

LocMoE: A Low-Overhead MoE for Large Language Model Training

Jing Li[†], Zhijie Sun^{†,*}, Xuan He[†], Li Zeng, Yi Lin, Entong Li, Binfan Zheng, Rongqian Zhao and Xin Chen

Huawei Technologies Co., Ltd

{lijing473, sunzhijie3, hexuan22, zengli43, linyi11, lientong, zhengbinfan1, zhaorongqian, chenxin}@huawei.com

Abstract

The Mixtures-of-Experts (MoE) model is a widespread distributed and integrated learning method for large language models (LLM), which is favored due to its ability to sparsify and expand models efficiently. However, the performance of MoE is limited by load imbalance and high latency of All-to-All communication, along with relatively redundant computation owing to large expert capacity. Load imbalance may result from existing routing policies that consistently tend to select certain experts. The frequent inter-node communication in the All-to-All procedure also significantly prolongs the training time. To alleviate the above performance problems, we propose a novel routing strategy that combines load balance and locality by converting partial inter-node communication to that of intra-node. Notably, we elucidate that there is a minimum threshold for expert capacity, calculated through the maximal angular deviation between the gating weights of the experts and the assigned tokens. We port these modifications on the PanGu- Σ model based on the MindSpore framework with multi-level routing and conduct experiments on Ascend clusters. The experiment results demonstrate that the proposed LocMoE reduces training time per epoch by 12.68% to 22.24% compared to classical routers, such as hash router and switch router, without impacting the model accuracy.

1 Introduction

Large Language Models (LLM), such as GPT [Brown *et al.*, 2020] and LLaMA [Touvron *et al.*, 2023], have recently gone viral due to their distinguished capabilities in word processing and data analysis. The architectures of these LLMs are mostly derived from the Transformer, which is on the basis of the self-attention mechanism [Vaswani *et al.*, 2017]. Since the predictive ability of the Transformer-based model correlated strongly with the model size [Kenton and Toutanova, 2019], the parameter scales of existing LLMs have increased dramatically to assure accuracy. The complex construction,

along with the large parameter scale, triggers the rapid surge in demand for computing resources, resulting in escalating training and inference costs that hinder the development of LLMs [Lewis *et al.*, 2021]. Aiming at the problem, Mixtures-of-Experts (MoE) [Jacobs *et al.*, 1991] provide an effective way to extend the model capacity at a fixed computational overhead [He *et al.*, 2021], thus emerging as the preferred option for some renowned LLMs.

A typical MoE framework consists of a gated network and several expert networks that selectively activate a portion of parameters for various inputs to participate in computation [Clark *et al.*, 2022]. Owing to such a structure, the computational complexity remains relatively invariant when the scale of parameters increases [Puigcerver *et al.*, 2020]. Since each token activates only one or a few experts, sparse routing of the gated network delivers the token to the most appropriate expert(s) [Zuo *et al.*, 2021]. If the routing strategy is not well-designed, it may lead to the overtraining of a few experts and under-training of others, ultimately evolving into inefficient learning and uneven load distribution [Shazeer *et al.*, 2016]. To address this shortcoming, Switch Transformer [Fedus *et al.*, 2022] simplifies the routing mechanism of MoE while adding an auxiliary loss that encourages a balanced load across experts. Moreover, the frequent All-to-All communication delay has also limited the performance of MoE [Rajbhandari *et al.*, 2022]. It is estimated that the time-consuming ratio of All-to-All under 8 A100 GPUs in a single node is about 31.18% and would be much higher in multiple nodes [Nie *et al.*, 2023]. HeTuMoE [Nie *et al.*, 2022] further puts forward a hierarchical All-to-All strategy, which fully utilizes the bandwidth of intra-node NVLink and inter-node Infiniband to cope with the problem of low bandwidth utilization due to frequent inter-machine transfers of small data volumes.

In this paper, we propose **LocMoE**, a low-overhead routing strategy and a communication optimization scheme, and it is applied in PanGu- Σ model [Ren *et al.*, 2023]. PanGu- Σ is a sparse model extended by the dense model PanGu- α [Zeng *et al.*, 2021]. With Ascend cluster [Liao *et al.*, 2021], it is measured that the All-to-All communication in PanGu- Σ takes 18.10% and 28.74% of the training time under 128 Ascend 910A Neural Network Processing Units (NPU) and 256 Ascend 910A NPUs, respectively. It still has potential for further reduction, and we make the following optimiza-

^{*}Corresponding author

tions based on:

- **Orthogonal gating weight with Grouped Average Pooling (GrAP) layer.** The GrAP layer is adopted in gating values computation. It provides a natural way to perform class activation mapping and reduce computational costs. Above all, the orthogonality of gating weight facilitates the explicit decisions of the router.
- **Locality-based expert regularization.** Redistribute on the basis of load balance, add the locality loss as the regularization term, and transform partial inter-node communication into intra-node communication with higher bandwidth. The local experts are encouraged to compete with skilled experts, and the time consumption of communication is reduced while avoiding the under-training of some experts.
- **Reduction of expert capacity without losing accuracy.** Our work proves and solves the critical value of MoE’s expert capacity in the NLP sector for the first time, and its relationship with input corpus features is also elucidated. Furthermore, we find fewer class-discriminative tokens need to be learned by experts than class-correlated ones. The experimental results also confirm that the model accuracy would not be affected after downward adjusting the expert capacity within the critical limit.

After applying the above improvements, the time consumption of All-to-All communication decreases by 5.13%. The elapsed time per epoch decreases by up to 22.24% with our cluster groups (containing 8, 16, and 32 node with 64, 128, and 256 Ascend 910A NPUs, abbreviated as 64N, 128N, and 256N in the following paragraphs).

The remainder of this paper is organized as follows: Section 2 displays related works of MoE in the field of NLP, the Ascend architecture, and the base model PanGu- Σ . Section 3 demonstrates the methodology details of the LocMoE and the theoretical bounds. Section 4 analyses the results of comparison experiments. Section 5 summarizes this work and the prospects for its future research orientation.

2 Related Work

MoE. MoE is a strategy for model designing, combining with several expert networks, to enhance the model capacity and efficiency. The concept of MoE was first proposed in 1991 and became the prototype of the existing MoE structure [Jacobs *et al.*, 1991]. Sparsely-gated MoE [Shazeer *et al.*, 2016] was proposed to expand the model capacity adequately under the same arithmetic power, and the gating is designed to allow TopK experts to be activated in an iteration. GShard [Lepikhin *et al.*, 2020] was the first work to migrate the MoE to Transformer, using the expert capacity to limit the tokens processed by each expert to a certain range. In addition, the auxiliary loss is proposed in GShard’s random routing to deal with the winner-take-all drawback of MoE. Regarding expert capacity, the work of pMoE [Chowdhury *et al.*, 2023] has proved for the first time that each expert can be fully trained even when dealing with samples much smaller

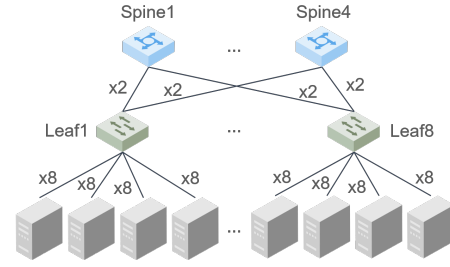


Figure 1: The networking scheme applied in the Ascend cluster.

than the number of tokens, but has a threshold. Switch Transformer [Fedus *et al.*, 2022] selects only the top expert to maximize MoE’s sparsity and proposes a corresponding auxiliary loss to achieve load balance. Facebook AI Research implements the Hash FFN layer [Roller *et al.*, 2021] with the balanced hash function, and the distribution of the experts’ load is close to the ideal state. Taking into consideration both convergence and accuracy, StableMoE [Dai *et al.*, 2022] adopts a two-stage training procedure. In the first stage, the imbalance of assignment and the cross-entropy of routing features are adopted as loss penalty terms, and the model directly learns with the routing strategy in the second stage. X-MoE [Chi *et al.*, 2022] rewrites the score function between the token and the expert by reducing dimensionality. Task-MoE [Kudugunta *et al.*, 2021] describes task-based routing at multiple granularities: token level, sentence level, and task level. HetuMoE [Nie *et al.*, 2022] proposes the hierarchical AlltoAll strategy, which combines hierarchical networks and aggregated information to improve transmission efficiency.

