MCUBERT: Memory-Efficient BERT Inference on Commodity Microcontrollers

Anonymous ACL submission

Abstract

In this paper, we propose MCUBERT to enable language models like BERT on tiny microcontroller units (MCUs) through network and scheduling co-optimization. We observe the embedding table contributes to the major storage bottleneck for tiny BERT models. Hence, at the network level, we propose an MCU-aware two-stage neural architecture search algorithm based on clustered low-rank approximation for embedding compression. To reduce the inference memory requirements, we further propose a novel fine-grained MCU-friendly scheduling strategy. Through careful computation tiling and re-ordering as well as kernel design, we drastically increase the input sequence lengths supported on MCUs without any latency or accuracy penalty. MCUBERT reduces the parameter size of BERT-tiny and BERT-mini by 5.7× and 3.0× and the execution memory by 3.5× and 4.3×, respectively. MCUBERT also achieves 1.5× latency reduction. For the first time, MCUBERT enables lightweight BERT models on commodity MCUs and processing more than 512 tokens with less than 256KB of memory.

1 Introduction

IoT devices based on microcontroller units (MCUs) are ubiquitous, enabling a wide range of speech and language applications on the edge, including voice assistant (Xu et al., 2019; Chung et al., 2021), real-time translation (Zhu et al., 2020; Yang et al., 2020), smart home (Shin et al., 2019), etc. Language models (LMs), e.g., BERT (Devlin et al., 2018), are fundamental to these applications. While cloud offloading is heavily employed for LM processing, it suffers from high latency overhead, privacy concerns, and a heavy reliance on WiFi or cellular networks (Banbury et al., 2021; Fedorov et al., 2022). Hence, there is a growing demand for deploying BERT-like LMs on MCUs.

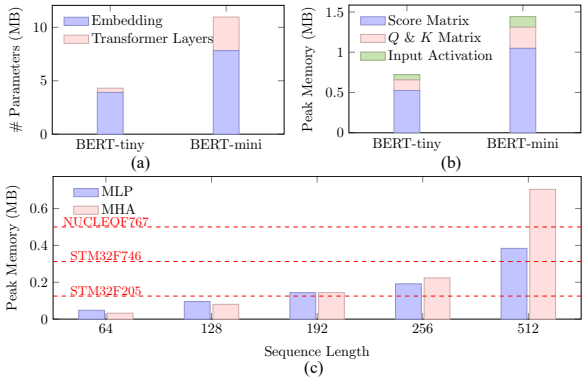

Figure 1: Enabling BERT on MCUs faces memory challenges: (a) the Flash storage limits the model size; (b) the SRAM memory limits the peak execution memory; (c) for long sequence lengths, memory requirements of both MHA and multi-layer perceptron (MLP) become bottleneck.

Though appealing, enabling BERT on MCUs is very challenging. On one hand, MCUs only have a very simple memory hierarchy with highly constrained memory budgets (Banbury et al., 2021; Burrello et al., 2021). For example, a state-of-the-art (SOTA) ARM Cortex-M7 MCU only has 320 KB static random-access memory (SRAM) to store intermediate data, e.g., activation, and 1 MB Flash to store program and weights, directly limiting the peak execution memory and the total parameter size of LMs (Lin et al., 2020; Banbury et al., 2021; Burrello et al., 2021; Lin et al., 2021). As shown in Figure 1, even with 8-bit quantization, the BERT-tiny model (Bhargava et al., 2021; Turc et al., 2019) exceeds the SRAM and Flash capacity by more than 2.3× and 4.3×, respectively.

On the other hand, the computation graph of a Transformer block in BERT is more complex when compared with convolutional neural networks (CNNs): each multi-head attention (MHA) block not only comprises more MCU-unfriendly tensor layout transformation operators, e.g., reshape, transpose, etc, but also organizes them with linear

operators in a more complex topology. Naively executing these operators can result in significant memory consumption and poor performance (Ivanov et al., 2021; Burrello et al., 2021).

Existing works, however, cannot resolve these challenges. Most SOTA BERT designs (Jiao et al., 2019; Lan et al., 2019; Sanh et al., 2019) and network optimization designs (Xu et al., 2021; Zhang et al., 2022) focus on large BERT models on GPUs or smartphones. As the memory and storage of these platforms are much more abundant than MCUs, these methods usually focus on optimizing the Transformer layers for better latency. Another class of works (Fedorov et al., 2019; Lin et al., 2020, 2021; Banbury et al., 2021) develop MCU-friendly CNNs by jointly optimizing the network architecture and scheduling, which, however, cannot be directly applied to BERT.

In this paper, we propose MCUBERT, a network and scheduling co-optimization framework to enable BERT on commodity MCUs for the first time. We observe for small BERT models, e.g., BERT-tiny, the embedding table accounts for the major storage bottleneck and thus, propose to leverage the MCU-friendly clustered low-rank approximation for embedding compression. We further propose an MCU-aware two-stage differentiable neural architecture search (NAS) algorithm to improve the accuracy of compressed models. To reduce the execution memory usage and latency, we observe the execution of a Transformer block can be carefully tiled without accuracy penalty and develop an MCU-friendly fine-grained scheduling algorithm. Our contributions can be summarized as follows:

- We propose an MCU-aware two-stage NAS algorithm based on clustered low-rank approximation for embedding compression to overcome the storage bottleneck.
- To reduce the execution peak memory and latency, we propose MCU-friendly scheduling optimization for the Transformer block.
- MCUBERT can reduce the model size by 5.7× and 3.0× as well as the peak memory by 3.5× and 4.3× for BERT-tiny and BERT-mini, respectively, enabling to process more than 512 tokens simultaneously with less than 256KB SRAM. MCUBERT also achieves 1.5× and 1.3× latency reduction compared to SOTA inference engines, CMSIS-NN and (Burrello et al., 2021), respectively.

Table 1: Comparison with prior-art methods (Opt represents Optimization).

Method	Platform	Model	Network Opt	Scheduling Opt		
				MLP	MHA	Kernel
(Banbury et al., 2021)	MCU	CNN	Conv	-	-	✗
(Lin et al., 2020, 2021)	MCU	CNN	Conv	-	-	✓
(Liang et al., 2023)	MCU	ViT	MLP	✓	✗	✓
(Sanh et al., 2019)	GPU	BERT	Linear	✗	✗	✗
(Dao et al., 2022)	GPU	BERT	-	✗	✓	✓
Ours	MCU	BERT	Embedding	✓	✓	✓

2 Related Works

While no existing works deploy BERTs on MCUs, there are BERT compression algorithms proposed for other platforms, e.g., GPUs, including pruning (Shi et al., 2021; Cui et al., 2019; Zafrir et al., 2021; Campos et al., 2022), low-precision quantization (Bai et al., 2020; Zhang et al., 2020; Shen et al., 2020; Qin et al., 2022), low-rank factorization (Baevski and Auli, 2018; Chen et al., 2018a; Hsu et al., 2022; Lan et al., 2019; Lioutas et al., 2019), etc. Adaptive embedding (Baevski and Auli, 2018) compresses the embedding table with clustered low-rank approximation: it divides the tokens in an embedding table into clusters first and applies low-rank approximation with different ratios to each cluster based on the cluster importance. Adaptive embedding leverages token frequency as the proxy metric for its importance, which leads to high accuracy degradation. Scheduling optimizations are also proposed for LMs on GPUs (Rasley et al., 2020; Fang et al., 2021; Dao et al., 2022; Zhai et al., 2023). Most existing works for deep learning on MCUs focus on CNNs (Rusci et al., 2020; Banbury et al., 2021; Lin et al., 2020). Recent works, MCUFormer (Liang et al., 2023) and (Burrello et al., 2021), also optimize Vision Transformers (ViT) for MCUs, the memory bottleneck of which is the MLP blocks. We provide a qualitative comparison with these prior-art methods in Table 1, and a more detailed review in Appendix F. As shown in Table 1, our proposed MCUBERT first deploys transformer model BERT on MCUs, which has more parameters. MCUBERT compresses the embedding table, which account for most of the parameters and can't be stored in MCU. Besides MLP block, We also carefully optimize the MHA block, which comprises more MCU-unfriendly operators and can also be the memory bottleneck of BERT inference. Compared with GPU-based methods, MCUBERT conduct network and scheduling optimization in a more MCU-friendly mode.

