

Flexible Communication for Optimal Distributed Learning over Unpredictable Networks

Sahil Tyagi

Indiana University Bloomington, USA

styagi@iu.edu

Martin Swany

Indiana University Bloomington, USA

swany@indiana.edu

Abstract—Gradient compression alleviates expensive communication in distributed deep learning by sending fewer values and its corresponding indices, typically via Allgather (AG). Training with high compression ratio (CR) achieves high accuracy like DenseSGD, but has lower parallel scaling due to high communication cost (i.e., parallel efficiency). Using lower CRs improves parallel efficiency by lowering synchronization cost, but degrades model accuracy as well (statistical efficiency). Further, speedup attained with different models and CRs also varies with network latency, effective bandwidth and collective op used for aggregation. In many cases, collectives like Allreduce (AR) have lower cost than AG to exchange the same amount of data. In this paper, we propose an AR-compatible Topk compressor that is bandwidth-optimal and thus performs better than AG in certain network configurations. We develop a flexible communication strategy that switches between AG and AR based on which collective is optimal in the current settings, and model the pareto-relationship between parallel and statistical efficiency as a multi-objective optimization (MOO) problem to dynamically adjust CR and accelerate training while still converging to high accuracy.

1. INTRODUCTION

Deep Neural Network (DNN) training requires numerous passes over the entire dataset (i.e., epochs) to achieve decent accuracy. Synchronous Data-Parallel training *scales-out* by processing different samples concurrently across multiple workers and computing local updates, which are then aggregated either via Parameter Server (PS) [1] or Allreduce (AR) [38]. Gradient synchronization cost is further exacerbated by the growing size of DNNs, which causes slowdowns and degrade the *parallel efficiency* of distributed training [2]. For e.g., storing gradients as single-precision (32-bit) floats in a billion parameter model implies transmitting 4GB of updates on each iteration. Communication collectives like Allgather (AG), Ring-based AR and Tree-based AR offer different latency-bandwidth trade-offs for the same message size [3], which affects overall communication cost. The bandwidth cost of sending-receiving data to-from PS increases linearly with cluster-size, while Ring-AR is bandwidth-optimal. However, latency cost is low for PS and high in Ring-AR. Further, latency or bandwidth itself can vary in a cluster over time due to various factors like resource sharing, QoS priorities, contention among multiple jobs, etc. Thus, parallel efficiency of distributed training is also affected by network availability.

Gradient compression expedites training by lowering the exchange volume to reduce communication time, and is measured by the *degree of compression* or *Compression Ratio*

(CR). High CR converge to high accuracy like DenseSGD (i.e., training *without* any compression), albeit with little to no speedup. Low CRs considerably reduce communication cost, but may degrade final accuracy or require more steps to converge, thus increasing training time [2]. A majority of compression methods use AG to aggregate values and corresponding indices of compressed tensors across workers [3]. In some network configurations, cost of AG can be high than using ring or tree AR to transmit the same message-size. Thus, distributed training exhibits varying communication costs depending on DNN and cluster size under constrained networks with variable latency or bandwidth. Even small models like ResNet18 [29] can suffer from frequent communication in high-latency, low-bandwidth settings. Additionally, distributed DNN training suffers from *statistical inefficiency* from noisy gradients generated by training with smaller mini-batches [4], [5]. Noise is further aggravated when models are updated by a compressed representation of the original gradients [2].

We propose a novel compression technique called **AR-Topk**¹ that is compatible with both ring and tree-AR, and describe its two variants: *staleness-based (STAR-Topk)* and *gradient variance-based (VAR-Topk)*, followed by their communication cost analysis. We then develop heuristics to choose the optimal collective between AR-Topk (ring or tree) and AG for a given network configuration. To balance the pareto-relationship between parallel and statistical efficiency, we approach gradient compression as a multi-objective optimization problem to calculate optimal CR at different training stages.

2. BACKGROUND AND RELATED WORK

We first discuss the synchronization cost of data-parallel training as we scale from intra to inter-node devices, followed by studying how intra-node parallel efficiency is affected by latency, bandwidth, collective op used, model and cluster-size. We also gloss over prior art covering statistical efficiency in distributed training and its impact on convergence quality/time.

A. Distributed Data-Parallel Training

Synchronous distributed training improves throughput by performing more work per-iteration. With weak scaling, N workers concurrently process mini-batch $d_{(i,n)}$ of size b from dataset \mathcal{D}_n to update model parameters w_{i+1} at step i by

¹Code available at <https://github.com/sahiltyagi4/ARTopk>

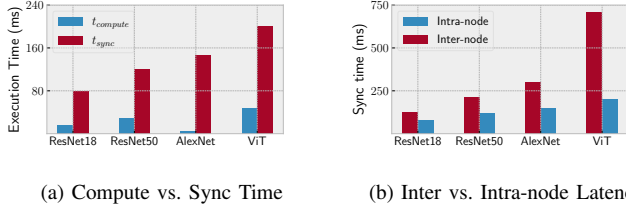


Fig. 1. (a) Intra-node computation-communication times for 8 workers. (b) Latency of gradient aggregation over 8 GPUs connected as: 8 GPUs/node vs. 1 GPU/node. Inter-node GPUs connected over 10Gbps network.

step-size η to minimize loss function $\mathcal{L}(\cdot)$. Eqn. (1a) shows how the local updates are aggregated. The compute cost of each iteration is identical, comprised of mini-batch sampling, computing and averaging gradients, followed by model update. It is thus mainly comprised of computation, communication and IO overhead, shown in Eqn. (1b).

$$w_{i+1} = w_i - \eta \frac{1}{N} \sum_{n=1}^{n=N} \frac{\partial}{\partial w_i} \left(\frac{1}{|b|} \sum_{d_{(i,n)} \in \mathcal{D}_n} \mathcal{L}(x_{(i,n)}, w_i) \right) \quad (1a)$$

$$t_{step} = t_{compute} + t_{sync} + t_{IO} \quad (1b)$$

Fig. (1a) plots the computation and communication time of ResNet18, ResNet50 [29], AlexNet [30] and Vision Transformer (ViT) [31] on single node with 8 NVIDIA V100 GPUs. As DNN size increases (left to right), so does the communication overhead. Compared to compute time, communication tends to be the primary bottleneck on modern ML accelerators.

1) *Intra vs. Inter-node Scaling*: In data-parallel training, synchronization cost to aggregate updates between nodes depends on the physical distance between them. GPU devices on a single node are connected via high-speed interconnects like PCIe Gen4 (64GB/s), Gen5 (128GB/s) or NVLink 4.0 [7] with 900GB/s bandwidth to enable *intra-node scaling*. This type of scale-up on a single system is limited by finite PCIe lanes, for e.g., NVIDIA DGX systems can accommodate up to 8 GPUs/node [8] and NVSwitch can connect up to 256 GPUs in a cluster. *Inter-node scaling* spans out to multiple servers, but suffers from high communication cost to transmit massive tensors over networks with comparatively lower bandwidths. The typical bandwidth in a data-center is about 10Gbps [9] and even lower at the edge. We compare gradient synchronization costs for various DNNs in intra and inter-node settings. In the former, 8 V100s are co-located on a single node (i.e., 8 GPUs/node) with peer-to-peer (P2P) transport enabled to allow CUDA direct access, while the latter is comprised of 8 machines with 1 GPU each connected via 10Gbps interface. Inter-node communication cost is much higher while aggregating updates via Allreduce on NVIDIA’s NCCL communication backend (Fig. (1b)).

Table I. Bandwidth Complexity and Communication Cost of Primitives

Operation	BW Complexity	Communication cost
PS (Star)	$\mathcal{O}(MN)$	$2\alpha + 2(N-1)M\beta$
Ring-AR	$\mathcal{O}(M)$	$2(N-1)\alpha + 2\frac{(N-1)}{N}M\beta$
Tree-AR	$\mathcal{O}(M \log(N))$	$2\alpha \log(N) + 2\log(N)M\beta$
Broadcast	$\mathcal{O}(M \log(N))$	$\alpha \log(N) + \log(N)M\beta$
Allgather	$\mathcal{O}(MN)$	$\alpha \log(N) + (N-1)M\beta$

2) *Cost of Communication Collectives*: Gradients can be aggregated in various ways depending on the topology and algorithm used. Centralized mechanisms exchange updates via Parameter servers (PS) based on a star-topology, while decentralized approaches use different flavors of Allreduce or AR (e.g., reduce-broadcast, scatter-reduce-allgather, ring, tree, etc.). Table (I) lists down the α - β communication cost [28], [38] and bandwidth complexity of collectives commonly used in distributed training. The α term is latency, $1/\beta$ is bandwidth and M is the model/gradient size exchanged among N workers. In PS, synchronization cost increases with both cluster and model-size. Ring-AR is dominated by latency cost as its bandwidth-optimal, i.e., β term is nearly independent of N and increases only with M . Although Tree-AR is *logarithmic* bandwidth-optimal, it has lower latency cost than ring; α term is proportional to $\log(N)$ in the former vs. $(N-1)$ in the latter. Hence, communication cost of different collectives varies for a given network configuration, cluster-size and model-size, as well as their implementation. For e.g., AR over CUDA GPUs is faster with NCCL over Gloo, while Gloo performs better than MPI on CPUs.

