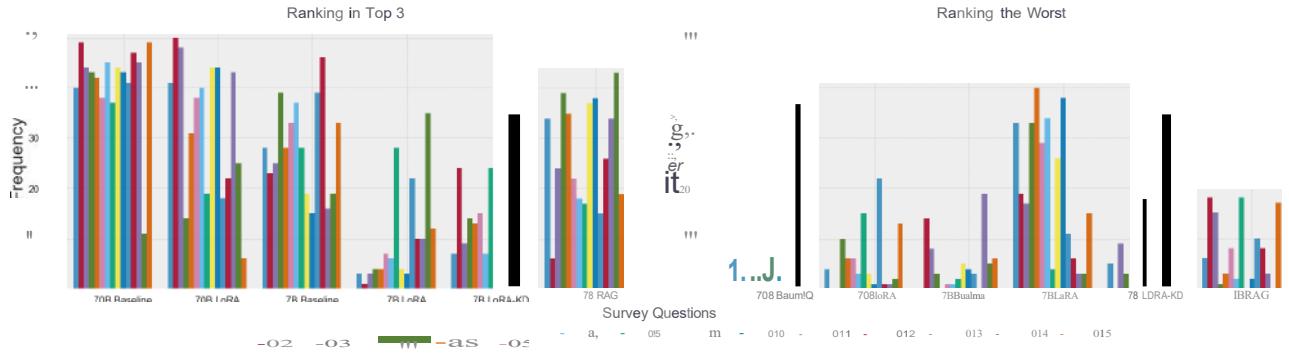


Fig. 2. These charts show histograms of which configurations were ranked in the top half and declared the worst according to third-year microelectronics students. Survey participants had to order the outputs of each configuration on IS questions. A total of 51 rankings were considered after filtering for quality.

II. PRIOR WORK ON LLMS FOR EDA

In the burgeoning field of EDA, early explorations into the applications of LLMs have already returned promising results, particularly in the nuanced areas of chip design, debugging, and script generation. Recently developed LLM-powered ChatEDA is capable of streamlining the IC design flow from RTL to GDSII [3]. ChatEDA integrates Automage, a fine-tuned LLM based on Llama-2-70B architecture, with an EDA tool. Automage serves as an interface that accepts human requests and manipulates the EDA tool through API for task completion. ChatEDA was tested on performance evaluation, parameter grid search, parameter tuning, customized optimization and clock period minimization.

On the other hand, Nvidia took a slightly different approach with their ChatNeMo, a Llama-2 based LLM for chip design which contributes greatly to improving productivity as an engineering chatbot assistant. It is also capable of generating EDA scripts and, bug summarization and analysis [13]. ChatNeMo outperforms GPT-4 at engineering assistant chatbot and EDA script generation tasks while showing comparable performance at bug summarization and analysis whereas ChatEDA has shown comparable or better performance than GPT-4 in all its evaluated cases.

Another work [2] explores the possibility of LLM applications in conversational hardware design by having a hardware engineer co-architect a microprocessor architecture with GPT-4 and this design was sent to tapeout. In addition to these, the possibility of LLM applications in generating VLSI design specifications has also been explored. SpecLLM has shown significant proficiency in assisting engineers in generating and reviewing architecture specifications [11].

Hardware security assessment is one more field which has studied the feasibility of language models. The authors of [8] present an automated flow to identify suitable modules in large HDL databases for hardware trojan insertion using a general-purpose LLM. The model's ability to pinpoint candidate modules for the attack can be indicative of its significant comprehension of RTL codes and system design.

III. ADAPTATION TECHNIQUES FOR LLMS

A. Low-Rank Adaptation

Low-rank adaptation (LoRA) addresses many issues associated with adapting LLMs for domain-specific usage [6]. This method bypasses the expensive backpropagation of gradients across all parameters by keeping the backbone model frozen. This is done by assuming that the update to the model's weights have low-rank. In other words, instead of updating the backbone, we learn parameters A and B which learn the required changes to the output of the backbone. More explicitly, if we write the parameter update equation, LoRA makes the following approximation:

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{v}^T \mathbf{C}(\mathbf{w}_t) \quad (1)$$

$$\mathbf{w}_{t+1} = \mathbf{w}_t - \eta \mathbf{v}^T \mathbf{A} \mathbf{B} \mathbf{w}_t \quad (2)$$

$$\text{i.e. } \eta \mathbf{v}^T \mathbf{C}(\mathbf{w}_t) = \eta \mathbf{v}^T \mathbf{A} \mathbf{B} \mathbf{w}_t \quad (3)$$

Where $\mathbf{v} \in \mathbb{R}^{d \times k}$, $\mathbf{A} \in \mathbb{R}^{d \times r}$, and $\mathbf{B} \in \mathbb{R}^{r \times k}$. Note that $d \times k$ represents the size of the backbone model's (very large) parameter shapes. By selecting $r \ll \min(d, k)$, LoRA provides a resource-efficient update to the backbone.