3 MCUBERT: MCU-friendly Network/Scheduling Co-Optimization

3.1 Motivations and Overview

We now discuss our observations on key challenges that prevent running BERTs directly on MCUs and introduce the overall flow of MCUBERT. The notations used in the section is summarized in Table 2.

Table 2: Notations used in the paper.

Notations	Meanings
v	Vocabulary size
s, h, d	Sequence length, # heads, and embedding dimension
c	# clusters
t	Sequence length each tile
U, V	Unitary matrices generated by SVD
Σ	Vector of singular values
α, β	NAS parameters for embedding compression
i	Loop variables for clusters, $0 \leq i \leq c-1$
j	Loop variables for tokens in vocabulary, $0 \leq i \leq v-1$
l	Loop variables for singular values, $0 \leq i \leq d-1$

Observation 1: the tight MCU Flash storage limits model size and forces to use small BERT models, for which the embedding table becomes the major bottleneck. To satisfy the Flash storage constraints of MCUs, which is often less than 2 MB (Lin et al., 2020; Banbury et al., 2021), we are forced to consider only small BERT models, e.g., BERT-tiny and BERT-mini (Devlin et al., 2018), which have lower embedding dimensions and fewer Transformer layers. However, there still exists a 4.3× gap between the model size and the MCU Flash, even for BERT-tiny with 8-bit quantization. As shown in Figure 1(a), the embedding table contributes to more than 90% of the parameters of BERT-tiny and becomes the major bottleneck. Hence, embedding compression is required to enable BERT on MCUs.

Observation 2: the MCU SRAM size limits the execution peak memory of both MHA and MLP, especially for long sequence lengths. During BERT inference, all the activations need to be stored in the MCU SRAM. Although Flash paging and re-materialization has been proposed (Patil et al., 2022) to reduce the SRAM requirements, the introduced Flash access and re-computation incur high power and latency overhead. As shown in Figure 1(c), with the increase of the sequence length, the peak memory of both MHA and MLP increases significantly and quickly exceed the MCU SRAM limit. Depending on the sequence length, both MLP and MHA can be the memory bottleneck. Meanwhile, as shown in Figure 1(b), for a long input sequence, the score matrix in MHA incurs highest memory consumption. Therefore, both MLP and MHA, especially the score matrix, needs to be optimized to enable processing long sequences.

Observation 3: the naive design of computation kernels introduces non-negligible latency overhead. Both MHA and MLP involve many linear/batched matrix multiplication operators with large shapes. MHA further complicates the inference with extensive tensor shape transformation operators, e.g., reshape, transpose, etc. Naive kernel implementation, e.g., CMSIS-NN, cannot well utilize the hardware resource and suffers from over 30% latency overhead. To improve the resource utilization, it is important to leverage the memory locality and the instruction-level parallelism (ILP) of MCU. Memory access patterns can also be carefully designed to fuse the tensor shape transformation operators with the linear operators. All of these require dedicated kernel optimization.

MCUBERT overview Based on these observations, we propose MCUBERT, an MCU-friendly network/scheduling co-optimization framework to enable BERT models on tiny devices. The overview of MCUBERT is shown in Figure 2. MCUBERT leverages MCU-friendly clustered low-rank approximation for embedding compression and features a MCU-aware two-stage NAS to explore the trade-off between accuracy and model sizes, enabling to reduce the model size and fit BERT models into the MCU Flash storage. To enable processing long sequences, MCUBERT proposes MCU-friendly scheduling optimization. By carefully tiling and re-ordering the computation as well as designing efficient kernels, inference peak memory and latency can be reduced without accuracy penalty. Such optimization enables to support more than 512 tokens on a small MCU with less than 256 KB SRAM.

Note while we focus on optimizing BERT-like encoder LMs for MCUs, our proposed techniques, including embedding compression and scheduling optimization, can benefit GPT-like decoder LMs (Radford et al., 2019) on MCUs as well. Decoder LMs face more challenges on the storage and dynamic shape induced by the KV cache, which we leave for future research.

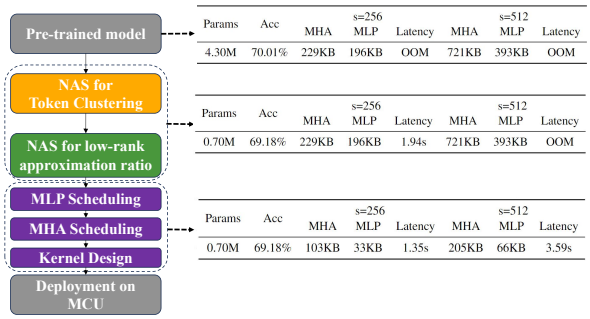

Figure 2: MCUBERT overview. (Params stands for parameters, Acc stands for MNLI accuracy, and OOM stands for out of memory.)

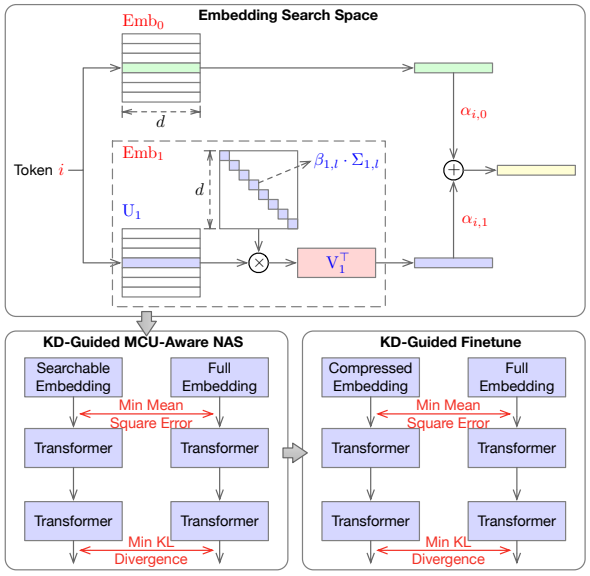

Figure 3: Our proposed NAS search formulation for embedding compression.

3.2 MCU-aware NAS for Embedding Compression

MCU-aware NAS Formulation To reduce the parameter size and satisfy the tight Flash storage constraint, we propose embedding compression based on clustered low-rank approximation following (Baevski and Auli, 2018), which is more MCU-friendly as MCUs do not natively support sub-8-bit quantization or sparse computation. Following (Baevski and Auli, 2018), we set the number of clusters to 4 and show its impact in ablation study. Two important questions remain to be answered: 1) how to cluster the tokens and 2) how to determine the approximation ratio for each cluster. We propose differentiable NAS to guide the compression directly with task loss.

