Protocol Models: Scaling Decentralized Training with Communication-Efficient Model Parallelism

Sameera Ramasinghe, Thalaiyasingam Ajanthan, Gil Avraham, Yan Zuo, and Alexander Long

Pluralis Research

Scaling models has led to significant advancements in deep learning, but training these models in decentralized settings remains challenging due to communication bottlenecks. While existing compression techniques are effective in data-parallel, they do not extend to model parallelism. Unlike data-parallel training, where weight gradients are exchanged, model-parallel requires compressing activations and activation gradients as they propagate through layers, accumulating compression errors. We propose a novel compression algorithm that compresses both forward and backward passes, enabling up to 99% compression with no convergence degradation with negligible memory/compute overhead. By leveraging a recursive structure in transformer networks, we predefine a low-dimensional subspace to confine the activations and gradients, allowing full reconstruction in subsequent layers. Our method achieves up to $100 \times$ improvement in communication efficiency and enables training billion-parameter-scale models over low-end GPUs connected via consumer-grade internet speeds as low as 80Mbps, matching the convergence of centralized datacenter systems with 100Gbps connections with model parallel.

Keywords: Protocol Models, Low-rank Compression, Pipeline Parallelism, Decentralized Training

1. Introduction

Scaling models and datasets has been pivotal in driving deep learning advancements, with model sizes expanding from millions of parameters (Krizhevsky et al., 2012) to billions (Dubey et al., 2024; Kolesnikov et al., 2020) and even trillions (Ren et al., 2023). These larger models exceed the memory capacity of a single device, requiring distributed training approaches to manage computation across multiple devices.

A common solution is distributed data parallelism (DDP) (Li et al., 2020) or its more advanced variant, fully sharded data parallelism (FSDP) (Zhao et al., 2023), which distributes data across nodes while replicating the model on each device. This enables larger batch sizes and higher throughput, but constrains the model size to the memory of a single device. Model parallelism (MP) addresses this limitation by distributing parameters across devices (Huang et al., 2019; Krizhevsky et al., 2012), enabling the training of models that surpass single-node memory constraints. MP includes tensor parallelism, which splits individual layers, and pipeline parallelism, which distributes layers across devices; the latter is the focus of this work. Modern large-scale training combines DDP and MP to achieve scalability. Despite these strategies, all approaches require the transfer of large amounts of data between devices (Narayanan et al., 2021), limiting training to high-performance computing clusters with fast interconnects. These infrastructures are costly and accessible only to resource-rich organizations (Power and Progress; Vergara Larrea et al., 2019), creating disparities that restrict broader research and risk centralizing innovation.

 $Corresponding \ author(s): \ same era@puralis.ai$

^{© 2025} Pluralis Research. All rights reserved.

Decentralized training provides an alternative by leveraging consumer-grade devices, enabling individuals with small devices to participate in large-scale model training, reducing reliance on major corporations. This approach democratizes access to large-scale model training by leveraging underutilized GPUs in personal computers and volunteer networks (Ryabinin et al., 2023; Yuan et al., 2022). However, decentralized training faces significant challenges due to limited bandwidth and high latency in heterogeneous networks (Yuan et al., 2022), necessitating communication-efficient compression algorithms to minimize data transfer while preserving training performance.

Most existing communication-efficient techniques focus on DDP (Douillard et al., 2023; Lian et al., 2017; Peng et al., 2024; Ryabinin et al., 2021), where *weight gradients* are computed independently on each node (for each model replica) and then compressed before peer-to-peer communication. To this end, techniques such as sparsification (Lin et al., 2017; Wang et al., 2017; Wangni et al., 2018), quantization (Alistarh et al., 2017; Bernstein et al., 2018; Wang et al., 2023), and low-rank approximations (Zhao et al., 2024) exploit the redundancy in weight gradients to reduce communication. However, in MP, information must be passed between *layers*, requiring the communication of *activations* and *activation gradients*. Unlike weight gradients, activations lack inherent redundancy and approximation errors accumulate across layers, leading to degraded convergence (Bian et al., 2024; Rudakov et al., 2023). These challenges prevent the straightforward application of DDP compression techniques to MP, and hence to date, MP decentralized training remains infeasible, resulting in massive slowdowns over centralized training.

To bridge this gap, we propose a novel compression algorithm tailored for MP. We show that as training progresses, weight matrices exhibit rank collapse, converging to low-rank subspaces. Thus, by *explicitly* constraining specific weight matrices to such low-rank subspaces and leveraging a recursive structure inherent in transformer networks, we demonstrate that transformer layer activations—despite their high rank—can be decomposed into a dynamic low-rank component and a static high-rank component. This decomposition enables efficient compression of information passed between layers during both the forward and backward passes, ensuring *lossless* reconstruction in subsequent layers.

We validate the practical effectiveness of our approach through extensive evaluations on billionparameter-scale models. Our compression method enables the distribution of large-scale models across consumer-grade GPUs with internet-grade connections (80Mbps) while matching the convergence performance of centralized setups with 100 Gbps interconnects. We achieve up to $100 \times$ improvement in communication efficiency without any degradation in convergence. Further, we successfully train an 8B-parameter LLaMA model (Dubey et al., 2024) with layers split across four different geographical regions, connected via the internet, and achieve convergence on par with baselines utilizing datacenter-grade connections. By addressing critical limitations in decentralized training, our method intend to remove significant barriers to scaling large models in resource-constrained environments, democratizing access to large-scale deep learning.

2. Related works

Decentralized training involves a group of *autonomous* devices (*i.e.*, no central orchestrator) collaborating to train a large-scale model by leveraging MP/DDP methods. These devices, often geographically distributed and heterogeneous, are connected via networks with varying delays and bandwidth constraints. Key advancements in this area encompass both theoretical insights (Koloskova et al., 2019; 2020; Lian et al., 2017) as well as practical approaches (Diskin et al., 2021; Ryabinin and Gusev, 2020). Despite the progress, current efforts predominantly focus on DDP (Diskin et al., 2021; Koloskova et al., 2019; 2020; Lian et al., 2017), which hinders scaling models beyond the memory capacity of local nodes. SWARM parallelism (Ryabinin et al., 2023) and Tasklets (Yuan

et al., 2022), treat the problem as a scheduling challenge, with the former addressing the stochasticity inherent in decentralized cluster settings. However, all methods to date still face significant scalability challenges due to communication bottlenecks inherent in MP. Our method overcomes this limitation by introducing an effective communication compression technique for MP (specifically pipeline parallel), mitigating a major barrier in scaling decentralized training.

Communication compression accelerates distributed training over bandwidth-limited networks by reducing data transfers. Key strategies include **sparsification**, which transmits only significant parameter updates (Lin et al., 2017; Wang et al., 2017; Wangni et al., 2018); **quantization**, which lowers communication by reducing parameter precision (Alistarh et al., 2017; Bernstein et al., 2018; Dettmers et al., 2021; Karimireddy et al., 2019; Tang et al., 2021; Wang et al., 2023; Wu et al., 2018); and **low-rank projection**, which compresses gradients via projection onto lower-dimensional subspaces (Vogels et al., 2019; Zhao et al., 2024). While successful in data-parallel settings (DDP), these techniques face difficulties in model-parallel (MP) setups, including error accumulation across layers and unstructured activations (Bian et al., 2024; Rudakov et al., 2023), causing degraded convergence. Recent work by (Wang et al., 2021) proposed low-rank MP communication, but required significant architectural changes, preventing training from scratch. In contrast, our approach involves minimal initialization changes without architectural modifications, enabling efficient MP training from scratch with improved scalability.

3. Transformer block

We provide a brief exposition of the transformer block and proceed to describe the proposed compression method. Let the input to the l^{th} layer be $\mathbf{X}^l \in \mathbb{R}^{b \times n \times d}$, where b, n, and d are the batch size, sequence length, and embedding dimension, respectively. Given the weight matrices of each attention head h as $\mathbf{W}_{Q,h}^l, \mathbf{W}_{K,h}^l, \mathbf{W}_{V,h}^l \in \mathbb{R}^{d \times d_H}$, with $d_H = d/H$ where H is the number of attention heads, the following computations are performed for each attention head: $\mathbf{X}_{Q,h}^l = \mathbf{X}^l \mathbf{W}_{Q,h}^l$, $\mathbf{X}_{K,h}^l = \mathbf{X}^l \mathbf{W}_{K,h}^l, \mathbf{X}_{V,h}^l = \mathbf{X}^l \mathbf{W}_{V,h}^l$. The rest of the computations are as follows:

$$\mathbf{X}_{\text{head, h}}^{l} = f_{\text{softmax}} \left(\frac{\mathbf{X}_{Q,h}^{l} \mathbf{X}_{K,h}^{\top}}{\sqrt{d_{H}}} \right) \mathbf{X}_{V,h}^{l}$$
(1)

$$\begin{aligned} \mathbf{X}_{\text{concat}}^{l} &= [\mathbf{X}_{\text{head}, 1}^{l}, \dots, \mathbf{X}_{\text{head}, H}^{l}] \\ \mathbf{X}_{\text{attn}}^{l} &= \mathbf{X}_{\text{concat}}^{l} \mathbf{W}_{p_{1}}^{l} + \mathbf{X}^{l} \\ \mathbf{X}_{\text{hidden}}^{l} &= f_{\text{relu}}(\mathbf{X}_{\text{attn}}^{l} \mathbf{W}_{1}^{l}) \\ \mathbf{X}^{l+1} &= \mathbf{X}_{\text{hidden}}^{l} \mathbf{W}_{p_{2}}^{l} + \mathbf{X}_{\text{attn}}^{l} \end{aligned}$$
(2)

where $\mathbf{W}_{p_1}^l \in \mathbb{R}^{d \times d}$, $\mathbf{W}_1^l \in \mathbb{R}^{d \times d_{\mathrm{ff}}}$, and $\mathbf{W}_{p_2}^l \in \mathbb{R}^{d_{\mathrm{ff}} \times d}$. We omit the layer norms for brevity, which does not affect any of our derivations. d_{ff} is usually an integer multiple of d. We will refer to $\mathbf{W}_{p_2}^l$ and $\mathbf{W}_{p_1}^l$ as projection matrices from here onward.

4. Subspace networks

4.1. Rank collapse of projection matrices

We leverage the observation that weight gradients of projection matrices inherently reside in a lowdimensional subspace, as widely reported across various architectures (Cosson et al., 2023; Gur-Ari



Figure 1: **Rank collapse in projection matrices.** Consistent with Statements 7.2 and 7.3 (and Theorems C.2, D.1 - Appendix), we empirically observe a natural rank collapse in the projection matrices (of non-compressed models). Shown is an 8-layer, 2B-parameter model, with the stable (effective) ranks of the projection matrices for the 4th (middle) and 7th (penultimate) layers plotted over training steps.

et al., 2018; Li et al., 2024b; Vogels et al., 2019; Wang et al., 2018; Yang et al., 2023; Zhang et al., 2024; Zhao et al., 2021) (see Appendix C.1 for empirical validation). Together with AdamW's decoupled weight decay—which suppresses negligible gradient components—this drives projection matrices to converge toward a subspace spanned by dominant gradients, resulting in effectively low-rank structures. Formal treatment is provided in Statements 7.2, 7.3, and their associated theorems.

To validate this phenomenon, we train an 8-layer, 2B-parameter model on WikiText (Merity et al., 2016) and track the stable rank of its projection layers during training (hidden dimension of 4096 and a context length of 2048. Stable rank is computed as $\sum_i \sigma_i^2 / \max_i(\sigma_i^2)$, where σ_i denotes the *i*th singular value. As shown in Fig. 1, the stable rank of projection matrices sharply declines, consistent with our theory. This structure, combined with the recursive nature of transformers, allows us to design a nearly-lossless low-rank compression scheme. While prior works observe similar rank collapse in self-attention matrices (Abbe et al., 2024; Dong et al., 2021; Sanyal et al.), we focus specifically on projection matrices as the basis for compression. We also validate this using official checkpoints of fully-trained large scale models (Appendix G).

4.2. Investigating the activation structure

We utilize the natural rank collapse we discussed thus far for compression. Observing the transformer block in Eq. 1 and 2, we identify that a recursive structure emerges for layer outputs due to the skip connections:

$$\mathbf{X}^{l+1} = \mathbf{X}_{\text{hidden}}^{l} \mathbf{W}_{p_2}^{l} + \mathbf{X}_{\text{concat}}^{l} \mathbf{W}_{p_1}^{l} + \mathbf{X}^{l}.$$
(3)

which can be expanded as:

$$\mathbf{X}^{l+1} = \sum_{i=1}^{l} (\mathbf{X}^{i}_{\text{hidden}} \mathbf{W}^{i}_{p_{2}} + \mathbf{X}^{i}_{\text{concat}} \mathbf{W}^{i}_{p_{1}}) + \mathbf{X}^{0}$$
(4)

Here,

$$\mathbf{X}^0 = \mathrm{PE} + \mathrm{TE},\tag{5}$$

where $PE, TE \in \mathbb{R}^{b \times n \times d}$ represent the positional and token embeddings, respectively. Let $[p_1, p_2, \ldots, p_n]$ be the sequence of positional indices, and $[t_1, t_2, \ldots, t_n]$ be the corresponding token indices. We

denote embedding matrices $\mathbf{P} \in \mathbb{R}^{n \times d}$ for positional embeddings and $\mathbf{T} \in \mathbb{R}^{v \times d}$ for token embeddings, where v is the vocabulary length. Thus the tokens and positional indices are embedded via a lookup: $PE = \mathbf{P}[p_{1:n}, :]$ and $TE = \mathbf{T}[t_{1:n}, :]$. Observing Eq. 4, we note that $\mathbf{X}^0 = PE + TE$ contributes as a common additive term to all layer outputs. Therefore, we consider the rank of the residual activations when PE and TE are subtracted:

$$\hat{\mathbf{X}}^{l+1} = \mathbf{X}^{l+1} - \text{PE} - \text{TE} = \sum_{i=1}^{l} (\mathbf{X}^{i}_{\text{hidden}} \mathbf{W}^{i}_{p_{2}} + \mathbf{X}^{i}_{\text{concat}} \mathbf{W}^{i}_{p_{1}})$$
(6)

Recall that $\operatorname{Row}(\mathbf{AB}) \subseteq \operatorname{Row}(\mathbf{B})$, for any two matrices \mathbf{A}, \mathbf{B} where $\operatorname{Row}(\cdot)$ denotes the row space. Thus, it is clear that if the rows of the projection weights $(\mathbf{W}_{p_2}, \mathbf{W}_{p_1})$ up to layer l span a common low-dimensional subspace, then the rows of $\hat{\mathbf{X}}^{l+1}$ is also restricted to the same subspace, since vector spaces are closed under addition.

4.3. Compressing the forward pass

Recall that our analysis so far indicates that the projection matrices naturally confine themselves to a smaller subspace as training progresses. Consequently, the residual activations $\hat{\mathbf{X}}^{l+1}$ are also restricted to a smaller subspace, if the union of those subspaces is low-dimensional. Based on this insight, it is intriguing to explore the feasibility of *explicitly forcing* the rows of projection matrices to vary within a common low-dimensional subspace \mathcal{S} , throughout training, to facilitate activation compression. As discussed in Section 4.2, this forces the rows of activation outputs ($\hat{\mathbf{X}}^{l+1}$) to span the same subspace \mathcal{S} . Surprisingly, we find that even with extreme low-dimensional \mathcal{S} , the networks can achieve almost the same convergence rates as in the unaltered ones. This allows us to significantly reduce the communication between blocks, which we show next.