B. Knowledge Distillation

Knowledge Distillation (KD) [5] is a knowledge-transfer technique where a larger (teacher) network produces soft targets for a smaller (student) model. This can play a pivotal role in reducing the performance gap between larger and smaller models. Fine-tuning a smaller (student) model through KD can show improved performance compared to a normally fine-tuned small model. As an example, the authors of DistilBERT show that they can retain 97% of the original BERT's performance despite a significantly smaller parameter count [15]. This indicates that a smaller model that can be deployed on weaker hardware, e.g. personal computer, can maintain feasibility in handling complex tasks related to EDA. Written explicitly, the loss used for KD is:

$$\begin{aligned} \text{KD Loss} = & (1 - \alpha) \cdot \text{CrossEntropy}(\text{student}(x), y) \\ & + \alpha \cdot \text{MSE}(\text{student}(x), \text{teacher}(x)) \end{aligned} \quad (4)$$

TABLE I
CONFIGURATIONS' PERFORMANCE ON REASONING AND ACCURACY QUESTIONS

RAQ: Reasoning	Ground Truth	70B Baseline	70B LoRA	7B Baseline	7B LoRA	7B LoRA-KD	7B RAG
1a	Increase	Increase	Increase	Increase	Increase	Increase	Decrease
1b	$3 \mu\text{m}$	330nm					
2a	0.775 mA	5.58 mA	5.58 mA	1.28 A	1.395 A	0.06 A	x
2b	1.55 V	11.16 V	8.6 V	0.83 V	0.647 V	0.7 V	x
3a	2V	2V	2V	SV	0.625 V	SV	6V
3b	3 kn	8 k!1	4 kn	to k!1	13 kn	4 k!1	x
4	2	2	2	2	2	2	2
Sa	0.66 k!1	666.67 !1	2/3 kn	0.5 k!1	3 kn	0.5 kn	x
Sb	0.667 kn	666.67 !1	0.5 k!1	0.25 kn	I kn	1.6 kn	x
RAQ: T/F Accuracy		84%	84%	72%	76%	76%	80%

Evaluated on the reasoning and T/F questions from the RAQ benchmark. Note that some of the reasoning questions required multi-step thinking, eg. *tased on your answer to part a, what is part b?* The x symbol denotes that the model refused to answer the question due to fallacious ethical reasons.

Where $.Cy$ is the typical loss incurred between the student predictions and the target. $.CDist$ is the loss between what the student predicted and the teacher predicted on an input. More elaborate distillation techniques exist, for example, *patient* KD aims at having the student mimic the teacher's intermediate layers in addition to the teacher's outputs [16]. We refer readers to [21] for further exploration.

C. Low-Rank Knowledge Distillation (LoRA-KD)

Although not entirely unprecedented, the combination of low-rank approximations and knowledge distillation is far less explored in the context of **LLM** fine-tuning. The authors of LoSparse [12] introduce a new compression scheme for transformers [18] based on a truncated singular value decomposition. In their experiments, they find that combining this parameter compression scheme with knowledge distillation further improves performance.

In our work, we reformulate this concept in accordance with figure I. We begin by fine-tuning the teacher (Llama-2-70B) using LoRA. Afterwards, we fine-tune the student (Llama-2-7B) via LoRA using $.CKD$. We hypothesize that, if the updates to the teacher can be done in a low-rank fashion, then the underlying *knowledge* being learned is also low-rank; therefore, the knowledge to be distilled to the student is also low-rank.

There are several advantages to doing this:

- As with ordinary RAG, a pre-trained model can be repurposed via hot-swapping the adaptation layer. For example, EDA educators can use LoRA-KD to learn separate (small) adaptation layers for English and Spanish in the context of a bilingual classroom.
- KO has been used to enhance domain-adaptation tasks. We hypothesize that the *dark knowledge* distilled from the teacher to the student will facilitate enhanced reasoning capabilities [20].
- The training process remains fast. In our experiments, fine-tuning the student via LoRA-KD did not take much more time than ordinary LoRA.