Let c denote the number of clusters and Emb_i denote the embedding table for i-th cluster, where $0 \leq i \leq c-1$. To enable NAS for token clustering, as shown in Figure 3, we introduce trainable architecture parameters α ($0 \leq \alpha \leq 1$) for each token in the vocabulary. Let v denote the vocabulary size. Then, for the j-th token ($0 \leq j \leq v-1$), we define $\{\alpha_{j,0}, \ldots, \alpha_{j,c-1}\}$, where $\sum_{i=0}^{c-1} \alpha_{j,i} = 1$. Then, the embedding for the j-th token becomes:

$$\alpha_{j,0}\text{Emb}_0[j] + \sum_{i=1}^{c-1} \alpha_{j,i}\text{Emb}_i[j],$$

where $\text{Emb}_i[j]$ denotes the embedding vector of j-th token from the i-th embedding table.

To leverage NAS for low-rank approximation, instead of directly searching the reduced dimension for each embedding table, we apply singular value decomposition (SVD) and search the singular values that can be pruned, which enables to better inherit the pretrained weights and is crucial for BERT compression. Specifically, we first decompose Emb_i with SVD as $U_i\text{Diag}(\Sigma_i)V_i^\top$, where U_i and V_i are unitary matrices, Σ_i is the vector of singular values, and $\text{Diag}(\Sigma_i)$ represents the diagonal matrix with singular values in Σ_i filled in the diagonal. We also introduce architecture parameters β ($0 \leq \beta \leq 1$) for each cluster. For the i-th cluster, we define $\{\beta_{i,0}, \ldots, \beta_{i,d-1}\}$, where d is the embedding dimension, to indicate the importance of each singular value. Then, we have

$$\text{Emb}_i = U_i\text{Diag}(\beta_{i,0}\Sigma_{i,0}, \ldots, \beta_{i,d-1}\Sigma_{i,d-1})V_i^\top.$$

Based on α and β, the embedding size for the i-th cluster can be approximated as

$$\text{Size}(\text{Emb}_i) = \sum_l \beta_{i,l}(d + \sum_j \alpha_{j,i}),$$

where $\sum_j \alpha_{j,i}$ approximates the number of tokens for i-th cluster. $\sum_l \beta_{i,l} \sum_j \alpha_{j,i}$ and $\sum_l \beta_{i,l} d$ represent the size of low-rank approximated U_i and $V_i^\top$, respectively. We introduce ℓ_1-penalty of the embedding table size into the objective function to encourage a lower parameter size:

$$\min_{\alpha,\beta} \min_{w} \ell_{w,\alpha,\beta} + \lambda \sum_i \text{Size}(\text{Emb}_i).$$

where w denotes the model parameters and λ is the hyperparameter to balance model size and accuracy. α, β, and w are trained together. Upon convergence, for the j-th token, we determine its cluster by $\text{argmax}_i \alpha_{j,i}$, while for the i-th cluster, we set a threshold β_i^* and only keep the l-th singular values with $\beta_{i,l} > \beta_i^*$.

Improving NAS Convergence In practice, we find that searching token clustering and low-rank approximation ratio, i.e., α and β, simultaneously makes the differentiable NAS training unstable and directly degrades the accuracy of the searched models. As shown in Figure 6, the accuracy of the searched models is even lower compared to the baseline adaptive embedding (Baevski and Auli, 2018). We hypothesize this is because of the large discrepancy of optimal low-rank approximation ratios for different clusters. Because the token clustering keeps changing during NAS, some singular values in a certain cluster may be incorrectly preserved due to tokens that eventually belong to a different cluster. These incorrectly preserved singular values in turn may impact the token clustering, leading to sub-optimal results.

To encourage the NAS convergence, we decompose the search space and propose a two-stage NAS strategy: in the first stage, we fix different low-rank approximation ratios for each cluster, and only search the token clustering, i.e., α. The approximation ratios are chosen following (Baevski and Auli, 2018). Then, in the second stage, given the fixed token clustering, we search the best approximation ratio for each cluster, i.e., β. While we can iterate back to the first stage with updated approximation ratios, we empirically find it is not necessary. We hypothesize this is because the token clustering reflects the token importance, which should remain stable for different low-rank approximation ratios. The algorithm details are shown in Appendix C. We guide the NAS with knowledge distillation (KD) to improve convergence following (Sanh et al., 2019). We simply use the trained full model as the teacher.

Our proposed two-stage NAS incurs low training cost as only a few training epochs (less than 2) are needed for each stage. After the two-stage NAS, we quantize both the weights and activations to 8 bits and fine-tune the compressed model with KD to get the final deployable model.

3.3 MCU-friendly Scheduling Optimization

We now introduce our MCU-friendly scheduling optimization to reduce the inference memory and latency. We optimize the scheduling of both MHA and MLP since and also optimize the kernel implementation for MHA for better efficiency.

MLP Scheduling The memory bottleneck of MLP comes from the second linear layer. As shown in Figure 4(a), 3 tensors need to be stored, whose sizes add up to $6 \cdot s \cdot d$. We observe the computation of different tokens in an MLP is independent. This enables us to divide the input activation into tiles along the token dimension. Assume each tile has t tokens. Then, we can re-order the computation to finish all the MLP computations for one tile before moving to the next tile. Moreover, we can also perform an in-place addition with the residual and directly overwrite the t input tokens as shown in Figure 4(b). Thereby, we can reduce the execution memory to $s \cdot d + t \cdot 5d$. As t is usually much smaller than s, $6\times$ memory reduction can be achieved. In practice, t is carefully chosen based on the memory constraints and the computation resources. Hence, such scheduling optimization incurs negligible latency overhead on commodity MCUs.

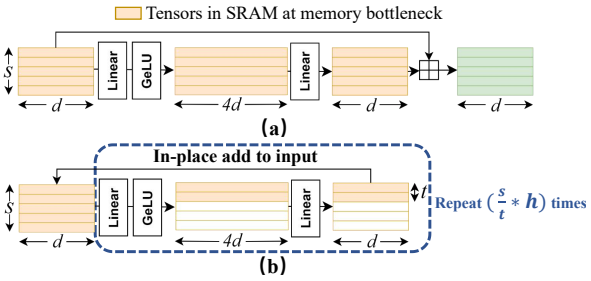

Figure 4: MLP scheduling to reduce peak execution memory. The tensor with yellow color will be saved in SRAM at the memory bottleneck.

MHA Scheduling Unlike MLP, the computation of different tokens depends on each other, making it hard to fully tile the MHA computation. We observe the following optimization opportunities: 1) the computation of each head is independent, and 2) the score tensor accounts for the major bottleneck, shown in Figure 1, and its computation can be tiled along token dimension.

Based on the observation, we propose a new MHA scheduling in Figure 5(b). We first tile the computation along the head dimension, i.e., h, and then, further tile the query tensor along the token dimension. This indicates we compute the attention between t tokens and all the s tokens per head each time. It enables us to reduce the memory of the score matrix, breaking the quadratic increase of execution memory into a linear increase with s. Again, we carefully choose t considering both memory constraints and computation parallelism to minimize latency overhead.