Ascend Architecture. The pivot architecture of the Ascend mainly consists of multilevel on-chip memory, load/storage units, and instruction management units [Liao *et al.*, 2021]. System-on-Chip (SoC) adopts the Mesh Network-on-Chip (NoC) [Kumar *et al.*, 2022] architecture to provide a unified and scalable communication network, realizing a high bandwidth of 256GB/s [Li *et al.*, 2022b]. In Ascend 910A server, every eight NPUs are divided into two groups on the board. The intra-group connection is based on the Huawei Cache Coherence System (HCCS) [Xia *et al.*, 2021]. The Ascend 910A chip delivers 320 Tera FLOPS at semi-precision (FP16) and 640 Tera OPS at integer precision (INT8). Our cluster is built based on a two-tier Fat-tree networking scheme on the single plane, with each Leaf switch connecting to 4 NPU servers (model Atlas 800 9000), as in Figure 1. The algorithm bandwidth of each communication operator in Huawei Collective Communication Library (HCCL) is displayed in Figure 2.

PanGu Series Model. The fields of PanGu series large models are mainly divided into NLP, computer vision, multimodality, graph network, and scientific computing [Mi *et al.*, 2022; Shen *et al.*, 2023; Bi *et al.*, 2023]. Thereinto, the models in the field of NLP focus primarily on text generation and semantic understanding. The most representative NLP model in the PanGu series is the PanGu- α [Zeng *et al.*, 2021], which is an LLM in the Chinese domain with up to 200 billion parameters. It also applies the auto-parallel framework based

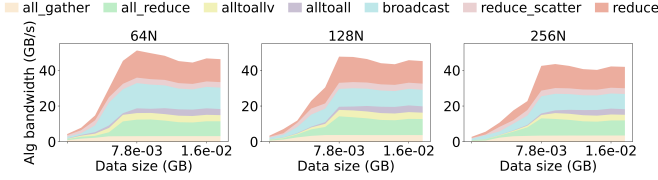


Figure 2: The algorithm bandwidth of each communication operator in HCCL under 64N, 128N, and 256N, respectively.

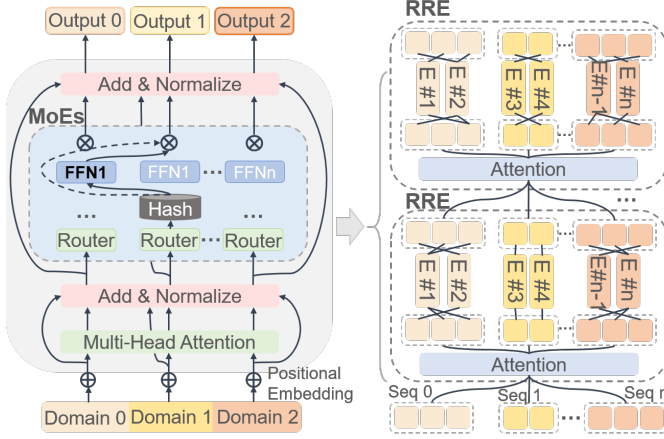


Figure 3: The architecture of sparse Transformer layers in PanGu- Σ .

on the MindSpore [Tong *et al.*, 2021]. PanGu- π [Wang *et al.*, 2023] mitigates feature collapse in the Transformer architecture by introducing more nonlinearities in the feed-forward networks (FFN) and MSA modules. Utilizing the intrinsic parameters of PanGu- α , PanGu- Σ [Ren *et al.*, 2023] is extended to a sparse model containing 1.085 trillion parameters by the conception of MoE.

3 Methodology

3.1 PanGu- Σ

The PanGu- Σ architecture consists of both dense and sparse Transformer encoder layers, stacked Transformer decoder layers modeled in the autoregressive language, and a query layer. The sparse Transformer layer of PanGu- Σ , with several conditionally activated feedforward sublayers, incorporates the MoE principle, as displayed in Figure 3. The RRE module is responsible for routing the token to the appropriate expert. It contains two levels of routing: in the first level, the experts are grouped by domains, and the token is assigned to one of the groups. In the second level, the token is routed to a particular expert of this group homogeneously. The second level of routing can be viewed as random hash routing, which does not contain learnable parameters.

3.2 MoE With Local Routing Strategy

MoE in Encoder Layers of Transformers

Similar to the classic MoE skeleton applied to Transformer structures such as GShard, the MoE layer in our model mainly consists of a MSA layer, a gating network, a routing module,

and several expert FFNs. The output of the MoE layer can be depicted as follows:

$$y_m = \sum_{i=1}^n \mathcal{R}_{m,E_i} \cdot W_{E_i,\text{out}} \cdot \text{GeLU}(W_{E_i,\text{in}} \cdot x_m) \quad (1)$$

Assume that the MoE layer contains n experts, \mathcal{R}_{m,E_i} denotes the expert score acquired by the gating network when expert i provides the largest gating value. The expert network of expert i consists of two linear transformations with a Gaussian Error Linear Unit (GeLU) activation, which is the product of input and the standard Gaussian cumulative distribution function. Thereinto, the gating function \mathcal{G} is the critical component of router \mathcal{R} . Typically, it is designed to be a dense layer extracting the feature of input tensor:

$$i^* = \arg \max_{i \in [n]} (\text{softmax}(\mathcal{G}_{m,E_i})) \quad (2)$$

$$\mathcal{R}_{m,E_i} = \mathbb{1}\{i = i^*\} (\text{softmax}(\mathcal{G}_{m,E_i})) \quad (3)$$

where i^* is the index of the most appropriate expert, and $\mathcal{G}_{m,E_i} = \text{ReLU}(\omega_i \cdot x_m + \epsilon_i)$. The token would be sent to the Top-1 expert with the largest expert score screened by Softmax. To reduce the parameter scale and the computational overhead, the gating value is obtained via the GrAP layer instead of the dense layer [Li *et al.*, 2022a]. The feature extraction with the GrAP layer is delineated in Figure 4. It can be regarded as the dense layer with the fixed weight ω_i :

$$\omega_i = \frac{n}{d} (\omega_{i,j} = \mathbb{1}\{i \frac{d}{n} \leq j < (i+1) \frac{d}{n}\}, 0 \leq j < d) \quad (4)$$

where d denotes the dimension of activation. Notably, the gating weights of the GrAP layer are orthogonal. From a perspective of semantic, irrelevant tokens are inclined to be routed to experts of different domains, which is conducive to convergence and accuracy [Guo *et al.*, 2020]. Besides, the GrAP layer has greater computation efficiency.