D. Retrieval Augmented Generation (RAG)

RAG [10] operates by integrating a neural retriever with a sequence-to-sequence (seq2seq) generator. The retriever pro-

duces a distribution, $Pr(z|x)$ from a dense vector index the fine-tuning dataset based on the input query. These documents then serve as additional context for the seq2seq generator, enabling it to produce outputs that are informed by the retrieved information, z , in addition to the user's input, x . RAG's seq2seq probability distribution is defined as:

$$p(y_t|x) \approx \sum_{z \in \text{top-}k(p_r(\cdot|x))} p_r(z|x) p_{\text{Llama}}(y_t|x, z, y_{[0:t-1]}) \quad (5)$$

This method combines the strengths of pre-trained parametric models with non-parametric external knowledge sources. For our work, we use the pre-trained **MinILM** model [19] as the retriever and the pre-trained Llama-2-7B as the generator.

IV. FINE-TUNING DATASET AND THE RAQ BENCHMARK

Our fine-tuning dataset consists of several well-known textbooks on microelectronics, VLSI circuit design, and fabrication technologies. In addition to these, we also included some recently published works related to DDR5 design and its corresponding JEDEC standard. After filtering the data, the number of tokens was calculated using Llama-2 tokenizer. The dataset contains 3,168,414 tokens and 12,988 unique tokens. Due to copyright reasons, we cannot release the fine-tuning dataset. However, we list all the components of the dataset within the appendix so that readers can assemble it themselves.

We created a benchmark to evaluate the performance of the different models which includes 70 carefully-curated domain-specific questions. Among them, there are 40 qualitative questions and 25 true/false questions. The 65 aforementioned questions are meant to evaluate the *accuracy* and *quality* of the LLM's responses on domain knowledge. Furthermore, 5 questions are designed to evaluate the models' capabilities to *reason* upon circuit design decisions based on given specifications. Hence, we name it the *Reasoning-Accuracy-Quality* (RAQ) benchmark.

¹Readers should note that our equation for $p(y|x)$ differs from [10]. We omit the multiplication across all document sources because we concatenate all of the fine-tuning texts into a single file. Therefore, in this scenario RAG sequence and RAG token are the same.

TABLE II
COMPARISON OF MODEL/ADAPTATION COMBINATIONS EVALUATED BY HUMAN EXPERT AND GPT-4.5 TURBO VIA LIKERT SCALE.

Configuration	Human Expert		GPT-4.5 Turbo		Pearson Correlation	
	Accuracy	Quality	Accuracy	Quality	Accuracy	Quality
70B Baseline	4.2±1.96	3.975±1.96	4.625±2.00	4.2±1.95	0.51	0.43
70B LoRA	4.35±1.84	4.275±1.84	5.7±1.65	5.475±1.67	0.47	0.51
7B Baseline	3.4±1.92	3.275±1.87	4.15±1.90	3.9±1.93	0.53	0.57
7B LoRA	3.5±1.82	3.3±1.73	4.2±1.93	3.95±1.90	0.66	0.73
7B LoRA-KD	3.525±1.99	3.3±1.83	4.475±1.70	4.175±1.60	0.60	0.59
7B RAG	4.2±1.87	3.7±1.76	4.55±1.96	4.15±1.94	0.12	0.13

Each response to the 40 qualitative questions was evaluated on a 7-point Likert scale by a human expert and GPT-4.5 Turbo. 7 denotes "strongly agreed with" and 1 denotes "strongly disagreed with." The subcolumns correspond to how much the evaluator agreed/disagreed with the accuracy/quality of the response. The standard deviations across the questions are written in sub-scripts. The correlation quantifies the consistency between the human expert and GPT-4.5 Turbo.