Our proposed MHA scheduling shares similarities with FlashAttention (Dao et al., 2022), which

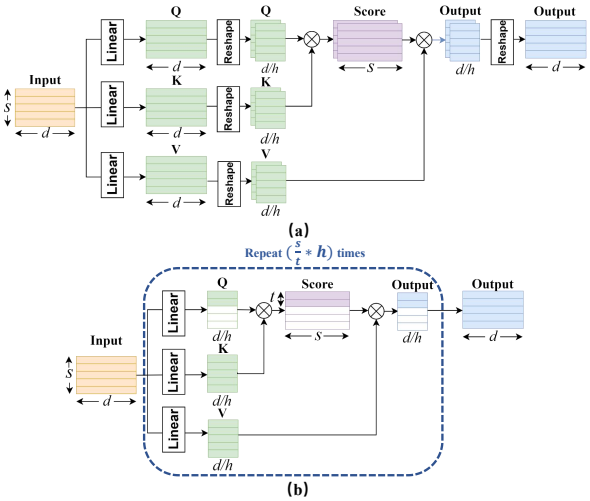

Figure 5: MHA scheduling to reduce the tensor transformation latency and peak memory.

tiles the key and value matrix. Hence, FlashAttention requires updating the score matrix repetitively. Either the score matrix has to be stored in high precision, leading to a higher peak execution memory, or quantization and de-quantization operations are needed during each output accumulation step, bringing heavy computation pressure and accuracy loss. Our method is more MCU-friendly and will not face such issues.

Kernel Design Optimization The scheduling optimization reduces the peak memory for MLP and MHA. For each tile of computation, we also design optimized kernels to reduce inference latency. First, we apply a two-level loop blocking to better fit into the MCU memory hierarchy. Compared to CMSIS-NN (Lai et al., 2018), we carefully design the loop dimensions based on MCU configurations to achieve higher locality and lower latency. We provide more details in Appendix E. Second, we unroll the reduction loop with a factor of 64 to harness the ILP of MCU and better utilize the hardware instruction pipeline. Moreover, we also carefully design the memory access patterns for the linear operators to fuse all the tensor shape transformation operators and avoid dedicated memory accesses. The kernel design optimization enables to fully leverage the hardware characteristics of MCUs for more efficient BERT processing.

4 Experiments

4.1 Experiment Setup

Dataset We search and evaluate our models on the General Language Understanding Evaluation (GLUE) benchmark (Wang et al., 2018), which is a collection of text classification tasks. For most evaluations and comparisons, we leverage the Multi-Genre Natural Language Inference Corpus (MNLI) dataset, which is the largest dataset in GLUE.

Table 3: Accuracy comparison on MNLI. The ratio of the first cluster is not shown as it equals to embedding dimension (Emb stands for embedding, params stands for parameters, and Acc stands for accuracy).

Model	Clustering Cutoffs	Ratios	Emb Params (M)	Total Params (M)	Acc (%)
BERT-tiny	-	-	3.907M	4.300M	70.01%
FWSVD	-	-	0.368M	0.761M	60.83%
Adaptive Emb	1000,4000,10000	32,8,2	0.318M	0.712M	67.01%
MCUBERT	510,1065, 1915	109, 18, 2	0.231M	0.624M	68.33%
	1218, 2534, 4061	54, 2, 2	0.307M	0.700M	69.18%
BERT-mini	-	-	7.814M	10.960M	74.80%
Adaptive Emb	1000,4000,10000	64, 16, 4	0.648M	3.794M	69.25%
MCUBERT	1014, 3348, 4912	41, 41, 2	0.492M	3.638M	71.89%
	1354, 4301, 6094	41, 42, 2	0.613M	3.759M	72.59%
	1871, 5591, 7788	52, 18, 2	0.779M	3.925M	74.22%

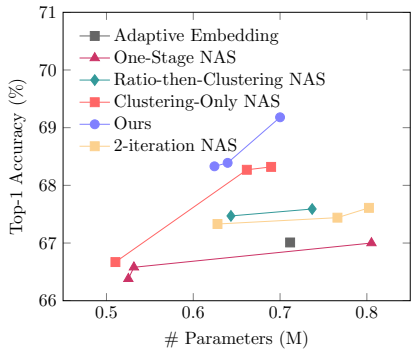

Figure 6: Importance of two-stage NAS.

Searching setting We select lightweight BERT models, i.e., BERT-tiny and BERT-mini for our experiments, and the pre-trained models are adopted from (Bhargava et al., 2021; Turc et al., 2019). The model configuration is shown in Appendix A. For NAS, we use AdamW optimizer with a zero weight decay. We use a batch size of 32 for training and set the learning rate to 5×10^{-5}.

Model deployment We deploy our model on different MCUs, i.e., NUCLEO-F746 with 320KB SRAM and 1MB Flash, NUCLEO-F767 with 512KB SRAM and 2MB Flash as well as NUCLEO-H7A3ZI-Q with 1.4 MB SRAM and 2MB Flash, to measure the latency and the peak memory usage. The batch size is fixed to 1.

4.2 Importance of two-stage NAS

As described in Section 3, we propose a two-stage NAS strategy, searching token clustering in the first stage and approximation ratios in the second stage. To demonstrate the importance of such two-stage

formulation, we compare MCUBERT with other baselines based on BERT-tiny and MNLI dataset, as shown in Figure 6, including 1) baseline adaptive embedding; 2) one-stage NAS that searches the token clustering and approximation ratios together; 3) searching approximation ratios first followed by token clustering (denoted as Ratio-then-Clustering NAS); 4) clustering-only NAS that only searches the token clustering; and 5) 2-round NAS that repeats the two-stage NAS for 2 rounds. As shown in Figure 6, our two-stage NAS outperforms all the other strategies and achieves the best Pareto front. One-stage NAS that ideally has the largest search space produces the worst Pareto front. This is because of the large discrepancy among optimal low-rank approximation ratios for different token clustering, which brings convergence difficulties as discussed in Section 3.

4.3 Accuracy Comparison on GLUE

Comparison on MNLI We compare MCUBERT with the baseline FWSVD (Hsu et al., 2022) and adaptive embedding (Baevski and Auli, 2018) on the MNLI dataset, as shown in Table 3.

Results and Analysis For both BERT-tiny and BERT-mini, MCUBERT can simultaneously achieve better accuracy and smaller parameter size compared to the baseline FWSVD and adaptive embedding. Specifically, for BERT-tiny, MCUBERT can achieve 1.4× and 1.6× embedding parameter reduction with 1.3% and 7.5% better accuracy compared to adaptive embedding and FWSVD, respectively. With the same model size, MCUBERT improves the accuracy by 2.2% and 8.4%. For BERT-mini, MCUBERT achieves 3.3% better accuracy with a smaller parameter size compared to adaptive embedding. We also observe for BERT-tiny, MCUBERT tries to assign a higher approximation ratio for the second cluster but assign more tokens to the fourth cluster.

Comparison on Other Datasets We then compare MCUBERT with adaptive embedding on other GLUE datasets. Note we focus on the MRPC, SST2, and QQP datasets as even the full BERT-tiny and BERT-mini suffers from accuracy issues on COLA and RTE datasets (Devlin et al., 2018). As shown in Table 4, MCUBERT consistently outperforms adaptive embedding on all three tasks by 3.43%, 0.69%, 0.88% better accuracy for BERT-tiny, respectively. For different datasets, the number of tokens and the approximation ratio of each cluster are different, indicating the necessity of task loss-guided compression.

