

An All-Reduce Compatible Top- K Compressor for Communication-Efficient Distributed Learning

Chuyan Chen*
Peking University
Beijing, China
chuyanchen@stu.pku.edu.cn

Chenyang Ma*
Peking University
Beijing, China
2300010754@stu.pku.edu.cn

Zhangxin Li*
Peking University
Beijing, China
2200011085@stu.pku.edu.cn

Yutong He
Peking University
Beijing, China
yutonghe@pku.edu.cn

Yanjie Dong†
Shenzhen MSU-BIT University
Shenzhen, P. R. China
ydong@smbu.edu.cn

Kun Yuan†
Peking University
Beijing, P. R. China
kunyuan@pku.edu.cn

Abstract—Communication remains a central bottleneck in large-scale distributed machine learning, and gradient sparsification has emerged as a promising strategy to alleviate this challenge. However, existing gradient compressors face notable limitations: **Rand- K** discards structural information and performs poorly in practice, while **Top- K** preserves informative entries but loses the contraction property and requires costly **All-Gather** operations. In this paper, we propose **ARC-Top- K** , an **All-Reduce-Compatible Top- K** compressor that aligns sparsity patterns across nodes using a lightweight sketch of the gradient, enabling index-free **All-Reduce** while preserving globally significant information. **ARC-Top- K** is provably contractive and, when combined with momentum error feedback (EF21M), achieves linear speedup and sharper convergence rates than the original EF21M under standard assumptions. Empirically, **ARC-Top- K** matches the accuracy of **Top- K** while reducing wall-clock training time by up to 60.7%, offering an efficient and scalable solution that combines the robustness of **Rand- K** with the strong performance of **Top- K** .

Index Terms—Distributed Optimization, Communication Compression, Error Feedback, Top- K , All-Reduce

I. INTRODUCTION

The rise of large-scale machine learning has established distributed training as a fundamental paradigm in modern AI systems. In centralized cloud environments, massive datasets are partitioned and sharded across N nodes (e.g., GPUs) within a data center to enable data-parallel training [1] [2]. In contrast, federated learning keeps data on N client devices (e.g., smartphones) for privacy and regulatory compliance, while a coordinating server aggregates model updates instead of raw data [3]–[5]. Both paradigms give rise to the following distributed stochastic optimization problem:

$$\min_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) := \left\{ \frac{1}{N} \sum_{i=1}^N [f_i(\mathbf{x}) := \mathbb{E}_{\xi \sim \mathcal{D}_i} F_i(\mathbf{x}, \xi)] \right\}. \quad (1)$$

*Equal contributions. †Co-Corresponding author

The work is supported by the NSF China under Grants 92370121, the National Key Research and Development Program of China under Grant 2024YFA1012902, and Guangdong Provincial Key Research Project for Regular Universities (2025ZDZX3048)

Here, N denotes the number of computing nodes, ξ represents a random data sample drawn from the local distribution $\mathcal{D}_i$, and $f_i : \mathbb{R}^d \rightarrow \mathbb{R}$ is the local loss function at node i , which may be non-convex. While each node possesses its local data and loss function, they collaborate to minimize the global loss function $f(\mathbf{x})$ defined in problem (1).

In such distributed settings, each node i can only access its local data and calculate local gradient $\nabla F_i(\mathbf{x}; \xi)$ of its loss function during optimization. However, communication is required to obtain information from other nodes, which frequently becomes the primary bottleneck and dominates end-to-end training time. The associated cost scales with both the model dimension d and the number of nodes N , rendering it prohibitive for large-scale learning tasks. Consequently, minimizing the volume of communicated bits per round is a critical problem for efficient and scalable distributed learning.

A. Limitations of Existing Communication-Saving Approaches

Gradient sparsification is a prominent strategy to reduce communication by transmitting only a small fraction of gradient coordinates. The two most common gradient sparsification compressors are **Rand- K** and **Top- K** [6] [7]. **Rand- K** samples K coordinates of the gradient uniformly at random to form an unbiased estimate, whereas **Top- K** preserves the K entries with the largest magnitudes. Despite the apparent trade-off between structure-awareness and statistical simplicity, both methods suffer from significant drawbacks.

Rand- K employs uniform sampling that disregards gradient structure, often discarding informative coordinates and inducing substantial compression error, which in turn leads to degraded accuracy and slower convergence in practice. In contrast, **Top- K** leverages gradient magnitudes but introduces two key challenges: algorithmically, independent **Top- K** selection across nodes misaligns entry indices, breaking error contraction properties of the global gradient and undermining theoretical convergence guarantees; systemically, the absence of shared indices requires transmitting both values and indices, precluding efficient **All-Reduce** and forcing slower

TABLE I
COMPARISON BETWEEN DIFFERENT GRADIENT SPARSIFICATION COMPRESSORS. THE SYMBOL “↓” INDICATES THAT A SMALLER VALUE OF THE QUANTITY IS PREFERABLE. THE SYMBOL 😊 IMPLIES SUPERIOR PERFORMANCE

EF21M with Compressor	Global Contractive Property	Primitive	Communication per Iteration (↓)	Asymptotic Convergence Rate (↓)	Transient Complexity (↓)	Empirical Performance
No Compression	-	All-Reduce	$2mn$	$\mathcal{O}(\frac{\sigma}{\sqrt{NT}})$	$\mathcal{O}(\frac{N}{\sigma^2})$	😊
Top- K	✗	All-Gather	$(N-1)(nK+K)$	$\mathcal{O}(\frac{\sigma}{\sqrt{NT}})$	$\mathcal{O}(\frac{N^3}{\sigma^2})$	😊
Rand- K	✓	All-Reduce	$2Kn$	$\mathcal{O}(\frac{\sigma}{\sqrt{NT}})$	$\mathcal{O}(\frac{N}{\sigma^2})$	😞
ARC-Top- K	✓	All-Reduce	$2Kn+2mr$	$\mathcal{O}(\frac{\sigma}{\sqrt{NT}})$	$\mathcal{O}(\frac{N}{\sigma^2})$	😊
ARC-Top- K ($r=1$)	✓	All-Reduce	$2Kn+2m$	$\mathcal{O}(\frac{\sigma}{\sqrt{NT}})$	$\mathcal{O}(\frac{N}{\sigma^2})$	😊

Gather/Scatter primitives that increase latency and underutilize bandwidth. These motivate the research question:

*(Question) Can we design a communication-efficient sparsification method that not only **selects informative gradient components** but also enables the use of **index-free collectives** (e.g., All-Reduce), achieving both superior convergence properties and strong empirical performance?*

B. Main Results

To address the fundamental question, we propose a novel **All-Reduce-Compatible Top- K (ARC-Top- K)** compressor in this paper. **ARC-Top- K** has three key features:

- **Informative gradient entry selection.** **ARC-Top- K** adapts to gradient structure and selects statistically informative coordinates. This design ensures performance comparable to **Top- K** , while achieving higher accuracy and faster convergence than **Rand- K** which uses random sparsification.
- **Superior convergence guarantees.** **ARC-Top- K** maintains the contractive property of the compressed global gradient. Consequently, communication-efficient methods using **ARC-Top- K** as the gradient compressor can achieve faster theoretical convergence rates than using **Top- K** .
- **High-performance implementation.** By aligning entry indices across nodes, **ARC-Top- K** enables the use of index-free All-Reduce primitives. In contrast, the absence of shared indices in **Top- K** necessitates slower Gather/Scatter operations, which increase latency and underutilize bandwidth. Our implementation demonstrates that **ARC-Top- K** matches the accuracy of **Top- K** while reducing wall-clock training time by up to 60.69%.