Localized Bias Weighting Loss

A general observation on the original two-level routing strategy of PanGu- Σ reveals that the router is devoid of the learning process. Although meeting load balance requirements, it lacks interpretability for distinguishing experts by domain. LocMoE rewrites the loss, consisting of two parts: auxiliary loss and locality loss. The auxiliary loss is first proposed in the sparsely-gated MoE [Shazeer *et al.*, 2016] and is also applied in Switch Transformer [Fedus *et al.*, 2022]:

$$L_{\text{aux}} = \alpha n \sum_{i=1}^n f_i P_i \quad (5)$$

$$f_i = \frac{1}{T} \sum_{x \in \beta} \mathbb{1}\{\arg \max p(x) = i\}, P_i = \frac{1}{T} \sum_{x \in \beta} p_i(x)$$

where n denotes the number of experts, β denotes the batch containing T tokens. f_i denotes the proportion of tokens assigned to expert i , and P_i denotes the average probability that

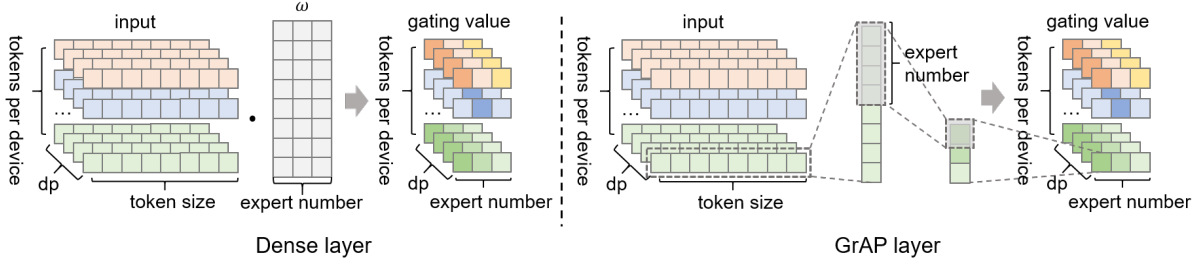


Figure 4: Difference between feature extraction via the dense layer and the GrAP layer.

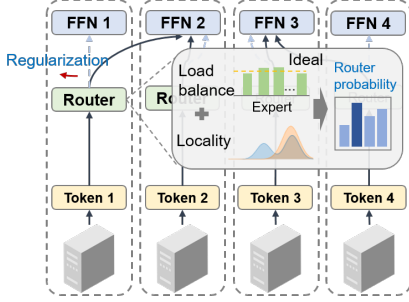


Figure 5: The action principle of locality loss.

the router chooses expert i . The auxiliary loss has substantiated that it can cause the balance of routing, as the loss would achieve its minimum under a uniform distribution. The hyperparameter α is set to 0.01, which refers to the value in the previous work [Fedus *et al.*, 2022].

The second part is locality loss, in line with the expectation that tokens are more likely to be assigned to local experts under the premise of load balance. The loss function can be measured by the difference between the current distribution and the fully localized distribution. The current distribution reflects the assignment distribution of all experts in the current batch, and the difference can be described using Kullback-Leibler (KL) divergence:

$$L_{loc} = \mu \text{KL}(D_c || D_l) = -\mu \int D_c(x) \ln \left[\frac{D_l(x)}{D_c(x)} \right] dx \quad (6)$$

where D_c denotes the current distribution and D_l denotes the fully localized distribution, and μ is the hyperparameter. The locality, along with the auxiliary loss, acts as the soft constraint that impels the tokens in the same domain to be trained by local experts, as shown in Figure 5. The blue dashed arrows are the contribution of the locality loss, and the final assignment of tokens considers the synthetic effect of gating value, auxiliary loss, and locality loss. The task loss is the sum of the above loss items and cross-entropy:

$$L_{task} = L_{aux} + L_{loc} + L_{cross} \quad (7)$$

where $L_{cross} = - \sum_{t=1}^T \log \frac{\exp(c_t^*)}{\sum_{i=1}^N \exp(c_{t,i})}$

Critical Value of Expert Capacity

The introduction of expert capacity aims to avoid the training block caused by the assignment imbalance of tokens. In general, an empirical expert capacity factor c_f is set to limit the scale of the expert capacity: $ec = \lceil \frac{b_s * c_f}{ep * n} \rceil$. b_s , c_f , ep , and n denote the batch size, the capacity factor, the degree of expert parallelism, and the number of experts, respectively. The computational workload of experts can be equalized in this way. However, the work of pMoE reveals that the sample size each expert needs to process has its lower bound [Chowdhury *et al.*, 2023], which can provably reduce the training costs. Inspired by the work on pMoE, we migrated the assumptions of data distribution in MoE from CV to the NLP domain, in conjunction with the network structure, while further discovering some significant conclusions:

Assumption 1. *The magnitudes of gating weight $\|\omega\|$ are equivalent for all experts.*

Lemma 1 (Minimum Angle of Expert). *The top-1 router is essentially the mechanism to select the expert i^* with the minimum angle $\theta_{i^*,j}$ to the gating weight ω_i .*

Assumption 2. *Suppose all tokens with size of d uniformly distributed on the unit sphere, that is, $\|x_m\| = 1$.*

Lemma 2 (Equivalent Probability for Assignment). *Suppose i_j denotes the expert to which the token j is routed. On account of the spherical symmetry, the probabilities for $i_j = i'$ are equivalent for all experts under the conditions of orthogonal gating weights (brought from the GrAP layer). That is, $P\{i_j = i'\} = \frac{1}{n}$.*

Assumption 3. *Assume that if $\delta_{i,j} \geq \delta$, the token j should be assigned to the expert i , where $\delta_{i,j} = \cos(\theta_{i,j})$*

Lemma 3 (Assignment Probability for Unit Vector). *For the uniformly distributed unit vector j , the probability it should be assigned to the expert i is:*

$$p_\delta = 1 - I_{\delta^2} \left(\frac{1}{2}, \frac{d-1}{2} \right) \quad (8)$$

where $I_x(a, b) = \frac{1}{B(a, b)} \int_0^x t^{a-1} (1-t)^{b-1} dt$, $0 \leq x \leq 1, a > 0, b > 0$

As for the activation size, for large d , when $\delta = \Theta(\frac{1}{\sqrt{d}})$, $p_\delta \approx 0.3$. When δ is larger than $\frac{1}{\sqrt{d}}$, p_δ declines to 0 rapidly.

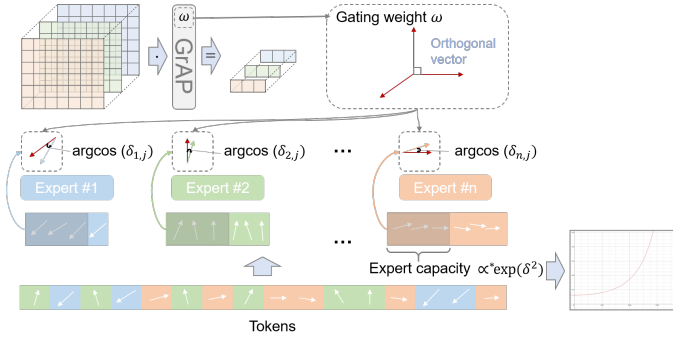


Figure 6: The correlation between expert capacity and the angle between token and gating weight.