Table 4: Accuracy comparison on other GLUE datasets (Emb stands for embedding, and params stands for parameters).

Model	Metrics	MRPC	SST2	QQP
BERT-tiny	Emb Params	3.91M	3.91M	3.91M
	Accuracy	74.02%	82.45%	87.16%
Adaptive Emb	Emb Params	0.32M	0.32M	0.32M
	Accuracy	70.34%	81.42%	84.32%
MCUBERT-tiny	Emb Params	0.31M	0.26M	0.32M
	Accuracy	73.77%	82.11%	85.20%
Model	**Metrics**	**MRPC**	**SST2**	**QQP**
BERT-mini	Emb Params	7.81M	7.81M	7.81M
	Accuracy	78.90%	85.32%	89.25%
Adaptive Emb	Emb Params	0.65M	0.65M	0.65M
	Accuracy	75.49%	83.60%	87.93%
MCUBERT-mini	Emb Params	0.65M	0.49M	0.52M
	Accuracy	77.94%	83.83%	88.12%

Table 5: Latency (ms) comparison on NUCLEO-F767 (F7) and NUCLEO-H7A3ZI-Q (H7) for different sequence lengths (OOM stands for out of memory).

MCU	Methods	64	128	192	256	512
F7	CMSIS-NN	386	838	1355	1939	OOM
	(Burrello et al., 2021)	366	780	1263	1812	OOM
	MCUBERT	264	578	942	1354	3591
H7	CMSIS-NN	320	697	1138	1638	4233
	(Burrello et al., 2021)	284	625	1016	1471	3866
	MCUBERT	213	469	768	1116	2940

4.4 Peak Memory and Latency Comparison

Peak memory comparison We compare the inference latency and peak memory with the baseline CMSIS-NN (Lai et al., 2018) and (Burrello et al., 2021). We re-implement (Burrello et al., 2021) based on the original paper. As shown in Figure 7, both CMSIS-NN and (Burrello et al., 2021) do not consider the execution memory in the scheduling, leading to a fast memory growth and out-of-memory issue for a long sequence length, i.e., 512. For BERT-tiny, MCUBERT achieves more than 1.9× and 3.5× peak memory reduction compared to the baseline when the sequence lengths are 64 and 512, respectively. This indicates MCUBERT can support 3×, 2×, and 2× longer input sequences compared to CMSIS-NN and (Burrello et al., 2021) on MCUs with 512 KB, 320 KB and 128 KB SRAM, respectively.

Latency comparison The latency comparison for different sequence lengths is shown in Table 5. MCUBERT achieves around 1.5× and 1.3× la-

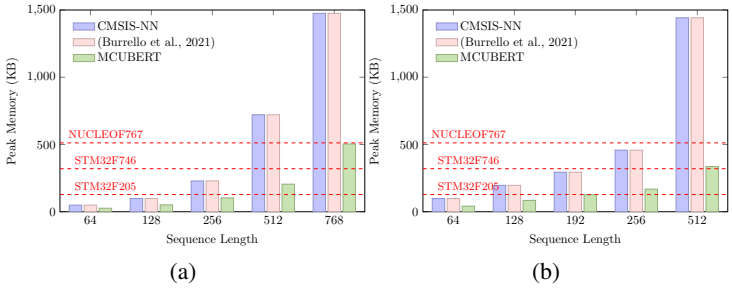

Figure 7: Peak memory comparison for different sequence lengths for (a) BERT-tiny and (b) BERT-mini.

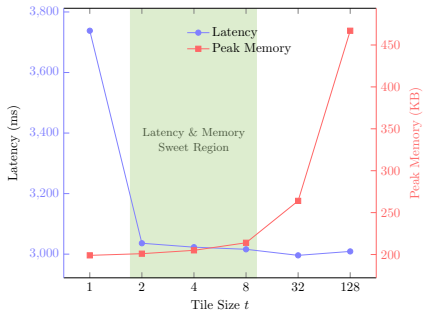

Figure 8: Comparison of latency and memory with different tile sizes t.

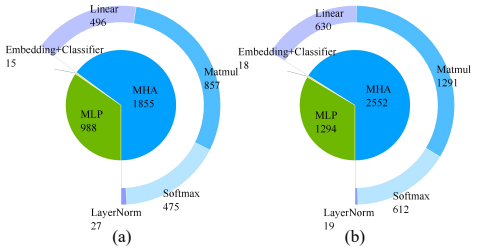

Figure 9: Latency (ms) breakdown for (a) MCUBERT and (b) (Burrello et al., 2021) with 512 sequence length (Matmul stands for matrix multiplication.).

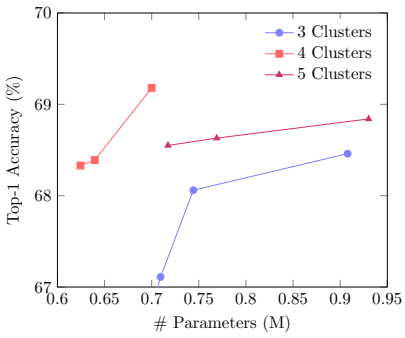

Figure 10: Impact of the number of clusters.

tency reduction compared to CMSIS-NN and (Burrello et al., 2021) consistently for different sequence lengths and different MCUs. We visualize the latency breakdown for (Burrello et al., 2021) and MCUBERT. As shown in Figure 9, MCUBERT reduces the latency of both the linear and matrix multiplication operators and fuses all tensor shape transformation operators, demonstrating the high efficiency of our kernel design. *With the model and scheduling optimization, MCUBERT enables to deploy BERT-tiny on NUCLEO-F746 for the first time.*

4.5 Ablation Study

Impact of the number of clusters MCUBERT follows adaptive embedding and empirically fix the number of clusters to 4. We now evaluate its impact on the accuracy-parameter Pareto front. We use the BERT-tiny model and MNLI dataset for comparison. As shown in Figure 10, the Pareto front for $c = 4$ indeed outperforms that for $c = 3$ or $c = 5$. We observe $c = 3$ performs much worse when the model parameter size is small. We hypothesize this is because when $c = 3$, important tokens are forced to have a small dimension, leading to low accuracy.

Impact of fine-grained scheduling To evaluate the impact of fine-grained scheduling on both latency and memory, we change the number of tokens in each tile, i.e., t. The change in latency and peak memory is plotted in Figure 8. As we can observe, the peak execution memory reduces consistently with the decrease of t while the latency remains roughly the same for $t \leq 2$. The high latency for $t = 1$ is because of hardware under-utilization.

5 Conclusion

This work proposes MCUBERT, a network/scheduling co-optimization framework enabling BERT on MCUs. For network optimization, MCUBERT proposes an MCU-aware two-stage NAS algorithm with clustered low-rank approximation for embedding compression. For scheduling optimization, we leverage tiling, in-place computation, and kernel optimization to simultaneously reduce peak memory and latency. MCUBERT overcomes all existing baselines and enables to run the lightweight BERT models on commodity MCUs for the first time.

6 Limitations

In this section, we discuss the current limitations and potential avenues for future research.