Table I compares **Rand- K** , **Top- K** , and **ARC-Top- K** . The results show that **ARC-Top- K** enables communication-efficient algorithms to achieve the same theoretical convergence rate and communication overhead as **Rand- K** , while matching the empirical performance of **Top- K** . All reported convergence rates are obtained in combination with EF21M [8], and the transient complexity reflects the number of iterations required for the asymptotic rate term to dominate. In the table, N is the number of nodes, T is the number of iterations, σ is the standard deviation of the stochastic gradient noise. Tensors in all algorithm are compressed in 2-

dimensional view of $m \times n$ while K represents the number of selected rows in each tensor.

Notation. We introduce the set $[N] := \{1, \dots, N\}$. Given $\mathbf{A} \in \mathbb{R}^{m \times n}$, operator $\text{vec}(\cdot)$ returns column-wise vectorization $\text{vec}(\mathbf{A}) \in \mathbb{R}^{mn}$ stacking the rows of $\mathbf{A}$. $\text{diag}(\cdot)$ returns the vector of diagonal entries $\text{diag}(\mathbf{M}) \in \mathbb{R}^n$ for a square matrix $\mathbf{M} \in \mathbb{R}^{n \times n}$. $\arg \text{top}_K(\mathbf{z})$ returns the index set of the K largest entries of $\mathbf{z} \in \mathbb{R}^n$. Given an index set $\mathcal{I}$ and a matrix $\mathbf{U}$, $[\mathbf{U}]_{\mathcal{I},:}$ selects the rows of $\mathbf{U}$ indexed by $\mathcal{I}$ and zero other rows.

II. RELATED WORKS

Distributed Learning. Distributed optimization underpins large-scale training in both data-center clusters and cross-device networks [1] [2] [9]–[12]. In the parameter server architecture, a central server aggregates gradients from N nodes and broadcasts the averaged update. Federated learning (FL) adapts this worker-server template under privacy and availability constraints, which exacerbate client drift and statistical heterogeneity [3] [13] [14]. Techniques such as client subsampling, proximal regularization, and adaptive aggregation have been proposed to mitigate these issues [5] [15] [16]. By contrast, synchronous data-parallel training in data centers employs collective communication (e.g., All-Reduce) to compute global averages without a central server, enabling the training of frontier large-scale models [17]–[20]. Our work focuses on the communication bottleneck common to both centralized and federated regimes.

Communication Compression. Communication compression reduces data volume by transmitting compact representations of variables, thereby mitigating communication overhead and improving the scalability of distributed training. Existing approaches fall into three main categories: low-rank projection, quantization, and sparsification. Low-rank projection sketches gradients into lower-dimensional subspaces [21]–[24]. Quantization encodes coordinates with fewer bits; stochastic schemes remain unbiased and integrate well with error feedback [25] [26]. Sparsification transmits only a subset of entries. **Rand- K** is unbiased but structure-agnostic, often discarding principal components and incurring large compression errors [6] [7]. **Top- K** retains large-magnitude entries and typically yields better training performance, but it breaks All-Reduce compatibility and introduces additional overhead [27] [28].




Fig. 1. Workflow of the ARC-Top-K algorithm, detailing the process of gradient compression and aggregation across two nodes in a distributed system.

Error Feedback. Error feedback (EF) mitigates the adverse effects of compression by incorporating past residuals into subsequent gradient updates, thereby preserving more informative signals [29] [30]. EF21 [31] extends this principle by maintaining a local gradient tracker for each node, which alleviates the impact of data heterogeneity and improves convergence rates. Building on this theoretical foundation, NEOLITHIC and its variants [32] [33] establish lower bounds for distributed learning under communication compression. Notably, EF21M [8] achieves these bounds through momentum, attaining nearly optimal convergence rates. Our **ARC-Top-K** compressor can be seamlessly integrated with EF21M to further enhance convergence in distributed settings.

III. LIMITATIONS OF THE TOP-K COMPRESSOR

Contractive Compressor. Contractive compressors are widely used in communication-efficient distributed optimization. Their key property is that the compression error diminishes in expectation as the underlying variable approaches zero. A formal definition of a contractive compressor is given below:

Definition 1 (Contractive Compressor). A compressor $\mathcal{C}(\cdot)$ is contractive if and only if

$$\mathbb{E}_{\mathcal{C}} [\|\mathcal{C}(\mathbf{g}) - \mathbf{g}\|_2^2] \leq (1 - \alpha) \|\mathbf{g}\|_2^2, \quad \forall \mathbf{g} \in \mathbb{R}^d, \quad (2)$$

where $0 < \alpha \leq 1$ is the contractive factor. The expectation is taken over the randomness of the operator $\mathcal{C}$.

Non-Contraction of Top-K in Distributed Learning. Although the classical **Top-K** operator is contractive when applied to a single vector, this property does not extend to distributed learning where the operator is applied independently across multiple nodes and the results are subsequently averaged. Formally, let $\mathcal{C}_{\text{local}}(\cdot)$ denote the local **Top-K** compressor, $\mathbf{g} := (1/N) \sum_{i=1}^N \mathbf{g}^{(i)}$ the uncompressed average, and $\mathcal{C}(\mathbf{g}) := (1/N) \sum_{i=1}^N \mathcal{C}_{\text{local}}(\mathbf{g}^{(i)})$ the average of compressed vectors. The following proposition demonstrates that $\mathcal{C}(\mathbf{g})$ may fail to be contractive in terms of the globally averaged gradient $\mathbf{g}$. This non-contractive property leads to unfavorable

worst-case behavior, preventing algorithms using **Top-K** from achieving convergence rates as fast as those of **Rand-K**.

Proposition 1. Let $N = 2$, $d = 2$, and $K = 1$. Consider $\mathbf{g}^{(1)} = [-1, 0.1]^\top$, $\mathbf{g}^{(2)} = [1, 0.1]^\top$, and $\mathbf{g} = \frac{1}{2}(\mathbf{g}^{(1)} + \mathbf{g}^{(2)})$. Let $\mathcal{C}(\cdot)$ be **Top-K** compressor defined above. It holds that

$$\|\mathcal{C}(\mathbf{g}) - \mathbf{g}\|_2^2 = \|\mathbf{g}\|_2^2,$$

which implies that $\mathcal{C}(\mathbf{g})$ is a non-contractive compressor.

Proof. Applying $\mathcal{C}_{\text{local}}$ locally with $K=1$ yields $\mathcal{C}_{\text{local}}(\mathbf{g}^{(1)}) = [-1, 0]^\top$, $\mathcal{C}_{\text{local}}(\mathbf{g}^{(2)}) = [1, 0]^\top$. Thus $\mathcal{C}(\mathbf{g}) = \frac{1}{2}(\mathcal{C}_{\text{local}}(\mathbf{g}^{(1)}) + \mathcal{C}_{\text{local}}(\mathbf{g}^{(2)})) = [0, 0]^\top$, whereas $\mathbf{g} = \frac{1}{2}(\mathbf{g}^{(1)} + \mathbf{g}^{(2)}) = [0, 0.1]^\top$. Thus, it holds $\|\mathcal{C}(\mathbf{g}) - \mathbf{g}\|_2^2 = \|\mathbf{g}\|_2^2$. $\square$

System-Level Limitations of Top-K. **Top-K** also imposes significant communication overhead. Because the selected supports differ across nodes, indices must be transmitted together with values, increasing communication sizes and necessitating gather-and-merge operations. In parameter-server systems, aggregation $\mathbf{g} := (1/N) \sum_{i=1}^N \mathbf{g}^{(i)}$ typically densifies the global gradient, requiring full-size communication from the server to all nodes. In collective-based clusters, heterogeneous supports preclude bandwidth-optimal All-Reduce, forcing implementations to fall back to variable-length All-Gather followed by a merge step, which underutilizes high-throughput collectives and increases memory traffic.