V. SETUP AND EXPERIMENTS

To assess the suitability of various adaptation methods, we performed four experiments using the RAQ Benchmark:

- 1) **Student Survey.** We selected 15 questions which would be most relevant for a third-year undergraduate microelectronics classroom. We recorded the responses from each configuration and asked students to provide the ordinal rankings in terms of what they preferred. To ensure quality, we kept the configurations anonymous and asked students to explain why they ranked the best/worst models as they did. After pruning low-quality submissions, we had 51 rankings.
- 2) **True/False Q&A.** We prompted each configuration to answer true or false to determine accuracy. This portion was taken from the T/F section.
- 3) **Likert Test.** Each configuration was asked to answer all 40 qualitative questions. Using a 7-point Likert scale, the responses were scrutinized in terms of *accuracy* and subjective *quality*. We² (human expert) and ChatGPT-4.5 Turbo were the evaluators.
- 4) **Reasoning Test.** We tested each configuration with 5 reasoning questions. These questions have unambiguous or numerical answers. Generated responses were compared against ground truth values.

For all experiments, we use $\text{LoRA_Rank} = 4$, the Adam($\beta_1 = 10^{-4}$) optimizer [7], a sequence length of 128 and a batch size of 16. All the models underwent fine-tuning with LoRA for a total of 20 epochs. Regarding the selection of checkpoints for the models: the 16th epoch checkpoint was chosen for the 7B LoRA model, the 17th epoch checkpoint was utilized for the 70B LoRA (teacher) model, and the 14th epoch checkpoint was selected for the 7B LoRA-KD (student) model. These checkpoints were all selected via early stopping. We set $\alpha = 0.80$ and temperature = 2.0 for KD.

VI. RESULTS AND CONCLUSION

In this work, we try to explore the feasibility of using language models in EDA education. Table I investigates the

configurations' capabilities to reason based on the given information and optimize a given design. A few interesting observations were made while evaluating the models on reasoning/optimization questions.

- 1) All the models had difficulty with numerical calculations and assigning proper units to a calculated value.
- 2) In our experiments, the models performed better when they were asked the different sections of the questions one by one in separate prompts, rather than putting all the questions in a single prompt
- 3) RAG tended to refuse answering due to dubious ethical reasons regarding the "danger of transistors."

An analysis of the data presented in tables I and II gives insights into the strengths and weaknesses of the configurations. For the Likert test and true/false accuracy, 7B RAG performs strongly but for reasoning/optimization, it exhibits a sharp decline in performance. This indicates that RAG alone cannot improve performance across all areas. On the contrary, the various fine-tuned versions manage to perform well on the reasoning portion while remaining within half a standard deviation from RAG on the Likert test.

The responses collected from the students underscore each configuration's communication skills and human expectations which can serve as an important guideline when fine-tuning an LLM. An improvement of LoRA-KD over LoRA can be observed in figure 2 where the responses generated by 7B LoRA-KD were far less likely to be ranked last. Another interesting facet is the agreement between the students with respect to the question (i.e. the entropy). For example, for Q14, there was high agreement that the 70B Baseline did the worst. On the other hand, for Q15, the students seemed split between whether 7B RAG, 7B LoRA, or 70B LoRA was the worst.

While the existing works are significant milestones of the application of LLMs in EDA, its potential in this field has yet to be fully realized. Development of specialized large language models capable of understanding the intricacies of domain-specific EDA tasks is crucial for its continued applications in EDA [4]. This study highlights some strengths and weaknesses of different open-source offline LLM configurations.

²We recognize there could be bias if we, the authors, evaluate these models. To promote transparency, we release the model responses on our GitHub