3) *Communication Scheduling*: Efficient communication scheduling accelerates training by hiding synchronization costs. Although data-parallel training can be scaled out by processing data on many workers simultaneously, backpropagation itself is inherently sequential. Input is fed to the innermost layer while outputs are generated on the outermost layer of DNNs. By chain rule, gradients are first computed on the outer layers and propagate back to the inner layers. ML frameworks improve iteration time by hiding communication cost and overlapping it with computation [10], although this can be inefficient in DNNs with many layers of non-uniform size. Thus, updates of current layer are sent while gradients of subsequent layer are still being computed. With *tensor-fusion* [13], smaller layers are batched together to perform fewer communication operations. Frameworks like PyTorch allow *bucketing* to specify message-size for each allreduce call, with the default being 25MB. Gradient computation time has decreased with the development of optimized GPU architectures, while size of DNNs continues to grow exponentially. Thus, gain of interleaving computation with communication is saturating as the former tends to be lower than the latter.

B. Statistical Efficiency of Data-Parallel Training

Although deep learning is an iterative-convergent process, features learned by a model over the epochs are not uniformly attributed to each training step due to the stochastic nature

of gradient descent, i.e., iterations do not contribute equally towards the overall learning of a model. Gradients calculated from a randomly sampled mini-batch tend to be noisier than those computed over the full-batch (i.e., entire training data) [4]. The divergence between local and averaged gradients represents the signal-to-noise ratio in the approximated gradients, which has prior been studied for fine-tuning cluster and batch-size in distributed training [5], [6], [14], [15]. The rate of gradient change throughout training is not uniform either. Gradients computed in early stages tend to be highly volatile change drastically as a model adjusts its parameters towards the minima aggressively [2], [16]. Thus, gradients are large initially and get smaller as a model converges and eventually saturates. Additionally, there are critical regions in training where gradients are extremely sensitive [2], [11], [17], [18]. This volatility is influenced by training hyperparameters like step-size or learning-rate schedule, gradient clipping, SGD variant (e.g., Nesterov, Adam, Adagrad, etc.). First-order gradients or *variance* approximates second-order information and thus helps detect critical regions and measure the statistical efficiency in data-parallel training [11], [18], [19].

C. Gradient Compression to Mitigate Communication Cost

Lossy gradient compression attenuates communication cost by reducing the volume of gradients exchanged over the network. Sparsification methods like layerwise Topk, or LWTopk [20] and MSTopk [21] retrieve the largest k% gradients and transmit their respective values and indices, while setting every other tensor element to 0. The effective CR in this case is thus $k/100$. Other sparsifiers like DGC [43] and SIDCo [44] compress by modeling gradients over some distribution and estimating threshold for a target compression ratio. Quantization techniques like signSGD [24] and TernGrad [25] limit or quantize the gradients to a fixed subset of values, while methods like 1-bit SGD [26] reduce the bit-width of single-precision gradients. Low-rank approximators like PowerSGD [27] decompose $\mathbb{R}^{p \times q}$ dimensional gradients into two lower r -rank matrices with dimensions $\mathbb{R}^{p \times r}$ and $\mathbb{R}^{r \times q}$. Additionally, lossy gradient compression methods rely on *error-feedback* to avoid losing significant updates when using a high degree of compression [20], [21]. In Eqn. (2), gradients from step $(i - 1)$ that are ineligible for communication after applying compression are *not* discarded, but appended to gradients computed in the next iteration, $g_o^{(i)}$. A compression operator \mathcal{C} thus reduces the error-fed gradients $g_e^{(i)}$ to $g_c^{(i)}$ and updates residual gradients accordingly.

$$g_e^{(i)} = g_o^{(i)} + \text{residual}^{(i-1)} \quad (2a)$$

$$g_c^{(i)} = \mathcal{C}(g_e^{(i)}) \quad \text{and} \quad \text{residual}^{(i)} = g_e^{(i)} - g_c^{(i)} \quad (2b)$$

1) *Parallel Efficiency in Gradient Compression*: Although it reduces the synchronization cost and speeds up training, compression itself is *not a zero-cost operation*. Parsing and compressing gradients before transmission, and deparsing

them before applying SGD-update incurs an additional, non-negligible overhead. With gradients compressed to CR c , the per-step time for distributed training in Eqn. (1b) reduces to Eqn. (3).

$$t_{step}^{(c)} = t_{compute} + t_{sync}^{(c)} + t_{IO} + t_{comp-decomp}^{(c)} \quad (3)$$

The computation time in Eqn. (1b) and (3) irrespective of compression since it depends on batch-size and model used. However, only a fraction of the total updates are sent instead of the entire gradient tensor so as to reduce synchronization time from t_{sync} in DenseSGD to $t_{sync}^{(c)}$. For compression to be viable, its cumulative compression and communication overhead needs to be lower than the synchronization cost of DenseSGD: $t_{sync}^{(c)} + t_{comp-decomp}^{(c)} < t_{sync}$. Interleaving gradient computation (on subsequent layers) and communication (on current layer) may not be efficient due to contention between them for the same compute resources.

Collective operation also influences the overall speedup in distributed training. In DenseSGD, gradients are reduced either via PS or AR (as described in §2-A2). PS and Ring/Tree-AR have low-latency and low-bandwidth cost respectively. However, the resulting gradients post-compression operation are not compatible with PS or AR-like reductions. Allgather (AG) is typically used to synchronize compressed updates by transmitting select values and its indexes from each worker, since an index chosen on one worker may not have necessarily been chosen for communication on a different worker. A gradient tensor of size G compressed to top- $k\%$ values requires transmitting $\lceil 2kG/100 \rceil$ data-points (i.e., $\lceil kG/100 \rceil$ values and $\lceil kG/100 \rceil$ indices). Although some compressors like Randomk and PowerSGD are allreduce-friendly, the former has a poor convergence quality while the latter has a considerable compression overhead and achieves lower compression rate compared to others [27].

2) *Latency and Bandwidth Variability*: Variable and unpredictable networks are prevalent on edge, cloud and HPC due to a variety of reasons. Different classes of VMs have varying networking tiers available to them in terms of ingress/egress bandwidth [33], [34]. In mobile computing and IoT, edge devices tend to have slower networking compared to fog nodes/servers. Network heterogeneity can also arise from the topology (e.g., star, mesh, ring, tree, etc.), different links (ethernet, infiniband) and physical mediums (e.g., fiber, copper). Both latency and bandwidth can vary even over the same link with time due to a variety of factors. Network bandwidth, which measures the data-transfer rate of a connection can fluctuate over time because:

- *Network Congestion*: The effective bandwidth can decrease over high network traffic, increasing data-transfer time on account of saturated links.
- *Quality of Service (QoS) prioritization*: Jobs deployed on a shared cluster may have different QoS priorities associated with them, such that high-priority jobs are

Table II. Topk Compression + Communication Cost for CRs 0.1 and 0.001 via Allgather (AG) vs. Ring-AR on uncompressed data for different α (ms) and $1/\beta$ (Gbps) on tensors with 100 Million and 1 Billion parameters.

Configuration		Time cost (ms)		
Tensor size	$(\alpha, 1/\beta)$	AG 0.1	AG 0.001	Ring-AR
10^8	(10, 10)	525	70	716
	(10, 5)	976	74	1271
	(10, 1)	4568	111	5773
	(100, 10)	798	340	1975
	(100, 5)	1248	345	2530
	(100, 1)	4830	380	7028
	10^9	(10, 10)	5010	482
(10, 5)		9507	534	11380
(10, 1)		45355	898	56190
(100, 10)		5280	745	7024
(100, 5)		9805	791	12621
(100, 1)		45645	1154	57442

allocated a bigger chunk of the total bandwidth. As a result, low-priority jobs suffer from high-communication cost due to lower available bandwidth.