IV. ALL-REDUCE COMPATIBLE TOPK

Compressor Design. The key idea of **ARC-Top-K** is to align sparsity patterns across all nodes while retaining the most significant entries of the compressed vector. To achieve this, **ARC-Top-K** first reshapes the gradient vector into a matrix:

$$\mathbf{g}_t^{(i)} \in \mathbb{R}^d \rightarrow \mathbf{G}_t^{(i)} \in \mathbb{R}^{m \times n}, \quad n = d/m. \quad (3)$$

At iteration t , all N nodes synchronize a random seed s , which deterministically generates a shared Gaussian projection

matrix $\mathbf{V} \in \mathbb{R}^{n \times r}$, with $\text{vec}(\mathbf{V}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{nr})$. Next, a row-wise sketch of the global gradient can be achieved through:

$$\mathbf{P}_t^{(i)} = \frac{1}{\sqrt{r}} \mathbf{G}_t^{(i)} \mathbf{V} \in \mathbb{R}^{m \times r}, \quad \mathbf{P}_t = \frac{1}{N} \sum_{i=1}^N \mathbf{P}_t^{(i)}.$$

Row importance is then estimated as

$$\boldsymbol{\Sigma}_t = \text{diag}(\mathbf{P}_t \mathbf{P}_t^\top) \in \mathbb{R}^m, \quad \mathcal{I}_t = \arg \text{top}_K(\boldsymbol{\Sigma}_t), \quad (4)$$

and $\mathcal{I}_t$ is the index set of the top- K important rows. Next we aggregate these important rows from each node

$$[\mathbf{G}_t]_{\mathcal{I}_t, :} = \frac{1}{N} \sum_{i=1}^N [\mathbf{G}_t^{(i)}]_{\mathcal{I}_t, :} \in \mathbb{R}^{m \times n}. \quad (5)$$

All unselected rows are set to zero in $\mathbf{G}_t$, and the result is reshaped back into $\hat{\mathbf{g}}_t = \text{vec}(\mathbf{G}_t) \in \mathbb{R}^d$. Such $\hat{\mathbf{g}}_t$ is regarded as a compressed estimate of the globally averaged gradient $\mathbf{g}_t = (1/N) \sum_{i=1}^N \mathbf{g}_t^{(i)}$ using **ARC-Top- K** , see Algorithm 1 for more implementation details.

ARC-Top- K Captures Important Entries. Our key idea is that each p -th element of $\boldsymbol{\Sigma}_t \in \mathbb{R}^m$ provides an estimate of the importance of the p -th row of the global gradient matrix $\mathbf{G}_t = \frac{1}{N} \sum_{i=1}^N \mathbf{G}_t^{(i)}$ in expectation. To see it, let $\mathbf{u}^{(i)} \in \mathbb{R}^{1 \times n}$ be the p -th row of $\mathbf{G}_t^{(i)}$ and define $\mathbf{u} = \frac{1}{N} \sum_{i=1}^N \mathbf{u}^{(i)}$:

$$\mathbb{E}[\boldsymbol{\Sigma}_{t,p}] = \frac{1}{r} \mathbb{E}[\|\mathbf{u} \mathbf{V}\|_2^2] = \frac{1}{r} \sum_{j=1}^r \mathbb{E}[(\mathbf{u} \mathbf{v}_j)^2] = \|\mathbf{u}\|_2^2, \quad (6)$$

where $\boldsymbol{\Sigma}_{t,p}$ is the p -th element in $\boldsymbol{\Sigma}_t$. The expectation is taken over the random projection matrix $\mathbf{V} = [\mathbf{v}_1, \dots, \mathbf{v}_r] \in \mathbb{R}^{n \times r}$, whose columns are sampled i.i.d. from $\mathcal{N}(\mathbf{0}, \mathbf{I}_n)$. Equality (6) shows that $\boldsymbol{\Sigma}_{t,p}$ provides an estimate of the squared norm of the p -th row of $\mathbf{G}_t$, while the hyperparameter r controls the variance of this estimate, with larger values of r yielding smaller estimate variance. For this reason, expressions (4)–(6) imply that **ARC-Top- K** preserves the most significant rows while zeroing out the less significant ones.

ARC-Top- K is a Contractive Compressor. As shown in Proposition 1, the **Top- K** compressor is not contractive in distributed learning with respect to the globally averaged gradient. In contrast, the following proposition establishes that **ARC-Top- K** is globally contractive, providing the foundation for its superior convergence properties.

Proposition 2. Let $\mathbf{g}_t := \frac{1}{N} \sum_{i=1}^N \mathbf{g}_t^{(i)}$ be the global variable and let $\mathcal{C}(\mathbf{g}_t)$ be the output of Algorithm 1. Then

$$\mathbb{E}_C[\|\mathcal{C}(\mathbf{g}_t) - \mathbf{g}_t\|_2^2] \leq (1 - \alpha) \|\mathbf{g}_t\|_2^2, \quad \alpha = \frac{K}{m},$$

where the expectation is taken over the randomness of $\mathcal{C}$.

Proof. By the reshaping definition of $\mathbf{G}_t$ in (3), we have

$$\|\mathbf{g}_t\|_2^2 = \|\mathbf{G}_t\|_F^2, \quad (7)$$

$$\|\mathcal{C}(\mathbf{g}_t) - \mathbf{g}_t\|_2^2 = \|[\mathbf{G}_t]_{\mathcal{I}_t, :} - \mathbf{G}_t\|_F^2. \quad (8)$$

Algorithm 1: ARC-Top- K ($\{\mathbf{g}_t^{(i)}\}_{i=1}^N, r, m, \mu$).

Input: N nodes; projection dimension r ; number of rows m ; compressor ratio μ with selected row number $K = \lceil \mu m \rceil$; local variable $\mathbf{g}_t^{(i)} \in \mathbb{R}^d$ at iteration t for $i \in [N]$.

Output: Compressed local gradient $\mathcal{C}_{\text{local}}(\mathbf{g}_t^{(i)})$ and global gradient $\mathcal{C}(\mathbf{g}_t)$ with $\mathbf{g}_t := \sum_{i=1}^N \mathbf{g}_t^{(i)}$.

(On i -th node)

Sample random seed s and synchronize it across nodes. Reshape $\mathbf{G}_t^{(i)} \in \mathbb{R}^{m \times n}$ s.t. $\text{vec}(\mathbf{G}_t^{(i)}) = \mathbf{g}_t^{(i)}$, $n = d/m$. Generate $\mathbf{V} \in \mathbb{R}^{n \times r}$ with s s.t. $\text{vec}(\mathbf{V}) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_{nr})$.