Let $\operatorname{Row}(\mathbf{W}_{p_2}^l)$, $\operatorname{Row}(\mathbf{W}_{p_1}^l) \subseteq \mathcal{S}$. Further, Let $\mathbf{U}_k \in \mathbb{R}^{d \times k}$ be a matrix with orthonormal columns and $\operatorname{Col}(\mathbf{U}_k) = \mathcal{S}$. Then, the following holds:

$$\hat{\mathbf{X}}^{l+1} = \hat{\mathbf{X}}^{l+1} \mathbf{U}_k \mathbf{U}_k^{\top} \tag{7}$$

In other words, \hat{X}^{l+1} remains unaltered by the projection, since it is already in S. The above formulation introduces a property that can be leveraged for compression during the forward pass. Specifically, the dimensionality of $\hat{\mathbf{X}}^{l+1}\mathbf{U}_k \in \mathbb{R}^{b \times n \times k}$, is substantially smaller than that of $\hat{\mathbf{X}}^{l+1} \in \mathbb{R}^{b \times n \times d}$ since $k \ll d$. If each node in a distributed system is initialized with a shared copy of \mathbf{U}_k , this matrix does not need to be transmitted repeatedly. Instead, $\hat{\mathbf{X}}^{l+1}\mathbf{U}_k$ can be transmitted to the next node, and the original \mathbf{X}^{l+1} is reconstructed as:

$$\mathbf{X}^{l+1} = \hat{\mathbf{X}}^{l+1} \mathbf{U}_k^{\top} + \text{PE} + \text{TE}.$$

This approach ensures exact recovery of \mathbf{X}^{l+1} without approximation.

4.3.1. Decomposition of high-rank components

For practical utilization of above approach, we still need to subtract PE and TE from \mathbf{X}^{l+1} to compute $\hat{\mathbf{X}}^{l+1}$ (and add them back in the next layer). While PE is deterministic and can be computed locally within each node, TE varies depending on the batch, making it impossible to do so.

One potential solution is restricting the embedding table \mathbf{T} also to \mathcal{S} . However, we observed that this degrades network performance due to severely limiting the representation capacity of the token embeddings. Instead, we propose modeling TE as a composition of a fixed high-rank component and a trainable low-rank component:

$$TE = \mathbf{T}_{fixed}[t_{1:n}, :] + \mathbf{T}_{\mathcal{S}}[t_{1:n}, :]$$

where we obtain a trainable low rank embedding table $\mathbf{T}_{\mathcal{S}} = \mathbf{T}_{\text{fixed}} \mathbf{U}_k \mathbf{U}_k^T$. At the beginning of training, $\mathbf{T}_{\text{fixed}}$ is transmitted to all nodes and stored. During the forward pass, we compress the activations as:

$$\mathbf{X}_{\text{compressed}}^{l+1} = \left(\hat{\mathbf{X}}^{l+1} + \mathbf{T}_{\mathcal{S}}[t_{1:n}, :] \right) \mathbf{U}_{k}$$
$$= \left(\mathbf{X}^{l+1} - \text{PE} - \mathbf{T}_{\text{fixed}}[t_{1:n}, :] \right) \mathbf{U}_{k}, \tag{8}$$

Note that in Eq. 8 both PE and $\mathbf{T}_{\text{fixed}}[t_{1:n},:]$ —which are generally high-rank—are subtracted from X^{l+1} so the remaining $(\mathbf{X}^{l+1} - \text{PE} - \mathbf{T}_{\text{fixed}}[t_{1:n},:])$ is already in S and is low-rank. Further, since $\text{Row}(\mathbf{T}_{S}[t_{1:n},:]) \subseteq \text{Col}(\mathbf{U}_{k})$, it is implicitly captured in the compressed activations $\mathbf{X}^{l+1}_{\text{compressed}}$ Reconstruction at the next node is then performed as:

$$\mathbf{X}_{\text{recovered}}^{l+1} = \mathbf{X}_{\text{compressed}}^{l+1} \mathbf{U}_{k}^{\top} + \text{PE} + \mathbf{T}_{\text{fixed}}[t_{1:n}, :] = \mathbf{X}^{l+1}$$

ensuring a lossless recovery of \mathbf{X}^{l+1} while not compromising its high ranked-ness.

A natural question arises: would explicitly restricting the projection matrices to a fixed S, instead of allowing this property to emerge organically, adversely affect convergence? Note that this is a form of constraint optimization and there are well known convergence guarantees. However, in Sec. 7, we provide a convergence proof (to at least a first-order stationary point) for completeness for the above partial projection case. Furthermore, we conduct extensive experiments across a variety of settings to empirically validate the convergence. Further, since the fixed embeddings are ephemeral, they have a negligible effect on the peak GPU memory (Appendix 8.8)

4.4. Compression in backpropagation

In the previous section, we showed that constraining the rows of projection matrices to a shared low-dimensional subspace S, coupled with decomposing the embedding table into low-rank and highrank components, facilitates compression of activations in the forward pass. This same constraint naturally facilitates lossless gradient compression in the backward pass. Specifically, let $\nabla_L(\mathbf{X}^{l+1})$ denote the gradient of \mathbf{X}^{l+1} , with respect to the loss, that needs to be propagated to the previous layer. This gradient can be compressed as:

$$\left(\nabla_L(\mathbf{X}^{l+1})\right)_{\text{compressed}} = \nabla_L(\mathbf{X}^{l+1})\mathbf{U}_k \in \mathbb{R}^{b \times n \times k},\tag{9}$$

and subsequently fully recovered in the previous layer l as:

$$\left(\nabla_L(\mathbf{X}^{l+1})\right)_{\text{recovered}} = \left(\nabla_L(\mathbf{X}^{l+1})\right)_{\text{compressed}} \mathbf{U}_k^\top = \nabla_L(\mathbf{X}^{l+1}).$$
(10)

Remarkably, this formulation ensures that the gradient flow to the computational graph prior to \mathbf{X}^{l+1} remains lossless, with no approximation error. Intuitively, after backpropagation through the parameter matrix \mathbf{W}_{p_2} , the gradient has the form $\nabla_L(\mathbf{X}^{l+1})\mathbf{W}_{p_2}^{\top}$. Because $\operatorname{Row}(\mathbf{W}_{p_2}) \subseteq \operatorname{Col}(\mathbf{U}_k) = \mathcal{S}$, the projection $\nabla_L(\mathbf{X}^{l+1})\mathbf{U}_k\mathbf{U}_k^{\top}\mathbf{W}_{p_2}^{\top}$ does not alter the resulting gradient flow. Full derivation is provided in Appendix A.

4.5. Subspace updates using Grassmann manifold

We observe that restricting the column spaces of projection layers to a fixed subspace, even at highratios, is able to maintain surprisingly adequate convergence. To further improve the convergence, we allow the subspaces to slowly drift. To align the subspace with the gradient directions, we minimize the norm of the gradient components that lie outside the subspace, as measured at the last Transformer layer. Let $\nabla_L(\mathbf{X}_t^{\text{final}}) \in \mathbb{R}^{b \times n \times d}$ denote the activation gradients at the last compressed transformer layer. The leftover gradient, which lies outside the subspace, is $\hat{\mathbf{X}}_t^{\text{final}} = \nabla_L(\mathbf{X}_t^{\text{final}})(\mathbf{I} - \mathbf{U}_k\mathbf{U}_k^{\top})$, where $\mathbf{I} - \mathbf{U}_k\mathbf{U}_k^{\top}$ is the projection operator onto \mathcal{S}^{\perp} . We accumulate $\hat{\mathbf{X}}_t^{\text{final}}$ over K iterations to obtain the metric $\mathcal{L}_{\text{Grassmann}} = \frac{1}{K} \sum_{t=k}^{k+K} \|\hat{\mathbf{X}}_t^{\text{final}}\|_F^2$, where $\|\cdot\|_F$ is the Frobenius norm. We aim to minimize $\mathcal{L}_{\text{Grassmann}}$ over all possible \mathbf{U}_k .

A straightforward way to minimize $\mathcal{L}_{\text{Grassmann}}$ is by performing SVD on it and updating the subspace using the left singular vectors. However, abrupt changes to the subspace can disrupt convergence. Thus, we perform smooth updates by taking steps on the Grassmann manifold. The Grassmann manifold $\mathcal{G}(k, n)$ is the set of all k-dimensional subspaces of \mathbb{R}^n . A point on $\mathcal{G}(k, n)$ is represented by an orthonormal matrix $\mathbf{U}_k \in \mathbb{R}^{n \times k}$, where the columns of \mathbf{U}_k form a basis for the subspace. Thus, defining \mathcal{S} as a point on the Grassmann manifold enables taking smooth steps on the manifold. To minimize $\mathcal{L}_{\text{Grassmann}}$, we employ gradient descent on $\mathcal{G}(k, n)$. First, the Euclidean gradient of $\mathcal{L}_{\text{Grassmann}}$ with respect to \mathbf{U}_k , denoted $\nabla_{\mathcal{L}_{\text{Grassmann}}}(\mathbf{U}_k)$, is projected onto the tangent space to obtain the Riemannian gradient:

$$(\nabla_{\mathcal{L}} \mathbf{U}_k)_{\text{tangent}} = \nabla_{\mathcal{L}_{\text{Grassmann}}} (\mathbf{U}_k) - \mathbf{U}_k \mathbf{U}_k^\top \nabla_{\mathcal{L}_{\text{Grassmann}}} (\mathbf{U}_k).$$
(11)

Then, we perform a gradient descent step $\mathbf{U}_k^{\text{new}} = \mathbf{U}_k - \eta \, (\nabla_{\mathcal{L}} \mathbf{U}_k)_{\text{tangent}}$ where η is the step size. Then, to map $\mathbf{U}_k^{\text{new}}$ back to the manifold, we apply a retraction by orthonormalizing the columns of $\mathbf{U}_k^{\text{new}}$ using QR decomposition $\mathbf{U}_k^{\text{new}}$, $\mathbf{R} = \text{QR}(\mathbf{U}_k^{\text{new}})$, where $\text{QR}(\cdot)$ denotes the QR decomposition. In practice, we perform this subspace update on \mathbf{U}_k very infrequently (per every 500 iterations), and transmit to all the layers.

5. Modifying AdamW for subspace networks

Our approach can be framed as a constrained optimization problem in which the projection matrices must remain within the subspace S. To ensure they stay in this feasible set throughout training, it requires projecting them onto S at each training iteration. We propose a modified version of the AdamW optimizer that preserves the row space of $\mathbf{W}_{p_2}^l$ in S. That means, once initialized within S, the modified optimizer ensures that the $\mathbf{W}_{p_2}^l$ no longer require iterative projection steps, guaranteeing they remain in S without incurring any approximation error. Note that however, we still need to project $\mathbf{W}_{p_1}^l$ onto S due to the nonlinearity of activations as discussed in Appendix. A.

AdamW defines the momentum term as $\mathbf{M}_t = \beta_1 \mathbf{M}_{t-1} + (1 - \beta_1) \nabla_L(\mathbf{W}_{p_2}(t))$ and the second order momentum $\mathbf{V}_t = \beta_2 \mathbf{V}_{t-1} + (1 - \beta_2) (\nabla_L(\mathbf{W}_{p_2}(t)))^2$. Then, the weight update is defined as

$$\mathbf{W}_{p_2}(t+1) = \mathbf{W}_{p_2}(t) - \eta \alpha_t \hat{\mathbf{M}}_t - \lambda \mathbf{W}_{p_2}(t), \qquad (12)$$

where $\alpha_t = \frac{1}{\sqrt{\hat{\mathbf{v}}_t + \epsilon}}$, $\hat{\mathbf{V}}_t = \frac{\mathbf{V}_t}{\beta_2 + \epsilon}$, $\hat{\mathbf{M}}_t = \frac{\mathbf{M}_t}{\beta_1 + \epsilon}$, and η is the fixed learning rate. λ is a hyperparameter controlling the adaptive weight decay. Note that in the second term on the right hand side in Eq. 12, α_t is not a constant scaling factor across coordinates. Thus, it changes the direction of the rows of $\hat{\mathbf{M}}_t$. This causes the $\mathbf{W}_{p_2}(t+1)$ to drift from \mathcal{S} . Therefore, we make α_t constant row-wise so that the resulting update $\eta \alpha_t \hat{\mathbf{M}}_t - \lambda \mathbf{W}_p(t)$ is closed within \mathcal{S} . Specifically, we alter $\hat{\mathbf{V}}_t$ as below. We first take

$$\mu_{\text{row}}(\hat{\mathbf{V}}_t) = \frac{1}{m} \sum_{i=1}^m \hat{\mathbf{V}}_t(:,i) \in \mathbb{R}^{n \times 1}.$$
(13)

where $\mu_{\rm row}(\cdot)$ is the row-wise mean. Then,

$$\mathbf{V}_t = \mathbf{1}_m \cdot \mu_{\text{row}}(\mathbf{V}_t),\tag{14}$$

where $\mathbf{1}_m \in \mathbb{R}^{1 \times m}$ is a column vector of ones, and $\tilde{\mathbf{V}} \in \mathbb{R}^{n \times m}$ has the same dimensions as $\hat{\mathbf{V}}_t$. Then we substitute $\hat{\mathbf{V}}_t$ with $\tilde{\mathbf{V}}$ in Eq. 12. We only apply the above modification to \mathbf{W}_{p_2} and keep the usual updates as it is for other weights.

6. Computational Overhead of subspace updates

Our compression introduces two additional computational components: *weight projection* and *Grassmann updates*. Here, we provide an analysis of the overhead associated with these operations, demonstrating their minimal impact on memory usage and computational efficiency.

Overhead of Weight Projection. We first assess the computational cost incurred by the weight projection step. For a practical evaluation, we consider a model with approximately 2 billion parameters (8 layers, 4k model dimension, 16 attention heads), pipelined across 8 A10G GPUs. In our experiments, a single forward pass takes approximately 4.61s, while the weight projection step adds only about 0.05s of computation time. Thus, the relative overhead introduced by weight projection is approximately 1%, indicating a negligible computational burden.

Overhead of Grassmann Updates. Updating the subspace via Grassmann manifold optimization initially appears complex; however, by deriving closed-form Euclidean gradients, the practical implementation becomes computationally inexpensive. Specifically, given a Grassmann loss defined as:

$$\mathcal{L}_{\text{Grassmann}} = \frac{1}{K} \sum_{t=k}^{k+K} \left\| \nabla_L \left(\mathbf{X}_t^{\text{final}} \right) \left(\mathbf{I} - \mathbf{U}_k \mathbf{U}_k^{\top} \right) \right\|_F^2,$$

and using the Frobenius norm expansion, we obtain:

$$\mathcal{L}_{\text{Grassmann}} = \text{const} - \frac{1}{K} \sum_{t} \text{Tr} \left(\mathbf{G}_{t} \mathbf{U}_{k} \mathbf{U}_{k}^{\top} \mathbf{G}_{t}^{\top} \right),$$

where $\mathbf{G}_t = \nabla_L(\mathbf{X}_t^{\text{final}})$. By defining the symmetric accumulation matrix $\mathbf{S} = \frac{1}{K} \sum_t \mathbf{G}_t^\top \mathbf{G}_t$, the closed-form Euclidean gradient simplifies elegantly to:

$$\nabla_{\mathcal{L}_{\text{Grassmann}}}(\mathbf{U}_k) = -2\mathbf{SU}_k.$$

Practically, implementing this gradient involves straightforward matrix operations in PyTorch:

- Accumulation: Compute and accumulate $\mathbf{G}_t^{\top} \mathbf{G}_t$ once per iteration via standard matrix multiplication.
- Gradient Computation: Multiply S by \mathbf{U}_k once per a fixed interval (e.g., every 500 iterations).
- **Riemannian Projection:** Project back onto the Grassmann manifold via a basic linear algebra operation, again executed at low frequency.