- *Resource Sharing/Contention*: Even without explicit prioritization, concurrent jobs in a shared cluster share finite networking resources, resulting in lower-bandwidth available per-job.
- *Network Scheduling*: Switches and routers may use different scheduling algorithms to allocate bandwidth for different traffic flows, which influences how networking resources are shared among applications.

The communication delay over a network, or its latency, can fluctuate for jobs running on the edge or even data-centers. Local area network (LAN) has very low latency, ranging from micro to milliseconds, while inter-rack latency tends to be in the order of tens of milliseconds, while inter-node latency and across data-centers is even worse (see Fig. (1b)). Network load and congestion can also cause latency-surge under high traffic flows due to queuing delays as multiple jobs perform simultaneous data transfers. Further, network routing that decides the data-path of a packet also influences its latency. Although routing protocols dictate the optimal path for transmission, different routes can result in varying latency due to network congestion and physical distance. As a result, fluctuating latency and bandwidth over a cluster of nodes can affect a job’s completion time.

On 8 inter-node V100s connected via 40Gbps network, we vary α (latency in ms) and $1/\beta$ (bandwidth in Gbps) using the `linux traffic control` or ‘tc’ module [35] and measure the time to synchronize tensors with 100 million and 1 billion parameters represented as single-precision floats. Latency is controlled via `netem` (network emulation) `qdisc` (queuing discipline), while bandwidth is adjusted via ‘htb’ (hierarchical token bucket) `qdisc`. We compare the time taken to reduce the original, uncompressed tensors using allreduce (Ring-AR) collective vs. time to communicate tensors compressed to different CRs using allgather (AG) with different network configurations. For e.g., 1B parameter tensor using CR 0.001 with AG exchanges 2 million data-points, while Ring-AR com-

Table III. Step-time, Accuracy and Diff. (w.r.t. DenseSGD) on a 4ms, 20Gbps network. DenseSGD uses Ring-AR, LWTopk/MSTopk use AG.

Model	Method	t_{step} (ms)	Acc.	Diff.
ResNet18	DenseSGD	98.7	90.8%	0.0%
	LWTopk 0.1	62	90.15%	-0.65%
	LWTopk 0.01	38.4	88.49%	-2.31%
	LWTopk 0.001	36.8	87.2%	-3.6%
	MSTopk 0.1	83.22	90.59%	-0.21%
	MSTopk 0.01	60.22	89.83%	-0.97%
	MSTopk 0.001	58	88.69%	-2.11%
ResNet50	DenseSGD	152.5	98.75%	0.0%
	LWTopk 0.1	130.24	98.65%	-0.1%
	LWTopk 0.01	79.4	98.3%	-0.45%
	LWTopk 0.001	72.7	96.55%	-2.2%
	MSTopk 0.1	124.35	98.64%	-0.11%
	MSTopk 0.01	72.63	98.49%	-0.26%
	MSTopk 0.001	67.33	97.73%	-1.02%
AlexNet	DenseSGD	230.6	82.41%	0.0%
	LWTopk 0.1	156.8	81.68%	-0.71%
	LWTopk 0.01	33.45	81.45%	-0.96%
	LWTopk 0.001	21.3	78.9%	-3.51%
	MSTopk 0.1	214.5	81.8%	-0.61%
	MSTopk 0.01	91.4	81.7%	-0.71%
	MSTopk 0.001	79	78.92%	-3.49%
ViT	DenseSGD	475	80.4%	0.0%
	LWTopk 0.1	362.4	79%	-1.4%
	LWTopk 0.01	94.64	78.75%	-1.65%
	LWTopk 0.001	67.7	78.25%	-2.15%
	MSTopk 0.1	543.6	80.47%	+0.07%
	MSTopk 0.01	275.24	79.87%	-0.53%
	MSTopk 0.001	248.8	79.76%	-0.64%

munication all 1 billion values. For AG, Table (II) shows the cumulative time, i.e., sum of compression and communication costs. Although AG at say, CR c performs better than Ring-AR for the same (α, β) due to smaller message-size, Ring-AR does *not* take $(1/c)\times$ more time to communicate. This is because the communication cost of different collectives varies with message-size, cluster-size and network parameters. Ring-AR is bandwidth-optimal but vulnerable to high-latency, as also seen from the wider gap between AG and Ring-AR in 100ms latency case. Synchronization overhead in AG increases considerably on low-bandwidth networks even when latency stays the same. These results corroborate the α - β communication cost model described in Table (I).

3) *Statistical Efficiency in Gradient Compression*: As explained in §2-B, distributed training has a statistical efficiency aspect associated with it due to gradient *variance* or *noise*. This is further exacerbated in lossy compression where model weights are updated with an even less accurate gradient representation [2]. To see the impact of varying degrees of compression on convergence, we train DNNs with *LWTopk* [20] and *MSTopk* [21] at CRs [0.1, 0.01, 0.001]. We compare final model accuracy with that of DenseSGD and observe differences in their convergence quality. Layerwise or LWTopk retrieves top- $k\%$ gradient values and indexes on a layer-by-layer basis, while MSTopk approximates top- k on the entire gradient tensor (i.e., not on layerwise basis) via multi-sampling and uses binary search to find the threshold corresponding to target CR. Threshold estimation in MSTopk involves multiple trials or searches for a configurable number of rounds to pre-

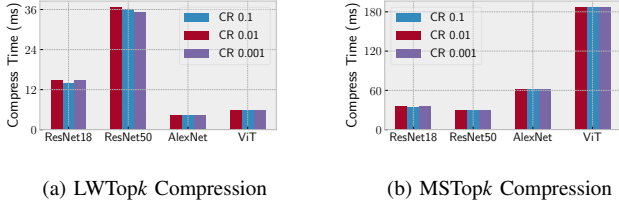


Fig. 2. Compression overhead of different techniques. For the same CR, MSTopk has higher compression cost due to multi-round threshold estimation.

dict an accurate threshold; we use 25 rounds in our evaluation. Training routine and hyperparameters are described in §3-C1.

Table (III) tabulates the step-time, test accuracy and convergence difference (w.r.t. DenseSGD) while training on the same cluster of 8 inter-node GPUs when the latency and bandwidth are limited to 4ms and 20Gbps via t.c. . Step-time includes gradient computation, IO overhead, compression and communication cost for different CRs using AG-collective, while DenseSGD has no compression cost but transmits entire gradients via AR-collective. Further, the step-time of MSTopk is higher than LWTopk irrespective of the model or CR. This is because the former has higher compression cost on account of multiple rounds needed to compute the threshold, as also seen in Fig. (2). In some cases, like ViT at CR 0.1, MSTopk has higher iteration time than even DenseSGD because of high compression overhead. However, MSTopk achieves similar or *even better* accuracy than LSTopk since its estimates threshold over the entire model, while the latter compresses on a layerwise basis. Thus, in LWTopk, models with non-uniform layers and skewed gradients may lose more critical updates since the fraction of gradients selected from each layer is fixed (in proportion to the layer’s size). At the same time, more significant updates from larger layers may not even be eligible for communication if most of the critical updates are cluttered in only a few layers.

ResNet18 converges to 90.8% accuracy with DenseSGD, while LSTopk achieves 88.49% and 87.2% accuracy for CRs 0.01 and 0.001 respectively. MSTopk on ResNet18 attains 89.83% and 88.69% for similar CRs. We observe a similar pattern for all other DNNs; as more compression is applied and CR is lowered, model convergence degrades as well. This happens because we lose considerably more information at smaller CRs. GraVAC [2] defines a heuristic called *Compression gain* to measure this loss by comparing prior and post-compression gradients. It compares the variance between errorfed and compressed gradients (i.e., $g_e^{(i)}$ and $g_c^{(i)}$ from Eqn. (2)) at step i as: $E[\|g_c^{(i)}\|^2] / E[\|g_e^{(i)}\|^2]$.