Theorem 1 (Lower Bound of Expert Capacity). *From Lemma 3, assume that p_i denotes the probability of the token routed to the expert i is class-discriminative, thus, $p_i \leq n[1 - I_{\delta^2}(\frac{1}{2}, \frac{d-1}{2})]$. The lower bound of the expert capacity can be described as:*

$$ec_{\min} = \frac{1}{p_i} \geq \frac{1}{n[1 - I_{\delta^2}(\frac{1}{2}, \frac{d-1}{2})]}$$

for large d ,

$$ec_{\min} \geq \frac{1}{n \cdot \operatorname{erfc}(\sqrt{\frac{\delta^2 d}{2 - \delta^2}})} > \frac{1}{n} \exp\left(\frac{\delta^2 d}{2 - \delta^2}\right), \quad (9)$$

where erfc is the complementary error function. Figure 6 portrays the schematic diagram of our discovery, which describes the correlation between the expert capacity and the minimum angle of experts. The expert capacity correlates negatively with the minimum angle between token and gating weight, and it grows exponentially with the decrease of the angle θ .

Group-Wise All-to-All and Communication Overlap

Since the All-to-All is an aggregate communication operator, other operations would be performed until the data is completely transmitted, leading to low hardware utilization efficiency. Our model applies the group-wise exchange algorithm embedded in MindSpore to split and rearrange the All-to-All operations. In the tensor parallel (TP) domain, each device is responsible for a portion of the All-to-All data transmission in its respective expert parallel (EP) domain. Then, the All-Gather operation is conducted to synchronize tokens on all devices in the TP domain. The communication volume is diverted to the TP domain with high-speed bandwidth, which reduces the overall All-to-All communication time. In addition, FFN computation and communication are sliced and overlapped to mask the delay caused by communication, eventually reducing the time of communication.

4 Experiment Results and Analysis

We conduct experiments on the Ascend cluster groups (see environment configuration in Appendix C) to verify the effect of LocMoE. The existing classical MoEs, such as HashMoE

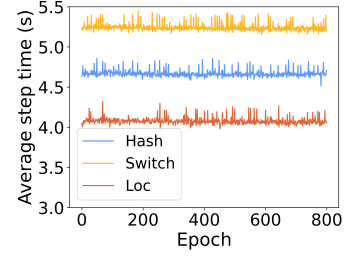


Figure 7: The average time consumption of steps in each epoch with multiple MoEs under 64N.

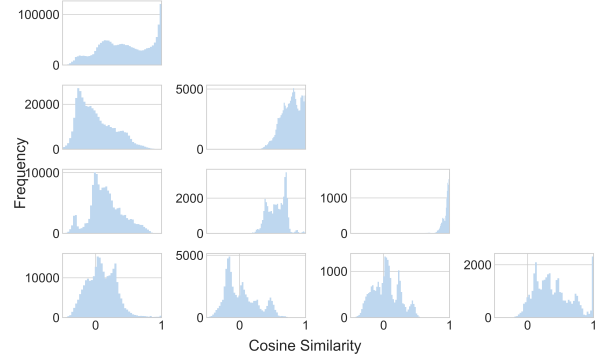


Figure 8: The histograms of cosine similarities between tokens routed to two experts.

and SwitchMoE, are implemented in PanGu- Σ and made contrasts. The average training time with these MoEs under 64N is displayed in Figure 7. It can be seen that LocMoE has an average speedup of $1.15\times$ to $1.29\times$ compared to HashMoE and SwithMoE, respectively.

4.1 Analysis for Expert Capacity

To verify our proof on Lemma 1, the angle between tokens, as well as the angle between the token and the gating weight, are explored, shown in Figure 8 and 9, respectively. Figure 8 at row i and column j plots the distribution of cosine similarities between every pair of tokens routed to the expert i and j , respectively. Specifically, the diagonal ones denote the distributions of tokens from the same expert. It can be seen that the tokens routed to the same expert are more alike, with the cosine similarity closer to 1.

Then, we select a representative expert to discuss the phenomenon further. Figure 9 illustrates the frequency of cosine similarity between tokens and gating weights. The orange bar stands for the cosine distribution of the angle between the token and the gating weight, which corresponds to the expert where the token is routed. The blue bar denotes the distribution of the cosine similarity between the above tokens and another expert. Obviously, most tokens close to the specific expert are indeed routed to it, and the distribution has wide differences with other experts. From our experiments, the δ in Formula 8 is about 0.03.

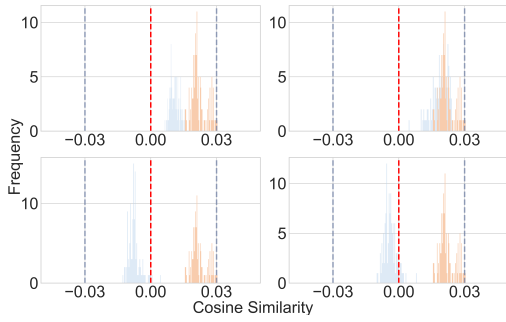


Figure 9: The histograms of cosine similarities between angles between token and the gating weight.

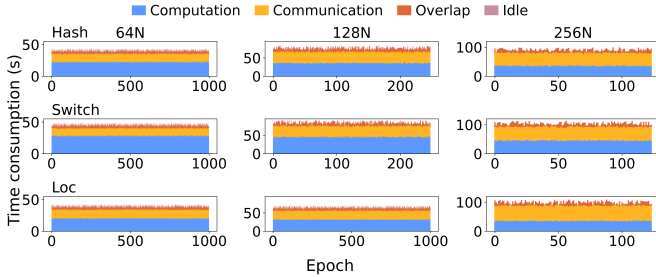


Figure 10: The composition of training time in each epoch under our cluster groups.

4.2 Ablation Analysis

The ablation study is built around aspects including the proportion of computation and communication time, load equalization, and astringency. Moreover, in order to prove that the modifications do not affect the model accuracy, the results for inference are also evaluated for verification.

Proportion of Computation and Communication

We record the total elapsed time per epoch as well as the time consumption for computation, communication, overlapping, and idle with MoEs under different cluster configurations, as shown in Figure 10. Following the analysis of the average time consumption per step, LocMoE has both minimal computation overhead and communication overhead under 64N and 128N. However, under 256N, although LocMoE still has the lowest computation costs, its performance does not surpass the HashMoE. The reason is that load balance is more critical than locality when some devices may not have experts. Due to some of the aforementioned engineering optimizations, the proportion of elapsed computation time of LocMoE slightly fluctuates when the amount of devices increases. Meanwhile, the proportion of communication also rises, and the degree of overlapping becomes deeper.