These matrix operations are inexpensive and infrequent, resulting in minimal computational overhead. Consequently, our method's overall cost of maintaining and updating the subspace is negligible compared to the cost of the forward and backward passes of transformer models, affirming the efficiency of our approach.

7. Theoretical insights

This section provides several theoretical insights for the proposed method. We structure our analysis into four key statements (and put the corresponding formal theorems in the Appendix). The first statement indicates that if there is a lossy compression between the layers, as the model depth increases, the approximation error of compression can grow exponentially. The second and third statements demonstrate that, even in an uncompressed network, if the weight gradients are confined to a particular subspace, the weight matrices also naturally converge to a low-dimensional subspace with AdamW. The fourth observation establishes that explicitly enforcing a subset of weights onto a low-dimensional subspace does not harm convergence.

Statement 7.1

If the compression of activations and activation gradients between layers in a model-parallel setting introduces approximation errors, these errors can accumulate exponentially with increasing depth, provided the weight and activation norms are sufficiently large. Refer to Theorem B.1 for a formal proof.

The above result suggests that extending compression techniques from DDP (which are lossy) to MP (which requires compressing information passed between adjacent layers) leads to the accumulation of approximation errors. This occurs because the compression at one layer directly impacts downstream layers, a phenomenon not present in DDP. Additionally, the lack of exploitable structure in activations and activation gradients (Bian et al., 2024; Rudakov et al., 2023) typically results in larger approximation errors compared to the gradients of weights. This limitation makes such compression methods unsuitable for MP in large-scale models.

Statement 7.2

If the gradients of a particular weight matrix in a network are constrained to a fixed subspace, then under AdamW, the weight *updates* asymptotically converge to that subspace over a sufficiently large number of training steps. For a formal proof, see Theorem C.2.

This provides a critical insight: if the gradients of an unconstrained network predominantly lie within a specific low-dimensional subspace S—a property we empirically validate (see Appendix A)—then the AdamW optimizer asymptotically restricts updates outside of S. While this behavior is straightforward for vanilla stochastic gradient descent (SGD), it is non-trivial for AdamW due to its adaptive learning rate mechanism.

Statement 7.3

If Statement 7.2 holds, then the corresponding weight matrices asymptotically converge to the same subspace, irrespective of their initialization. For a formal proof, see Theorem D.1.

Intuitively, this result indicates that the decoupled weight decay mechanism in AdamW systematically suppresses components of the weight matrix that receive negligible gradient updates. Consequently, the learned weights converge to a low-dimensional subspace defined by the gradient updates.

Statement 7.4

A network in which a subset of weights is constrained to a low-dimensional subspace converges to a first-order stationary point with a convergence rate of O(1/T). For a formal proof, see Proposition E.1.

This result is a straightforward extension of the the convergence rate guarantees for constrained optimization using proximal gradient descent on non-convex functions. It shows that even when a *subset* of parameters is restricted to a lower-dimensional subspace, the standard convergence rate of O(1/T) in terms of stationarity remains intact.

8. Experiments

8.1. Experimental Setup

We evaluate decoder-only models (based on Llama 3 (Dubey et al., 2024)) across four large-scale datasets: WikiText (WT) (Merity et al., 2016), BookCorpus (BC) (Zhu et al., 2015), OpenWebText (OWT) (Gokaslan et al., 2019), and C4 (Raffel et al., 2019). For WT, we use the standard splits; for BC and OWT, we randomly select 10% of training data as validation; for C4, due to computational constraints, we report training loss only. The base model has a context length of 1024, embedding dimension 4096, 24 heads, and 8 layers (~2B parameters); larger models (up to 8B parameters) are noted explicitly in ablation sections. We use a base learning rate $\eta = 3e$ -4 (with warmup and linear decay), weight decay 0.01, and batch size 32, unless otherwise specified. We use GPipe (Huang et al., 2019) via torch.distributed.pipelining, integrating our compression into all but the final transformer layer.



Figure 2: Convergence in low-bandwidth settings. From left to right: OpenWebText, WikiText, and BookCorpus. In each plot, the training curves are presented against wall-clock time for an 8-layer (2B) model. Decentralized models utilize 80Mbps connections while the centralized model has datacenter-grade 100Gbps links. Our compressed model achieves on-par convergence to the centralized model with 80Mbps bandwidth budget. In contrast, the non-compressed decentralized model with 80Mbps links suffers from significantly slower convergence due to the communication bottleneck.

We initialize \mathbf{U}_k with isotropic Gaussian noise and set k = 40, achieving $100 \times$ compression. Bandwidth simulations sample from $\mathcal{N}(\mathcal{B}, 0.2\mathcal{B})$ per pass, defining 'centralized' as 100Gbps or 16Gbps setups, with all others as 'decentralized'. Experiments (except the 8B Llama run on L4 GPUs with internet-based decentralized connections) use A10g GPUs (24GB VRAM) with one layer per GPU. **Our method's effectiveness increases with faster accelerators, as slower GPUs allow more computation-communication overlap.**

Square-Cube Law. Square-cube law (Ryabinin et al., 2023) states that in distributed training, computation scales cubically with model size per node, while communication grows only quadratically.



Figure 3: **Performance against depth.** Two key observations: (1) Our compression matches (even slightly exceeds) with centralized baseline as the # layers increases. Fig. 5, further validates this. (2) For the 16-layer model, we fit two layers per GPU (using A100s), increasing computation per block. As shown, the performance gap between centralized and decentralized models slightly decreases, as the ratio between computation and bandwidth bottlenecks is reduced—consistent with the square-cube law.



Figure 4: **Throughput Gain.** As bandwidth becomes increasingly limited, the compressed models achieve a significantly higher throughput gain in both inference (left) and training (right). Results are shown for 8-layer (2B) models.

This partially offsets communication bottlenecks with computational overhead. Thus, c-times slower communication does not lead to a c-times slower convergence. Hence, by improving communication efficiency by $100\times$, we achieve convergence speeds comparable to 100Gbps setups, even with 80Mbps links.

8.2. Convergence in low-bandwidth settings

Our method enables training models over extremely low-bandwidth connections. We trained networks on both 80 Mbps and datacenter-grade 100 Gbps connections. Fig. 2 illustrates the train curves of an 8-layer (2B) model against wall-clock time. As expected, training over 80 Mbps links in the decentralized setting significantly degrades convergence. In contrast, with our compression, the decentralized model achieves on par convergence to the network trained over 100 Gbps connections. To demonstrate the generalization, the perplexity scores of each model over validation sets is shown in Table. 1. As evident, the decentralized model with our compression even surpasses the performance of the centralized model for the same training time. To further validate test-time performance, we train models to convergence using the compute-optimal 1:20 model-to-token ratio from the Chinchilla scaling law (Hoffmann et al.), reaching compute-optimality at 12B training tokens with superior performance (Appendix 8.4).





Figure 5: Training convergence and throughput on an 8B LLaMA model. All runs use 64 L4 GPUs distributed over 8 instances. Centralized instances reside within one region (min bandwidth 16 Gbps), while decentralized instances span 4 regions (min bandwidth 60 Mbps), highlighting pipeline parallel bottlenecks due to reduced internode bandwidth.



Table 1: **Perplexity scores**. Models trained for 12 hours on OpenWebText (OWT), BookCorpus (BC), and WikiText (WT). Bandwidth (B/W) and tokens per second (TPS) are reported. Our method outperforms even the centralized model, achieving significantly higher TPS compared to the non-compressed decentralized baseline.

Model	$ \mathbf{B} / \mathbf{W}$	OWT↓	$\mathbf{BC}{\downarrow}$	$\mathbf{WT}{\downarrow}$	$\mathbf{TPS}\!\!\uparrow$
Decentralized Decentralized Compressed (Ours)	80Mbps 80Mbps	925.19 46.75	108.85 17.63	601.84 23.01	$36.12 \\ 592.41$
Centralized	100Gbps	47.22	18.35	23.08	602.57

8.3. Throughput gain

Our compression also significantly accelerates inference. As inference requires less computation than training, bandwidth becomes the dominant bottleneck; hence, our compression yields substantial gains. Fig. 4 illustrates gains in both training and inference: at inference, we achieve almost a $100 \times$ speedup at 80 Mbps. Although this advantage diminishes at higher bandwidths (e.g., 100 Gbps), we still observe about a $3 \times$ improvement. This indicates that even centralized systems benefit from reduced inference latency using our approach, which can translate into considerable cost savings with the recent trend of inference time scaling of large language models (Bi et al., 2024; Snell et al., 2024) A similar trend holds for training throughput as well.

8.4. Validation at compute-optimal

We conduct additional training experiments to demonstrate the effectiveness of our proposed compression method at the compute-optimal point of model training. The compute-optimal point, as defined by prior work such as Chinchilla, represents an ideal trade-off between model size and token count, optimizing performance within a fixed computational budget. Typically, for architectures like LLaMA, this optimal point corresponds to approximately a 1:20 model-to-token ratio. Consequently, we trained a 640M-parameter model for around 12 billion tokens, aligning closely with the compute-optimal recommendation. The training was executed using 8 FSDP workers, achieving approximately 22k tokens per second throughput. The results, summarized in Table 2, demonstrate that our compression method not only matches but marginally outperforms the non-compressed baseline in terms of validation perplexity.

Table 2: Validation perplexity (lower is better) at the compute-optimal point (12 billion tokens) for a 640M-parameter model. Both ours and centralized model are trained for a same number of iterations. Training the decentralized uncompressed model to compute optimal is infeasible (estimated over 200 days) so we only report the TPS.

Model	$\mathbf{C4}$	BookCorpus	TPS
Decentralized Decentralized Compressed (Ours) Centralized	- 12.53 12.61	12.67 12.79	871 22k 22.5k

These results confirm that our proposed compression method maintains or slightly improves model performance at the crucial compute-optimal training regime.

8.5. Scaling across globally distributed GPUs

To further explore the scalability, we trained a LLaMA 8B parameter variant (Dubey et al., 2024) with 2048 context length, using TorchTitan (Liang et al., 2024) on the C4 dataset across 64 L4 GPUs distributed across 8 instances. We use a pipeline parallel setup with 32 stages running in 2 FSDP dimensions, where the 32 transformer layers are distributed one layer per stage. We evaluated two environment configurations: Centralized and Decentralized. In the Centralized setting all instances were located in the same cloud region and the bandwidth spans between 16Gbps-27Gbps. For the Decentralized case the 8 instances were distributed across 4 distinct regions (North America, Europe, and Asia). Additionally, no two consecutive stages were placed in the same region for the decentralized setup, hence, bandwidth spans from 60Mbps-350Mbps. As shown in Fig. 5, our compression method in the decentralized configuration matches the wall-clock time (even slightly improving) and TPS with the centralized setting. In contrast, the decentralized setting w/o compression was 13x slower.

8.6. Ablations

We conduct ablations on the C4 dataset to evaluate the robustness of our method. Fig. 3 compares performance across model depths. If our compression were lossy, deeper models would accumulate errors, degrading performance relative to the centralized baseline (Bian et al., 2024) (see also Theorem B.1). However, our results show that even as depth increases from 8 (2B) to 16 (3.5B) layers, convergence remains on par with centralized baselines. Further, our large-scale experiment (Fig. 5) confirms that 32-layer models scale effectively, demonstrating decentralized training of large models with MP for the first time. Fig. 3 also highlights that in the 16-layer model, assigning two layers per GPU (A100, 40GB VRAM) increases per-GPU computation, reducing bandwidth bottlenecks and narrowing the gap between decentralized and centralized models. This validates the square-cube law (Ryabinin et al., 2023), showing how computation-to-communication balance impacts decentralized training. However, note that decentralized training primarily targets low-end GPUs rather than high-end hardware. Ablations over other design choices and the negligible memory overhead are discussed in Appendix 8.8.

8.7. Comparison against lossy compressions

As per Statement 7.1, standard compression methods used in DDP do not effectively extend to MP. We train an 8-layer model on the WikiText, comparing our compression method with TopK, quantization, and low-rank projection. As shown in Fig. 6, with an aggressive compression rate of $\times 100$, such compression schemes fail to converge. In contrast, our method achieves convergence on par with the non-compressed model. Note that for quantization, the best compression rate we can achieve is $16 \times$ for 16bit precision.

8.8. Memory Overhead Analysis

As our method requires addition of fixed high-rank and dynamic low-rank embeddings, one potential concern is the memory overhead, as sequence length L grows. We empirically and theoretically demonstrate that the absolute memory overhead introduced by our approach remains constant and neglegible, while the relative overhead decreases with increasing sequence lengths.

We first empirically validate this and then give a theoretical explanation. To this end, we performed experiments using a 2B-parameter Transformer model (8 layers, 4k model dimension, 16 attention heads) distributed across eight NVIDIA H100 GPUs under varying sequence lengths L. Table 3 summarizes these results. Remarkably, the absolute overhead consistently remains around 400 MB, irrespective of sequence length. As L grows from 8k to 24k, the relative memory overhead correspondingly drops from 4.0% to 0.6%.

Table 3: Peak memory usage comparison between baseline and our subspace method as sequence length (L) scales.

L	Baseline (GB)	Ours (GB)	Overhead	Relative Overhead
8k	9.66	10.06	$\sim \! 400 \mathrm{MB}$	$\sim 4.0\%$
16k	36.51	36.91	${\sim}400 {\rm MB}$	$\sim 1.1\%$
24k	76.00	76.46	${\sim}400 {\rm MB}$	${\sim}0.6\%$

This consistent and negligible overhead can be explained through PyTorch's memory management behavior. Firstly, fixed embedding lookups in our method are ephemeral since they are non-trainable; PyTorch does not store activation gradients for them. After they are added to the activations, PyTorch's caching allocator reuses the memory (instead of invoking cudaMalloc/cudaFree), allowing temporary tensors to be released prior to attention. Thus, their memory footprint doesn't persist into later stages. Secondly, embedding lookups ($\mathcal{O}(B \times L \times D)$ memory usage) are inherently minor relative to attention ($\mathcal{O}(B \times L^2 \times D)$) and MLP layers ($\mathcal{O}(B \times L \times D^2)$), making their temporary storage impact negligible in peak memory usage.

Further, the cached embedding tables themselves do not pose a meaningful threat to scalability. For instance, even a large embedding table (10k-dimensional embedding and 50k vocabulary size) would consume approximately 0.93 GB in 16-bit precision, requiring only impractically large embedding dimensions (e.g., 100k+) to surpass 10 GB. State-of-the-art frontier models (e.g., DeepSeek (DeepSeek AI Team, 2023), LLAMA-405B (Dubey et al., 2024)) remain far below such thresholds (7K and 16K, respectively). Therefore, even with significantly increased embedding dimensions, memory overhead remains dominated by MLP and attention computations, not embedding storage.

In conclusion, our method demonstrates robust and scalable memory behavior, effectively managing memory overhead even at extremely large sequence lengths, validating its practicality for decentralized training at scale.

8.9. Memory Overhead with Increasing Workers

Another key consideration is how our compression scheme scales in terms of memory overhead with an increasing number of workers. One might suspect that distributing long contexts across more workers could lead to increased overhead per worker. We clarify that this concern is unfounded due to the way embeddings and attention computations are handled in our setup.

To elaborate, in a context-parallel training setting, each worker handles a distinct segment of the input sequence. Embedding lookups are performed *locally on each worker* before attention computations. The resulting key-value (KV) tensors maintain the same shape and size as a standard, non-compressed model, ensuring no extra storage requirement. Importantly, the ephemeral embedding additions are discarded immediately after their local usage, well before the memory-intensive attention and MLP layers, thus never contributing to peak memory usage.