The overhead of computing compression gain is low as it only needs first-order gradients already computed in backpropagation. We plot gains of LWTopk and MSTopk at various CRs in Fig. (3) and see that gain tends to be smaller for lower CRs as more gradient information is lost there. Gain also changes in early training stages and certain critical regimes [2], [16]–

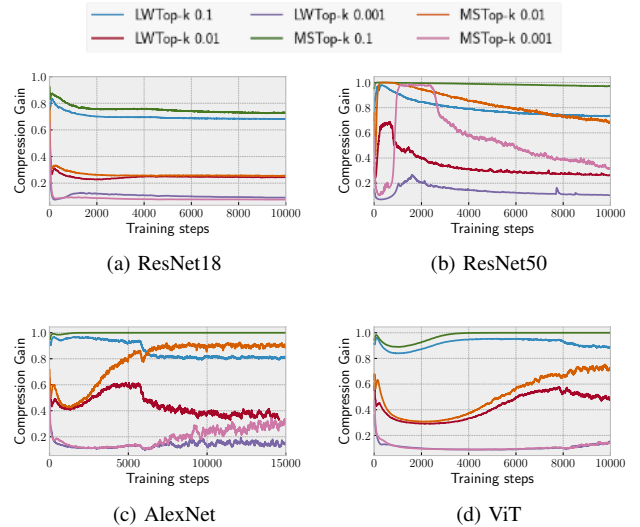


Fig. 3. Compression gain measures statistical efficiency in lossy compressors.

[18] (like step-size decay in our case) before it eventually saturates. Additionally, we see a correlation between compression gain and model convergence for each compressor and CR configuration. From Table (III), ResNet50 at CR 0.001 achieves 96.55% accuracy with LWTopk and a higher 97.73% with MSTopk. At the same time, MSTopk has a higher compression gain than LWTopk, seen in Fig. (3b). Further, the generalization performance of different compressors at different CRs can also be inferred by their compression gains. For e.g., AlexNet with LWTopk at CR 0.1 and MSTopk with CR 0.01 have roughly the same 81.68% and 81.7% accuracy, while their gains also saturate to the same level (Fig. (3c)). Thus, we see how compression gain works as a practical and effective heuristic to evaluate the efficacy of a compressor and measure statistical efficiency in lossy gradient compression.

3. AR-TOPk COMPRESSION

In this paper, we propose a new compression technique compatible with AR-collective called *AR-Topk* along with its two variants, followed by evaluation and communication cost analysis of the same. However, as we will see, aggregating updates via AR collective may not necessarily be better than AG, depending on network and distributed training parameters. We thus propose a flexible compression-communication scheme that utilizes the faster collective based on the α - β communication model. To determine the ideal CR and collective to use (either AG, AR-Topk with ring or tree-AR) that best balances parallel and statistical efficiency when training with lossy compression, we model gradient compression as a multi-objective optimization (MOO) problem.

A. Allreduce-friendly Topk Compression (AR-Topk)

We implement Top- k sorting using max-heap and retrieve top $k\%$ values over fused gradient tensors, i.e., compression is applied over the entire model and not just layerwise (as with LWTopk). Compared to MSTopk which is based on threshold estimation, the compression overhead of max-heap

Algorithm 1: AR-Topk Compression

```

1 Input: CR  $c$ , step-size  $\eta$ , cluster-size  $N$ , worker-rank  $r$ 
2 function train():
3   residual(0,r) = 0
4   for  $i = 1, 2, \dots$  do
5      $G_{(i,r)} = \nabla \mathcal{F}(x_{(i,r)}, w_{(i,r)}) + \text{residual}_{(i-1,r)}$ 
6      $\mathbf{g}_{(i,r)}, \mathbf{ix}_{(i,r)} = \text{Topk}(G_{(i,r)}, c)$ 
7     if STAR-Topk then
8        $\tilde{r}_{(i)} = i \% N$ 
9     else if VAR-Topk then
10      var [N] = 0
11      var [r] =  $\|g_{(i,r)}\|^2$ 
12      var = AllGather(var)
13       $\tilde{r}_{(i)} = \text{var.indexOf}(\max(\text{var}))$ 
14      $\tilde{\mathbf{ix}}_{(i,r)} = \text{Broadcast}(\mathbf{ix}_{(i,r)}, \text{src} = \tilde{r}_{(i)})$ 
15      $\tilde{\mathbf{g}}_{(i,r)} = G_{(i,r)}[\tilde{\mathbf{ix}}_{(i,r)}]$ 
16     residual(i,r) =  $G_{(i,r)} - \tilde{\mathbf{g}}_{(i,r)}$ 
17      $\tilde{\mathbf{g}}_{(i,r)} = \text{AllReduce}(\tilde{\mathbf{g}}_{(i,r)})$ 
18      $w_{(i+1,r)} = w_{(i,t)} - \eta \cdot \tilde{\mathbf{g}}_{(i,r)}$ 
19     i++

```

based sorting is much lower on GPUs. Between AG and AR, the collective that incurs the least communication overhead in distributed training depends on the message-size to transmit, cluster-size, target CR and network-specific parameters like latency and bandwidth. The message size is doubled with AG as both values and indices are exchanged, while AR is not out-of-the-box compatible with lossy compression. AR-Topk takes the compressed gradient values and indexes from Top- k sorting and initiates a Broadcast from a worker to disperse its corresponding indices, followed by AllReduce to aggregate values on indices that were broadcasted in the first step. The AR call in second step can use either ring or tree-based reduction (covered in §3-D).

B. AR-Topk Worker Selection Mechanism

After broadcasting the compressed indices from a worker, gradient values reduced at those positions may not necessarily correspond to top- k gradients across all workers. *Additionally, the natural question that then arises is, ‘How do we choose which worker gets to broadcast its indices to all other workers in the cluster for aggregating its local top- k values?’* We explore two mechanisms to determine worker selection in AR-Topk, one *variance-based* and the other *staleness-based*:

- 1) **Variance-based AR-Topk (VAR-Topk)** selects worker with largest gradient variance and disperses the indexes of its local top- k values, followed by calling AR on those elements. As discussed in §2-B and §2-C3, the significance of updates can be measured by their gradient variance. Thus, worker with maximum variance sets the tone for calling AR-collective to aggregate values at indices corresponding to its local top- k values.

Table IV. Convergence and t_{step} (ms) in DenseSGD and STAR/VAR-Topk. DenseSGD uses Tree-AR on 4ms latency and 20Gbps bandwidth network.

Model	Method	t_{step} (ms)	Acc.	Diff.
ResNet18	DenseSGD	146.21	90.79%	0.0%
	STAR-Topk 0.1	64.83	90.19%	-0.6%
	STAR-Topk 0.01	49.68	89.89%	-0.9%
	STAR-Topk 0.001	48.17	88.81%	-1.98%
	VAR-Topk 0.1	77.2	90.15%	0.64%
	VAR-Topk 0.01	62.1	89.5%	-1.29%
ResNet50	VAR-Topk 0.001	60.5	88.7%	-2.09%
	DenseSGD	294.35	98.72%	0.0%
	STAR-Topk 0.1	100.8	98.68%	-0.04%
	STAR-Topk 0.01	67.68	98.57%	-0.15%
	STAR-Topk 0.001	64.37	98.38%	-0.34%
	VAR-Topk 0.1	113.4	97.65%	-1.07%
AlexNet	VAR-Topk 0.01	80	98.2%	-0.52%
	VAR-Topk 0.001	76.5	97.33%	-1.39%
	DenseSGD	614	82.4%	0.0%
	STAR-Topk 0.1	129.5	81.74%	-0.66%
	STAR-Topk 0.01	50.3	80.41%	-1.99%
	STAR-Topk 0.001	42.38	79.57%	-2.83%
ViT	VAR-Topk 0.1	142	83.15%	+0.75%
	VAR-Topk 0.01	62.6	81.78%	-0.62%
	VAR-Topk 0.001	54.4	78.06%	-4.34%
	Dense-SGD	1348.5	80.42%	0.0%
	STAR-Topk 0.1	276.32	79.72%	-0.7%
	STAR-Topk 0.01	104.13	79.18%	-1.24%
ViT	STAR-Topk 0.001	86.91	79.23%	-1.19%
	VAR-Topk 0.1	289.2	80.12%	-0.3%
	VAR-Topk 0.01	117	79.2%	-1.22%
	VAR-Topk 0.001	99.7	79.8%	-0.62%

- 2) **Staleness-based AR-Topk (STAR-Topk)** chooses top- k indices from workers in a round-robin fashion, which are then broadcasted to call AR on gradients at those indices across all workers. In this way, each worker gets to transmit its compressed update once every N iterations (in a cluster of N workers). While error-feedback effectively works as *delayed* or *stale* updates, STAR-Topk further adds to it by adding an implicit staleness of N steps.