The overall time consumption proportions are shown in Figure 11, and the impact of our innovations on these operations can be visually detected. LocMoE always has a relatively smaller time proportion of computation and a higher overlapping proportion compared to SwitchMoE. The actual computation time of LocMoE approaches or is a bit lower than HashMoE, which has no extra computation for token features. When the resource increases, the time proportions

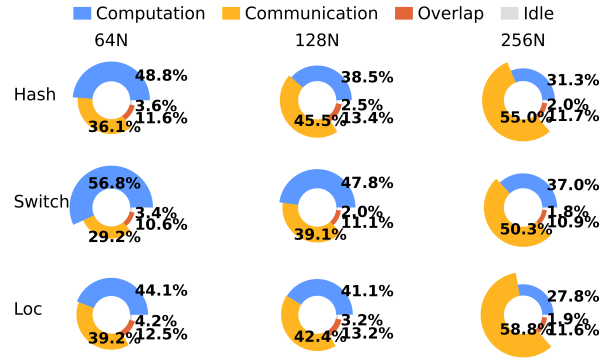


Figure 11: The time consumption ratio with different MoEs under our cluster groups.

of communication for these MoEs all reflect an increasing trend. Specifically, from 64N to 128N, the increasing communication proportion in LocMoE is not as tangible as HashMoE and SwitchMoE. It is shown that the communication time of LocMoE is markedly elevated under 256N with 32 nodes as was expected. The phenomenon indicates that LocMoE is more appropriate for cases whose number of experts is larger than that of nodes. The locality would lose efficacy when the local expert does not exist. Overall, LocMoE offers more notable enhancements in computation.

Distribution of Expert Assignment

Taking the cluster of 64N as an instance, Figure 12 portrays the distribution of expert assignments in different MoEs during the training process. Since HashMoE adopts an absolute balance strategy, the allocation of tokens is quite balanced at initialization. However, SwitchMoE and LocMoE initialize from the allocation to a single expert. To avoid misinterpretation, the analysis begins at epoch 200. The vertical axis indicates the number of tokens assigned to each communication group, and the horizontal axis indicates the index of experts. There are 16 experts in our experiments, and the index range is from 0 to 15. The cumulative number of tokens routed to expert i can be observed along the specific vertical axis corresponding to the expert. Each occurrence of a non-zero value means a new token being assigned to this expert. Successive color bars indicate that shuffled tokens are continuously assigned to the same experts, thus causing imbalances to arise. As can be seen from the figure, the rigid constraints in HashMoE ensure that its assignment is even. However, almost no token is routed to experts with an index of 9 to 15 in the subfigure of SwitchMoE. It results in nearly 40% of the experts' invalidation; to make matters worse, the phenomena of "winner-take-all" is pronounced in expert number 5 and expert number 6. LocMoE, due to the localized bootstrapping, can allocate the token to these experts evenly during the training process, indicating that the dual constraints of auxiliary and locality loss can steadily enhance resource utilization.

Astringency and Accuracy

The astringency is measured by the valid perplexity throughout the process, and the comparison of the convergence speed under different MoEs is depicted in Figure 13. The overall

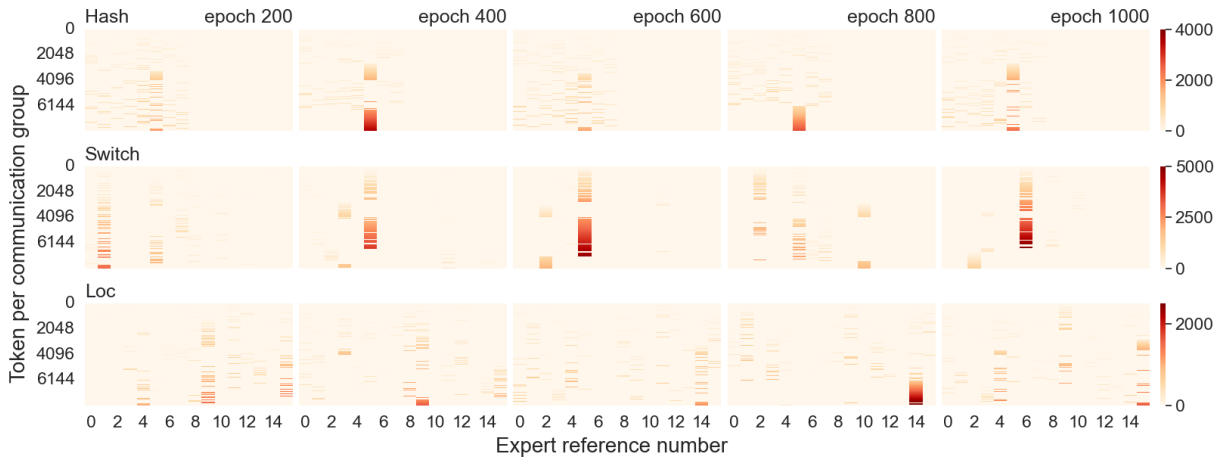


Figure 12: The allocation of tokens with different routing strategies.

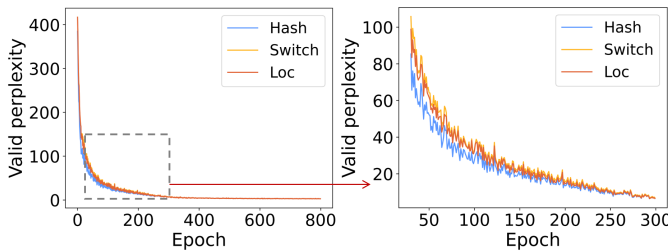


Figure 13: The valid perplexity throughout the training stage of multiple MoEs.

convergence speed of LocMoE is between that of HashMoE and SwitchMoE in the early stage, and they have an analogical tendency of convergence after a certain amount of epochs.

HashMoE exhibits better convergence performance due to the fixed and uniformly assignment of RRE. This phenomenon may be perverse because the unlearned routers make it hard to distinguish experts and converge rapidly. The reason for such a situation may be the relatively small angle between tokens in corpora. Concretely, from Appendix B, the dataset contains the fine-grained classification of materials in a specific domain, and the similarities between these items are inherently high. Thus, the composition of the dataset needs to be ameliorated. As for LocMoE, more experts participate in the early training process due to locality. Compared to SwitchMoE, whose routing probability relies only on the token feature, it may promote astringency using LocMoE.

The performance on multiple NLP tasks (see Appendix E) compared with the original PanGu- Σ is illustrated in Figure 14. All models are pre-trained with the corpora introduced in Appendix B from scratch. The LocMoE and the baseline (original PanGu- Σ) are both more adept at the query type, while they have difficulties with tasks of the fault tree. The samples of query type are displayed in Appendix F. Due to the enhancement of discrimination for experts and tokens, the comprehension and expressive ability of semantics in various tasks is generally improved.

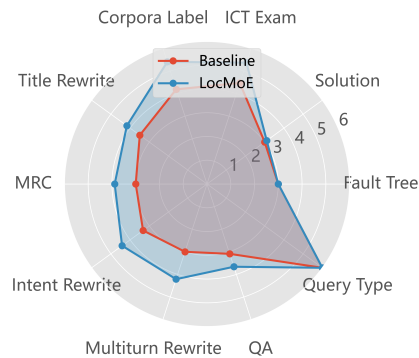


Figure 14: The scores of inference tasks compared with the baseline.

5 Conclusion

In this paper, we propose a low overhead structure named LocMoE to relieve the performance bottleneck of existing MoE. The modifications mainly revolve around the mechanism of token assignment. The locality loss, which can be delineated as the distribution difference of token assignments, is proposed to promote locality computation on the premise of load balance. We also provide the theoretical demonstration for the lower bound of the expert capacity to achieve the same effect by training fewer tokens. To meet the assumption of orthogonal gating weight, the GrAP layer is adopted instead of the dense layer to calculate the gating values, and it can also reduce the overhead of computation. Incorporating group-wise All-to-All and communication overlapping features, the elapsed time of communication is further reduced. The experiments are performed on Ascend clusters with 64, 128, and 256 NPUs. Compared with current state-of-the-art MoEs, the performance improvement of training is up to 22.24%. Evaluating multiple NLP tasks, it is detected that the interactive capability of our model is also enhanced. From the results that explore the relationship between the scale of expert capacity and the token features, we find that the dataset construction still needs to be improved. In future work, we will further organize the multilingual corpora from more fields.