To empirically validate our claims, we extend our experiments to include varying numbers of workers using Ring Attention, a state-of-the-art context-parallel mechanism where each worker exchanges KV tensors only with its immediate neighbors. As embedding operations occur locally and KV tensors remain unchanged in size, our compression scheme integrates seamlessly.

Table 4 illustrates that, irrespective of the increase in sequence length (L) and corresponding increase in the number of workers, our per-worker memory overhead remains constant (around 400MB) and does not scale with either sequence length or worker count.

L	Num. Workers	Baseline (GB)	Ours (GB)	Overhead/Worker	Relative Overhead
8k	1	9.66	10.06	$\sim 400 \mathrm{MB}$	$\sim \!\! 4.0\%$
16k	1	36.51	36.91	$\sim 400 \mathrm{MB}$	$\sim 1.1\%$
24k	1	76.00	76.46	$\sim 400 \mathrm{MB}$	${\sim}0.6\%$
50k	2	76.13	76.55	$\sim 400 \mathrm{MB}$	$\sim 0.55\%$
65k	3	78.19	79.62	$\sim 400 \mathrm{MB}$	$\sim 0.54\%$

Table 4: Peak memory per worker with increasing sequence lengths and workers using our compression scheme.

Thus, our design effectively maintains low, constant memory overhead per worker, demonstrating excellent scalability in decentralized contexts. The embedding memory per-worker overhead remains negligible and constant, independent of the scaling of sequence length or the number of distributed workers, validating our approach for large-scale decentralized training scenarios.

9. Broader impact

Our work presents what we believe to be the first compression method that allows practical, large-scale, decentralized model-parallel training, enabling efficient scaling in low-bandwidth environments. This marks a significant step toward democratizing access to large-scale AI by reducing dependence on costly datacenter infrastructure and expanding opportunities for deep learning research. Our approach has far-reaching implications for making AI more accessible, cost-effective, and environmentally sustainable across academia, startups, and industry.

10. Conclusion

We propose a novel compression technique that, for the first time, enables aggressive compression in MP without harming convergence. By leveraging structured subspace constraints, we achieve up to $100 \times$ communication efficiency while preserving convergence. Our compression enhances both inference and training efficiency, reducing latency even in centralized settings. Extensive experiments across varying model depths, bandwidth conditions, and large-scale deployments validate the effectiveness and scalability of our approach. Notably, we demonstrate its real-world applicability by successfully training an 8B-parameter LlaMa model on low-end GPUs distributed across multiple global regions, connected solely via internet-grade (60 Mbps) links, while achieving convergence comparable to a centralized setup. Our results establish that decentralized training of large-scale models using model parallelism—previously hindered by severe communication bottlenecks—is now practical.

References

- Emmanuel Abbe, Samy Bengio, Enric Boix-Adsera, Etai Littwin, and Joshua Susskind. Transformers learn through gradual rank increase. *Advances in Neural Information Processing Systems*, 36, 2024.
- Dan Alistarh, Demjan Grubic, Jerry Li, Ryota Tomioka, and Milan Vojnovic. Qsgd: Communicationefficient sgd via gradient quantization and encoding. *Advances in neural information processing* systems, 30, 2017.
- Lukas Balles and Philipp Hennig. Dissecting adam: The sign, magnitude and variance of stochastic gradients. In *International Conference on Machine Learning*, pages 404–413. PMLR, 2018.
- Jeremy Bernstein, Yu-Xiang Wang, Kamyar Azizzadenesheli, and Animashree Anandkumar. signsgd: Compressed optimisation for non-convex problems. In *International Conference on Machine Learning*, pages 560–569. PMLR, 2018.
- Zhenni Bi, Kai Han, Chuanjian Liu, Yehui Tang, and Yunhe Wang. Forest-of-thought: Scaling test-time compute for enhancing llm reasoning. arXiv preprint arXiv:2412.09078, 2024.
- Song Bian, Dacheng Li, Hongyi Wang, Eric Xing, and Shivaram Venkataraman. Does compressing activations help model parallel training? *Proceedings of Machine Learning and Systems*, 6:239–252, 2024.
- Romain Cosson, Ali Jadbabaie, Anuran Makur, Amirhossein Reisizadeh, and Devavrat Shah. Low-rank gradient descent. *IEEE Open Journal of Control Systems*, 2023.
- DeepSeek AI Team. DeepSeek LLMs, 2023. URL https://deepseek.com/.
- Alexandre Défossez, Léon Bottou, Francis Bach, and Nicolas Usunier. A simple convergence proof of adam and adagrad. arXiv preprint arXiv:2003.02395, 2020.
- Tim Dettmers, Mike Lewis, Sam Shleifer, and Luke Zettlemoyer. 8-bit optimizers via block-wise quantization. arXiv preprint arXiv:2110.02861, 2021.
- Michael Diskin, Alexey Bukhtiyarov, Max Ryabinin, Lucile Saulnier, Anton Sinitsin, Dmitry Popov, Dmitry V Pyrkin, Maxim Kashirin, Alexander Borzunov, Albert Villanova del Moral, et al. Distributed deep learning in open collaborations. Advances in Neural Information Processing Systems, 34:7879–7897, 2021.

- Yihe Dong, Jean-Baptiste Cordonnier, and Andreas Loukas. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. In *International Conference on Machine Learning*, pages 2793–2803. PMLR, 2021.
- Arthur Douillard, Qixuan Feng, Andrei A Rusu, Rachita Chhaparia, Yani Donchev, Adhiguna Kuncoro, Marc'Aurelio Ranzato, Arthur Szlam, and Jiajun Shen. Diloco: Distributed lowcommunication training of language models. arXiv preprint arXiv:2311.08105, 2023.
- Abhimanyu Dubey, Abhinav Jauhri, Abhinav Pandey, Abhishek Kadian, Ahmad Al-Dahle, Aiesha Letman, Akhil Mathur, Alan Schelten, Amy Yang, Angela Fan, et al. The llama 3 herd of models. arXiv preprint arXiv:2407.21783, 2024.
- Aaron Gokaslan, Vanya Cohen, Ellie Pavlick, and Stefanie Tellex. Openwebtext corpus. http://Skylion007.github.io/OpenWebTextCorpus, 2019.
- Guy Gur-Ari, Daniel A Roberts, and Ethan Dyer. Gradient descent happens in a tiny subspace. arXiv preprint arXiv:1812.04754, 2018.
- Jordan Hoffmann, Sebastian Borgeaud, Arthur Mensch, Elena Buchatskaya, Trevor Cai, Eliza Rutherford, Diego de Las Casas, Lisa Anne Hendricks, Johannes Welbl, Aidan Clark, et al. Training compute-optimal large language models.
- Yanping Huang, Youlong Cheng, Ankur Bapna, Orhan Firat, Dehao Chen, Mia Chen, HyoukJoong Lee, Jiquan Ngiam, Quoc V Le, Yonghui Wu, et al. Gpipe: Efficient training of giant neural networks using pipeline parallelism. Advances in neural information processing systems, 32, 2019.
- Sai Praneeth Karimireddy, Quentin Rebjock, Sebastian Stich, and Martin Jaggi. Error feedback fixes signsgd and other gradient compression schemes. In *International Conference on Machine Learning*, pages 3252–3261. PMLR, 2019.
- Diederik P Kingma. Adam: A method for stochastic optimization. arXiv preprint arXiv:1412.6980, 2014.
- Alexander Kolesnikov, Lucas Beyer, Xiaohua Zhai, Joan Puigcerver, Jessica Yung, Sylvain Gelly, and Neil Houlsby. Big transfer (bit): General visual representation learning. In Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16, pages 491–507. Springer, 2020.
- Anastasia Koloskova, Sebastian Stich, and Martin Jaggi. Decentralized stochastic optimization and gossip algorithms with compressed communication. In *International Conference on Machine Learning*, pages 3478–3487. PMLR, 2019.
- Anastasia Koloskova, Nicolas Loizou, Sadra Boreiri, Martin Jaggi, and Sebastian Stich. A unified theory of decentralized sgd with changing topology and local updates. In *International Conference* on Machine Learning, pages 5381–5393. PMLR, 2020.
- Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. Advances in neural information processing systems, 25, 2012.
- Haochuan Li, Alexander Rakhlin, and Ali Jadbabaie. Convergence of adam under relaxed assumptions. Advances in Neural Information Processing Systems, 36, 2024a.
- Hongkang Li, Meng Wang, Shuai Zhang, Sijia Liu, and Pin-Yu Chen. Learning on transformers is provable low-rank and sparse: A one-layer analysis. In 2024 IEEE 13rd Sensor Array and Multichannel Signal Processing Workshop (SAM), pages 1–5. IEEE, 2024b.

- Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke, Jeff Smith, Brian Vaughan, Pritam Damania, et al. Pytorch distributed: Experiences on accelerating data parallel training. arXiv preprint arXiv:2006.15704, 2020.
- Xiangru Lian, Ce Zhang, Huan Zhang, Cho-Jui Hsieh, Wei Zhang, and Ji Liu. Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent. Advances in neural information processing systems, 30, 2017.
- Wanchao Liang, Tianyu Liu, Less Wright, Will Constable, Andrew Gu, Chien-Chin Huang, Iris Zhang, Wei Feng, Howard Huang, Junjie Wang, et al. Torchtitan: One-stop pytorch native solution for production ready llm pre-training. arXiv preprint arXiv:2410.06511, 2024.
- Yujun Lin, Song Han, Huizi Mao, Yu Wang, and William J Dally. Deep gradient compression: Reducing the communication bandwidth for distributed training. arXiv preprint arXiv:1712.01887, 2017.
- Mingrui Liu, Wei Zhang, Francesco Orabona, and Tianbao Yang. Adam⁺: A stochastic method with adaptive variance reduction. arXiv preprint arXiv:2011.11985, 2020.
- I Loshchilov. Decoupled weight decay regularization. arXiv preprint arXiv:1711.05101, 2017.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. Pointer sentinel mixture models, 2016.
- Deepak Narayanan, Mohammad Shoeybi, Jared Casper, Patrick LeGresley, Mostofa Patwary, Vijay Korthikanti, Dmitri Vainbrand, Prethvi Kashinkunti, Julie Bernauer, Bryan Catanzaro, et al. Efficient large-scale language model training on gpu clusters using megatron-lm. In Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, pages 1–15, 2021.
- Bowen Peng, Jeffrey Quesnelle, and Diederik P Kingma. Decoupled momentum optimization. arXiv preprint arXiv:2411.19870, 2024.
- How Much Longer Can Computing Power and Drive Artificial Intelligence Progress. Ai and compute.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *arXiv e-prints*, 2019.
- Xiaozhe Ren, Pingyi Zhou, Xinfan Meng, Xinjing Huang, Yadao Wang, Weichao Wang, Pengfei Li, Xiaoda Zhang, Alexander Podolskiy, Grigory Arshinov, et al. Pangu-{\Sigma}: Towards trillion parameter language model with sparse heterogeneous computing. arXiv preprint arXiv:2303.10845, 2023.
- Mikhail I Rudakov, Aleksandr Nikolaevich Beznosikov, Ya A Kholodov, and Alexander Vladimirovich Gasnikov. Activations and gradients compression for model-parallel training. In *Doklady Mathematics*, volume 108, pages S272–S281. Springer, 2023.
- Max Ryabinin and Anton Gusev. Towards crowdsourced training of large neural networks using decentralized mixture-of-experts. Advances in Neural Information Processing Systems, 33:3659–3672, 2020.
- Max Ryabinin, Eduard Gorbunov, Vsevolod Plokhotnyuk, and Gennady Pekhimenko. Moshpit sgd: Communication-efficient decentralized training on heterogeneous unreliable devices. Advances in Neural Information Processing Systems, 34:18195–18211, 2021.

- Max Ryabinin, Tim Dettmers, Michael Diskin, and Alexander Borzunov. Swarm parallelism: Training large models can be surprisingly communication-efficient. In *International Conference on Machine Learning*, pages 29416–29440. PMLR, 2023.
- Sunny Sanyal, Ravid Shwartz-Ziv, Alexandros G Dimakis, and Sujay Sanghavi. Inheritune: Training smaller yet more attentive language models, 2024. URL https://arxiv. org/abs/2404.08634.
- Charlie Snell, Jaehoon Lee, Kelvin Xu, and Aviral Kumar. Scaling llm test-time compute optimally can be more effective than scaling model parameters. arXiv preprint arXiv:2408.03314, 2024.
- Hanlin Tang, Shaoduo Gan, Ammar Ahmad Awan, Samyam Rajbhandari, Conglong Li, Xiangru Lian, Ji Liu, Ce Zhang, and Yuxiong He. 1-bit adam: Communication efficient large-scale training with adam's convergence speed. In *International Conference on Machine Learning*, pages 10118–10129. PMLR, 2021.
- Verónica G Vergara Larrea, Wayne Joubert, Michael J Brim, Reuben D Budiardja, Don Maxwell, Matt Ezell, Christopher Zimmer, Swen Boehm, Wael Elwasif, Sarp Oral, et al. Scaling the summit: deploying the world's fastest supercomputer. In *High Performance Computing: ISC High Performance 2019 International Workshops, Frankfurt, Germany, June 16-20, 2019, Revised Selected Papers 34*, pages 330–351. Springer, 2019.
- Thijs Vogels, Sai Praneeth Karimireddy, and Martin Jaggi. Powersgd: Practical low-rank gradient compression for distributed optimization. Advances in Neural Information Processing Systems, 32, 2019.
- James Vuckovic. Kalman gradient descent: Adaptive variance reduction in stochastic optimization. arXiv preprint arXiv:1810.12273, 2018.
- Chong Wang, Xi Chen, Alexander J Smola, and Eric P Xing. Variance reduction for stochastic gradient optimization. Advances in neural information processing systems, 26, 2013.
- Hongyi Wang, Scott Sievert, Shengchao Liu, Zachary Charles, Dimitris Papailiopoulos, and Stephen Wright. Atomo: Communication-efficient learning via atomic sparsification. Advances in neural information processing systems, 31, 2018.
- Hongyi Wang, Saurabh Agarwal, and Dimitris Papailiopoulos. Pufferfish: Communication-efficient models at no extra cost. *Proceedings of Machine Learning and Systems*, 3:365–386, 2021.
- Jialei Wang, Mladen Kolar, Nathan Srebro, and Tong Zhang. Efficient distributed learning with sparsity. In International conference on machine learning, pages 3636–3645. PMLR, 2017.
- Jue Wang, Yucheng Lu, Binhang Yuan, Beidi Chen, Percy Liang, Christopher De Sa, Christopher Re, and Ce Zhang. Cocktailsgd: Fine-tuning foundation models over 500mbps networks. In International Conference on Machine Learning, pages 36058–36076. PMLR, 2023.
- Jianqiao Wangni, Jialei Wang, Ji Liu, and Tong Zhang. Gradient sparsification for communicationefficient distributed optimization. Advances in Neural Information Processing Systems, 31, 2018.
- Jiaxiang Wu, Weidong Huang, Junzhou Huang, and Tong Zhang. Error compensated quantized sgd and its applications to large-scale distributed optimization. In *International conference on machine learning*, pages 5325–5333. PMLR, 2018.
- Greg Yang, James B Simon, and Jeremy Bernstein. A spectral condition for feature learning. arXiv preprint arXiv:2310.17813, 2023.