Alg. (1) illustrates how workers are selected in AR-Topk compression. Error-fed gradients $G_{(i,r)}$ at step ‘ i ’ are compressed to CR ‘ c ’ on each worker ‘ r ’ in the first step. STAR-Topk selects a worker based on the current iteration number and cluster-size (line 8 in Alg. (1)). VAR-Topk involves additional steps/overhead, such that an array of size ‘ N ’ holds the gradient variance of worker ‘ r ’ at index ‘ r ’ (line 11). An AG call on the array synchronizes variance across all workers (line 12), and the worker with maximum variance is chosen for communication (line 13). *The overhead of this step is low as the message-exchange size is just ‘ N ’ 32-bit floats.* Once a worker (denoted by rank $\tilde{r}_{(i)}$) is selected either by STAR or VAR-Topk, its corresponding indices are dispersed via broadcast (line 14) The gradients residing at those indexes on the remaining workers are then selected, residual gradients are updates and finally aggregated via AR-collective (lines 15-17). Lastly, model parameters are updated and training proceeds to the next iteration (lines 18-19).

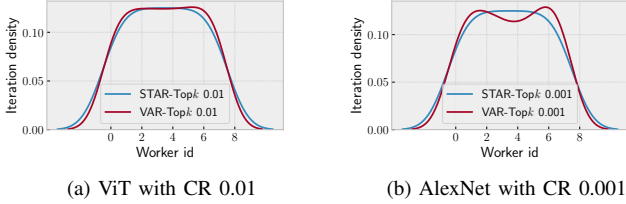


Fig. 4. Iteration density of the 8 workers (ranked 0-7) that broadcast their topk indices with STAR and VAR-Topk compression. STAR-Topk has similar density across all workers due to its round-robin nature, while some workers may have higher density of updates over others in VAR-Topk.

C. Training Speedup and Convergence in AR-Topk

1) *DNNs and Hyperparameters*: The following models are deployed with PyTorch version v1.12.1 [12], using NCCL v2.10.3 for communication between 8 inter-node GPUs. ResNet18 trains on CIFAR100 [32] for 50 epochs with per-worker batch-size 32, momentum 0.1, weight decay 0.0005, step-size 0.1 and decay factor 0.1 at epoch 15, 30 and 45. ResNet50 runs on Food101 [39] dataset for 100 epochs with batch-size 64, momentum 0.1, step-size 0.1 and weight decay 0.0001. AlexNet trains on Caltech101 [36] for 50 epochs with batch-size 32, momentum 0.9, weight decay 0.0005, step-size 0.01 and decay factor 0.1 at epoch 25. ViT runs on Caltech256 [37] for 50 epochs with batch-size 32, momentum 0.9, weight decay 0.0005, step-size 0.01 and decay rate 0.2 on epoch 40. We report the top-1 test accuracy and time to complete one training step (in ms).

2) *Training with STAR-Topk and VAR-Topk Compression*: Table (IV) compares the model convergence and per-iteration speedup in STAR and VAR-Topk with DenseSGD. We train on a network limited to 4ms latency and 20Gbps bandwidth via t_c , and configure NCCL to use Tree-AR in DenseSGD by setting environment variable `NCCL_ALGO` to ‘tree’. The general trend is that accuracy decreases as CR is lowered in both STAR and VAR-Topk. For the same CR, STAR-Topk performs better than VAR-Topk at the current cluster-size and network configuration. For ResNet50 CR 0.001, STAR-Topk had 0.34% lower accuracy than DenseSGD, while VAR-Topk converged to 1.39% lower accuracy. To understand the deviation in generalization quality between STAR and VAR-Topk even at the same CR, we look at the iteration kernel density estimates (KDE) of each worker that broadcasts its topk indices over the entire training, with worker ids ranked from 0 to $(N - 1)$. We see similar iteration densities of STAR and VAR-Topk in ViT CR 0.01 (Fig. (4a)), which also have close accuracy as shown in Table (IV); STAR-Topk achieves 79.18% while VAR-Topk saturates to 79.2%. With AlexNet at CR 0.001, STAR and VAR-Topk converged to 79.57% and 78.06% respectively, a considerable 1.51% difference. In this case, iteration density of VAR-Topk is more skewed and non-uniform compared to STAR-Topk, as seen in Fig. (4b). Worker ranks 1 and 6 broadcast considerably more updates than others, while rank 4 transmits its topk gradients much less frequently.

Table V. Comparing STAR-Topk, VAR-Topk and LWTopk on ResNet18/50, AlexNet and ViT. AR-Topk variants use Allreduce, LWTopk uses Allgather.

DNN config		t_{step} (ms)			Accuracy (%)		
Model	CR	STAR	VAR	LW	STAR	VAR	LW
Res.18	0.1	64.83	77.2	62	90.19	90.15	90.15
	0.01	49.68	62.1	38.4	89.89	89.5	88.49
	0.001	48.17	60.5	36.8	88.81	88.7	87.2
Res.50	0.1	100.8	113.4	130.24	98.68	97.65	98.65
	0.01	67.68	80	79.4	98.57	98.2	98.3
	0.001	64.37	76.5	72.7	98.38	97.33	96.55
Alex.	0.1	129.5	142	156.8	81.74	83.15	81.68
	0.01	50.3	62.6	33.45	80.41	81.78	81.45
	0.001	42.38	54.4	21.3	79.57	78.06	78.9
ViT	0.1	276.32	289.2	362.4	79.72	80.12	79
	0.01	104.13	117	94.64	79.18	79.2	78.75
	0.001	86.91	99.7	67.7	79.23	79.8	78.25

The infrequent updates of rank 4 are thus more stale which in turn degrades model quality. On the other hand, STAR-Topk has a much more even and uniform iteration density due to its round-robin approach.

Additionally, iteration time is higher in VAR-Topk over STAR-Topk. This is because the former involves additional steps over the latter, as illustrated in lines (10-13) of Alg. (1). VAR-Topk makes an additional AG call to collect gradient variance from all workers, exchanging $4N$ bytes of data in the process. The overhead of this AG call is negligible due to its small message-size. *Although STAR-Topk performs better both in terms of parallel (due to its lower compute overhead) and statistical efficiency (as it achieves better accuracy with its round-robin approach), we conjecture that VAR-Topk compression is also a viable approach, especially when training on unbalanced and non-i.i.d. (independent and identically distributed) data or when training over a massive number of devices, as commonly seen in federated learning. With unbalanced data, collecting updates from workers with large variance prioritizes critical updates. When cluster-size N is huge, the round-robin updates of STAR-Topk can get too stale and degrade model convergence.*

3) *Performance of AR-Topk over LWTopk Compression*: In this section, we directly compare AR-Topk (both variants) with LWTopk in Table (V) where the CR is fixed throughout training. STAR/VAR-Topk use AR collective communication, while LWTopk uses AG. The network latency is limited to 4ms, while bandwidth is capped at 20Gbps. We observe that STAR-Topk has a lower iteration time than LWTopk in some cases, and higher in others. For e.g., STAR-Topk is faster than LWTopk at all CRs of ResNet50, and CR 0.1 in AlexNet and ViT. Inversely, LWTopk is faster at all CRs in ResNet18, and CRs (0.01, 0.001) in AlexNet and ViT. VAR-Topk has higher iteration cost across all models due to its additional variance compute operation. Thus, AG works better than AR in some cases, and vice versa, determined by network topology, latency, bandwidth, CR, model and cluster-size. AR-Topk applies tensor fusion prior compression, i.e., we compress gradients as a whole across all layers while LWTopk does so on a layer-by-layer basis. This appeal of the

former approach can be seen from model accuracy in Table (V); STAR-Topk achieves higher accuracy across all CRs of ResNet18 and ResNet50, while VAR-Topk attains higher accuracy for majority of the CRs in AlexNet and ViT.