Appendix

A. Proof Sketch in 3.2

A.1 Proof for Lemma 1

Proof. According to the previous definition, $\delta_{i^*,j} = \cos(\theta_{i^*,j})$, where i^* is the expert that the token j routed to. $\theta_{i^*,j}$ is the angle between token j and the gating weight ω_{i^*} corresponding to the expert i^* . Combined with Formula (3) and (4) in Section 3.2, we have:

$$\begin{aligned} i^* &= \arg \max_{i \in [n]} \langle \omega_i, x_m \rangle \\ \text{where } \langle \omega_i, x_m \rangle &= \|\omega_i\| \cdot \|x_m\| \cdot \cos(\theta_{i^*,j}) \\ i^* &= \arg \max_{i \in [n]} \langle \omega_i, x_m \rangle \\ &= \arg \max_{i \in [n]} (\delta_{i^*,j}) \end{aligned} \quad (10)$$

□

A.2 Proof for Lemma 3

Proof. The area of a hyperspherical cap in a n -sphere of radius r can be obtained by integrating the surface area of an $(n-1)$ -sphere of radius $r \sin \theta$ with arc element $r d\theta$ over a great circle arc, that is:

$$\begin{aligned} A_n^{\text{cap}}(r) &= \int_0^\phi A_{n-1}(r \sin \theta) r d\theta \\ &= \frac{2\pi^{(n-1)/2}}{\Gamma(\frac{n-1}{2})} r^{n-1} \int_0^\phi \sin^{n-2} \theta d\theta \\ &= \frac{2\pi^{(n-1)/2}}{\Gamma(\frac{n-1}{2})} r^{n-1} J_{n-2}(\phi) \\ &= \frac{2\pi^{(n-1)/2}}{\Gamma(\frac{n-1}{2})} r^{n-1} \frac{1}{2} B\left(\frac{n-1}{2}, \frac{1}{2}\right) I_{\sin^2 \phi} \left(\frac{n-1}{2}, \frac{1}{2}\right) \\ &= \frac{1}{2} \frac{2\pi^{(n-1)/2}}{\Gamma(\frac{n-1}{2})} r^{n-1} \frac{\Gamma(\frac{n-1}{2}) \Gamma(\frac{1}{2})}{\Gamma(\frac{n}{2})} I_{\sin^2 \phi} \left(\frac{n-1}{2}, \frac{1}{2}\right) \\ &= \frac{1}{2} \frac{2\pi^{n/2}}{\Gamma(\frac{n}{2})} r^{n-1} I_{\sin^2 \phi} \left(\frac{n-1}{2}, \frac{1}{2}\right) \\ &= \frac{1}{2} A_n(r) I_{\sin^2 \phi} \left(\frac{n-1}{2}, \frac{1}{2}\right) \end{aligned} \quad (11)$$

where $A_n(r)$ denotes the area of the high-dimensional sphere. p_δ can be viewed as the proportion of the symmetrical areas formed by θ to that of the entire sphere, shown as Figure 15:

$$\begin{aligned} p_\delta &= \frac{2A_n^{\text{cap}}(r, \theta)}{A_n(r)} \\ &= I_{1-\delta^2} \left(\frac{d-1}{2}, \frac{1}{2}\right) \\ &= 1 - I_{\delta^2} \left(\frac{1}{2}, \frac{d-1}{2}\right) \end{aligned} \quad (12)$$

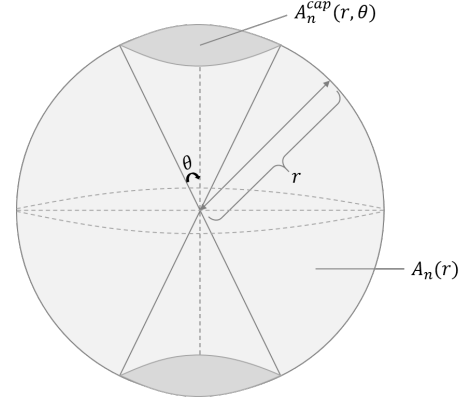


Figure 15: The schematic of p_δ

Suppose $\delta = \sqrt{\frac{1}{d-\frac{3}{2}}}$, when d is large, δ approximates to $\sqrt{\frac{1}{d}}$, then:

$$\begin{aligned} I_{\delta^2} \left(\frac{1}{2}, \frac{d-1}{2}\right) &\approx I\left(\frac{\delta^2(d-1+\frac{1}{2}-1)}{2-\delta^2}, \frac{1}{2}\right) + \Theta\left[\left(\frac{d-1}{2}\right)^{-2}\right] \\ &\approx I\left(\frac{\frac{1}{d-\frac{3}{2}}(d-\frac{3}{2})}{2-\frac{1}{d-\frac{3}{2}}}, \frac{1}{2}\right) \\ &= I\left(\frac{1}{2} \left(\frac{d-\frac{3}{2}}{d-2}\right), \frac{1}{2}\right) \\ &= \frac{1}{\Gamma(\frac{1}{2})} \int_0^{\frac{1}{2} \left(\frac{d-\frac{3}{2}}{d-2}\right)} \exp(-t) t^{\frac{1}{2}} dt \\ &\approx \frac{1}{\Gamma(\frac{1}{2})} \int_0^{\frac{1}{2}} e^{-t} t^{-\frac{1}{2}} dt \\ &= \frac{1}{\Gamma(\frac{1}{2})} \gamma\left(\frac{1}{2}, \frac{1}{2}\right) \\ &= \text{erf}\left(\frac{\sqrt{2}}{2}\right) \end{aligned} \quad (13)$$

where γ is the incomplete gamma function. Combined with Formula (2) in Section 3.2, $\text{erf}\left(\frac{\sqrt{2}}{2}\right) \approx 0.68$, then:

$$p_\delta = 1 - I_{\delta^2} \left(\frac{1}{2}, \frac{d-1}{2}\right) \approx 0.3 \quad (14)$$

□

A.3 Proof for Theorem 1

Proof. Refer to the assumption about distributions of class-discriminative and class-irrelevant patterns in pMoE [Chowdhury *et al.*, 2023], with analogy, the tokens satisfy $\delta_{i,j} \geq \delta$ can be regarded as the class-discriminative token. Then, the problem we need to explore can be converted to find the minimum amount of tokens that make at least one class-discriminative token routed to expert i .