REFERENCES

- [1] Josh Achiam, Steven Adler, Sandhini Agarwal, Lama Ahmad, Ilge Akkaya, Florencia Leoni Aleman, Diogo Almeida, Janko Altenschmidt, Sam Altman, Shyamal Anadkat, et al. Gpt-4 technical report. *arXiv preprint arXiv:2303.08774*, 2023.
- [2] Jason Blocklove, Siddharth Garg, Ramesh Karri, and Hammond Pearce. Chip-chat Challenges and opportunities in conversational hardware design. In *2023 ACM/IEEE 5th Workshop on Machine learning for CAD (MLCAD)*, pages 1-6. IEEE, 2023.
- [3] Zhuolun He, Haoyuan Wu, Xinyun Zhang, Xufeng Yao, Su Zheng, Haisheng Zheng, and Bei Yu. Chateda: A large language model powered autonomous agent for eda, 2024.
- [4] Zhuolun He and Bei Yu. Large language models for eda: Future or mirage? In *Proceedings of the 2024 International Symposium on Physical Design*, ISPD '24, page 65-66, New York, NY, USA, 2024. Association for Computing Machinery.
- [5] Geoffrey Hinton, Oriol Vinyals, and Jeff Dean. Distilling the knowledge in a neural network, 2015.
- [6] Edward J Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.
- [7] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [8] Georgios Kokolakis, Athanasios Moschos, and Angelos D Keromytis. Harnessing the power of general-purpose llms in hardware trojan design.
- [9] Ehsan Latif, Luyang Fang, Ping Ma, and Xiaoming Zhai. Knowledge distillation of um for automatic scoring of science education assessments, 2024.
- [10] Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Kilttler, Mike Lewis, Wen-tau Yih, Tim Rocklischel, et al. Retrieval-augmented generation for knowledge-intensive nlp tasks. *Advances in Neural Information Processing Systems*, 33:9459-9474, 2020.
- [11] Mengming Li, Wenji Fang, Qijun Zhang, and Zhiyao Xie. Specllm: Exploring generation and review of vlsi design specification with large language model, 2024.
- [12] Yixiao Li, Yifan Yu, Qingru Zhang, Chen Liang, Pengcheng He, Weizhu Chen, and Tuo Zhao. Losparse: Structured compression of large language models based on low-rank and sparse approximation. In *International Conference on Machine Learning*, pages 20336-20350. PMLR, 2023.
- [13] Mingjie Liu, Teodor-Dumitru Ene, Robert Kirby, ChrisCheng, Nathaniel Pinckney, Rongjian Liang, Jonah Alben, Himyanshu Anand, Sanmitra Banerjee, Ismet Bayraktaroglu, Bonita Bhaskaran, Bryan Catanzaro, Arjun Chaudhuri, Sharon Clay, Bill Dally, Laura Dang, Parikshit Deshpande, Siddhanth Dhadhi, Sameer Halepete, Eric Hill, Jiashang Hu, Sumit Jain, Ankit Jindal, Brucek Khailany, George Kokai, Kishor Kunal, Xiaowei Li, Charley Lind, Hao Liu, Stuart Oberman, Sujeev Omar, Sreedhar Pratty, Jonathan Raiman, Ambar Sarkar, Zhengjiang Shao, Hanfei Sun, Pratik P Suthar, Varun Tej, Walker Turner, Kaizhe Xu, and Haoxing Ren. Chipnemo: Domain-adapted llms for chip design, 2024.
- [14] Mingjie Liu, Nathaniel Pinckney, Brucek Khailany, and Haoxing Ren. Invited paper: Verilogeval: Evaluating large language models for verilog code generation. In *2023 IEEE/ACM International Conference on Computer Aided Design (ICCAD)*, pages 1-8, 2023.
- [15] Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distmed version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.
- [16] Siqi Sun, Yu Cheng, Zhe Gan, and Jingjing Liu. Patient knowledge distillation for bert model compression. *arXiv preprint arXiv:1908.09355*, 2019.
- [17] Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023.
- [18] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.
- [19] Wenhui Wang, Furu Wei, Li Dong, Hangbo Bao, Nan Yang, and Ming Zhou. Minilm: Deep self-attention distillation for task-agnostic compression of pre-trained transformers, 2020.
- [20] Yufei Wang, Haoliang Li, Lap-pui Chau, and Alex C. Kot. Embracing the dark knowledge: Domain generalization using regularized knowledge distillation. In *Proceedings of the 29th ACM International Conference on Multimedia*, MM '21, page 2595-2604, New York, NY, USA, 2021. Association for Computing Machinery.
- [21] Xiaohan Xu, Ming Li, Chongyang Tao, Tao Shen, Reynold Cheng, Jinyang Li, Can Xu, Dacheng Tao, and Tianyi Zhou. A survey on knowledge distillation of large language models, 2024.

VII. APPENDIX

A. Fine-tuning Sources

The following sources were used in fine-tuning. An enumerated list is also available on our github:

- 1) Fundamentals of Microelectronics - 2nd Edition - Behzad Razavi
- 2) Electronic Devices and Circuit Theory - 11th Edition - Robert L. Boylestad and Louis Nashelsky
- 3) CMOS VLSI Design - 4th Edition - Neil H. E. Weste and David M. Harris
- 4) Fundamentals of Semiconductor Manufacturing and Process Control - Gary S. May and Costas J. Spanos
- 5) Fabrication Engineering at the Micro and Nanoscale - 3rd Edition - Stephen A. Campbell
- 6) JEDEC Standard - Graphjcs Double Data Rate (GDDR5) SGRAM Standard
- 7) JEDEC Standard - Compression Attached Memory Module (CAMM2) Common Standard
- 8) JEDEC Standard - DDR5 Clocked Small Outline Dual Inline Memory Module (CSODIMM) Common Standard
- 9) DDR5 Clocked Unbuffered Dual Inline Memory Module (CUDIMM) Common Specification
- 10) JEDEC Standard - DDR5 262 Pin SODIMM Connector Performance Standard
- 11) JEDEC Standard - DDR5 Unbuffered Dual Inline Memory Module (UDIMM) Common Standard
- 12) JEDEC Standard - DDR5 288 Pin U/R/LR DIMM Connector Performance Standard
- 13) JEDEC Standard - DDR5 Load Reduced (LRDIMM) and Registered Dual Inline Memory Module (RDIMM) Common Specification
- 14) JEDEC Standard - DDR5 Clock Driver Definition (DDR5CKD0I)
- 15) JEDEC Standard - DDR5 Small Outline Dual InLine Memory Module (SODIMM) Common Standard
- 16) JEDEC Standard - DDR5 Registering Clock Driver Definition (DDR5RCD03)
- 17) JEDEC Standard - DDR5 DIMM Labels
- 18) JEDEC Standard - GDDR5 Measurement Procedures
- 19) JEDEC Standard - DDR5 Serial Presence Detect (SPD) Contents
- 20) JEDEC Standard - Graphics Double Data Rate (GDDR5X) SGRAM Standard
- 21) JEDEC Standard - DDR5 SDRAM
- 22) Improving Memory Reliability by Bounding DRAM Faults - KJERSTEN CRISS, KULJIT BAINS, RAJAT AGARWAL, TANJ BENNETT, TERRY GRUNZKE, JANGRYUL KEITH KIM, HOEJU CHUNG, MUNSEON JANG
- 23) Optimizing DDR5 address signal integrity using stochastic learning algorithms - Nitin Bhagwath, Daniel DeAraujo, Jayaprakash Balachandran, BaekKyu Choi
- 24) DDR5 Electrical Challenges in High-Speed Server Design - Douglas Winterberg, Vijender Kumar, Tom Chen, Bhyrav Mutnury
- 25) Modeling of DDR5 Signaling from Jitter Sequences to Accurate Bit Error Rate (BER) - Alaeddin A. Aydiner, Yunhui Chu, Oleg Mikulchenko, Jin Yan, Robert J. Friar, Ellen Yan Fu
- 26) LPDDR5 (6.4 Gbps) 1-tap DFE Optimal Weight Determination - Sunil Gupta, Ph.D.
- 27) Far-End Crosstalk Mitigation for Transmission Lines in DDR5 Using Glass-Weave Coating Structure - Xiao-Bo Yu, Qiang-Ming Cai, Liang Zhang, Chao Zhang, Lin Zhu, Xin Cao, and Jun Fan
- 28) Simulating DDR5 Systems with Clocked Receivers - Matthew Leslie, Justin Butterfield, Randy Wolff
- 29) Design and Analysis of Power Integrity of DDR5 Dual In-Line Memory Modules - Shinyoung Park, Vinod Arjun Huddar
- 30) Deterministic Policy Gradient-based Reinforcement Learning for DDR5 Memory Signaling Architecture Optimization considering Signal Integrity - Daehwan Lho, Hyunwook Park, Keunwoo Kim, Seongguk Kim, Boogyo Sim, Kyungjune Son, Keeyoung Son, Jihun Kim, Seonguk Choi, Joonsang Park, Haeyeon Kim, Kyubong Kong, Joungho Kim
- 31) Advancing DDR5 Test and Measurements: Fine-tuning a Large Language Model AI Expert in DDR5 Protocols - Xinran Li
- 32) DDR5 Design Challenges - Nitin Bhagwath, Randy Wolff, Shinichiro Ikeda, Eiji Fujine, Ryo Shibata, Yumiko Sugaya, Megumi Ono
- 33) Advanced Measurement and Simulation Approach for DDR5 On-chip SI/PI with the Probing Package - WonSuk Choi, SangKeun Kwak, Jaeseok Park, Jiyoung Do, Byeongseon Yun, Yoo-jeong Kwon, Dongyeop Kim, Kyudong Lee, Tae young Kim, Wonyoung Kim, Kyungsun Kim, Sung Joo Park, Jeonghyeon Cho and Hoyoung Song