D. Communication Cost Analysis of AR-Topk

Communication cost of AG to exchange $2Mc$ bytes of gradient values and indices is: $\alpha \log(N) + 2Mc\beta(N-1)$, where the network latency is denoted by α and $1/\beta$ is the available bandwidth to transmit M bytes of gradient information that are compressed to CR c . With AR-Topk, communication is a two-step process: one Broadcast to distribute top- k indexes from a single worker, followed by AR to reduce gradient values at those specific indexes. AR-Topk can either use Ring or Tree-AR during reduction phase, which have different costs associated with each of them (see Table (I)). Thus, communication cost of AR-Topk depends on the specific flavor of AR algorithm used, shown in Eqn. (4a) and (4b) for AR-Topk with Ring and Tree reduction respectively.

$$t_{ARTring}^{(c)} = \alpha[2(N-1) + \log(N)] + Mc\beta[2\frac{(N-1)}{N} + \log(N)] \quad (4a)$$

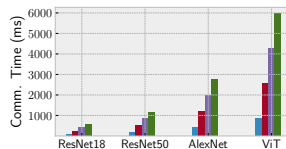
$$t_{ARTtree}^{(c)} = 3\alpha \log(N) + 3Mc\beta \log(N) \quad (4b)$$

We limit α to 1 ms and vary $1/\beta$ as 10, 5 and 1Gbps to measure communication time between AG, AR-Topk (ART) with Ring and Tree-AR for various DNNs and CRs on a cluster of 8 inter-node V100s. We fuse tensors into 64MB buckets with PyTorch gradient-bucketing and list the communication cost of each collective in Table (VI). At a moderately-high 10Gbps bandwidth and CR 0.1, ART-Ring has the least communication overhead across all DNNs. On the other hand, AG-collective has the least communication cost at very low CRs like 0.001 and high-to-moderate bandwidth (like 10 and 5Gbps). In low-bandwidth settings, AR-Topk had the advantage over AG, since the AR call made in the former is bandwidth-optimal. ART-Ring was the optimal collective in low-latency, high-bandwidth (10Gbps) and high CR (0.1) scenarios as well. Thus, the optimal collective operation to use in distributed training with gradient compression involves a large trade-off space and depends on network latency, bandwidth, model-size, cluster-size and CR. AG generally performs better when model size is relatively small (like ResNet18) with moderately-high network bandwidth or very low CRs. In contrast, AR-Topk with ring or tree performs better larger models in low-bandwidth cases, as well as high CRs on high-bandwidth networks.

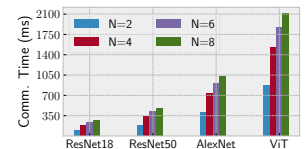
Moreover, we compare the *scale-out* cost between AR-Topk and AG as cluster-size increases. Fig. (5) shows the results for CR 0.1 as N scales from 2 to 8 on a network with an average latency of 5ms and a low 1Gbps bandwidth. The cost of ART-Ring has a gradual incline, while cost of AG collective increases much more sharply when N is large. This is because the bandwidth-cost of both broadcast increases as logarithm

Table VI. Communication cost of varying bandwidths with fixed latency of 1ms for AR-Topk (ART) with Ring and Tree-AR, and with AG-collective.

Configuration		Communication Time (ms)			
Model	$(\alpha, 1/\beta)$	CR	AG	ART-Ring	ART-Tree
ResNet18	(1,10)	0.1	54	35	43.2
		0.01	7.66	18.1	12.2
		0.001	3.28	16.7	9
	(1,5)	0.1	107.76	52.5	76.3
		0.01	13.83	20.8	16.1
		0.001	4.25	17.9	10.1
	(1,1)	0.1	526.3	194.7	345.6
		0.01	51.93	34.1	41.9
		0.001	8.86	19.5	12.8
ResNet50	(1,10)	0.1	115.1	52.9	83.4
		0.01	14.35	20.3	15.9
		0.001	4.65	18.1	10
	(1,5)	0.1	232	94.7	156.2
		0.01	28.1	26.1	24.2
		0.001	5.3	17.8	10.5
	(1,1)	0.1	1148	405.5	745
		0.01	126.5	58.8	83.7
		0.001	14.35	21	16.1
AlexNet	(1,10)	0.1	271.8	106.8	180.4
		0.01	32.73	25.2	25.8
		0.001	6	18.6	11.1
	(1,5)	0.1	544.5	200.4	354.8
		0.01	61.75	34.8	42.6
		0.001	8.92	19.3	13.1
	(1,1)	0.1	2718.7	964.4	1778
		0.01	282.7	111.8	186.8
		0.001	31.33	27	27.3
ViT	(1,10)	0.1	592.77	238.6	401.2
		0.01	68.48	36.2	46.2
		0.001	9.15	19.2	12.9
	(1,5)	0.1	1206	424.3	779.1
		0.01	127.45	58	86.2
		0.001	15.3	21.4	16.9
	(1,1)	0.1	5973	2047	3852
		0.01	601.8	222.8	385.2
		0.001	59.68	36.7	44.4



(a) AllGather communication



(b) AR-Topk communication

Fig. 5. Communication time increases more steeply with N in Allgather vs. AR-Topk CR 0.1. Network has an average 5ms latency and 1Gbps bandwidth.

of N , while Ring-AR is near-independent of N . On the other hand, bandwidth cost of AG increases on larger clusters.

As cost of AR-Topk can vary based on the AR collective used, we develop a rough heuristic to determine which AR reduction is apt based on current network configuration, model and CR. Looking at the α - β costs of ART-Ring and ART-Tree in Eqn. (4), the former should be used if its cost is lower than the latter, i.e., $t_{ring}^{(c)} < t_{tree}^{(c)}$. From solving, ART-Ring should be used over ART-Tree if the α - β relationship follows Eqn. (5a). Additionally, we earlier saw that at times its better to use AG for communication over AR-Topk. For a DNN with M

bytes of gradients compressed to CR c , ART-Ring or ART-Tree should be used over AG if the latency-bandwidth ratio satisfies Eqn. (5b) or (5c) respectively, otherwise communication is faster with AG collective.

$$\left(\frac{\alpha}{\beta}\right)_{\text{ART-Ring_over_ART-Tree}} < \left(\frac{\log(N) - (N-1)/N}{N-1 - \log(N)}\right) Mc \quad (5a)$$

$$\left(\frac{\alpha}{\beta}\right)_{\text{ART-Ring_over_AG}} < \left(1 - \frac{1}{N} - \frac{\log(N)}{2(N-1)}\right) Mc \quad (5b)$$

$$\left(\frac{\alpha}{\beta}\right)_{\text{ART-Tree_over_AG}} < \left(\frac{N-1}{\log(N)} - \frac{3}{2}\right) Mc \quad (5c)$$

E. Compression as a Multi-Objective Optimization Problem

The parallel and statistical efficiency of gradient compression with distributed training is pareto-related, i.e., one improves at the detriment of the other. This is seen from the step-times and accuracy numbers in both Tables (III) and (IV)) where smaller CRs lead to lower step-times (thus improving parallel efficiency) but at the cost of lower final accuracy (degrading the statistical efficiency). As gradient sensitivity changes throughout training, it is intuitive to change CR as training progresses as well. For e.g., we could use high CRs in early stages when gradients are volatile and lower CRs later when gradients have saturated and stop changing aggressively. In order to accelerate training while still achieving DenseSGD-level convergence, we propose an adaptive compression scheme on top of AR-Topk that models compression as a *multi-objective optimization (MOO)* problem by identifying three key metrics that affect the pareto-relationship between parallel and statistical efficiency:

- 1) **Compression time:** To first compress and then decompress the gradients, and depends on the compression method itself and may vary with CR. E.g., top- k sorting based on max-heap for G gradients has complexity $\mathcal{O}(G + k \log G)$. In this case, time decreases with CR.
- 2) **Communication time:** Between AG, ART-Ring and ART-Tree, the collective with least communication cost depends on α , β , M , N and CR (described in Eqn. (5)).
- 3) **Compression Gain:** Depends on factors like compression technique, CR, stage of training, and indicates statistical efficiency by measuring gradient information lost in compression.

$$c_{\text{optimal}} = \underset{c}{\operatorname{argmin}} \{ \mathcal{F}(t_{\text{comp}}^{(c)}, t_{\text{sync}}^{(c)}, \text{gain}^{-1(c)}) \} \quad (6)$$

The optimal CR in Eqn. (6), c_{optimal} can be determined as a solution to a MOO problem that aims to minimize compression time (t_{comp}) and communication time (t_{sync}), while maximizing compression *gain* (or minimizing $1/\text{gain}$). As both network parameters and gain can vary throughout training, so does c_{optimal} . Here, we bound candidate CRs in a min-max exploration space $[c_{\text{low}}, c_{\text{high}}]$ so as to avoid losing important updates at low CRs or incur considerable

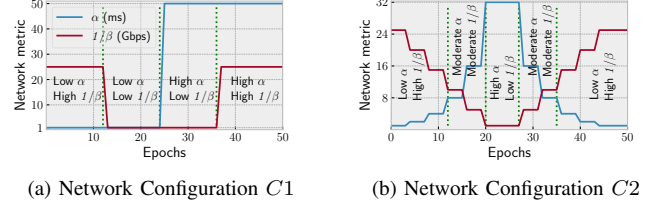


Fig. 6. Latency and Bandwidth is adjusted via ‘traffic control’ over the epochs to emulate different networking scenarios of low/moderate/high α and β .

communication cost at high CRs. The search for ideal CR in the current state is triggered whenever either the average latency or bandwidth changes beyond a certain threshold. A background process measures bandwidth using `iperf` and latency via `traceroute` between the nodes. The same process also uses `tc` to control α and β to emulate different network settings in our evaluation as described later. Based on α , β , M , N and CR c , the collective with least communication overhead is chosen (by Eqn. (5)). Compression gains for candidate CRs are re-evaluated only in sensitive periods if the inter-iteration gain with the current CR changes beyond a specified threshold. To do this, we first preserve the current model state via checkpoint-restore, select a candidate CR and run it for a few iterations, log the average gain and compression time, and then reload the saved state from checkpoint to evaluate the next candidate CR. The checkpoint-restore technique ensures model quality is not degraded from evaluating very low CRs.