First of all, in our method, we compress the embedding table of language models and successfully

deploy small models such as BERT-tiny and BERT-mini on MCUs. In fact, our method also benefits the compression for larger language models such as BERT-medium, the embedding table of which contributes to around 39% of the parameters. While for larger models, to deploy them on devices with tight storage and memory constraints, we need to further compress the following encoder blocks. This can be improved in future work.

Also, in our experiments, we deploy encoder models on MCUs. For decoder models such as GPT, our embedding compression method and scheduling optimization such as tiling and in-place calculation still benefit the deployment process. While for higher efficiency, its generative inference can be further optimized by special kernel design, which we will explore in the future.

7 Ethics Statement

In our paper, our method enables deploying small BERT models on MCUs, overcoming the tight storage and memory constraints of the device by careful model compression and scheduling design. Our method democratized AI to IoT devices and can enable AI applications such as intelligent wearing and smart home to areas with poor connections, protecting privacy and saving energy at the same time. However, as the target model is relatively small, even the full models have relatively limited capacity compared with large language models. Therefore, for complex tasks or high-risk domains such as biomedicine, MCUBERT should be carefully used.

References

Alexei Baevski and Michael Auli. 2018. Adaptive input representations for neural language modeling. *arXiv preprint arXiv:1809.10853*.

Haoli Bai, Wei Zhang, Lu Hou, Lifeng Shang, Jing Jin, Xin Jiang, Qun Liu, Michael Lyu, and Irwin King. 2020. Binarybert: Pushing the limit of bert quantization. *arXiv preprint arXiv:2012.15701*.

Colby Banbury, Chuteng Zhou, Igor Fedorov, Ramon Matas, Urmish Thakker, Dibakar Gope, Vijay Janapa Reddi, Matthew Mattina, and Paul Whatmough. 2021. Micronets: Neural network architectures for deploying tinyml applications on commodity microcontrollers. *Proceedings of Machine Learning and Systems*, 3:517–532.

Prajjwal Bhargava, Aleksandr Drozd, and Anna Rogers. 2021. Generalization in nli: Ways (not) to go beyond simple heuristics.

Alessio Burrello, Moritz Scherer, Marcello Zanghieri, Francesco Conti, and Luca Benini. 2021. A microcontroller is all you need: Enabling transformer execution on low-power iot endnodes. In *2021 IEEE International Conference on Omni-Layer Intelligent Systems (COINS)*, pages 1–6. IEEE.

Daniel Campos, Alexandre Marques, Tuan Nguyen, Mark Kurtz, and ChengXiang Zhai. 2022. Sparse* bert: Sparse models are robust. *arXiv preprint arXiv:2205.12452*.

Patrick Chen, Si Si, Yang Li, Ciprian Chelba, and Cho-Jui Hsieh. 2018a. Groupreduce: Block-wise low-rank approximation for neural language model shrinking. *Advances in Neural Information Processing Systems*, 31.

Tianqi Chen, Thierry Moreau, Ziheng Jiang, Lianmin Zheng, Eddie Yan, Haichen Shen, Meghan Cowan, Leyuan Wang, Yuwei Hu, Luis Ceze, et al. 2018b. {TVM}: An automated {End-to-End} optimizing compiler for deep learning. In *13th USENIX Symposium on Operating Systems Design and Implementation (OSDI 18)*, pages 578–594.

Yu-An Chung, Yu Zhang, Wei Han, Chung-Cheng Chiu, James Qin, Ruoming Pang, and Yonghui Wu. 2021. W2v-bert: Combining contrastive learning and masked language modeling for self-supervised speech pre-training. In *2021 IEEE Automatic Speech Recognition and Understanding Workshop (ASRU)*, pages 244–250. IEEE.

Baiyun Cui, Yingming Li, Ming Chen, and Zhongfei Zhang. 2019. Fine-tune bert with sparse self-attention mechanism. In *Proceedings of the 2019 conference on empirical methods in natural language processing and the 9th international joint conference on natural language processing (EMNLP-IJCNLP)*, pages 3548–3553.

Tri Dao, Dan Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. 2022. Flashattention: Fast and memory-efficient exact attention with io-awareness. *Advances in Neural Information Processing Systems*, 35:16344–16359.

Robert David, Jared Duke, Advait Jain, Vijay Janapa Reddi, Nat Jeffries, Jian Li, Nick Kreeger, Ian Nappier, Meghna Natraj, Tiezhen Wang, et al. 2021. Tensorflow lite micro: Embedded machine learning for tinyml systems. *Proceedings of Machine Learning and Systems*, 3:800–811.

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*.

Jiarui Fang, Yang Yu, Chengduo Zhao, and Jie Zhou. 2021. Turbotransformers: an efficient gpu serving system for transformer models. In *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, pages 389–402.

Igor Fedorov, Ryan P Adams, Matthew Mattina, and Paul Whatmough. 2019. Sparse: Sparse architecture search for cnns on resource-constrained microcontrollers. *Advances in Neural Information Processing Systems*, 32.

Igor Fedorov, Ramon Matas, Hokchhay Tann, Chuteng Zhou, Matthew Mattina, and Paul Whatmough. 2022. Udc: Unified dnas for compressible tinyml models for neural processing units. *Advances in Neural Information Processing Systems*, 35:18456–18471.

Yen-Chang Hsu, Ting Hua, Sungen Chang, Qian Lou, Yilin Shen, and Hongxia Jin. 2022. Language model compression with weighted low-rank factorization. *arXiv preprint arXiv:2207.00112*.

Andrei Ivanov, Nikoli Dryden, Tal Ben-Nun, Shigang Li, and Torsten Hoefler. 2021. Data movement is all you need: A case study on optimizing transformers. *Proceedings of Machine Learning and Systems*, 3:711–732.

Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu. 2019. Tinybert: Distilling bert for natural language understanding. *arXiv preprint arXiv:1909.10351*.

Liangzhen Lai, Naveen Suda, and Vikas Chandra. 2018. Cmsis-nn: Efficient neural network kernels for arm cortex-m cpus. *arXiv preprint arXiv:1801.06601*.

Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Soricut. 2019. Albert: A lite bert for self-supervised learning of language representations. *arXiv preprint arXiv:1909.11942*.

Yinan Liang, Ziwei Wang, Xiuwei Xu, Yansong Tang, Zhou Jie, and Jiwen Lu. 2023. Mcuformer: Deploying vision tranformers on microcontrollers with limited memory. *arXiv preprint arXiv:2310.16898*.

Ji Lin, Wei-Ming Chen, Han Cai, Chuang Gan, and Song Han. 2021. Mcunetv2: Memory-efficient patch-based inference for tiny deep learning. *arXiv preprint arXiv:2110.15352*.

Ji Lin, Wei-Ming Chen, Yujun Lin, Chuang Gan, Song Han, et al. 2020. Mcunet: Tiny deep learning on iot devices. *Advances in Neural Information Processing Systems*, 33:11711–11722.

Vasileios Lioutas, Ahmad Rashid, Krtin Kumar, Md Akmal Haidar, and Mehdi Rezagholizadeh. 2019. Distilled embedding: non-linear embedding factorization using knowledge distillation.

Shishir G Patil, Paras Jain, Prabal Dutta, Ion Stoica, and Joseph Gonzalez. 2022. Poet: Training neural networks on tiny devices with integrated rematerialization and paging. In *International Conference on Machine Learning*, pages 17573–17583. PMLR.