$\mathbf{P}_t^{(i)} \leftarrow \mathbf{G}_t^{(i)} \mathbf{V} \in \mathbb{R}^{m \times r}$.

$\mathbf{P}_t \leftarrow \frac{1}{N} \sum_{i=1}^N \mathbf{P}_t^{(i)}$. (All-Reduce)

$\boldsymbol{\Sigma}_t \leftarrow \text{diag}(\mathbf{P}_t \mathbf{P}_t^\top) \in \mathbb{R}^m$, $\mathcal{I}_t \leftarrow \arg \text{top}_K(\boldsymbol{\Sigma}_t) \in \mathbb{Z}^K$.

$\mathcal{C}_{\text{local}}(\mathbf{g}_t^{(i)}) := \text{vec}([\mathbf{G}_t^{(i)}]_{\mathcal{I}_t, :})$.

$\mathcal{C}(\mathbf{g}_t) := \frac{1}{N} \sum_{i=1}^N \mathcal{C}_{\text{local}}(\mathbf{g}_t^{(i)})$. (All-Reduce)

return $\mathcal{C}(\mathbf{g}_t)$ and $\mathcal{C}_{\text{local}}(\mathbf{g}_t^{(i)})$ for $i \in [N]$

Let $p_j := \mathbb{P}[j \in \mathcal{I}_t]$ be the selection probability of the j -th row. According to the Lemma 1 in [23],

$$\|\mathbf{G}_t[l, :]\|_2^2 \leq \|\mathbf{G}_t[j, :]\|_2^2 \Rightarrow p_l \leq p_j, \quad \forall l, j \in [m]. \quad (9)$$

Then by Chebyshev inequality, it holds

$$\begin{aligned} \mathbb{E}_C[\|[\mathbf{G}_t]_{\mathcal{I}_t, :} - \mathbf{G}_t\|_F^2] &= \|\mathbf{G}_t\|_F^2 - \sum_{j=1}^m p_j \|[\mathbf{G}_t]_{j, :}\|_2^2 \\ &\leq \|\mathbf{G}_t\|_F^2 - \frac{1}{m} \left(\sum_{j=1}^m p_j \right) \left(\sum_{j=1}^m \|[\mathbf{G}_t]_{j, :}\|_2^2 \right) \\ &= \|\mathbf{G}_t\|_F^2 - \frac{K}{m} \|\mathbf{G}_t\|_F^2. \end{aligned} \quad (10)$$

Substituting (7) and (8) into (10) completes the proof. $\square$

ARC-Top- K can be Efficiently Implemented. By aligning sparsity patterns across nodes, **ARC-Top- K** eliminates the need to transmit indices and enables the use of bandwidth-optimal All-Reduce rather than the more costly All-Gather (see Algorithm 1). This design achieves both high-fidelity gradient preservation and low communication cost, making **ARC-Top- K** particularly effective for large-scale distributed learning. Moreover, in deep learning tasks where model weights and gradients are naturally stored in matrix form, the vectorization step required by **ARC-Top- K** can be omitted, further simplifying its implementation.

Communication Cost of ARC-Top- K . We analyze the communication overhead of row-sparsifying updates for an $m \times n$ parameter block replicated across N nodes. At iteration t , each node selects K rows, yielding sparsity $\mu := K/m$, and transmits compressed updates through a collective primitive. In **ARC-Top- K** , a compact rank- r sketch is additionally exchanged to align the selected row sets across nodes. ‘‘Communication per iteration’’ refers to the number of scalar entries transmitted per node per iteration. As summarized in Table I, the dense baseline requires $2mn$ entries using All-Reduce.

Independent **Top-K** instead relies on All-Gather to merge both values and indices, which incurs $(N-1)(nK+K)$ entries and fails to preserve global contractivity. By contrast, **Rand-K** achieves global contractivity in expectation while communicating $2Kn$ entries via All-Reduce. Finally, **ARC-Top-K** restores global contractivity and maintains index-free All-Reduce through the shared sketch, with total communication volume $2Kn + 2mr$, which further reduces to $2Kn + 2m$ when $r = 1$.

V. COMMUNICATION-EFFICIENT METHOD USING ARC-Top-K

The state-of-the-art communication-efficient method to solve problem (1) is **Momentum SGD with Error Feedback (EF21M)** [8], which corrects compression bias by recycling residual information. Equipping EF21M with **ARC-Top-K**, we achieve the following recursions:

$$\mathbf{h}_t^{(i)} = (1 - \eta) \mathbf{h}_{t-1}^{(i)} + \eta \nabla F_i(\mathbf{x}_t, \boldsymbol{\xi}_t^{(i)}), \quad (11a)$$

$$\mathbf{g}_t^{(i)} = \mathbf{g}_{t-1}^{(i)} + \mathcal{C}_{\text{local}}(\mathbf{h}_t^{(i)} - \mathbf{g}_{t-1}^{(i)}), \quad (11b)$$

$$\mathbf{x}_{t+1} = \mathbf{x}_t - \frac{\gamma}{N} \sum_{i=1}^N \mathbf{g}_t^{(i)}, \quad (11c)$$

where compressor $\mathcal{C}_{\text{local}}(\cdot)$ is realized through **ARC-Top-K**. Step (11a) updates the local gradient tracker, which serves as a momentum term to smooth stochastic gradient estimates. Step (11b) performs error feedback by compressing only the difference between this tracker and the previously transmitted vector using **ARC-Top-K**, thereby retaining critical gradient information while reducing communication overhead. Because all nodes employ the same synchronized sketch, the resulting sparsity patterns are aligned across nodes, enabling efficient index-free All-Reduce aggregation in step (11c).

Vanilla EF21M [8] employs **Top-K** as the compressor in step (11b). Since **ARC-Top-K** preserves the contraction property, we next establish that replacing **Top-K** with **ARC-Top-K** enables EF21M to achieve faster theoretical convergence rates than its vanilla counterpart.

Assumption 1 (Smoothness and lower boundedness). *We assume f is L -smooth, $\|\nabla f(\mathbf{x}) - \nabla f(\mathbf{y})\| \leq L\|\mathbf{x} - \mathbf{y}\|$, $\forall \mathbf{x}, \mathbf{y} \in \mathbb{R}^d$. Moreover, we assume that f is lower bounded, i.e., $f^* := \inf_{\mathbf{x} \in \mathbb{R}^d} f(\mathbf{x}) > -\infty$.*

Assumption 2 (Stochastic gradient oracle). *For each $i \in [N]$, the stochastic gradient oracle $\nabla F_i(\mathbf{x}; \boldsymbol{\xi})$ is unbiased and has bounded variance, i.e., $\mathbb{E}_{\boldsymbol{\xi} \sim \mathcal{D}_i}[\nabla F_i(\mathbf{x}; \boldsymbol{\xi})] = \nabla f_i(\mathbf{x})$, $\mathbb{E}_{\boldsymbol{\xi} \sim \mathcal{D}_i}[\|\nabla F_i(\mathbf{x}; \boldsymbol{\xi}) - \nabla f_i(\mathbf{x})\|^2] \leq \sigma^2$.*

We need several lemmas to facilitate the convergence analysis.