Suppose p_i is the probability that the token routed to the expert i is a class-discriminative token; we have:

Hyperparameter	Description	Value
adam_eps	Terms to increase the stability of numerical calculations	1e-6
batch_size	The size of data input to the model for training each time, related to the number of devices	32
expert_num_per_dp_dim	Number of experts per communication group	1
expert_parallel	Number of experts in parallel	16
moe_layer_num	Number of MoE layers	8
num_heads	Number of parallel heads	40
op_level_model_parallel_num	Number of parallel models	8
sink_size	The size of data executed per sink	16

 Table 1: The critical hyperparameters in configuration of PanGu- Σ .

$$p_i \leq \frac{p_\delta}{\frac{1}{n}} = n[1 - I_{\delta^2}(\frac{1}{2}, \frac{d-1}{2})] \quad (15)$$

where the first inequality holds since the token satisfies $\delta_{i,j} \geq \delta$ may not always be routed to expert i . Then, the minimum value of expert capacity under the circumstance of at least one class-discriminative token routed to expert i can be written as:

$$\begin{aligned}
 ec_{\min} &= \frac{1}{p_s} = \frac{1}{n[1 - I_{\delta^2}(\frac{1}{2}, \frac{d-1}{2})]} \\
 \text{For large } d, I_{\delta^2}(\frac{1}{2}, \frac{d-1}{2}) &\approx I(\frac{\delta^2(d-\frac{3}{2})}{2-\delta^2}, \frac{1}{2}) \\
 &\approx \frac{1}{\Gamma(\frac{1}{2})} \gamma(\frac{1}{2}, \frac{\delta^2 d}{2-\delta^2}) \\
 &= 1 - \operatorname{erfc}(\sqrt{\frac{\delta^2 d}{2-\delta^2}}) \\
 \text{thus, } ec_{\min} &\geq \frac{1}{n \cdot \operatorname{erfc}(\sqrt{\frac{\delta^2 d}{2-\delta^2}})} \\
 &> \frac{1}{n} \exp(\frac{\delta^2 d}{2-\delta^2})
 \end{aligned} \quad (16)$$

□

B. Datasets

PanGu- Σ has already demonstrated its ability to learn efficiently and independently from text corpus in various domains. In this work, we will evaluate the performance of PanGu- Σ in detailed knowledge of a specific area. The materials connected to mobile network operators' services are chosen as input corpora. Concretely, blogs and technical documents in the form of *iCase*, *Wiki*, core network/Man-Machine language (MML), configuration translations, feature documents, etc., are collected. These corpora are in Chinese, English, or bilingual (Chinese-English).

C. Experimental Environment

The experiments are conducted on Ascend clusters, and the environment falls into three groups: 64, 128, and 256 Ascend 910A NPUs. The Ascend 910A series NPU has 32 AI Cores, with a maximum memory capacity of 2TB and a maximum

memory bandwidth of 1.07TB/s. The collective communication function on high-speed links such as PCI-E, HCCS, and RoCE is realized by HCCL, a high-performance collective communication library based on the Ascend. It provides communication primitives on single-node-multi-card and multi-node-multi-card, and it also supports various communication algorithms such as ring, mesh, HD, ring + HD, and mesh + HD.

The versions of the Compute Architecture for Neural Networks (CANN) suite (toolkit, CANN, driver) are 5.1.RC2.1, 1.84, and 23.0.rc2, respectively. The CANN is the heterogeneous computing architecture developed by Huawei, and it supports multiple AI frameworks, including MindSpore, PyTorch, TensorFlow, etc., providing interfaces to build AI applications on the Ascend platform. Our model runs on the MindSpore framework with version 2.0.0.

D. Model Configuration

The hyperparameter configuration of our model is listed in Table 1. Thereinto, *batch_size* and *sink_size* are relevant to the number of devices, and the values in the table are under 128N. The total number of experts can be obtained by *expert_num_per_dp_dim* * *expert_parallel*.

E. Measurement Metric

We design multiple NLP tasks to systematically evaluate the knowledge understanding and semantic expression capabilities of our model. These tasks are extracted from 10 business perspectives in the field of carrier networks, such as fault tree nodes, solutions, ICT certification exams, and title rewriting. Among them, taking the recognition fault tree node as an example, the construction of the NLP task is divided into two steps: firstly, the text is differentiated into difficulty levels (L1 to L3) according to the logical complexity of concepts and inter-conceptual relationships, and the samples are selected according to the hierarchies. L1 contains single and a group of connection parameters with integrity and independence, complete temporal connection parameters; L2 represents quantitative relationship parameters, referential relationship parameters, and combination parameters; L3 denotes the sample that cannot be intellectualized. In the next step, after the classification is completed, the prompt, derived from the structured specification of the fault discriminative approach, is applied to generate structured parameters and restores the discrimination logic.

Acknowledgments

We thank all anonymous reviewers for their valuable feedback. We would like to express our appreciation to Dachao Lin (Dr. Lin) for the improvements in theoretical proof in this paper. We are grateful for the guidance from the Noah's Ark Lab on model training. Moreover, the contributions of the MindSpore team and employees participating in model evaluation are greatly appreciated.

Contribution Statement

Jing Li, Zhijie Sun, and Xuan He wrote this paper and contributed equally to this work. Li Zeng, Yi Lin, Entong Li, and Binfan Zheng provided experimental analysis and offered suggestions for this paper. Rongqian Zhao and Xin Chen are the project leaders and provide support for this work.

References

- [Bi *et al.*, 2023] Kaifeng Bi, Lingxi Xie, Hengheng Zhang, Xin Chen, Xiaotao Gu, and Qi Tian. Accurate medium-range global weather forecasting with 3d neural networks. *Nature*, 619(7970):533–538, 2023.
- [Brown *et al.*, 2020] Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901, 2020.
- [Chi *et al.*, 2022] Zewen Chi, Li Dong, Shaohan Huang, Damai Dai, Shuming Ma, Barun Patra, Saksham Singhal, Payal Bajaj, Xia Song, Xian-Ling Mao, et al. On the representation collapse of sparse mixture of experts. *Advances in Neural Information Processing Systems*, 35:34600–34613, 2022.
- [Chowdhury *et al.*, 2023] Mohammed Nowaz Rabbani Chowdhury, Shuai Zhang, Meng Wang, Sijia Liu, and Pin-Yu Chen. Patch-level routing in mixture-of-experts is provably sample-efficient for convolutional neural networks. *arXiv preprint arXiv:2306.04073*, 2023.
- [Clark *et al.*, 2022] Aidan Clark, Diego De Las Casas, Aurelia Guy, Arthur Mensch, Michela Paganini, Jordan Hoffmann, Bogdan Damoc, Blake Hechtman, Trevor Cai, Sebastian Borgeaud, et al. Unified scaling laws for routed language models. In *International Conference on Machine Learning*, pages 4057–4086. PMLR, 2022.
- [Dai *et al.*, 2022] Damai Dai, Li Dong, Shuming Ma, Bo Zheng, Zhifang Sui, Baobao Chang, and Furu Wei. Stablemoe: Stable routing strategy for mixture of experts. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7085–7095, 2022.
- [Fedus *et al.*, 2022] William Fedus, Barret Zoph, and Noam Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *The Journal of Machine Learning Research*, 23(1):5232–5270, 2022.
- [Guo *et al.*, 2020] Yuanyuan Guo, Yifan Xia, Jing Wang, Hui Yu, and Rung-Ching Chen. Real-time facial affective computing on mobile devices. *Sensors*, 20(3):870, 2020.
- [He *et al.*, 2021] Jiaao He, Jiezhong Qiu, Aohan Zeng, Zhilin Yang, Jidong Zhai, and Jie Tang. Fastmoe: A fast mixture-of-expert training system. *arXiv preprint arXiv:2103.13262*, 2021.
- [Jacobs *et al.*, 1991] Robert A Jacobs, Michael I Jordan, Steven J Nowlan, and Geoffrey E Hinton. Adaptive mixtures of local experts. *Neural computation*, 3(1):79–87, 1991.
- [Kenton and Toutanova, 2019] Jacob Devlin Ming-Wei Chang Kenton and Lee Kristina Toutanova. Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of NAACL-HLT*, pages 4171–4186, 2019.
- [Kudugunta *et al.*, 2021] Sneha Kudugunta, Yanping Huang, Ankur Bapna, Maxim Krikun, Dmitry Lepikhin, Minh-Thang Luong, and Orhan Firat. Beyond distillation: Task-level mixture-of-experts for efficient inference. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3577–3599, 2021.
- [Kumar *et al.*, 2022] N Ashok Kumar, A Kavitha, P Venkatramana, and Durgesh Nandan. Architecture design: Network-on-chip. In *VLSI Architecture for Signal, Speech, and Image Processing*, pages 147–165. Apple Academic Press, 2022.
- [Lepikhin *et al.*, 2020] Dmitry Lepikhin, HyoukJoong Lee, Yuanzhong Xu, Dehao Chen, Orhan Firat, Yanping Huang, Maxim Krikun, Noam Shazeer, and Zhifeng Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. In *International Conference on Learning Representations*, 2020.
- [Lewis *et al.*, 2021] Mike Lewis, Shruti Bhosale, Tim Dettmers, Naman Goyal, and Luke Zettlemoyer. Base layers: Simplifying training of large, sparse models. In *International Conference on Machine Learning*, pages 6265–6274. PMLR, 2021.
- [Li *et al.*, 2022a] Jianjun Li, Yu Han, Ming Zhang, Gang Li, and Baohua Zhang. Multi-scale residual network model combined with global average pooling for action recognition. *Multimedia Tools and Applications*, pages 1–19, 2022.
- [Li *et al.*, 2022b] Yuan Li, Ke Wang, Hao Zheng, Ahmed Louri, and Avinash Karanth. Ascend: A scalable and energy-efficient deep neural network accelerator with photonic interconnects. *IEEE Transactions on Circuits and Systems I: Regular Papers*, 69(7):2730–2741, 2022.
- [Liao *et al.*, 2021] Heng Liao, Jiajin Tu, Jing Xia, Hu Liu, Xiping Zhou, Honghui Yuan, and Yuxing Hu. Ascend: a scalable and unified architecture for ubiquitous deep neural network computing: Industry track paper. In *2021 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*, pages 789–801. IEEE, 2021.