1) **MOO Configuration Details:** We assign $c_{\text{low}} = 0.001$, $c_{\text{high}} = 0.1$, and set candidate CRs by scaling c_{low} with a factor of 3. Thus, selected candidate CRs are $[0.1, 0.033, 0.011, 0.004, 0.001]$, where each one is launched for only 10 iterations in order to measure its *gain*, t_{comp} and t_{sync} to formulate the MOO problem. The overhead of this exploration phase is low as its trained for only a few iterations and checkpoint-restore is performed in system memory (thus avoiding expensive disk read/writes). Additionally, this evaluation is triggered only when the inter-iteration gain with current CR, or *gain-threshold* changes by 10% or more. Following this, we initiate the search for c_{optimal} only if the emulated network latency or bandwidth changes. This is done since changes in α and β influence the communication time of CRs (i.e., t_{sync}), while other factors like M and N stay the same. Based on current network and training configuration, the optimal collective with least communication cost is selected. If AR-Topk is chosen over AG, we save the current model state and select ART-Ring or ART-Tree (whichever is faster according to Eqn. (5a)) by setting `NCCL_ALGO` environment variable to `RING / TREE` and redeploying the model from the last saved state. We use the NSGA-II algorithm (Non-dominated Sorting Genetic Algorithm) [42] on top of a MOO framework called `pymoo` [41] to find c_{optimal} using data gathered from candidate CRs.

Fig. (6) shows configuration $C1$ and $C2$ where each epoch has a specific latency and bandwidth, set using `tc`. Specifi-

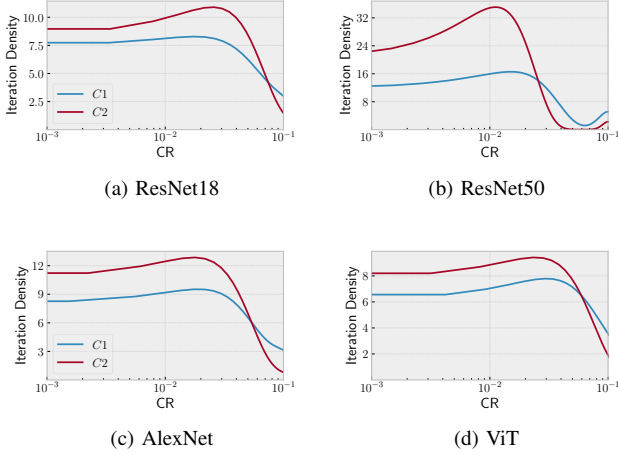


Fig. 7. Iteration density distributions of different CRs used over the course of training for $C1$ and $C2$ configurations when gain-threshold is set to 10%.

cally, we emulate 4 scenarios in $C1$: ($low-\alpha$, $high-1/\beta$) from epoch 1-12, ($low-\alpha$, $low-1/\beta$) from 13-24, ($high-\alpha$, $low-1/\beta$) from 25-36 and ($high-\alpha$, $high-1/\beta$) thereafter. We regard 1ms as low and 50ms as high-latency in this configuration, while inter-node $1/\beta$ of 25 and 1Gbps is considered high and low bandwidth respectively. Configuration $C2$ imitates settings with ($moderate-\alpha$, $moderate-1/\beta$) between epochs 12-19 and 28-35, ($high-\alpha$, $low-1/\beta$) between 20-27 and ($low-\alpha$, $high-1/\beta$) between epochs 0-11 and 36 onwards. This schedule applies to ResNet18, AlexNet and ViT that train for 50 epochs. For ResNet50 that runs for 100 epochs, we scale the number of epochs for each $(\alpha, 1/\beta)$ configuration by $2\times$. This way, $C1$ applies ($low-\alpha$, $high-1/\beta$) state from epoch 1-24 (instead of 1-12), etc. Similarly, $C2$ for ResNet50 exhibits ($high-\alpha$, $low-1/\beta$) between epochs 40-54, and so on.

2) *Performance of MOO-based Compression:* For $C1$ and $C2$, ResNet18 reached 89.88% and 90.1% accuracy, while ResNet50 converged to 98.56% and 98.62% respectively. AlexNet reached 82.4% in $C1$ and 83.6% with $C2$, while ViT converged to 80.75% and 81.2% with $C1$ and $C2$ respectively.

By comparing with convergence on *static* CRs in AR-Topk (Table (IV)) or AG (LW / MS-Topk in Table (III)), we see that dynamically adjusting CR during training while cognizant of a model’s *compression gain* achieves better test accuracy. In the larger AlexNet and ViT models, MOO-based training achieves slightly better accuracy than even DenseSGD as flexible CRs allow for better exploration of the local minimas on individual workers and thus, improves generalization. Fig. (7) plots the kernel density estimates (KDE) of training iterations with various CRs used throughout training for $C1$ and $C2$. The CRs selected via MOO are the knee-point or pareto-front of compression time, communication time and gain, and lie in the min-max range $[c_{low}, c_{high}]$. $C2$ has higher density than $C1$ as search for $c_{optimal}$ is triggered only when network latency or bandwidth changes over the epochs, and $C2$ changes

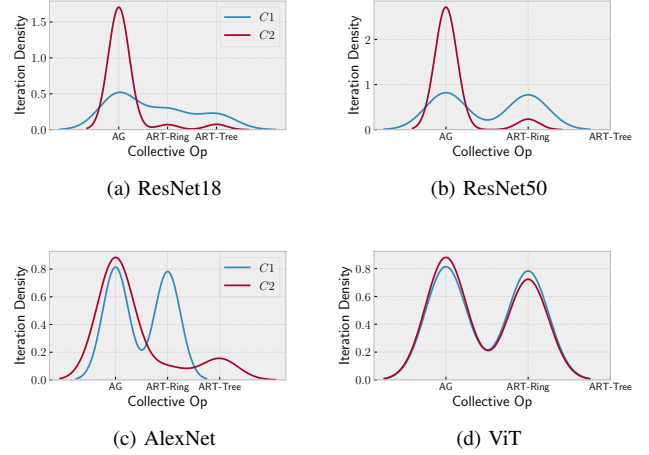


Fig. 8. Density distributions of AG, ART-Ring and ART-Tree collectives used throughout training for configurations $C1$ and $C2$ with gain-threshold 10%.

much more often than $C1$ (from Fig. (6b)). The density peaks lie between CR 0.01 and 0.1 in Fig. (7), implying that most iterations use a CR in that range for most of the training.

Based on α - β cost model, collectives have varying overheads depending on network configuration, CR, model and cluster-size. While optimizing for $c_{optimal}$, we dynamically choose the collective with least communication cost as well. Fig. (8) shows the iteration densities of different collectives used over the course of training as the network configuration changes. For smaller models like ResNet18 and ResNet50, configuration $C2$ uses AG over AR-Topk for most iterations. A major chunk of $C2$ comprises of ($low-\alpha$, $high-1/\beta$) and $moderate-(\alpha, 1/\beta)$ phase. In Table (VI), we see that both ResNet18 and ResNet50 use AG at CRs 0.01 and 0.001 and a 1ms, 10Gbps network; a low-latency, moderate-to-high bandwidth network setting. For ResNet50 and ViT, ART-Tree always had a higher synchronization overhead than ART-Ring. Thus, density of ART-Tree is 0 in Fig. (8b) and (8d). Larger models like AlexNet and ViT communicate with both AG and ART-Ring, as similarly seen from the communication time for different α - β settings in Table (VI).