Lemma 1. *Suppose Assumption 1 holds, and let the iterate be updated as $\mathbf{x}_{t+1} = \mathbf{x}_t - \gamma \mathbf{g}_t$ for some vector $\mathbf{g}_t \in \mathbb{R}^d$ and step size $\gamma > 0$. Then the following inequality holds:*

$$\begin{aligned} f(\mathbf{x}_{t+1}) &\leq f(\mathbf{x}_t) - \frac{\gamma}{2} \|\nabla f(\mathbf{x}_t)\|^2 - \frac{1}{4\gamma} \|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2 \\ &\quad + \frac{\gamma}{2} \|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2. \end{aligned} \quad (12)$$

Lemma 2. *Suppose Assumption 1 holds, and let $\mathcal{C}$ be a contractive compressor with parameter $\alpha \leq \frac{1}{2}$. Let the averaged sequence be $\mathbf{h}_t := \frac{1}{N} \sum_{i=1}^N \mathbf{h}_t^{(i)}$, $\mathbf{g}_t := \frac{1}{N} \sum_{i=1}^N \mathbf{g}_t^{(i)}$ and $\nabla F(\mathbf{x}_t, \boldsymbol{\xi}_t) := \frac{1}{N} \sum_{i=1}^N \nabla F_i(\mathbf{x}_t, \boldsymbol{\xi}_t^{(i)})$. Consider*

$$\begin{aligned} \mathbf{h}_t &= \mathbf{h}_{t-1} + \eta(\nabla F(\mathbf{x}_t, \boldsymbol{\xi}_t) - \mathbf{h}_{t-1}), \\ \mathbf{g}_t &= \mathbf{g}_{t-1} + \mathcal{C}(\mathbf{h}_t - \mathbf{g}_{t-1}). \end{aligned}$$

Then the following inequality holds:

$$\begin{aligned} \mathbb{E}[\|\mathbf{g}_t - \mathbf{h}_t\|^2] &\leq (1 - \frac{\alpha}{2}) \mathbb{E}[\|\mathbf{g}_{t-1} - \mathbf{h}_{t-1}\|^2] + \frac{\eta^2 \sigma^2}{N} \\ &\quad + \frac{4\eta^2}{\alpha} \mathbb{E}[\|\mathbf{h}_{t-1} - \nabla f(\mathbf{x}_{t-1})\|^2] \\ &\quad + \frac{4L^2 \eta^2}{\alpha} \mathbb{E}[\|\mathbf{x}_t - \mathbf{x}_{t-1}\|^2]. \end{aligned} \quad (13)$$

Lemma 3. *Suppose Assumption 1 holds, and let $0 < \eta \leq 1$. For each $i \in \{1, \dots, N\}$, define the sequence $\{\mathbf{h}_t^{(i)}\}_{t \geq 0}$ by the recursion $\mathbf{h}_t^{(i)} = \mathbf{h}_{t-1}^{(i)} + \eta(\nabla F_i(\mathbf{x}_t, \boldsymbol{\xi}_t^{(i)}) - \mathbf{h}_{t-1}^{(i)})$.*

Then, for all $t \geq 0$, it holds that

$$\begin{aligned} \mathbb{E}[\|\mathbf{h}_{t+1} - \nabla f(\mathbf{x}_{t+1})\|^2] &\leq (1 - \eta) \mathbb{E}[\|\mathbf{h}_t - \nabla f(\mathbf{x}_t)\|^2] \\ &\quad + \frac{3L^2}{\eta} \mathbb{E}[\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2] + \frac{\eta^2 \sigma^2}{N}. \end{aligned} \quad (14)$$

The proofs of Lemma 1 and Lemma 3 are provided in [8]. Lemma 2 is new; its analysis is straightforward, and we omit the details here due to space constraints. The following theorem establishes the convergence of recursion (11a)–(11c).

Theorem 1 (ARC-Top-K Convergence with EF21M). *Under Assumptions 1, 2, with the learning rate $\gamma \leq 1/(4L)$, momentum η and initial batch size B_{init} ,*

$$\begin{aligned} &\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\mathbf{x}_t)\|^2] \leq \\ &\mathcal{O}\left(\frac{\delta_0}{\gamma T} + \frac{\eta^3 \sigma^2}{N \alpha^2} + \frac{\eta^2 \sigma^2}{N \alpha} + \frac{\eta \sigma^2}{N} + \frac{\eta \sigma^2}{N \alpha^2 B_{\text{init}} T} + \frac{\sigma^2}{N \eta B_{\text{init}} T}\right), \end{aligned}$$

*where $\delta_0 := f(\mathbf{x}_0) - \inf_{\mathbf{x}} f(\mathbf{x})$. If we further choose γ , η and B_{init} properly, **ARC-Top-K** with EF21M converges as*

$$\begin{aligned} &\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\mathbf{x}_t)\|^2] \leq \\ &\mathcal{O}\left(\left(\frac{L \delta_0 \sigma^2}{NT}\right)^{\frac{1}{2}} + \left(\frac{L \delta_0 \sigma}{\alpha^{\frac{1}{2}} N^{\frac{1}{2}} T}\right)^{\frac{2}{3}} + \left(\frac{L \delta_0 \sigma^{\frac{2}{3}}}{\alpha^{\frac{2}{3}} N^{\frac{1}{3}} T}\right)^{\frac{3}{4}} + \frac{L \delta_0}{\alpha T}\right). \end{aligned}$$

Proof of Theorem 1. By applying Lemma 1 and decompose the error between $\mathbf{g}_t$ and $\nabla f(\mathbf{x}_t)$ into two terms by $\|\mathbf{g}_t - \nabla f(\mathbf{x}_t)\|^2 \leq 2\|\mathbf{g}_t - \mathbf{h}_t\|^2 + 2\|\mathbf{h}_t - \nabla f(\mathbf{x}_t)\|^2$ we obtain:

$$\begin{aligned} &\frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E}[\|\nabla f(\mathbf{x}_t)\|^2] \\ &\leq \frac{2\delta_0}{\gamma T} + \frac{2}{T} \sum_{t=0}^{T-1} V_t + \frac{2}{T} \sum_{t=0}^{T-1} P_t - \frac{1}{2\gamma^2 T} \sum_{t=0}^{T-1} R_t. \end{aligned} \quad (15)$$

It is noted that the quantities V_t , P_t , and R_t are defined as follows: $V_t = \mathbb{E} [\|\mathbf{g}_t - \mathbf{h}_t\|^2]$, $P_t = \mathbb{E} [\|\mathbf{h}_t - \nabla f(\mathbf{x}_t)\|^2]$, $R_t = \mathbb{E} [\|\mathbf{x}_{t+1} - \mathbf{x}_t\|^2]$. From Lemma 2, it follows that

$$\frac{1}{T} \sum_{t=0}^{T-1} V_t \leq \frac{8\eta^2}{\alpha^2 T} \sum_{t=0}^{T-1} P_t + \frac{8L^2\eta^2}{\alpha^2 T} \sum_{t=0}^{T-1} R_t + \frac{2\eta^2\sigma^2}{N\alpha} + \frac{2V_0}{\alpha T}. \quad (16)$$

From Lemma 3, it follows that

$$\frac{1}{T} \sum_{t=0}^{T-1} P_t \leq \frac{3L^2}{\eta^2} \cdot \frac{1}{T} \sum_{t=0}^{T-1} R_t + \frac{\eta\sigma^2}{N} + \frac{1}{\eta T} P_0. \quad (17)$$

Substituting the results of 16 and 17 into 15, and by choosing the step size γ appropriately, we obtain:

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\mathbf{x}_t)\|^2] &\leq \frac{2\delta_0}{\gamma T} + \left(\frac{16\eta^3}{\alpha^2} + \frac{4\eta^2}{\alpha} + 2\eta \right) \frac{\sigma^2}{N} \\ &\quad + \left(\frac{16\eta}{\alpha^2} + \frac{2}{\eta} \right) \frac{P_0}{T} + \frac{4V_0}{\alpha T}. \end{aligned} \quad (18)$$

given $P_0 = \mathbb{E} [\|\mathbf{h}_0 - \nabla f(\mathbf{x}_0)\|^2] \leq \frac{\sigma^2}{NB_{\text{init}}}$, $V_0 = 0$ we obtain

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\mathbf{x}_t)\|^2] &\leq \\ &\mathcal{O} \left(\frac{\delta_0}{\gamma T} + \frac{\eta^3\sigma^2}{N\alpha^2} + \frac{\eta^2\sigma^2}{N\alpha} + \frac{\eta\sigma^2}{N} + \frac{\eta\sigma^2}{N\alpha^2 B_{\text{init}} T} + \frac{\sigma^2}{N\eta B_{\text{init}} T} \right), \end{aligned}$$

To further simplify the convergence bound, we impose constraints such that

$$\frac{\eta^3\sigma^2}{N\alpha^2} \leq \frac{L\delta_0}{\eta T}, \quad \frac{\eta^2\sigma^2}{N\alpha} \leq \frac{L\delta_0}{\eta T}, \quad \frac{\eta\sigma^2}{N} \leq \frac{L\delta_0}{\eta T}, \quad \frac{\eta\sigma^2}{N\alpha^2 B_{\text{init}} T} \leq \frac{L\delta_0}{\eta T},$$

and

$$\frac{\sigma\sqrt{L\delta_0}}{\alpha\sqrt{B_{\text{init}}NT}} \leq \max \left\{ \frac{L\delta_0}{\alpha T}, \left(\frac{L\delta_0\sigma^{\frac{2}{3}}}{\alpha^{\frac{2}{3}}N^{\frac{1}{3}}T} \right)^{\frac{3}{4}}, \left(\frac{L\delta_0\sigma}{\alpha^{\frac{1}{2}}N^{\frac{1}{2}}T} \right)^{\frac{3}{4}}, \left(\frac{L\delta_0\sigma^2}{NT} \right)^{\frac{1}{2}} \right\}.$$

Putting all these together, we achieve the convergence rate

$$\begin{aligned} \frac{1}{T} \sum_{t=0}^{T-1} \mathbb{E} [\|\nabla f(\mathbf{x}_t)\|^2] &\leq \\ &\mathcal{O} \left(\left(\frac{L\delta_0\sigma^2}{NT} \right)^{\frac{1}{2}} + \left(\frac{L\delta_0\sigma}{\alpha^{\frac{1}{2}}N^{\frac{1}{2}}T} \right)^{\frac{3}{4}} + \left(\frac{L\delta_0\sigma^{\frac{2}{3}}}{\alpha^{\frac{2}{3}}N^{\frac{1}{3}}T} \right)^{\frac{3}{4}} + \frac{L\delta_0}{\alpha T} \right), \end{aligned}$$

which completes the proof. $\square$

Asymptotic Rate and Transient Iterations. According to Theorem 1, the asymptotic convergence rate is dominated by $\mathcal{O}(1/\sqrt{NT})$ as $T \rightarrow \infty$, which matches the rate achieved by algorithms using other compressors such as **Top-K** and **Rand-K**. However, to reach this asymptotic regime, the algorithm must run for a sufficiently large number of iterations, commonly referred to as the transient phase. Based on Theorem 1, we establish that the recursion (11a)–(11c) with **ARC-Top-K** has transient complexity on the order of $\mathcal{O}(N/\sigma^2)$. This is substantially smaller than the transient complexity of **Top-K**, which scales as $\mathcal{O}(N^3/\sigma^2)$ (see Table I). The advantage arises from the contractive property of **ARC-Top-K**, in contrast to the non-contraction of **Top-K**.

TABLE II
EVALUATION ACCURACY (%) OF TRAINING RESNET ON CIFAR-10 WITH MEAN $\pm$ STANDARD DEVIATION OVER 3 RUNS.

Method	Adam		MSGD	
	ResNet-18	ResNet-50	ResNet-18	ResNet-50
Dense	94.01 $\pm$ 0.42	94.02 $\pm$ 0.15	95.08 $\pm$ 0.08	94.96 $\pm$ 0.31
<i>(without EF)</i>				
Top-K	93.58 $\pm$ 0.10	93.86 $\pm$ 0.14	94.14 $\pm$ 0.10	94.58 $\pm$ 0.32
Rand-K	92.54 $\pm$ 0.09	92.83 $\pm$ 0.17	92.27 $\pm$ 0.27	91.63 $\pm$ 0.39
ARC-Top-K	93.00 $\pm$ 0.12	93.85 $\pm$ 0.10	93.79 $\pm$ 0.09	94.39 $\pm$ 0.06
<i>(with EF)</i>				
Top-K	93.98$\pm$0.05	94.20$\pm$0.14	94.95 $\pm$ 0.14	94.80 $\pm$ 0.20
Rand-K	93.90 $\pm$ 0.09	93.88 $\pm$ 0.41	94.81 $\pm$ 0.18	94.72 $\pm$ 0.38
ARC-Top-K	93.91 $\pm$ 0.05	93.97 $\pm$ 0.06	95.00$\pm$0.18	95.00$\pm$0.09

VI. EXPERIMENT

We evaluate **ARC-Top-K** on both vision and language tasks, and further measure wall-clock time and accuracy across varying node counts. Comparisons are made against uncompressed training, standard **Top-K**, and **Rand-K** baselines. Gradient compression begins after 1000 iterations, following [21]. For transformers, we compress only two-dimensional tensors, which account for the majority of gradients. For convolutional neural networks, we follow [34] and reshape four-dimensional convolutional kernels into two-dimensional matrices for compression.

A. Pre-Training on CIFAR

We evaluate **ARC-Top-K** by training ResNet on CIFAR-10 [35]. Each experiment is repeated 3 times with different random seeds. We report mean accuracy and standard deviation. ResNet-18 and ResNet-50 are trained for 200 epochs with a learning rate of 1×10^{-3} , weight decay of 5×10^{-4} and local batch size of 16. All algorithms use sparsity $\mu = 0.2$. For **ARC-Top-K**, we set projection dimension $r = 4$. Table II reports the results, with **Dense** denoting the uncompressed baseline. The highest accuracy in each model-optimizer pair is highlighted in bold. Results are shown both without EF and with EF enabled.

Across all settings, **ARC-Top-K** consistently surpasses **Rand-K** in accuracy. Moreover, **ARC-Top-K** attains accuracy comparable to **Top-K** and the uncompressed baseline while incurring lower communication cost. When combined with MSGD, **ARC-Top-K** achieves the highest accuracy for both ResNet-18 and ResNet-50, demonstrating an effective balance between communication efficiency and model performance.

TABLE III
RESULTS OF FINE-TUNING ROBERTA-BASE ON GLUE.