Haotong Qin, Yifu Ding, Mingyuan Zhang, Qinghua Yan, Aishan Liu, Qingqing Dang, Ziwei Liu, and Xianglong Liu. 2022. Bibert: Accurate fully binarized bert. *arXiv preprint arXiv:2203.06390*.

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.

Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. 2020. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 3505–3506.

Manuele Rusci, Alessandro Capotondi, and Luca Benini. 2020. Memory-driven mixed low precision quantization for enabling deep network inference on microcontrollers. *Proceedings of Machine Learning and Systems*, 2:326–335.

Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. 2019. Distilbert, a distilled version of bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*.

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*, volume 34, pages 8815–8821.

Han Shi, Jiahui Gao, Xiaozhe Ren, Hang Xu, Xiaodan Liang, Zhenguo Li, and James Tin-Yau Kwok. 2021. Sparsebert: Rethinking the importance analysis in self-attention. In *International Conference on Machine Learning*, pages 9547–9557. PMLR.

Joonbo Shin, Yoonhyung Lee, and Kyomin Jung. 2019. Effective sentence scoring method using bert for speech recognition. In *Asian Conference on Machine Learning*, pages 1081–1093. PMLR.

Iulia Turc, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. Well-read students learn better: The impact of student initialization on knowledge distillation. *CoRR*, abs/1908.08962.

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*.

Hu Xu, Bing Liu, Lei Shu, and Philip S Yu. 2019. Bert post-training for review reading comprehension and aspect-based sentiment analysis. *arXiv preprint arXiv:1904.02232*.

Jin Xu, Xu Tan, Renqian Luo, Kaitao Song, Jian Li, Tao Qin, and Tie-Yan Liu. 2021. Nas-bert: task-agnostic and adaptive-size bert compression with neural architecture search. In *Proceedings of the 27th ACM*

SIGKDD Conference on Knowledge Discovery & Data Mining, pages 1933–1943.

Jiacheng Yang, Mingxuan Wang, Hao Zhou, Chengqi Zhao, Weinan Zhang, Yong Yu, and Lei Li. 2020. Towards making the most of bert in neural machine translation. In *Proceedings of the AAAI conference on artificial intelligence*, volume 34, pages 9378–9385.

Ofir Zafrir, Ariel Larey, Guy Boudoukh, Haihao Shen, and Moshe Wasserblat. 2021. Prune once for all: Sparse pre-trained language models. *arXiv preprint arXiv:2111.05754*.

Yujia Zhai, Chengquan Jiang, Leyuan Wang, Xiaoying Jia, Shang Zhang, Zizhong Chen, Xin Liu, and Yibo Zhu. 2023. Bytetransformer: A high-performance transformer boosted for variable-length inputs. In *2023 IEEE International Parallel and Distributed Processing Symposium (IPDPS)*, pages 344–355. IEEE.

Wei Zhang, Lu Hou, Yichun Yin, Lifeng Shang, Xiao Chen, Xin Jiang, and Qun Liu. 2020. Ternarybert: Distillation-aware ultra-low bit bert. *arXiv preprint arXiv:2009.12812*.

Xiaofan Zhang, Zongwei Zhou, Deming Chen, and Yu Emma Wang. 2022. Autodistill: An end-to-end framework to explore and distill hardware-efficient language models. *arXiv preprint arXiv:2201.08539*.

Jinhua Zhu, Yingce Xia, Lijun Wu, Di He, Tao Qin, Wengang Zhou, Houqiang Li, and Tie-Yan Liu. 2020. Incorporating bert into neural machine translation. *arXiv preprint arXiv:2002.06823*.

A Model Configurations

Table 6 shows the detailed configurations of the pre-trained BERT-tiny and BERT-mini models, which are used in our experiments.

Table 6: Configurations of BERT pre-trained models.

Model	Heads	Layers	Emb dimension	Max sequence length	Vocab size
BERT-tiny	2	2	128	512	30522
BERT-mini	4	4	256	512	30522

B GLUE scores of pre-trained models

In our experiments, we focus on MNLI, MRPC, SST2, and QQP datasets in GLUE. We do not evaluate the searched models on COLA and RTE, as our baseline BERT-tiny and BERT-mini suffer from accuracy issues on these datasets (Devlin et al., 2018), shown in Table 7.

Table 7: GLUE scores of BERT pre-trained models. The evaluation metric for COLA is Mcc. The evaluation metric for other tasks is Acc.

Model	MNLI	MRPC	SST2	QQP	COLA	RTE
BERT-tiny	70.0%	74.0%	82.5%	87.2%	0.0%	57.2%
BERT-mini	74.8%	78.9%	85.3%	89.3%	0.0%	57.9%
BERT-base	84.6%	88.9%	93.5%	90.8%	52.1%	66.4%

C Algorithm details of hardware-aware NAS

We propose a two-stage NAS strategy: in the first stage, we fix different low-rank approximation ratios for each cluster, and only search the token clustering, i.e., α; then, in the second stage, given the fixed token clustering, we search the best approximation ratio for each cluster, i.e., β. Our assumption here is the token clustering reflects the token importance, which should remain stable for different low-rank approximation ratios. The algorithm details are shown in Algorithm 1, 2 and 3.

D Model Deployment on MCUs

There are two main approaches to deploy models on MCUs, interpretation (Chen et al., 2018b; David et al., 2021) and code generation (Lin et al., 2020, 2021; Burrello et al., 2021). The interpretation-based methods embed an on-device runtime interpreter to support flexible model deployment but require extra memory to store meta-information and extra time for runtime interpretation. Code-generation-based methods directly compile the given model into target code to save memory and reduce inference latency. MCUBERT uses a code generation method that is more specialized for our target model and device, leading to lower latency and lower memory usage.

E Kernel Design Optimization

Deploying model on MCUs is very challenging. MCUs have tight constraints on SRAM and Flash size, which saves activation and model weights respectively, limiting peak execution memory and the total parameter size of models. We solve this challenge using MCU-aware NAS and MCU-friendly scheduling, as shown in Section 3. While register locality mainly limits the latency of model inference, we conduct kernel design optimization for lower latency. In this section, we show more details of this optimization.

Figure 11 illustrates the kernel design for each tile of the matrix multiplication/linear operator. We apply a two-level loop-blocking based on the two-level memory hierarchy (SRAM, Register) of MCU. The innermost block is referred to as a micro-kernel, which computes entirely on registers. Properly setting its

Algorithm 1: Searching for token clustering

Input: Embedding matrix Emb; Factorized embedding U, V; architecture parameters α; vocabulary size v; # clusters c; embedding dimension d; Division value div; # training epochs $epochs$

Output: Clustering Results cls

for all $i = 1, 2, \cdots, v$ **do**
  for all $j = 1, 2, \cdots, c$ **do**
    $\alpha_{ij} = \frac{1}{c}$;
  end for
end for
$U_1 = Emb$
for all $i = 2, \cdots, c$ **do**
  $U_i, V_i = LowRankFactorization(Matrix = Emb, r = \frac{d}{div})$;
end for
for all $j = 1, 2, \cdots, v$ **do**
  Redefining the embedding of j_{th} token as: $\text{Emb}_j = \alpha_{j,1} U_{1,j} + \sum_{i=2}^{c} \alpha_{j,i} U_{i,j} V_i$
end for
for all $i = 1, 2, \cdots, epochs$ **do**
  Update weights w by descending $\nabla_w \ell_{w,\alpha}$
  Update architecture parameters α by descending $\nabla_\alpha(\ell_{w,\alpha} + \lambda \sum_i \text{Size}(\text{Emb}_i))$
end for
for all $i = 1, 2, \cdots, v$ **do**
  $cls_i = \arg\max_j(\alpha_{ij})$;
end for
return cls