- Xingyi Yang. Stochastic gradient variance reduction by solving a filtering problem. arXiv preprint arXiv:2012.12418, 2020.
- Binhang Yuan, Yongjun He, Jared Davis, Tianyi Zhang, Tri Dao, Beidi Chen, Percy S Liang, Christopher Re, and Ce Zhang. Decentralized training of foundation models in heterogeneous environments. *Advances in Neural Information Processing Systems*, 35:25464–25477, 2022.
- Qi Zhang, Yi Zhou, and Shaofeng Zou. Convergence guarantees for rmsprop and adam in generalizedsmooth non-convex optimization with affine noise variance. arXiv preprint arXiv:2404.01436, 2024.
- Jiawei Zhao, Florian Schäfer, and Anima Anandkumar. Zero initialization: Initializing neural networks with only zeros and ones. arXiv preprint arXiv:2110.12661, 2021.
- Jiawei Zhao, Zhenyu Zhang, Beidi Chen, Zhangyang Wang, Anima Anandkumar, and Yuandong Tian. Galore: Memory-efficient llm training by gradient low-rank projection. *arXiv preprint* arXiv:2403.03507, 2024.
- Yanli Zhao, Andrew Gu, Rohan Varma, Liang Luo, Chien-Chin Huang, Min Xu, Less Wright, Hamid Shojanazeri, Myle Ott, Sam Shleifer, et al. Pytorch fsdp: experiences on scaling fully sharded data parallel. arXiv preprint arXiv:2304.11277, 2023.
- Yukun Zhu, Ryan Kiros, Rich Zemel, Ruslan Salakhutdinov, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. Aligning books and movies: Towards story-like visual explanations by watching movies and reading books. In *The IEEE International Conference on Computer Vision (ICCV)*, December 2015.

Appendix

A. Gradient compression in backpropagation

In this section, we show that compressing the gradients by projecting them on to S does not induce any approximation error.

Let $\nabla_l(\mathbf{X}^{l+1})$ denote the gradient of the loss with respect to \mathbf{X}^{l+1} , the output of layer l. The gradient with respect to $\mathbf{W}_{p_2}^l$ is then

$$\nabla_{l}(\mathbf{W}_{p_{2}}^{l}) = (\mathbf{X}_{\text{hidden}}^{l})^{\top} \nabla_{l}(\mathbf{X}^{l+1}).$$
(15)

The residual gradient from the skip connection is given by

$$\left(\nabla_l(\mathbf{X}_{\text{attn}}^l)\right)^{\text{residual}} = \nabla_l(\mathbf{X}^{l+1}).$$
(16)

Hence, the update rule for $\mathbf{W}_{p_2}^l$ becomes

$$\mathbf{W}_{p_2}^l(t+1) = \mathbf{W}_{p_2}^l(t) - \gamma \left(\mathbf{X}_{\text{hidden}}^l\right)^\top \nabla_l(\mathbf{X}^{l+1}), \tag{17}$$

where γ is the learning rate, and t denotes the training timestep.

If we project $\nabla_l(\mathbf{X}^{l+1})$ onto $\operatorname{Col}(\mathbf{U}_k) = \mathcal{S}$, the update can be written as

$$\mathbf{W}_{p_2}^{l}(t+1) = \mathbf{W}_{p_2}^{l}(t) - \gamma \left(\mathbf{X}_{\text{hidden}}^{l}\right)^{\top} \left(\nabla_l(\mathbf{X}^{l+1}) \mathbf{U}_k \mathbf{U}_k^{\top}\right).$$
(18)

If γ acts as a row-wise constant, then it is straightforward to see that if $\operatorname{Row}(\mathbf{W}_{p_2}^l(t)) \subseteq \mathcal{S}$, it follows that $\operatorname{Row}(\mathbf{W}_{p_2}^l(t+1)) \subseteq \mathcal{S}$ as well, because vector spaces are closed under linear operations.

Recall that we modify AdamW so that its adaptive learning rate is constant across each row (see Section 5), ensuring this property holds. Therefore, by induction, if $\operatorname{Row}(\mathbf{W}_{p_2}^l(0)) \subseteq S$ at initialization, then $\operatorname{Row}(\mathbf{W}_{p_2}^l(t)) \subseteq S$ for all subsequent updates when the incoming gradients $\nabla_l(\mathbf{X}^{l+1})$ are projected onto S, removing the need for iterative projection of $\mathbf{W}_{p_2}^l$ on to S.

Continuing backpropagation, the gradient with respect to \mathbf{X}_{hidden}^{l} is

$$\nabla_{l} \left(\mathbf{X}_{\text{hidden}}^{l} \right) = \nabla_{l} \left(\mathbf{X}^{l+1} \right) \mathbf{U}_{k} \mathbf{U}_{k}^{\top} \left(\mathbf{W}_{p_{2}}^{l} \right)^{\top}.$$
(19)

Since $\operatorname{Row}(\mathbf{W}_{p_2}^l) = \mathcal{S}$, we can write $\mathbf{W}_{p_2}^l = \mathbf{W}_{p_2}^l \mathbf{U}_k \mathbf{U}_k^{\top}$. Substituting this in, we get

$$\nabla_l \left(\mathbf{X}_{\text{hidden}}^l \right) = \nabla_l \left(\mathbf{X}^{l+1} \right) \mathbf{U}_k \mathbf{U}_k^{\top} \left(\mathbf{W}_{p_2}^l \mathbf{U}_k \mathbf{U}_k^{\top} \right)^{\top}$$
(20)

$$= \nabla_l (\mathbf{X}^{l+1}) \mathbf{U}_k \mathbf{U}_k^\top \mathbf{U}_k \mathbf{U}_k^\top (\mathbf{W}_{p_2}^l)^\top.$$
(21)

Because \mathbf{U}_k is orthonormal, $\mathbf{U}_k^{\top} \mathbf{U}_k = \mathbf{I}$, so

$$\nabla_l (\mathbf{X}_{\text{hidden}}^l) = \nabla_l (\mathbf{X}^{l+1}) \mathbf{U}_k \mathbf{U}_k^\top (\mathbf{W}_{p_2}^l)^\top$$
(22)

$$= \nabla_l \left(\mathbf{X}^{l+1} \right) \left(\mathbf{W}_{p_2}^l \, \mathbf{U}_k \, \mathbf{U}_k^\top \right)^\top \tag{23}$$

$$=\nabla_l \left(\mathbf{X}^{l+1} \right) \left(\mathbf{W}_{p_2}^l \right). \tag{24}$$

This shows that projecting $\nabla_l(\mathbf{X}^{l+1})$ onto \mathcal{S} does not introduce any approximation error in $\nabla_l(\mathbf{X}^l_{\text{hidden}})$.

Similarly, the remaining gradients are computed as follows. First,

$$\nabla_{l} (\mathbf{W}_{1}^{l}) = (\mathbf{X}_{\text{attn}}^{l})^{\top} (\nabla_{l} (\mathbf{X}_{\text{hidden}}^{l}) \circ \nabla f_{\text{relu}}), \qquad (25)$$

where $(\cdot \circ \cdot)$ denotes elementwise (Hadamard) multiplication. The gradient for \mathbf{X}_{attn}^{l} then becomes

$$\nabla_{l} \left(\mathbf{X}_{\text{attn}}^{l} \right) = \left(\nabla_{l} \left(\mathbf{X}_{\text{hidden}}^{l} \right) \circ \nabla f_{\text{relu}} \right) \left(\mathbf{W}_{1}^{l} \right)^{\top} + \left(\nabla_{l} \left(\mathbf{X}_{\text{attn}}^{l} \right) \right)^{\text{residual}}$$
(26)

$$= \left(\nabla_l \left(\mathbf{X}_{\text{hidden}}^l\right) \circ \nabla f_{\text{relu}}\right) \left(\mathbf{W}_1^l\right)^\top + \nabla_l \left(\mathbf{X}^{l+1}\right).$$
(27)

Next,

$$\nabla_{l} (\mathbf{W}_{p_{1}}^{l}) = (\mathbf{X}_{\text{concat}}^{l})^{\top} \nabla_{l} (\mathbf{X}_{\text{attn}}^{l})$$

$$= (\mathbf{X}_{\text{concat}}^{l})^{\top} \left((\nabla_{\mathbf{X}} (\mathbf{X}_{\text{attn}}^{l}) \circ \nabla_{\mathbf{f}} (\mathbf{X}_{\text{ott}}^{l})^{\top} + \nabla_{\mathbf{f}} (\mathbf{X}_{\text{ott}}^{l+1}) \right)$$
(28)
$$(28)$$

$$= \left(\mathbf{X}_{\text{concat}}^{l}\right)^{\top} \left(\left(\nabla_{l} \left(\mathbf{X}_{\text{hidden}}^{l}\right) \circ \nabla f_{\text{relu}}\right) \left(\mathbf{W}_{1}^{l}\right)^{\top} + \nabla_{l} \left(\mathbf{X}^{l+1}\right) \right).$$
(29)

Recall that $\operatorname{Row}(\mathbf{W}_1^l) \subseteq \mathcal{S}$. Therefore, if $\nabla_l(\mathbf{X}^{l+1})$ is projected onto \mathcal{S} , it follows that $\operatorname{Row}(\nabla_l(\mathbf{W}_{p_1}^l)) \subseteq \mathcal{S}$.

Finally, the update rule for $\mathbf{W}_{p_1}^l$ is

$$\mathbf{W}_{p_{1}}^{l}(t+1) = \mathbf{W}_{p_{1}}^{l}(t) - \gamma \nabla_{l} \big(\mathbf{W}_{p_{1}}^{l}(t) \big),$$
(30)

The gradient updates for $\mathbf{X}_{\text{concat}}^{l}$ involve

$$\mathbf{X}_{\text{concat}}^{l} = \mathbf{X}_{\text{attn}}^{l} \left(\mathbf{W}_{p_{1}}^{l} \right)^{\top}.$$
 (31)

Because all gradients up to $\mathbf{X}_{\text{concat}}^{l}$ are preserved without loss, it follows via the chain rule that the gradient flow to the blocks below $\mathbf{X}_{\text{concat}}^{l}$ is also lossless. Consequently, when compressing $\nabla_{L}(\mathbf{X}^{l+1})$ as

$$\left(\nabla_L(\mathbf{X}^{l+1})\right)_{\text{compressed}} = \nabla_L(\mathbf{X}^{l+1}) \mathbf{U}_k \in \mathbb{R}^{b \times n \times k},$$
(32)

we can fully recover it in the previous layer:

$$\left(\nabla_L(\mathbf{X}^{l+1})\right)_{\text{recovered}} = \left(\nabla_L(\mathbf{X}^{l+1})\right)_{\text{compressed}} \mathbf{U}_k^\top \tag{33}$$

$$=\nabla_L (\mathbf{X}^{l+1}), \tag{34}$$

thereby incurring no approximation error. Notably, since $k \ll d$, the compressed gradient $(\nabla_L(\mathbf{X}^{l+1}))_{\text{compressed}}$ is significantly lower-dimensional than $\nabla_L(\mathbf{X}^{l+1})$, yielding substantial communication savings over low-bandwidth links. This shows that by projecting gradients onto the same subspace used for forward-pass compression, we can achieve *lossless gradient compression* during backpropagation. This strategy further eliminates the need for frequent projections of $\mathbf{W}_{p_2}^l$ maintaining its confinement to the intended subspace.

B. On the error accumulation of lossy compressions

A key difference between Distributed Data Parallel (DDP) training and Model Parallel (MP) training lies in how gradients are exchanged and how compression is applied. In DDP, gradients of model parameters are exchanged after each training step, enabling compression to be applied across the entire gradient vector in one operation. In contrast, PP training requires the exchange of activations and activation gradients between model partitions, leading to a layer-wise compression approach. Unlike DDP, where compression affects the gradient as a whole, MP training introduces independent compression errors at each layer. These errors accumulate progressively during training, impacting gradient computations layer by layer and posing unique challenges in maintaining model accuracy and stability.

Given these distinctions, it is critical to analyze MP training within the constraints of activation gradient compression. The layer-by-layer compression in MP can lead to compounded errors that significantly influence convergence behavior and overall model performance. The following result focuses on investigating the relative error with respect to the weight gradients, occurred by a compression error induced by a particular layer, on an arbitrary layer below it.

Theorem B.1. Consider a feedforward neural network with L layers, where layer l applies a (differentiable) function

$$x_{l+1} = f_l(x_l), \quad l = 1, \dots, L.$$

Let $\nabla_L(x_l)$ denote the gradient of the final loss \mathcal{L} with respect to the layer's input x_l . Suppose that:

- 1. The spectral norm of the Jacobian $\nabla f_l(x_l)$ is bounded above by $\nu > 0$ for all l, i.e. $\|\nabla f_l(x_l)\| \leq \nu$.
- 2. In backpropagation, an additional error e_l is introduced at each layer l, with $||e_l|| \le e$ for some constant e > 0.

Define ε_l to be the cumulative error in the gradient at layer l. Then for $\nu > 1$, ε_l can grow exponentially with the total number of layers L; in particular,

$$\|\varepsilon_l\| \leq e \frac{\nu^{L-l+1} - 1}{\nu - 1},$$

which is an exponential function of L when $\nu > 1$.

Proof. Recall the usual chain rule for the gradient of the final loss \mathcal{L} with respect to the input x_l of layer l:

$$\nabla_L(x_l) = \nabla_L(x_{l+1}) \nabla f_l(x_l).$$

Assume that at each layer we introduce an error in the gradient. Let

 $\varepsilon_l = (\text{true gradient at layer } l) - (\text{observed/propagated gradient at layer } l).$

When moving from layer l to layer l-1, the error recursion becomes:

$$\varepsilon_{l-1} = \varepsilon_l \nabla f_{l-1}(x_{l-1}) + e_{l-1},$$

where e_{l-1} is the *newly introduced* error at layer l-1. Unfolding this backwards gives a general expansion:

$$\varepsilon_l = e_l + \sum_{j=l+1}^{L} \left(\prod_{i=l}^{j-1} \nabla f_i(x_i) \right) e_j.$$

Taking the norm and using the assumption $\|\nabla f_i(x_i)\| \leq \nu$ and $\|e_j\| \leq e$, we get:

$$\|\varepsilon_l\| \leq \sum_{j=l}^{L} \left(\prod_{i=l}^{j-1} \|\nabla f_i(x_i)\| \right) \|e_j\| \leq \sum_{j=l}^{L} \nu^{j-l} e = e \sum_{k=0}^{L-l} \nu^k,$$

where k = j - l. This geometric sum is

$$\sum_{k=0}^{L-l} \nu^k = \frac{\nu^{L-l+1} - 1}{\nu - 1} \quad \text{(valid for } \nu \neq 1\text{)}.$$

Hence,

$$\|\varepsilon_l\| \leq e \frac{\nu^{L-l+1}-1}{\nu-1}.$$

Since $\nu > 1$, ν^{L-l+1} grows exponentially in L.

This is an important result. In particular, this result indicates that in cases where the spectral norm of the weight matrices is large enough, the upper bound for the error for the lower layers can grow exponentially as the depth of the network increases.

C. Vanishing of gradient updates outside a subspace

The result in this section highlights an important property of AdamW: the optimizer's weight updates progressively align with the subspace in which the gradients are constrained. In particular, if the gradients of a weight matrix lie within a specific low-dimensional subspace S, the updates made to the weight matrix are increasingly confined to S over time. This behavior ensures that any weight components orthogonal to S are asymptotically suppressed, with the weight matrix ultimately converging to S. To this end, we first prove Lemma C.1 and then Theorem C.2.

Lemma C.1 (Orthogonal Distortion Bound). Let $a = (a_1, \ldots, a_n) \in \mathbb{R}^n$ be a nonzero vector, and let $v_1, \ldots, v_n > 0$ be positive scalars. Define

$$m = \min_{1 \le i \le n} v_i, \quad M = \max_{1 \le i \le n} v_i, \quad \Gamma = \frac{M}{m}.$$

Form the vector

$$b = (v_1a_1, v_2a_2, \ldots, v_na_n) \in \mathbb{R}^n,$$

and let b_{\perp} be the component of b orthogonal to a. Then the following two bounds hold:

$$||b_{\perp}|| \leq \frac{M-m}{2} ||a|| = \frac{m(\Gamma-1)}{2} ||a|| \text{ and } ||b_{\perp}|| \leq \frac{\Gamma-1}{2} ||b||.$$

Proof. First, decompose b into its projection onto a plus its component orthogonal to a:

 $b = \operatorname{proj}_a(b) + b_{\perp}$, where $b_{\perp} \perp a$.