Method	CoLA	SST-2	MRPC	STS-B	QQP	MNLI	QNLI	RTE	Average
Dense	64.16	94.73	93.15	91.16	91.86	87.51	92.81	81.59	87.12
Top-K	65.29	94.27	92.76	91.13	92.02	87.45	92.84	80.51	87.03
Rand-K	62.26	94.61	92.47	90.93	91.88	87.34	93.04	80.51	86.63
ARC-Top-K	66.04	94.84	93.22	91.03	91.98	87.58	93.04	81.59	87.42

TABLE IV
VALIDATION PERPLEXITY OF PRE-TRAINING LLAMA ON C4.

Method	LLaMA-60M	LLaMA-130M
Dense	30.59	24.72
Rand-K	43.31	36.31
Top-K	32.20	29.19
ARC-Top-K	33.82	26.43

B. Fine-Tuning on GLUE

We fine-tune RoBERTa-base on GLUE using 4 NVIDIA RTX 4090 (24 GB) GPUs, employing Adam without weight decay and a maximum sequence length of 512 [36]. All tasks are fine-tuned for 30 epochs with a linear learning-rate schedule. We set sparsity $\mu=0.2$ for all compressors and fix the projection rank at $r=4$ for **ARC-Top- K** . For CoLA and MRPC we use a learning rate of 3×10^{-5} with local batch size 32, whereas all other tasks use 1×10^{-5} with size 16.

Table III reports accuracy on the 8 GLUE tasks, with the best score per column in **bold**. **Dense** denotes the baseline without gradient compression. Across tasks, **ARC-Top- K** consistently surpasses **Rand- K** , and it matches or exceeds **Top- K** and the uncompressed baseline while requiring substantially less communication, achieving the highest average score and a favorable accuracy–communication tradeoff.

C. Pre-Training on C4

We pre-train LLaMA models on the C4 corpus [20] [18] using standard Adam with a local batch size of 128 and a sequence length of 256. Following [22], we train LLaMA-60M and LLaMA-130M on 1.1B and 2.2B tokens respectively. All compressors use error feedback, retain $\mu = 0.2$ entries per tensor, and set $r=4$ for **ARC-Top- K** . We select the best result from learning rates $\{1 \times 10^{-3}, 2 \times 10^{-3}\}$.

Table IV reports the final validation perplexity. **ARC-Top- K** outperforms **Rand- K** in pre-training tasks consistently. **ARC-Top- K** also outperform **Top- K** when training large models of same compression ratio with less communication.

D. Further Study

Wall-clock performance. We measure per-iteration wall-clock time for pre-training LLaMA models (60M, 130M, 350M, and 1B) on C4 using 4 NVIDIA A100 (40 GB) GPUs. Each run uses Adam for 50 iterations. We report the mean time per iteration. The local batch size is fixed at 1. For sparsification, the sparsity is set to $\mu=0.2$. Gradients are communicated via the NCCL backend using shared memory (SHM) transport, and NVLink peer-to-peer is explicitly disabled to emulate a low-bandwidth, multi-machine setting. As summarized in Table V, sparsification yields limited time savings for small models because the compression overhead offsets communication gains. However, as model size increases and communication dominates, the relative overhead of compression diminishes and benefits become pronounced. For LLaMA-1B, **ARC-Top- K** reduces the average per-iteration time from

TABLE V
AVERAGE PER-ITERATION TRAINING TIME (IN SECONDS) OF SPARSIFICATION ALGORITHMS WHEN TRAINING LLAMA.

Method	LLaMA-60M	LLaMA-130M	LLaMA-350M	LLaMA-1B
Dense	0.1469	0.3243	0.8775	3.0854
Top-K	0.1526	0.3287	0.8497	2.9628
Rand-K	0.0752	0.1480	0.3544	0.9581
ARC-Top-K	0.1061	0.1798	0.3745	1.2130

TABLE VI
EVALUATION MATTHEWS CORRELATION OF FINE-TUNING PRE-TRAINED ROBERTA-BASE ON COLA WITH MEAN $\pm$ STANDARD DEVIATION OVER 3 RUNS.

Method	8 Nodes	16 Nodes	32 Nodes	64 Nodes
Dense	64.00 $\pm$ 0.80	63.67 $\pm$ 1.66	64.20 $\pm$ 1.49	63.78 $\pm$ 1.12
Top-K	64.01 $\pm$ 0.39	63.34 $\pm$ 0.90	63.50 $\pm$ 1.27	64.02$\pm$1.95
Rand-K	62.67 $\pm$ 1.48	62.60 $\pm$ 1.04	62.13 $\pm$ 1.15	62.42 $\pm$ 0.16
ARC-Top-K	64.50$\pm$1.30	64.15$\pm$0.47	64.53$\pm$0.64	63.70 $\pm$ 2.25

3.0854 s to 1.2130 s with a 60.69% reduction, and achieves the best wall-clock performance across all model sizes.

Scaling to 64 nodes. To validate the scalability of **ARC-Top- K** in multi-node distributed settings, we conduct experiments by fine-tuning RoBERTa-base under the description in Section VI-B, while fixing the global batch size to 64. The number of nodes scales from 8 to 64 for each method. Each experiment is repeated 3 times with independent random seeds, and the results are reported as the mean and standard deviation in Table VI. The results indicate that **ARC-Top- K** consistently matches the performance of the dense baseline across scales from 8 to 64 nodes, while substantially outperforming the **Rand- K** compressor under the same compression ratio. These findings underscore the strong scalability of **ARC-Top- K** , highlighting its promise for large-scale distributed training and federated learning scenarios involving massive edge devices. **Saving in communication bits.** We quantify communication bits by pre-training LLaMA-130M for 10,000 iterations with a global batch size of 256; all remaining hyperparameters follow Section VI-C. Figure 2 plots training loss against cumulative communicated bits. In the figure, **ARC-Top- K** reaches the target loss with fewer bits than the dense baseline. Furthermore, although **ARC-Top- K** involves slightly more communication per iteration than **Rand- K** , its faster convergence yields lowest total communication to attain the same loss.

VII. CONCLUSION

In this work, we propose **ARC-Top- K** , an All-Reduce-compatible sparsification scheme for communication-efficient distributed learning. By aligning sparsity patterns across nodes via a lightweight sketch, **ARC-Top- K** lowers communication while preserving convergence comparable to dense training. We establish convergence guarantees under distributed settings and experiments on RoBERTa fine-tuning and large-scale multi-node tasks show that **ARC-Top- K** matches dense base-




Fig. 2. Loss curves of pre-training LLaMA-130M on C4.

lines while significantly outperforming random sparsification at the same compression ratio. Finally, **ARC-Top- K** scales efficiently from 8 to 64 nodes, demonstrating its suitability for large-scale distributed and federated learning.