4. DISCUSSION

Unlike conventional distributed applications, not all work performed in DNN training holds equal contribution towards final model quality due to stochastic nature of SGD and critical periods in the learning phase. This opens up new avenues of adaptive communication for training speedups [2], [11], [22]. Gradient compression trade-offs speedup with convergence quality based on the chosen compression ratio. We model the same as a multi-objective optimization problem to minimize compression-compression cost (parallel efficiency), while maximizing compression gain (statistical efficiency) in distributed learning. Additionally, the optimal collective to use varies with a DNN and latency-bandwidth of a network.

The latter may fluctuate in data-center/edge due to network congestion, contention, scheduling, etc. Our work predicts the ideal compression configuration to use for fast learning over such unpredictable networks. Further, this work can improve convergence on unbalanced and skewed data in federated learning with its low-volume, high-frequency communication strategy (thus, routinely sharing important model features) compared to algorithms like FedAvg [23] that use a high-volume, low-frequency approach instead. An ML practitioner may not necessarily be aware of the network topology when deploying models over a cluster of nodes, or the network configuration may be different as the training infrastructure changes (for e.g., training in a private cluster vs. the cloud). Our system keeps track of the current network latency and available bandwidth to determine the ideal collective op with least communication cost based on the α - β model. Apart from top- k , our approach is compatible with other compressors (like DGC [43], SIDCo [44], etc.), and can be replaced easily.

5. CONCLUSION AND FUTURE WORK

The optimal communication collective depends on factors like network topology, latency, bandwidth, model-size, cluster-size and degree of compression (if any), among others. In this work, we propose a new compression-communication mechanism called AR-Top k to accelerate training when Allreduce is faster than Allgather. Further, we model compression as an optimization problem to balance the parallel and statistical efficiency, and see how CR changes with network state and stage of training. We dynamically choose the CR to attain better convergence quality than with static compression, and achieve maximal speedup with the appropriate communication collective. Currently, we propose two worker selection mechanisms in AR-Top k , one staleness-based and other variance-based. In our evaluation we see that STAR-Top k usually performs better than VAR-Top k on small-to-medium cluster-size. However, the former may degrade model convergence on large cluster-sizes due to stale updates from choosing workers in a round-robin fashion. In the future, we plan to combine the two approaches where AR-Top k automatically switches between the two based on the DNN test performance with each approach.

We also plan to deploy large language models (LLMs) with AR-Top k and evaluate trade-offs between convergence and overall speedup over massive attention-based models. As part of integration with other distributed training paradigms, we expect our adaptive compression technique to work even with model-parallelism, where gradient updates between subsequent layers can be compressed to different CRs over the course of training.

REFERENCES

- [1] Li, Mu et al. "Scaling Distributed Machine Learning with the Parameter Server." BigDataScience '14 (2014).
- [2] Tyagi, Sahil et al. "GraVAC: Adaptive Compression for Communication-Efficient Distributed DL Training." IEEE CLOUD (2023).
- [3] Agarwal, Saurabh et al. "On the Utility of Gradient Compression in Distributed Training Systems." MLSys (2021).
- [4] McCandlish, Sam et al. "An Empirical Model of Large-Batch Training." ArXiv abs/1812.06162 (2018).
- [5] Tyagi, Sahil et al. "Scavenger: A Cloud Service for Optimizing Cost and Performance of ML Training." IEEE/ACM CCGrid (2023).
- [6] Qiao, Aurick et al. "Pollux: Co-adaptive Cluster Scheduling for Goodput-Optimized Deep Learning." OSDI (2020).
- [7] NVLink: <https://www.nvidia.com/en-us/data-center/nvlink/>.
- [8] NVIDIA DGX: <https://www.nvidia.com/en-us/data-center/dgx-systems/>.
- [9] AWS <https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/ec2-instance-network-bandwidth.html>.
- [10] Sun, Peng et al. "Optimizing Network Performance for Distributed DNN Training on GPU Clusters: ImageNet/AlexNet Training in 1.5 Minutes." ArXiv abs/1902.06855 (2019).
- [11] Tyagi, Sahil and Swamy, Martin. "Accelerating Distributed ML Training via Selective Synchronization." IEEE CLUSTER (2023).
- [12] Li, Shen et al. "PyTorch distributed." VLDB Endowment 13 (2020).
- [13] Horovod Tensor Fusion: https://horovod.readthedocs.io/en/stable/tensorfusion_include.html.
- [14] Mai, Luo et al. "KungFu: Making Training in Distributed Machine Learning Adaptive." USENIX OSDI (2020).
- [15] Johnson, Tyler B. et al. "AdaScale SGD: A User-Friendly Algorithm for Distributed Training." ArXiv abs/2007.05105 (2020).
- [16] Frankle, Jonathan et al. "The Early Phase of Neural Network Training." ArXiv abs/2002.10365 (2020).
- [17] Achille, Alessandro et al. "Critical Learning Periods in Deep Neural Networks." ArXiv abs/1711.08856 (2017).
- [18] Agarwal, Saurabh et al. "Accordion: Adaptive Gradient Communication via Critical Learning Regime Identification." abs/2010.16248 (2020).
- [19] Yao, Zhewei et al. "Large batch training of neural networks with adversarial training and second-order information." abs/1810.01021 (2018).
- [20] Alistarh, Dan et al. "The Convergence of Sparsified Gradient Methods." NIPS (2018).
- [21] Shi, Shaohuai et al. "Towards Scalable Distributed Training of Deep Learning on Public Cloud Clusters." ArXiv abs/2010.10458 (2020).
- [22] Tyagi, S., and Swamy, M. (2022). ScaDLES: Scalable Deep Learning over Streaming data at the Edge. 2022 IEEE Big Data, 2113-2122.
- [23] McMahan, H. B. et al. "Communication-Efficient Learning of Deep Networks from Decentralized Data." AISTATS (2016).
- [24] Bernstein, Jeremy et al. "signSGD: compressed optimisation for non-convex problems." ArXiv abs/1802.04434 (2018).
- [25] Wen, Wei et al. "TernGrad: Ternary Gradients to Reduce Communication in Distributed Deep Learning." NIPS (2017).
- [26] Seide, Frank et al. "1-bit stochastic gradient descent and its application to data-parallel distributed training of speech DNNs." Interspeech (2014).
- [27] Vogels, Thijs et al. "PowerSGD: Practical Low-Rank Gradient Compression for Distributed Optimization." NIPS (2019).
- [28] Sarvotham, Shriram et al. "Connection-level analysis and modeling of network traffic." International Memory Workshop (2001).
- [29] He, Kaiming et al. "Deep Residual Learning for Image Recognition." 2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2015): 770-778.
- [30] Krizhevsky, Alex et al. "ImageNet classification with deep convolutional neural networks." Communications of the ACM 60 (2012): 84 - 90.
- [31] Dosovitskiy, Alexey et al. "An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale." ArXiv abs/2010.11929 (2020).
- [32] Krizhevsky, Alex. "Learning Multiple Layers of Features from Tiny Images." (2009).
- [33] GCP: <https://cloud.google.com/compute/docs/networking/configure-vm-with-high-bandwidth-configuration>.
- [34] Egress BW: <https://cloud.google.com/compute/docs/network-bandwidth>.
- [35] Linux traffic control: <https://man7.org/linux/man-pages/man8/tc.8.html>.
- [36] Fei-Fei, Li et al. "Learning Generative Visual Models from Few Training Examples: An Incremental Bayesian Approach Tested on 101 Object Categories." CVPR Workshop (2004).
- [37] G. Griffin, et al. "Caltech 256". doi: 10.22002/D1.20087.
- [38] Thakur, Rajeev et al. "Optimization of Collective Communication Operations in MPICH." IJHPCA (2005).
- [39] Bossard, Lukas et al. "Food-101 Mining Discriminative Components with Random Forests." ECCV (2014).
- [40] iPerf test tool for TCP, UDP and SCTP: <https://iperf.fr/>.
- [41] Blank, Julian and Kalyanmoy Deb. "Pymoo: Multi-Objective Optimization in Python." IEEE Access 8 (2020): 89497-89509.

- [42] Deb, Kalyanmoy, Samir Agrawal, Amrit Pratap and T. Meyarivan. "A fast and elitist multiobjective genetic algorithm: NSGA-II." *IEEE Trans. Evol. Comput.* 6 (2002): 182-197.
- [43] Lin, Yujun et al. "Deep Gradient Compression: Reducing the Communication Bandwidth for Distributed Training." *ICLR* (2018).
- [44] Abdelmoniem, Ahmed Mohamed et al. "An Efficient Statistical-based Gradient Compression Technique for Distributed Training Systems." *MLSys* (2021).