The norm of b_{\perp} can be expressed by the well-known identity

$$||b_{\perp}||^2 = ||b||^2 - \frac{(b \cdot a)^2}{||a||^2}.$$

Since $b = (v_1 a_1, \ldots, v_n a_n)$, we have

$$||b||^2 = \sum_{i=1}^n (v_i a_i)^2 = \sum_{i=1}^n v_i^2 a_i^2$$
, and $b \cdot a = \sum_{i=1}^n v_i a_i^2$.

Thus

$$||b_{\perp}||^2 = \sum_{i=1}^n v_i^2 a_i^2 - \frac{\left(\sum_{i=1}^n v_i a_i^2\right)^2}{\sum_{i=1}^n a_i^2}.$$

Let $A^2 = ||a||^2 = \sum_{i=1}^n a_i^2$. Define weights

$$s_i = \frac{a_i^2}{A^2}$$
 so that $s_1 + s_2 + \dots + s_n = 1$.

Then

$$\sum_{i=1}^{n} v_i^2 a_i^2 = A^2 \sum_{i=1}^{n} v_i^2 s_i, \quad \sum_{i=1}^{n} v_i a_i^2 = A^2 \sum_{i=1}^{n} v_i s_i$$

Hence

$$||b_{\perp}||^2 = A^2 \left[\sum_{i=1}^n v_i^2 s_i - \left(\sum_{i=1}^n v_i s_i \right)^2 \right]$$

Recognize that

$$\sum_{i=1}^{n} v_i^2 s_i - \left(\sum_{i=1}^{n} v_i s_i\right)^2$$

is precisely the variance $\operatorname{Var}_s(V)$ of the values v_i with respect to the probability distribution $\{s_i\}$. Since $m \leq v_i \leq M$ for all *i*, the variance of any random variable confined to [m, M] is at most $\frac{(M-m)^2}{4}$. (This is realized, for instance, by taking a two-point distribution at *m* and *M* with equal probabilities.) Therefore,

$$\sum_{i=1}^{n} v_i^2 s_i - \left(\sum_{i=1}^{n} v_i s_i\right)^2 \le \frac{(M-m)^2}{4}.$$

Consequently,

$$||b_{\perp}||^2 \leq A^2 \cdot \frac{(M-m)^2}{4} = \frac{(M-m)^2}{4} ||a||^2$$

and hence

$$\|b_{\perp}\| \leq \frac{M-m}{2} \|a\|$$

Substituting $M = m \Gamma$ (so $M - m = m(\Gamma - 1)$) gives

$$\|b_{\perp}\| \leq \frac{m(\Gamma-1)}{2} \|a\|.$$

For the alternative form in terms of ||b||, note that

$$||b|| = \sqrt{\sum_{i=1}^{n} (v_i a_i)^2} \ge \sqrt{\sum_{i=1}^{n} m^2 a_i^2} = m ||a||.$$

Hence $||a|| \leq \frac{1}{m} ||b||$, and

$$||b_{\perp}|| \le \frac{(M-m)}{2} ||a|| \le \frac{(M-m)}{2m} ||b|| = \frac{\Gamma-1}{2} ||b||.$$

This completes the proof.

Theorem C.2. Let $\mathbf{W}(t) \in \mathbb{R}^{m \times n}$ be the weight matrix of a neural network optimized using AdamW at the t-th iteration. Suppose the rows of the gradient matrix $\nabla_L(\mathbf{W}(t))$ lie within a subspace $S \subseteq \mathbb{R}^n$. Let Δ_t be defined as

$$\mathbf{W}(t+1) = \mathbf{W}(t) - \eta \left(\Delta_t + \alpha \mathbf{W}(t) \right)$$
(35)

where η is the learning rate and,

$$\Delta_t = \mathbf{M}_t \oslash (\sqrt{\mathbf{V}} + \epsilon) \tag{36}$$

where $(\cdot \otimes \cdot)$ denotes the element wise Hadamard division and $(\alpha, \epsilon) > 0$ where : α is the weight decay and ϵ is a constant added to avoid division by 0. \mathbf{M}_t is the first momentum matrix defined as

$$\mathbf{M}_t = \beta_1 \mathbf{M}_{t-1} + (1 - \beta_1) \nabla_L(\mathbf{W}(t))$$
(37)

and \mathbf{V} is the second moment,

$$\mathbf{V}_t = \beta_2 \mathbf{V}_{t-1} + (1 - \beta_2) \nabla_L (\mathbf{W}(t))^2 \tag{38}$$

Also, let Δ_t^{\perp} be the matrix whose rows consist of the components orthogonal to S from the rows of Δ_t . Assume that g_t^2 follows a sub-Gaussian distribution with a variance proxy $(\sigma(t))^2$ where g_t is any element of $\nabla_L(\mathbf{W}(t))$. Then, with probability of at least $1 - \delta$, $\lim_{t \to \infty} \frac{\|\Delta_t^{\perp}\|}{\|\Delta_t\|} = 0$ if $\lim_{t \to \infty} \sigma(t) = 0$.

Proof. Since the first-moment estimate \mathbf{M}_t in AdamW remains a (biased) linear combination of past gradients $\nabla_L(\mathbf{W}(t))$, its rows stays within \mathcal{S} (recall that vector spaces are closed under addition). However, the rows of second-moment estimate \mathbf{V}_t are not constant-value vectors, and thus, the element-wise division of \mathbf{M}_t by $(\sqrt{\mathbf{V}} + \epsilon)$ distorts and pushes the rows of \mathbf{M}_t outsied \mathcal{S} . Consequently, the rows of Δ_t are not confined to \mathcal{S} . Our goal is to show that $\frac{\|\Delta_t^{\perp}\|}{\|\Delta_t\|}$ asymptotically goes to zero with the variance decay.

Step 1: Bounding the second-moment ratio. Since the distribution of g_t^2 is sub-Gaussian, then for each g_t , we have

$$P(g_t^2 \ge x) \le 2 \exp\left(-\frac{x^2}{2\sigma^2}\right). \tag{39}$$

Note that we drop the coordinate indexes (i, j) from g_t for brevity. Using a union bound over $m \times n$ coordinates and at any given time t, we have

$$P(\max_{i,j} g_t^2 \ge x) \le 2mn \exp\left(-\frac{x^2}{2(\sigma(t))^2}\right)$$
(40)

Then, with probability at least $1 - \delta$, we get

$$\max_{i,j}(g_t^2) \le \sigma(t) \sqrt{2\ln\left(\frac{2mn}{\delta}\right)} \tag{41}$$

Since any element v_t of \mathbf{V}_t is a moving average of g_t^2 , we can expand it as,

$$v_t = \beta_2^{(t-1)} (1-\beta_2) g_1^2 + \beta_2^{(t-2)} (1-\beta_2) g_2^2 + \dots + \beta_2^2 (1-\beta_2) g_{t-2}^2 + \beta_2 (1-\beta_2) g_{t-1}^2 + (1-\beta_2) g_t^2$$
(42)

(assuming $v_0 = 0$). Note that for large t, we can have a $\tau(k)$ such that $v_t - \tau(k) < \mu$ for arbitrary small μ , where

$$\tau(k) = \beta_2^{(t-k)} (1-\beta_2) g_k^2 + \beta_2^{(t-k-1)} (1-\beta_2) g_{k+1}^2 + \dots \beta_2^2 (1-\beta_2) g_{t-2}^2 + \beta_2 (1-\beta_2) g_{t-1}^2 + (1-\beta_2) g_t^2$$
(43)

Plugging the above result to Eq. 41, with probability at least $1 - \delta$, we again have

$$\max_{i,j} [v_t]_{i,j} \le \max_{t \in [t-k,T]} \sigma(t) \sqrt{2 \ln\left(\frac{2mn}{\delta}\right)}$$
(44)

Then,

$$\frac{\sqrt{\max_{i,j}[v_t]_{i,j}} + \epsilon}{\sqrt{\min_{i,j}[v_t]_{i,j}} + \epsilon} \le 1 + \frac{\sqrt{\max_{i,j}[v_t]_{i,j}}}{\epsilon} \le 1 + \frac{\sqrt{\max_{t \in [t-k,T]} \sigma(t)}\sqrt{2\ln\left(\frac{2mn}{\delta}\right)}}{\epsilon}, \quad (45)$$

since $\sqrt{\min_{i,j} [v_t]_{i,j}} = 0$. Let $\kappa(t) = \frac{\sqrt{\max_{t \in [t-k,T]} \sigma(t)} \sqrt{2 \ln\left(\frac{2mn}{\delta}\right)}}{\epsilon}$. Then, the coordinate-wise learning-rate scaling $\frac{1}{\sqrt{v_t + \epsilon}}$ cannot differ among coordinates by more than $1 + \kappa(t)$ with a high probability.

Step 2: Relating off-subspace updates to the ratio. Consider any row vector of Δ_t to be $(\Delta_t)_{i,:}$. Then, by Applying the update Eq. 45 to Lemma 1, and by definition of Δ_t , we have

$$\|(\Delta_t)_{i,:}^{\perp}\| \le \frac{\kappa(t)}{2} \|(\Delta_t)_{i,:}\|$$
(46)

Step 3: Taking the limit $t \to \infty$. Because we have

$$\lim_{t \to \infty} \sigma(t) = 0, \tag{47}$$

trivially, we also have

$$\lim_{T \to \infty} \max_{t \in [T-k,T]} \sigma(t) = 0.$$
(48)

Then, it follows from the definition of $\kappa(t)$ that

$$\lim_{t \to \infty} \kappa(t) = 0. \tag{49}$$

Therefore, from Eq. 46, assuming $\|(\Delta_t)_{i,:}\| \neq 0$ we finally have

$$0 \le \lim_{t \to \infty} \frac{\|(\Delta_t)_{i,:}^{\perp}\|}{\|(\Delta_t)_{i,:}\|} \le 0$$
(50)

$$= \lim_{t \to \infty} \frac{\|(\Delta_t)_{i,:}^{\perp}\|}{\|(\Delta_t)_{i,:}\|} = 0$$
(51)

In simpler terms, the distortion of the rows of Δ_t outside S is ultimately goes to zero.

Discussion: Above theorem gives us an important insight. In particular, it ensures that if the gradients of an unrestricted network predominantly lie in a specific low-dimensional subspace S_{t} then the AdamW optimizer restricts updates outside that subspace. On the other hand, as training progresses, gradient norms decrease and become more uniform across coordinates, which has been rigorously proved in previous works within reasonable bounds (Défossez et al., 2020; Li et al., 2024a; Zhang et al., 2024). Consequently, the variance of the gradients also tends to diminish, justifying our assumption of $\lim_{t\to\infty} \sigma(t) \to 0$. This reduction further curtails the extent of any off-subspace updates, thereby indicating that AdamW naturally remains focused on the directions most important to learning.

Further, empirical evidence supports the notion that gradients of unconstrained networks often lie predominantly within a low-dimensional subspace (we also empirically validate this in Fig. 7). This observation aligns with the hypothesis that many deep learning problems exhibit significant redundancy in parameter updates, with key directions of optimization being confined to a smaller subspace. By leveraging this inherent structure, it becomes possible to guide optimization dynamics in a manner that preserves computational efficiency while maintaining convergence. We conduct an extended discussion next.

C.1. On the assumptions of Theorem C.2

For Theorem C.2, we employ two assumptions 1) the gradients of weight matrices are confined to a low-dimensional subspace and 2) the variance of the gradients diminish as the model converges. Next, we provide justifications for these assumptions.

C.1.1. The gradients of weight matrices are confined to a low-dimensional subspace

To empirically validate this behavior, we train an 8-layer network on C4 and track the stable rank progression of its projection matrices throughout training. As shown in Fig. 7, the gradients



Figure 7: Low dimensionality of the gradients of projection matrices. We measure the stable ranks of gradients of the projection matrices over the course of training. In each plot, the green dashed line indicates the maximum rank. As shown, all the weight matrices consistently maintain a significantly lower rank ($\sim < 5$) throughout the training. This validates the assumption of our theorem C.2 that the weight gradients of the projection matrices are confined to a low-dimensional subspace.

maintain an exceptionally low stable rank relative to the maximum possible rank from initialization to convergence. This confirms that the gradients primarily focus on a few dominant directions.

The assumption that weight gradients span a low-dimensional subspace is well-supported by prior research. (Gur-Ari et al., 2018) demonstrated that gradient descent in deep learning does not explore the full parameter space but instead operates within a constrained subspace, largely dictated by the spectral structure of the Hessian. The Hessian matrix features a dominant subspace associated with large eigenvalues, where most gradient updates occur, and a bulk subspace characterized by near-zero eigenvalues that correspond to flat directions in the loss landscape. Empirical evidence indicates that early in training, gradients align with this dominant Hessian subspace and remain confined within it throughout the optimization process.

This phenomenon extends to the learned representations themselves. (Zhao et al., 2021) showed that gradient descent naturally gravitates toward low-rank solutions, a behavior arising from implicit regularization within the optimization process. Instead of initializing with full-rank representations, training typically begins in a low-rank regime and expands rank as necessary. This trend has been consistently observed across multiple architectures, including MLPs, CNNs, and transformers. The concept of spectral bias further supports this observation, suggesting that neural networks initially learn broad, low-rank structures before refining finer details.

A practical application of this low-rank property is evident in (Cosson et al., 2023)'s introduction of Low-Rank Gradient Descent (LRGD), which explicitly leverages low-rank subspaces to expedite training. By identifying dominant gradient directions and confining updates to this reduced subspace, LRGD significantly lowers computational costs without sacrificing optimization effectiveness. Theoretical results demonstrate that in strongly convex settings, LRGD reduces oracle complexity from $O(p \log(1/\epsilon))$ to $O(r \log(1/\epsilon) + rp)$, where $r \ll p$. Similar improvements are observed in non-convex environments. Moreover, (Wangni et al., 2018) illustrated that in deep networks trained with cross-entropy loss, the Hessian spectrum exhibits a clear separation. The top subspace, whose dimensionality approximates the number of classes in the dataset, captures the majority of gradient activity. In contrast, the bulk subspace, filled with small eigenvalues, corresponds to nearly flat directions in the loss landscape. This structure implies that significant changes in the loss function occur primarily within the top subspace. Empirical findings also reveal that gradients quickly align with the top Hessian subspace early in training and persist within it. This effect has been documented across diverse architectures such as fully connected networks, convolutional networks, ResNets, and transformers. While individual Hessian eigenvectors evolve during training, the dominant subspace remains relatively stable, guiding the optimization path. The low-rank nature of gradients is further validated by (Vogels et al., 2019), who observed a rapid decay in the singular values of gradient matrices. This decay implies that the most meaningful updates are confined to a small, structured subspace. Additionally, the gradient covariance matrix in deep networks shows exponential spectral decay, suggesting that full-rank updates are not essential.

Above findings provide robust empirical and theoretical support for our low-rank gradient assumption, directly motivating our approach. By exploiting low-rank constraints, we achieve substantial reductions in both computation and communication overhead while preserving performance. Given that deep learning models inherently exhibit implicit regularization, explicitly enforcing a low-rank structure in our compression could even induce a beneficial regularization effect—potentially enhancing performance beyond that of non-compressed models, as indicated in Table 1.

C.1.2. The variance of the gradients diminish as the model converges.