REFERENCES

- [1] W. Shi, J. Cao, Q. Zhang, Y. Li, and L. Xu, "Edge computing: Vision and challenges," *IEEE internet of things journal*, vol. 3, no. 5, pp. 637–646, 2016.
- [2] M. Satyanarayanan, "The emergence of edge computing," *Computer*, vol. 50, no. 1, pp. 30–39, 2017.
- [3] J. Konečný, H. B. McMahan, D. Ramage, and P. Richtárik, "Federated optimization: Distributed machine learning for on-device intelligence," *arXiv preprint arXiv:1610.02527*, 2016.
- [4] J. Konečný, H. B. McMahan, F. X. Yu, P. Richtárik, A. T. Suresh, and D. Bacon, "Federated learning: Strategies for improving communication efficiency," *arXiv preprint arXiv:1610.05492*, 2016.
- [5] T. Li, A. K. Sahu, M. Zaheer, M. Sanjabi, A. Talwalkar, and V. Smith, "Federated optimization in heterogeneous networks," *Proceedings of Machine learning and systems*, vol. 2, pp. 429–450, 2020.
- [6] J. Wangni, J. Wang, J. Liu, and T. Zhang, "Gradient sparsification for communication-efficient distributed optimization," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [7] S. U. Stich, J.-B. Cordonnier, and M. Jaggi, "Sparsified sgd with memory," *Advances in neural information processing systems*, vol. 31, 2018.
- [8] I. Fatkhullin, A. Tyurin, and P. Richtárik, "Momentum provably improves error feedback!," *Advances in Neural Information Processing Systems*, vol. 36, pp. 76444–76495, 2023.
- [9] T. Chilimbi, Y. Suzue, J. Apacible, and K. Kalyanaraman, "Project adam: Building an efficient and scalable deep learning training system," in *11th USENIX symposium on operating systems design and implementation (OSDI 14)*, pp. 571–582, 2014.
- [10] A. Smola and S. Narayanamurthy, "An architecture for parallel topic models," *Proceedings of the VLDB Endowment*, vol. 3, no. 1-2, pp. 703–710, 2010.
- [11] M. Li, D. G. Andersen, A. J. Smola, and K. Yu, "Communication efficient distributed machine learning with the parameter server," *Advances in Neural Information Processing Systems*, vol. 27, 2014.
- [12] Z. Zhou, X. Chen, E. Li, L. Zeng, K. Luo, and J. Zhang, "Edge intelligence: Paving the last mile of artificial intelligence with edge computing," *Proceedings of the IEEE*, vol. 107, no. 8, pp. 1738–1762, 2019.
- [13] B. McMahan, E. Moore, D. Ramage, S. Hampson, and B. A. y Arcas, "Communication-efficient learning of deep networks from decentralized data," in *Artificial intelligence and statistics*, pp. 1273–1282, PMLR, 2017.
- [14] P. Kairouz, H. B. McMahan, B. Avent, A. Bellet, M. Bennis, A. N. Bhagoji, K. Bonawitz, Z. Charles, G. Cormode, R. Cummings, *et al.*, "Advances and open problems in federated learning," *Foundations and trends® in machine learning*, vol. 14, no. 1–2, pp. 1–210, 2021.
- [15] S. P. Karimireddy, S. Kale, M. Mohri, S. Reddi, S. Stich, and A. T. Suresh, "Scaffold: Stochastic controlled averaging for federated learning," in *International conference on machine learning*, pp. 5132–5143, PMLR, 2020.
- [16] J. Wang, Q. Liu, H. Liang, G. Joshi, and H. V. Poor, "Tackling the objective inconsistency problem in heterogeneous federated optimization," *Advances in neural information processing systems*, vol. 33, pp. 7611–7623, 2020.
- [17] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, *et al.*, "Language models are few-shot learners," *Advances in neural information processing systems*, vol. 33, pp. 1877–1901, 2020.
- [18] C. Raffel, N. Shazeer, A. Roberts, K. Lee, S. Narang, M. Matena, Y. Zhou, W. Li, and P. J. Liu, "Exploring the limits of transfer learning with a unified text-to-text transformer," *Journal of machine learning research*, vol. 21, no. 140, pp. 1–67, 2020.
- [19] J. Hoffmann, S. Borgeaud, A. Mensch, E. Buchatskaya, T. Cai, E. Rutherford, D. d. L. Casas, L. A. Hendricks, J. Welbl, A. Clark, *et al.*, "Training compute-optimal large language models," *arXiv preprint arXiv:2302.15556*, 2022.
- [20] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix, B. Rozière, N. Goyal, E. Hambro, F. Azhar, *et al.*, "Llama: Open and efficient foundation language models," *arXiv preprint arXiv:2302.13971*, 2023.
- [21] T. Vogels, S. P. Karimireddy, and M. Jaggi, "Powersgd: Practical low-rank gradient compression for distributed optimization," *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [22] J. Zhao, Z. Zhang, B. Chen, Z. Wang, A. Anandkumar, and Y. Tian, "Galore: Memory-efficient llm training by gradient low-rank projection," *arXiv preprint arXiv:2403.03507*, 2024.
- [23] C. Chen, Y. He, P. Li, W. Jia, and K. Yuan, "Greedy low-rank gradient compression for distributed learning with convergence guarantees," *arXiv preprint arXiv:2507.08784*, 2025.
- [24] Y. He, P. Li, Y. Hu, C. Chen, and K. Yuan, "Subspace optimization for large language models with convergence guarantees," *arXiv preprint arXiv:2410.11289*, 2024.
- [25] D. Alistarh, D. Grubic, J. Li, R. Tomioka, and M. Vojnovic, "Qsgd: Communication-efficient sgd via gradient quantization and encoding," *Advances in neural information processing systems*, vol. 30, 2017.
- [26] S. Horváth, C.-Y. Ho, L. Horváth, A. N. Sahu, M. Canini, and P. Richtárik, "Natural compression for distributed deep learning," in *Mathematical and Scientific Machine Learning*, pp. 129–141, PMLR, 2022.
- [27] Y. Lin, S. Han, H. Mao, Y. Wang, and W. J. Dally, "Deep gradient compression: Reducing the communication bandwidth for distributed training," *arXiv preprint arXiv:1712.01887*, 2017.
- [28] L. Zhang, L. Zhang, S. Shi, X. Chu, and B. Li, "Evaluation and optimization of gradient compression for distributed deep learning," in *2023 IEEE 43rd International Conference on Distributed Computing Systems (ICDCS)*, pp. 361–371, IEEE, 2023.
- [29] F. Seide, H. Fu, J. Droppo, G. Li, and D. Yu, "1-bit stochastic gradient descent and its application to data-parallel distributed training of speech dnns," in *Interspeech*, vol. 2014, pp. 1058–1062, Singapore, 2014.
- [30] S. P. Karimireddy, Q. Rejbock, S. Stich, and M. Jaggi, "Error feedback fixes signsgd and other gradient compression schemes," in *International Conference on Machine Learning*, pp. 3252–3261, PMLR, 2019.
- [31] P. Richtárik, I. Sokolov, and I. Fatkhullin, "Ef21: A new, simpler, theoretically better, and practically faster error feedback," *Advances in Neural Information Processing Systems*, vol. 34, pp. 4384–4396, 2021.
- [32] X. Huang, Y. Chen, W. Yin, and K. Yuan, "Lower bounds and nearly optimal algorithms in distributed learning with communication compression," *Advances in Neural Information Processing Systems*, vol. 35, pp. 18955–18969, 2022.
- [33] Y. He, X. Huang, and K. Yuan, "Unbiased compression saves communication in distributed optimization: When and how much?," *Advances in Neural Information Processing Systems*, vol. 36, pp. 47991–48020, 2023.
- [34] H. Wang, S. Sievert, S. Liu, Z. Charles, D. Papailiopoulos, and S. Wright, "Atomo: Communication-efficient learning via atomic sparsification," *Advances in neural information processing systems*, vol. 31, 2018.
- [35] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- [36] Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer, and V. Stoyanov, "Roberta: A robustly optimized bert pretraining approach," *arXiv preprint arXiv:1907.11692*, 2019.