- [Mi *et al.*, 2022] Fei Mi, Yitong Li, Yulong Zeng, Jingyan Zhou, Yasheng Wang, Chuanfei Xu, Lifeng Shang, Xin Jiang, Shiqi Zhao, and Qun Liu. Pangu-bot: Efficient generative dialogue pre-training from pre-trained language model. *arXiv preprint arXiv:2203.17090*, 2022.
- [Nie *et al.*, 2022] Xiaonan Nie, Pinxue Zhao, Xupeng Miao, Tong Zhao, and Bin Cui. Hetumoe: An efficient trillion-scale mixture-of-expert distributed training system. *arXiv preprint arXiv:2203.14685*, 2022.
- [Nie *et al.*, 2023] Xiaonan Nie, Xupeng Miao, Zilong Wang, Zichao Yang, Jilong Xue, Lingxiao Ma, Gang Cao, and Bin Cui. Flexmoe: Scaling large-scale sparse pre-trained model training via dynamic device placement. *Proceedings of the ACM on Management of Data*, 1(1):1–19, 2023.
- [Puigcerver *et al.*, 2020] Joan Puigcerver, Carlos Riquelme Ruiz, Basil Mustafa, Cedric Renggli, Andr’e Susano Pinto, Sylvain Gelly, Daniel Keysers, and Neil Houlsby. Scalable transfer learning with expert models. In *International Conference on Learning Representations*, 2020.
- [Rajbhandari *et al.*, 2022] Samyam Rajbhandari, Conglong Li, Zhewei Yao, Minjia Zhang, Reza Yazdani Aminabadi, Ammar Ahmad Awan, Jeff Rasley, and Yuxiong He. Deepspeed-moe: Advancing mixture-of-experts inference and training to power next-generation ai scale. In *International Conference on Machine Learning*, pages 18332–18346. PMLR, 2022.
- [Ren *et al.*, 2023] Xiaozhe Ren, Pingyi Zhou, Xinfan Meng, Xinjing Huang, Yadao Wang, Weichao Wang, Pengfei Li, Xiaoda Zhang, Alexander Podolskiy, Grigory Arshinov, et al. Pangu- σ : Towards trillion parameter language model with sparse heterogeneous computing. *arXiv e-prints*, pages arXiv–2303, 2023.
- [Roller *et al.*, 2021] Stephen Roller, Sainbayar Sukhbaatar, Jason Weston, et al. Hash layers for large sparse models. *Advances in Neural Information Processing Systems*, 34:17555–17566, 2021.
- [Shazeer *et al.*, 2016] Noam Shazeer, Azalia Mirhoseini, Krzysztof Maziarczyk, Andy Davis, Quoc Le, Geoffrey Hinton, and Jeff Dean. Outrageously large neural networks: The sparsely-gated mixture-of-experts layer. In *International Conference on Learning Representations*, 2016.
- [Shen *et al.*, 2023] Bo Shen, Jiabin Zhang, Taihong Chen, Daoguang Zan, Bing Geng, An Fu, Muhan Zeng, Ailun Yu, Jichuan Ji, Jingyang Zhao, et al. Pangu-coder2: Boosting large language models for code with ranking feedback. *arXiv preprint arXiv:2307.14936*, 2023.
- [Tong *et al.*, 2021] Zhihao Tong, Ning Du, Xiaobo Song, and Xiaoli Wang. Study on mindspore deep learning framework. In *2021 17th International Conference on Computational Intelligence and Security (CIS)*, pages 183–186. IEEE, 2021.
- [Touvron *et al.*, 2023] Hugo Touvron, Thibaut Lavril, Gautier Izacard, Xavier Martinet, Marie-Anne Lachaux, Timothée Lacroix, Baptiste Rozière, Naman Goyal, Eric Hambro, Faisal Azhar, et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023.
- [Vaswani *et al.*, 2017] 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.
- [Wang *et al.*, 2023] Yunhe Wang, Hanting Chen, Yehui Tang, Tianyu Guo, Kai Han, Ying Nie, Xutao Wang, Hailin Hu, Zheyuan Bai, Yun Wang, et al. Pangu- π : Enhancing language model architectures via nonlinearity compensation. *arXiv preprint arXiv:2312.17276*, 2023.
- [Xia *et al.*, 2021] Jing Xia, Chuanning Cheng, Xiping Zhou, Yuxing Hu, and Peter Chun. Kunpeng 920: The first 7-nm chiplet-based 64-core arm soc for cloud services. *IEEE Micro*, 41(5):67–75, 2021.
- [Zeng *et al.*, 2021] Wei Zeng, Xiaozhe Ren, Teng Su, Hui Wang, Yi Liao, Zhiwei Wang, Xin Jiang, ZhenZhang Yang, Kaisheng Wang, Xiaoda Zhang, et al. Pangu- α : Large-scale autoregressive pretrained chinese language models with auto-parallel computation. *arXiv preprint arXiv:2104.12369*, 2021.
- [Zuo *et al.*, 2021] Simiao Zuo, Xiaodong Liu, Jian Jiao, Young Jin Kim, Hany Hassan, Ruofei Zhang, Jianfeng Gao, and Tuo Zhao. Taming sparsely activated transformer with stochastic experts. In *International Conference on Learning Representations*, 2021.