The reduction in gradient variance as a model converges is a well-observed phenomenon in optimization literature. As training progresses, stochastic gradient estimates become more stable, leading to smoother updates and accelerated convergence. This effect is particularly pronounced in adaptive learning rate optimizers such as Adam (Kingma, 2014) and AdamW (Loshchilov, 2017). In particular, (Balles and Hennig, 2018) showed that these methods inherently reduce the variance of parameter updates over time, contributing to stable convergence. (Liu et al., 2020) further refined this idea by introducing a variant of Adam designed to enforce adaptive variance reduction, ensuring that the variance of the stochastic gradient estimator decreases as the model approaches optimality.

Beyond adaptive optimization, variance reduction techniques have been widely studied in the context of stochastic gradient descent (SGD). Traditional SGD suffers from high variance in gradient estimates, particularly in early training stages, which can hinder convergence. To mitigate this, methods such as Stochastic Variance Reduced Gradient (SVRG) and SAGA (Wang et al., 2013) have been proposed. These approaches reduce variance by incorporating control variates or maintaining a memory of past gradients, yielding more accurate updates and improved convergence rates. Additional techniques, including Kalman-based filtering (Vuckovic, 2018) and stochastic filtering methods (Yang, 2020), further enhance stability by adaptively reducing noise in gradient estimates.

Normalization layers have also been shown to smooth parameter spaces and gradients, effectively reducing variance. This smoothness, characterized by a lower Lipschitz constant, results in more predictive gradients, which can be interpreted as a form of implicit variance reduction. The effect is particularly beneficial in deep networks, where sharp loss landscapes can otherwise lead to erratic updates.

In summary, diminishing gradient variance is a natural outcome of optimization dynamics. As training progresses and model parameters approach optimality, stochastic gradients become increasingly stable, even in the absence of explicit variance reduction techniques. This effect can be attributed to the curvature of the loss landscape: near an optimal solution, the loss function tends to be flatter, leading to more consistent gradient estimates across mini-batches. Consequently, as the model converges, gradient variance inherently diminishes, reinforcing the validity of this assumption in our work.

D. Continuous-Time Gradient Flow with Weight Decay Converges to a Low-Dimensional Subspace

The result in this section reveals a fundamental property of optimization under AdamW: the optimizer inherently drives weight matrices to converge asymptotically to a shared subspace, regardless of their initial values. This behavior stems from the decoupled weight decay mechanism, which systematically suppresses components of the weight matrix that receive negligible or noisy gradient updates during training. Over time, the weight matrix becomes confined to a low-dimensional subspace spanned by the significant directions of the gradient updates, leading to an effectively low-rank representation.

Theorem D.1. Consider a continuous-time dynamics for $\mathbf{W}(t) \in \mathbb{R}^{m \times n}$ governed by

$$\frac{\mathrm{d}\mathbf{W}}{\mathrm{d}t} = -\eta \left(\Delta_t + \alpha \mathbf{W}(t)\right),\,$$

where η , $\alpha > 0$ are constants, and Δ_t represents the gradient-like update without weight decay (i.e. the "AdamW" part, ignoring the $\alpha \mathbf{W}(t)$ term). Assume that for large t, Δ_t has its rows confined predominantly to a subspace $S \subseteq \mathbb{R}^n$ (so that the orthogonal component of Δ_t vanishes asymptotically). Then, as $t \to \infty$, the component of $\mathbf{W}(t)$ lying in S^{\perp} decays exponentially to zero. Consequently, the rows of the final solution $\mathbf{W}(\infty)$ lie in (or arbitrarily close to) S.

Proof. Let $\Pi_{\mathcal{S}}$ denote the orthogonal projector onto \mathcal{S} (applied rowwise to any $m \times n$ matrix). Decompose:

$$\mathbf{W}(t) = \mathbf{W}_{\parallel}(t) + \mathbf{W}_{\perp}(t), \text{ where } \mathbf{W}_{\parallel}(t) = \Pi_{\mathcal{S}} \big[\mathbf{W}(t) \big], \quad \mathbf{W}_{\perp}(t) = \big(\mathbf{I} - \Pi_{\mathcal{S}} \big) \mathbf{W}(t).$$

By hypothesis (see also Theorem C.2), the rows of Δ_t become confined to \mathcal{S} as $t \to \infty$, i.e.

$$\lim_{t\to\infty} (\mathbf{I} - \Pi_{\mathcal{S}}) \Delta_t = 0.$$

Thus, for large t, the orthogonal component of Δ_t is negligible. Project the ODE onto \mathcal{S}^{\perp} :

$$\frac{\mathrm{d}\mathbf{W}_{\perp}(t)}{\mathrm{d}t} = \left(\mathbf{I} - \Pi_{\mathcal{S}}\right) \frac{\mathrm{d}\mathbf{W}}{\mathrm{d}t} = -\eta \left(\mathbf{I} - \Pi_{\mathcal{S}}\right) \left(\Delta_t + \alpha \mathbf{W}(t)\right).$$

For sufficiently large t, $(I - \Pi_S)\Delta_t$ is negligible, so asymptotically,

$$\frac{\mathrm{d}\mathbf{W}_{\perp}(t)}{\mathrm{d}t} \approx -\eta \,\alpha \left(\mathbf{I} - \Pi_{\mathcal{S}}\right) \left(\mathbf{W}(t)\right) = -\eta \,\alpha \,\mathbf{W}_{\perp}(t).$$

This is a linear ODE:

$$\frac{\mathrm{d}\mathbf{W}_{\perp}(t)}{\mathrm{d}t} = -\eta \,\alpha \,\mathbf{W}_{\perp}(t),$$

whose solution is

$$\mathbf{W}_{\perp}(t) = \exp(-\eta \,\alpha \,t) \,\mathbf{W}_{\perp}(0).$$

Hence $\mathbf{W}_{\perp}(t)$ decays exponentially to **0**. Therefore,

$$\mathbf{W}(t) \;=\; \mathbf{W}_{\parallel}(t) + \mathbf{W}_{\perp}(t) \;\longrightarrow\; \mathbf{W}_{\parallel}(t), \quad \text{with} \quad \mathbf{W}_{\parallel}(t) \;\in\; \mathcal{S}.$$

Thus, as $t \to \infty$, the rows of $\mathbf{W}(t)$ lie in \mathcal{S} (or arbitrarily close), proving the claim.

Discussion: Unlike standard stochastic gradient descent (SGD), which applies weight decay as a modification to the gradient updates, AdamW separates the weight decay term from the gradient computation. This decoupled weight decay has two key effects: 1.) **Suppression of Insignificant Components:** Components of the weight matrix orthogonal to the subspace spanned by significant gradient updates decay exponentially due to weight decay. This ensures that the optimization focuses on meaningful directions of the gradient. 2) **Low-Dimensional Convergence:** By amplifying the impact of directions where the gradient signal is strongest, the optimizer effectively projects the weight matrix onto a subspace defined by these directions, leaving other components asymptotically negligible.

To empirically validate this property, we analyzed the spectral structure of weight matrices trained under AdamW. As shown in Fig. 1, the spectrum of these matrices exhibits rapid decay, with the majority of the spectral mass concentrated in a few dominant directions.

E. Convergence of Partially-Projected Gradient Descent

This section establishes that constraining a subset of weights in a neural network to lie within a low-dimensional subspace does not degrade the theoretical convergence guarantees of constrained optimization. Specifically, the network still achieves convergence to a first-order stationary point at the rate of O(1/T), where T is the number of optimization steps. This is consistent with standard results in constrained optimization and provides a preliminary theoretical foundation for employing subspace constraints in practical applications.

Proposition E.1. Let $f(\mathbf{W}^{(s)}, \mathbf{W}^{(u)})$ be a possibly non-convex function that is L-smooth in all its parameters $\mathbf{W} = (\mathbf{W}^{(s)}, \mathbf{W}^{(u)})$, where $\mathbf{W}^{(s)} \in \mathbb{R}^k$ is a constrained block and $\mathbf{W}^{(u)} \in \mathbb{R}^d$ is unconstrained and k < d. Suppose we have a closed subspace $S \subseteq \mathbb{R}^k$, and we define the feasible set

$$\mathbf{W} \;=\; \mathcal{S} imes \mathbb{R}^{d_u}$$

Consider the following partially-projected gradient step:

$$\begin{split} \widetilde{\mathbf{W}}_{t+1}^{(s)} &= \mathbf{W}_{t}^{(s)} - \eta \, \nabla_{(s)} f(\mathbf{W}_{t}), \quad \widetilde{\mathbf{W}}_{t+1}^{(u)} &= \mathbf{W}_{t}^{(u)} - \eta \, \nabla_{(u)} f(\mathbf{W}_{t}), \\ \mathbf{W}_{t+1}^{(s)} &= \Pi_{\mathcal{S}} \Big(\widetilde{\mathbf{W}}_{t+1}^{(s)} \Big), \quad \mathbf{W}_{t+1}^{(u)} &= \widetilde{\mathbf{W}}_{t+1}^{(u)}, \end{split}$$

where $\Pi_{\mathcal{S}}(\cdot)$ is the Euclidean projection onto \mathcal{S} , $\nabla_{(S)}f$ and $\nabla_{(u)}f$ are the gradients of f w.r.t. the constrained and unconstrained blocks, respectively, and $\eta > 0$ is a fixed step size.

Then, for any $T \ge 1$ and $\eta \le \frac{1}{L}$, the iterates $\{\mathbf{W}_t\}$ satisfy the following stationary guarantee:

$$\min_{0 \le t < T} \left\| \nabla f(\mathbf{W}_t) \right\|^2 \le \frac{2 \left(f(\mathbf{W}_0) - f^* \right)}{\eta T}$$

where $f^* = \min_{\mathbf{W} \in \mathbf{W}} f(\mathbf{W})$ and \mathbf{W}_0 is the initial parameter vector.

Proof. Since $f(\mathbf{W}^{(s)}, \mathbf{W}^{(u)})$ is L-smooth in all parameters, for any \mathbf{W}, \mathbf{W}' we have

$$f(\mathbf{W}') \leq f(\mathbf{W}) + \nabla f(\mathbf{W})^{\top} (\mathbf{W}' - \mathbf{W}) + \frac{L}{2} \|\mathbf{W}' - \mathbf{W}\|^2.$$
(52)

Setting $\mathbf{W} = \mathbf{W}_t$ and $\mathbf{W}' = \mathbf{W}_{t+1}$, we get

$$f(\mathbf{W}_{t+1}) \leq f(\mathbf{W}_t) + \nabla f(\mathbf{W}_t)^{\top} (\mathbf{W}_{t+1} - \mathbf{W}_t) + \frac{L}{2} \|\mathbf{W}_{t+1} - \mathbf{W}_t\|^2.$$
(53)

Recall the update:

$$\mathbf{W}_{t+1}^{(s)} = \Pi_{\mathcal{S}} \Big(\mathbf{W}_{t}^{(S)} - \eta \,\nabla_{(S)} f(\mathbf{W}_{t}) \Big), \quad \mathbf{W}_{t+1}^{(u)} = \mathbf{W}_{t}^{(u)} - \eta \,\nabla_{(u)} f(\mathbf{W}_{t}).$$
(54)

$$\mathbf{W}_{t+1} = \left(\Pi_{\mathcal{S}}(\mathbf{W}_t^{(s)} - \eta \,\nabla_{(s)} f(\mathbf{W}_t)), \, \mathbf{W}_t^{(u)} - \eta \,\nabla_{(u)} f(\mathbf{W}_t) \right).$$
(55)

Projecting the $\widetilde{\mathbf{W}}_{t+1}^{(s)}$ block onto \mathcal{S} , is a 1-Lipschitz (nonexpansive) operation in the full Euclidean space. Concretely, consider $\|(\Pi_{\mathcal{S}}(\widetilde{\mathbf{W}}_{t+1}^{(s)}) - \mathbf{W}_{t}^{(s)})\|$. Also recall that $\mathbf{W}_{t}^{s} \in \mathcal{S}$. It follows from the classical projection lemma that,

$$\|(\mathbf{W}_{t+1}^{(s)}) - \mathbf{W}_{t}^{(s)})\| = \|\left(\Pi_{\mathcal{S}}(\widetilde{\mathbf{W}}_{t+1}^{(s)}) - \mathbf{W}_{t}^{(s)})\| \le \|\left(\widetilde{\mathbf{W}}_{t+1}^{(s)} - \mathbf{W}_{t}^{(s)})\|\right)\|$$
(56)

This happens because the projection minimizes the Frobenius norm to the subset, so the projected point is closer than the unprotected one.

We also leave the disjoint set $\mathbf{W}^{(u)}$ untouched. So we have,

$$\widetilde{\mathbf{W}}_{t+1}^{(u)} = \mathbf{W}_{t+1}^{(u)} \tag{57}$$

Since $\mathbf{W}^{(s)}, \mathbf{W}^{(u)}$ are disjoint blocks, We can combine the blocks (Pythagorean identity): The Frobenius norm is the Euclidean norm on the concatenated vector, so

$$\left\|\mathbf{W}_{t+1} - \mathbf{W}_{t}\right\|^{2} = \left\|\left(\mathbf{W}_{t+1}^{(s)} - \mathbf{W}_{t}^{(s)}\right)\right\|^{2} + \left\|\left(\mathbf{W}_{t+1}^{(u)} - \mathbf{W}_{t}^{(u)}\right)\right\|^{2}$$

From Eq. 56 and 57 we have,

$$\left\|\mathbf{W}_{t+1} - \mathbf{W}_{t}\right\|^{2} \le \left\|\left(\widetilde{\mathbf{W}}_{t+1}^{(s)} - \mathbf{W}_{t}^{(s)}\right)\right\|^{2} + \left\|\left(\widetilde{\mathbf{W}}_{t+1}^{(u)} - \mathbf{W}_{t}^{(u)}\right)\right\|^{2}$$
(58)

As $\mathbf{W}^{(s)}, \mathbf{W}^{(u)}$ are disjoint, we get,

$$\left\|\mathbf{W}_{t+1} - \mathbf{W}_{t}\right\|^{2} \le \left\|\widetilde{\mathbf{W}}_{t+1} - \mathbf{W}_{t}\right\|^{2}$$
(59)

$$\left\|\mathbf{W}_{t+1} - \mathbf{W}_t\right\| \le \left\|\widetilde{\mathbf{W}}_{t+1} - \mathbf{W}_t\right\|$$
(60)

 $(\widetilde{\mathbf{W}}_{t+1} - \mathbf{W}_t)$ is the update step which is equal to $\eta(\nabla f(\mathbf{W}_t))$. Thus,

$$\|\mathbf{W}_{t+1} - \mathbf{W}_t\| \leq \eta \|\nabla f(\mathbf{W}_t)\|.$$
(61)

Plugging this bound back into the smoothness inequality (Eq. 53), we can get

$$f(\mathbf{W}_{t+1}) \leq f(\mathbf{W}_t) + \nabla f(\mathbf{W}_t)^{\top} (\mathbf{W}_{t+1} - \mathbf{W}_t) + \frac{L}{2} \eta^2 \left\| \nabla f(\mathbf{W}_t) \right\|^2.$$
(62)

Further, from the optimality condition, it follows that,

$$\left(\mathbf{W}_{t+1} - \left(\mathbf{W}_{t} - \eta \nabla f(\mathbf{W}_{t})\right)\right)^{\top} (\mathbf{W}_{t+1} - \mathbf{W}_{t}) \le 0$$
(63)

$$\left(\eta \nabla f(\mathbf{W}_t)\right)^{\top} (\mathbf{W}_{t+1} - \mathbf{W}_t) \le -\|\mathbf{W}_{t+1} - \mathbf{W}_t\|^2$$
(64)

$$\nabla f(\mathbf{W}_t)^{\top} (\mathbf{W}_{t+1} - \mathbf{W}_t) \le -\frac{1}{\eta} \|\mathbf{W}_{t+1} - \mathbf{W}_t\|^2$$
(65)