Algorithm 2: $LowRankFactorization$

Input: Matrix to be factorized $Matrix$; Factorization ratio r;

Output: Factorization Results U and V

$U, \text{Diag}(\Sigma_0, \ldots, \Sigma_{d-1}), V = SVD(Matrix)$
$k = sort(\Sigma)_{d-r}$
for all $i = 0, 1, \cdots, d-1$ **do**
  if $\Sigma_i < k$ **then**
    $\Sigma_i = 0$;
  end if
end for
$U = U * \text{Diag}(\sqrt{\Sigma_0}, \ldots, \sqrt{\Sigma_{d-1}})$
$V = \text{Diag}(\sqrt{\Sigma_0}, \ldots, \sqrt{\Sigma_{d-1}}) * V$
return $U[:, NoneZeroRows], V[NoneZeroRows, :]$

shape helps exploit register-level locality. In our experiments, we set the micro-kernel shape to [M, N, K] = [4, 2, 4]. In comparison, CMSIS-NN (Lai et al., 2018) employs a shape of [1, 2, 4], which exhibits lower locality and decreases performance.

F Exsiting Works of Efficient NN Design

In this section, we focus on reviewing the model optimizations targeting at MCUs and for LMs. Deep learning models need to meet the tight storage and memory constraints to run on MCUs. Previous works propose network and scheduling optimization as summarized in Table 1. (Rusci et al., 2020; Banbury et al., 2021; Lin et al., 2020) compress CNNs with NAS to meet the storage constraints. MCUNet (Lin et al., 2020) also uses tiling and in-place updating for scheduling optimization. (Burrello et al., 2021) deploys small Transformer models which already satisfy the constraints of MCUs and conduct scheduling

Algorithm 3: Searching for low-rank approximation ratio

Input: Embedding matrix Emb; Factorized embedding U, V; architecture parameters β; vocabulary size v; # clusters c; embedding dimension d; # training epochs $epochs$; importance threshold β^*; searched token clustering cls

Output: Low-rank approximation ratio $ratios$

for all $i = 1, 2, \cdots, c$ **do**
  for all $j = 0, 1, \cdots, d-1$ **do**
    $\beta_{ij} = \beta_i^*$;
  end for
end for
Divide the Emb into c parts $\{Emb_1, \ldots, Emb_c\}$ according to token clustering cls
$U_1 = Emb_1$
for all $i = 2, \cdots, c$ **do**
  $U_i, \mathrm{Diag}(\Sigma_{i,0}, \ldots, \Sigma_{i,d-1}), V_i^\top = LowRankFactorization(matrix = Emb_i, dim = d)$;
end for
for all $j = 1, 2, \cdots, v$ **do**
  Redefining the embedding of j_{th} token in i_{th} cluster as:
  $\mathrm{Emb}_j = U_{i,j}\mathrm{Diag}(\beta_{i,0}\Sigma_{i,0}, \ldots, \beta_{i,d-1}\Sigma_{i,d-1})V_i^\top$
end for
for all $i = 1, 2, \cdots, epochs$ **do**
  Update weights w by descending $\nabla_w \ell_{w,\beta}$
  Update architecture parameters β by descending $\nabla_\beta(\ell_{w,\beta} + \lambda \sum_i \mathrm{Size}(\mathrm{Emb}_i))$
end for
for all $i = 1, 2, \cdots, c$ **do**
  $ratios_i = 0$
  for all $j = 0, 1, \cdots, d-1$ **do**
    if $\beta_{ij} > \beta_i^*$ **then**
      $ratios_i += 1$;
    end if
  end for
end for
return $ratios$

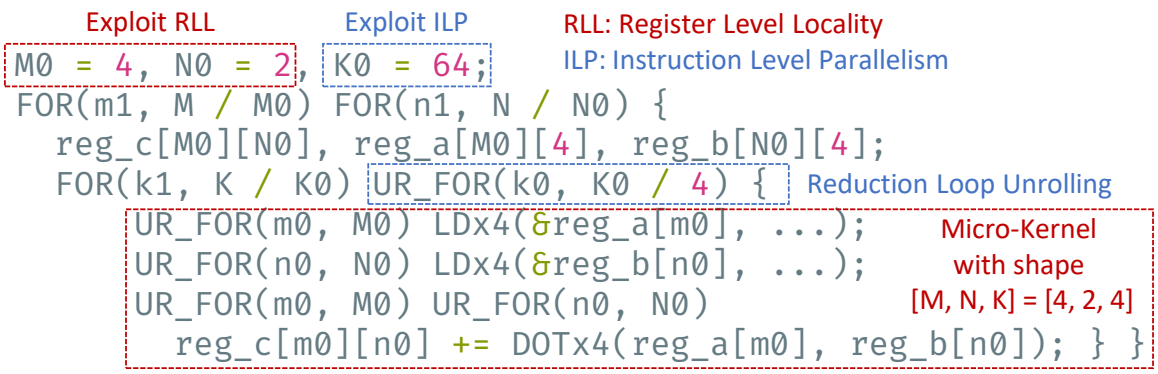

Figure 11: Matrix multiplication kernel design of MCUBERT. `UR_FOR` means fully unrolled for-loop. `LDx4` means loading 4 consecutive elements from SRAM to the register. `DOTx4` means computing dot-product of two 4-element vectors using SIMD instruction (e.g., `SMLAD` for ARM Cortex-M MCU). The quantization and writing-back operations are omitted here.

optimization to accelerate model inference. (Liang et al., 2023) deploys vision transformer (ViT) on MCUs mainly by compressing the MLP layers and searching the token numbers. However, there are new challenges for deploying BERT on MCUs. First, compared to tiny CNNs and small Transformer models,

BERT models have more parameters. Second, the memory bottleneck of BERT inference mainly lies in the MHA block when sequence length is long, which is shown in Figure 1(c). The computation of MHA block is more complex compared with CNN blocks and MLP layers, bringing new challenges for scheduling optimization. We propose MCUBERT to first deploy BERT, the most representative encoder transformer model, on MCUs.

Though no existing works deploy BERT models on MCUs, there are network optimization algorithms proposed for LMs on other platforms, e.g., GPUs. Distilled Embedding (Lioutas et al., 2019) and Albert (Lan et al., 2019) leverage low-rank factorization based on singular value decomposition (SVD) for embedding compression and prune singular values with small magnitudes. Adaptive embedding (Baevski and Auli, 2018) further divides the tokens into different clusters according to the token frequency and applies low-rank factorization separately to each cluster. However, adaptive embedding adapts empirical methods to cluster tokens and prunes singular values, leading to a large accuracy degradation.

There are also scheduling optimization proposed for Transformer models on GPUs. FlashAttention (Dao et al., 2022) only computes the attention score tensor partially each time and repetitively updates the accumulation of the partial sum to drastically reduce the execution memory. There are also previous works like (Zhai et al., 2023; Fang et al., 2021; Rasley et al., 2020) that use kernel fusion to reduce memory usage and accelerate inference. However, these methods usually do not consider quantization and targets at both training and inference. Our scheduling optimization is inspired by these methods but is more MCU-friendly.