Substituting above in Eq. 62,

$$f(\mathbf{W}_{t+1}) \leq f(\mathbf{W}_{t}) - \frac{1}{\eta} \|\mathbf{W}_{t+1} - \mathbf{W}_{t}\|^{2} + \frac{L}{2} \eta^{2} \|\nabla f(\mathbf{W}_{t})\|^{2}.$$
 (66)

$$f(\mathbf{W}_{t+1}) \leq f(\mathbf{W}_t) - \frac{1}{\eta} \|\mathbf{W}_{t+1} - \mathbf{W}_t\|^2 + \frac{L}{2} \|\mathbf{W}_{t+1} - \mathbf{W}_t\|^2.$$
(67)

$$f(\mathbf{W}_{t+1}) \leq f(\mathbf{W}_t) + \|\mathbf{W}_{t+1} - \mathbf{W}_t\|^2 \Big(\frac{L}{2} - \frac{1}{\eta}\Big).$$
 (68)

$$f(\mathbf{W}_{t+1}) - f(\mathbf{W}_t) \leq \|\mathbf{W}_{t+1} - \mathbf{W}_t\|^2 \Big(\frac{L}{2} - \frac{1}{\eta}\Big).$$
 (69)

$$f(\mathbf{W}_{t+1}) - f(\mathbf{W}_t) \leq \eta^2 \|\nabla f(\mathbf{W}_t)\|^2 \Big(\frac{L}{2} - \frac{1}{\eta}\Big).$$
 (70)

Choosing $\eta < \frac{2}{L}$

$$f(\mathbf{W}_{t+1}) - f(\mathbf{W}_t) \leq -\gamma \|\nabla f(\mathbf{W}_t)\|^2.$$
(71)

Taking a telescoping summation, we have,

$$\sum_{t=0}^{T-1} \left(f(\mathbf{W}_{t+1}) - f(\mathbf{W}_t) \right) \leq -\gamma \sum_{t=0}^{T-1} \|\nabla f(\mathbf{W}_t)\|^2$$
(72)

$$\frac{\left(f(\mathbf{W}_0) - f(\mathbf{W}_{t+1})\right)}{\gamma} \ge \sum_{t=0}^{T-1} \|\nabla f(\mathbf{W}_t)\|^2.$$
(73)

Let us set $\mathbf{W}_{t+1} = \mathbf{W}^*$, where \mathbf{W}^* is a minimizer of f. Then, we have $\left(f(\mathbf{W}_0) - f(\mathbf{W}_{t+1})\right) \leq \left(f(\mathbf{W}_0) - f(\mathbf{W}^*)\right)$, which gives,

$$\frac{\left(f(\mathbf{W}_0) - f(\mathbf{W}^*)\right)}{\gamma} \ge \sum_{t=0}^{T-1} \|\nabla f(\mathbf{W}_t)\|^2.$$
(74)

Dividing both sides by T, we get

$$\frac{\left(f(\mathbf{W}_0) - f(\mathbf{W}^*)\right)}{\gamma T} \ge \frac{1}{T} \sum_{t=0}^{T-1} \|\nabla f(\mathbf{W}_t)\|^2.$$
(75)

If the gradient norm is upper bounded as below, then, trivially we have,

$$\min_{0 \le t < T} \left\| \nabla f(\mathbf{W}_t) \right\|^2 \le \frac{2 \left(f(\mathbf{W}_0) - f(\mathbf{W}^*) \right)}{\gamma T}$$
(76)

This shows that the squared minimum gradient norm goes to zero at a O(1/T) rate, implying convergence to a stationary point in the non-convex sense.

Discussion: The result is a direct extension of the convergence guarantees for proximal gradient descent methods on non-convex functions. Proximal gradient methods are widely used in constrained optimization problems, where the update rule involves projecting the parameters onto a feasible set at each iteration. In the case of subspace-constrained networks, the feasible set corresponds to the predefined low-dimensional subspace, and the projection operator ensures that the parameters *partially* remain within this subspace. The convergence rate of O(1/T) in terms of stationarity reflects that, even in the presence of non-convexity and subspace constraints, the optimization algorithm efficiently reduces the gradient norm over time, driving the network parameters toward a stationary point. A possible implication of subspace restriction could be that it acts as an implicit regularizer, encouraging the network to learn more compact and robust representations. This can lead to improved generalization performance on unseen data. We hypothesize that this might be the reason for improved performance of the compressed models over their centralized counterparts (see Fig. 5).

Next result extends the above theorem to the stochastic mini-batch gradient descent optimization.

Theorem E.2. Let $f(\mathbf{W}^{(s)}, \mathbf{W}^{(u)})$ be *L*-smooth in $\mathbf{W} = (\mathbf{W}^{(s)}, \mathbf{W}^{(u)})$ with $\mathbf{W}^{(s)} \in \mathbb{R}^k$ (constrained block), $\mathbf{W}^{(u)} \in \mathbb{R}^d$ (unconstrained block), and let $S \subseteq \mathbb{R}^k$ be a closed subspace. Define the feasible set $\mathbf{W} := S \times \mathbb{R}^d$.

At each iteration t draw a mini-batch and compute an unbiased stochastic gradient $\mathbf{g}_t = (\mathbf{g}_t^{(s)}, \mathbf{g}_t^{(u)})$ satisfying

$$\mathbb{E}[\mathbf{g}_t \mid \mathbf{W}_t] = \nabla f(\mathbf{W}_t), \qquad \mathbb{E}[\|\mathbf{g}_t - \nabla f(\mathbf{W}_t)\|^2 \mid \mathbf{W}_t] \le \sigma^2 \quad (bounded \ variance). \tag{77}$$

Update the parameters with the stochastic partially-projected gradient (SPPG) step

$$\widetilde{\mathbf{W}}_{t+1}^{(s)} = \mathbf{W}_t^{(s)} - \eta_t \mathbf{g}_t^{(s)}, \quad \widetilde{\mathbf{W}}_{t+1}^{(u)} = \mathbf{W}_t^{(u)} - \eta_t \mathbf{g}_t^{(u)},$$
$$\mathbf{W}_{t+1}^{(s)} = \Pi_{\mathcal{S}}(\widetilde{\mathbf{W}}_{t+1}^{(s)}), \qquad \mathbf{W}_{t+1}^{(u)} = \widetilde{\mathbf{W}}_{t+1}^{(u)},$$

where each stepsize satisfies $\eta_t \in (0, 1/L]$.

Let $f^* := \min_{\mathbf{W} \in \mathbf{W}} f(\mathbf{W})$ and denote $\bar{\eta} := \frac{1}{T} \sum_{t=0}^{T-1} \eta_t$. Then for every horizon $T \ge 1$

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \left[\|\nabla f(\mathbf{W}_t)\|^2 \right] \leq \frac{2 \left(f(\mathbf{W}_0) - f^* \right)}{\bar{\eta} T} + L \bar{\eta} \sigma^2.$$
(78)

In particular, with the constant stepsize $\eta_t \equiv \eta \leq 1/L$ it holds that

$$\min_{0 \le t < T} \mathbb{E} \left[\|\nabla f(\mathbf{W}_t)\|^2 \right] \le \frac{2 \left(f(\mathbf{W}_0) - f^* \right)}{\eta T} + L \eta \sigma^2,$$

which attains the optimal $\mathcal{O}(1/\sqrt{T})$ rate when $\eta = \min\{1/L, \sqrt{2(f(\mathbf{W}_0) - f^*)/(L\sigma^2 T)}\}$.

Proof. The argument follows the deterministic proof but replaces exact gradients with the stochastic estimator and takes conditional expectations at the appropriate points.

Because the Euclidean projection onto a closed subspace is 1-Lipschitz, for every $t \|\mathbf{W}_{t+1}^{(s)} - \mathbf{W}_{t}^{(s)}\| \leq \|\widetilde{\mathbf{W}}_{t+1}^{(s)} - \mathbf{W}_{t}^{(s)}\| = \eta_{t} \|\mathbf{g}_{t}^{(s)}\|$. The unconstrained coordinates remain unprojected, hence $\|\mathbf{W}_{t+1}^{(u)} - \mathbf{W}_{t}^{(u)}\| = \eta_{t} \|\mathbf{g}_{t}^{(u)}\|$. Using the Pythagorean theorem for the two disjoint blocks,

$$\|\mathbf{W}_{t+1} - \mathbf{W}_t\| \leq \eta_t \|\mathbf{g}_t\|. \tag{79}$$

L-smoothness of f implies

$$f(\mathbf{W}_{t+1}) \le f(\mathbf{W}_t) + \nabla f(\mathbf{W}_t)^{\top} (\mathbf{W}_{t+1} - \mathbf{W}_t) + \frac{L}{2} \|\mathbf{W}_{t+1} - \mathbf{W}_t\|^2$$

Introduce the zero-mean "noise" term $\epsilon_t := \mathbf{g}_t - \nabla f(\mathbf{W}_t)$. Conditioned on \mathbf{W}_t , $\mathbb{E}[\epsilon_t] = 0$ and $\mathbb{E}[\|\epsilon_t\|^2] \leq \sigma^2$. Using equation 79 and $\mathbf{g}_t = \nabla f(\mathbf{W}_t) + \epsilon_t$,

$$\mathbb{E}[f(\mathbf{W}_{t+1}) \mid \mathbf{W}_t] \leq f(\mathbf{W}_t) - \eta_t \|\nabla f(\mathbf{W}_t)\|^2 + \frac{L\eta_t^2}{2} \mathbb{E}[\|\mathbf{g}_t\|^2 \mid \mathbf{W}_t] \\ \leq f(\mathbf{W}_t) - \eta_t \left(1 - \frac{L\eta_t}{2}\right) \|\nabla f(\mathbf{W}_t)\|^2 + \frac{L\eta_t^2}{2} \sigma^2,$$

where the last line used $\mathbb{E}[\|\mathbf{g}_t\|^2] = \|\nabla f(\mathbf{W}_t)\|^2 + \mathbb{E}[\|\epsilon_t\|^2] \le \|\nabla f(\mathbf{W}_t)\|^2 + \sigma^2$. Because $\eta_t \le 1/L$, $1 - L\eta_t/2 \ge 1/2$. Hence

$$\mathbb{E}[f(\mathbf{W}_{t+1})] \leq \mathbb{E}[f(\mathbf{W}_t)] - \frac{\eta_t}{2} \mathbb{E}\left[\|\nabla f(\mathbf{W}_t)\|^2\right] + \frac{L\eta_t^2}{2}\sigma^2.$$
(80)

Sum equation 80 over t = 0, ..., T - 1 and rearrange:

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} \Big[\|\nabla f(\mathbf{W}_t)\|^2 \Big] \leq \frac{2 \Big(\mathbb{E} [f(\mathbf{W}_0)] - f^* \Big)}{\sum_{t=0}^{T-1} \eta_t} + L \sigma^2 \frac{\sum_{t=0}^{T-1} \eta_t^2}{\sum_{t=0}^{T-1} \eta_t}$$

Noting that $\sum_t \eta_t = T\bar{\eta}$ and $\sum_t \eta_t^2 \leq T\bar{\eta}^2$ (by Jensen's inequality) yields

$$\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\mathbf{W}_t)\|^2] \leq \frac{2(f(\mathbf{W}_0) - f^*)}{\bar{\eta} T} + L \,\bar{\eta} \,\sigma^2$$

which is precisely equation 78.

F. Ablations

To validate the utility of the proposed method, we stress-test it with different network and training configurations. All the models are trained on C4.



Figure 8: **Convergence with different batch sizes.** Each block shows the training curves (with respect to training steps) and throughput for an 8-layer, 2B-parameter model on C4. Decentralized configurations use 80 Mbps connections, while the centralized configuration uses datacenter-grade 100 Gbps links. Even under a 80 Mbps bandwidth budget, our compressed model achieves convergence comparable to—and sometimes exceeding—the centralized configuration across varying batch sizes. Note that as the batch size increases, the compressed model achieves increasingly better results compared to the centralized model. The iteration-wise dynamics of the uncompressed decentralized model match those of the centralized model; hence, we omit its curve for clarity. Despite severe bandwidth constraints, our compressed model attains throughput on par with the centralized setting. In contrast, the uncompressed decentralized setup suffers significantly lower throughput due to communication bottlenecks. All models share the same network architecture.



Figure 9: **Convergence with different batch sizes.** Each block shows the training curves (with respect to wall-clock time) and throughput for an 8-layer, 2B-parameter model on C4. Decentralized configurations use 80 Mbps connections, while the centralized configuration uses datacenter-grade 100 Gbps links. Even under a 80 Mbps bandwidth budget, our compressed model achieves convergence comparable to—and sometimes exceeding—the centralized configuration across varying batch sizes while the decentralized non compressed model demonstrates extremely slow convergence.

G. Analysis of pretrained checkpoints

We investigate the stable ranks of the output projection layers of the official checkpoints of frontier open-weight LLMs (LlaMA, Qwen, Olmo, Phi). The observations are reported in Fig. 16. As shown, the weights demonstrate extremely low ranks across all the layers and models.



Figure 10: Convergence with different context lengths. Each block shows the training curves (with respect to training steps) and throughput for an 8-layer, 2B-parameter model on C4. Decentralized configurations use 80 Mbps connections, while the centralized configuration uses datacenter-grade 100 Gbps links. Even under a 80 Mbps bandwidth budget, our compressed model achieves convergence comparable to—and sometimes exceeding—the centralized configuration across varying context lengths. Note that as the context length increases, the compressed model achieves increasingly better results compared to the centralized model. The iteration-wise dynamics of the uncompressed decentralized model match those of the centralized model; hence, we omit its curve for clarity. Despite severe bandwidth constraints, our compressed model attains throughput on par with the centralized setting. In contrast, the uncompressed decentralized setup suffers significantly lower throughput due to communication bottlenecks.



Figure 11: Convergence with different context lengths. Each block shows the training curves (with respect to wall-clock time) and throughput for an 8-layer, 2B-parameter model on C4. Decentralized configurations use 80 Mbps connections, while the centralized configuration uses datacenter-grade 100 Gbps links. Even under a 80 Mbps bandwidth budget, our compressed model achieves convergence comparable to—and sometimes exceeding—the centralized configuration across varying context lengths while the decentralized non compressed model demonstrates extremely slow convergence.



Figure 12: **Convergence with increasing number of layers.** Note that as the number of layer increase, our model consistently matches (even slightly exceeds) the centralized model. This is in stark contrast to lossy compression schemes, where the model convergence severely degrades as the model depth increases (Bian et al., 2024).



Figure 13: Throughput across bandwidth constraints. We limit network bandwidth between GPUs and measure throughput during inference and training. The compressed model consistently achieves significantly higher throughput than the non-compressed model. Notably, even at 100Gbps, compression improves inference throughput by $3\times$, demonstrating benefits for even centralized setups.



Figure 14: Effect of the (Grassman) subspace updates. An 8-layer model is trained on C4 for this experiment. Since this performance gap seems to keep increasing towards the end of training, we emphasize the importance of infrequent subspace updates.



Figure 15: Effect of fixed token embedding. We decompose the token embeddings to a fixed high rank embedding and a dynamic low rank embedding. As shown in the figure, we observe inferior convergence when there is no such decomposition. An 8-layer model is trained on C4 for this experiment.



Figure 16: Stable ranks of output projection matrices (normalized by the maximum possible rank) across different frontier models. Statistics are computed on official fully pretrained checkpoints. All the models demonstrate extremely low ranks, solidifying our theoretical arguments.