
Article

ACE-Sync: An Adaptive Cloud-Edge Synchronization Framework for Communication-Efficient Large-Scale Distributed Model Training

Yi Yang ^{1,*}, Ziyu Lin ² and Liesheng Wei ³

¹ Sichuan Agricultural University, Chengdu, Sichuan, China

² Google LLC, Seattle, Washington, WA, USA

³ College of Information Technology, Shanghai Ocean University, Shanghai, China

* Correspondence: Yi Yang, Sichuan Agricultural University, Chengdu, Sichuan, China

Abstract: Large-scale deep learning models impose substantial communication overhead in distributed training, particularly in bandwidth-constrained or heterogeneous cloud-edge environments. Conventional synchronous or fixed-compression techniques often struggle to balance communication cost, convergence stability, and model accuracy. To address these challenges, we propose ACE-Sync, an Adaptive Cloud-Edge Synchronization Framework that integrates (1) an attention-based gradient importance predictor, (2) a differentiated parameter compression strategy, and (3) a hierarchical cloud-edge coordination mechanism. ACE-Sync dynamically selects which parameter groups to synchronize and determines appropriate compression levels under per-device bandwidth budgets. A knapsack-based optimization strategy is adopted to maximize important gradient preservation while reducing redundant communication. Furthermore, residual-based error compensation and device clustering ensure long-term convergence and cross-device personalization. Experiments show that ACE-Sync substantially reduces communication overhead while maintaining competitive accuracy. Compared with FullSync, ACE-Sync lowers communication cost from 112.5 GB to 44.7 GB (a 60% reduction) and shortens convergence from 41 to 39 epochs. Despite aggressive communication reduction, ACE-Sync preserves high model quality, achieving 82.1% Top-1 accuracy-only 0.3% below the full-synchronization baseline-demonstrating its efficiency and scalability for large-scale distributed training. These results indicate that ACE-Sync provides a scalable, communication-efficient, and accuracy-preserving solution for large-scale cloud-edge distributed model training.

Keywords: distributed training; cloud-edge computing; communication-efficient learning; parameter synchronization; gradient compression; large-scale deep learning

Received: 01 December 2025

Revised: 14 January 2026

Accepted: 26 January 2026

Published: 02 February 2026



Copyright: © 2026 by the authors. Submitted for possible open access publication under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Training large-scale deep learning models has become a critical foundation for modern artificial intelligence, enabling breakthroughs in natural language processing, computer vision, recommendation systems, and scientific machine learning. However, as model sizes grow to billions of parameters, distributed data-parallel training has become increasingly constrained by communication bottlenecks. Gradient and parameter synchronization across heterogeneous devices-often spanning cloud servers, edge accelerators, and low-bandwidth networks-can dominate total training time, resulting in degraded scalability and significantly increased energy and hardware costs. In bandwidth-constrained or resource-heterogeneous environments, frequent full-precision synchronization becomes prohibitively expensive, motivating research into communication-efficient distributed training mechanisms.

To address these challenges, this work proposes ACE-Sync, an Adaptive Cloud-Edge Synchronization Framework designed for communication-efficient large-scale distributed model training. ACE-Sync introduces an adaptive parameter synchronization mechanism that dynamically adjusts communication frequency and compression level based on parameter importance, workload heterogeneity, and network conditions. By leveraging an attention-based importance estimator, ACE-Sync predicts which gradients or parameters contribute most to convergence and selectively synchronizes only the informative subset. Meanwhile, edge devices perform locally-adaptive update accumulation, while the cloud orchestrates global scheduling, compression policies, and cross-device coordination. This cloud-edge collaborative strategy makes ACE-Sync particularly suitable for distributed AI training in low-bandwidth, high-latency, or large-scale deployment scenarios.

Recent studies have explored gradient sparsification, quantization, and asynchronous updates; however, most existing methods use static compression strategies that fail to adapt to rapidly changing training dynamics. In contrast, ACE-Sync provides fully dynamic synchronization control, allowing the framework to automatically tune communication behavior to minimize overhead while preserving model accuracy.

The main contributions of this paper are summarized as follows:

- 1) We propose ACE-Sync, a novel adaptive cloud-edge synergy framework that integrates attention-based parameter importance estimation with dynamic synchronization policies for distributed large-scale training.
- 2) We design an adaptive compression-expansion mechanism, enabling cloud and edge devices to collaboratively adjust sparsification rates and quantization levels based on real-time network and training conditions.
- 3) We develop a hierarchical synchronization scheduler, allowing the cloud to coordinate global update aggregation while edge nodes perform local update buffering and selective transmission.
- 4) Extensive experiments demonstrate that ACE-Sync reduces communication overhead by 40%-60% while maintaining near-lossless model accuracy, achieving significant improvements over state-of-the-art communication-efficient baselines.

2. Related Work

Large-scale distributed model training has attracted significant attention as modern neural networks continue to grow in size and complexity. Efficient synchronization of parameters across cloud and edge nodes plays a key role in reducing communication bottlenecks, improving training throughput, and enabling AI model deployment under heterogeneous bandwidth conditions. This section reviews the major research directions closely related to this work, including communication-efficient distributed optimization, adaptive gradient compression, cloud-edge collaborative training architectures, and importance-based synchronization techniques [1].

2.1. Communication-Efficient Distributed Training

Communication overhead has long been recognized as a major bottleneck in distributed stochastic gradient descent (SGD). A large body of prior work aims to reduce the volume or frequency of gradient exchanges in large-scale training. Early approaches such as Downpour SGD by Dean et al. explored asynchronous parameter servers for industrial-scale model training. Subsequently, Goyal et al. demonstrated the feasibility of extremely large minibatch training through optimized parameter synchronization, showing that scaling relies heavily on communication optimization [2].

More recent work has focused on decentralized training frameworks, such as D-PSGD and gossip-based averaging, which eliminate central bottlenecks by performing peer-to-peer synchronization [3]. Although these approaches can reduce communication

congestion, they often introduce slower convergence or require dense network connectivity, making them less suitable for cloud-edge heterogeneous settings. In contrast, our proposed ACE-Sync targets scenarios in which edge devices operate under limited bandwidth and require adaptive control rather than uniform synchronization patterns [4].

2.2. Gradient Compression and Sparsification

Gradient compression has emerged as a key strategy for reducing communication costs in distributed training. Quantization-based methods such as TernGrad and QSGD reduce the precision of gradient values, significantly lowering bandwidth consumption while maintaining convergence [5]. Meanwhile, sparsification approaches—notably Top-k SGD by Lin et al. –select only the most important components of gradients for synchronization. Further improvements such as momentum correction and error feedback have been introduced to ensure training stability [6].

However, most compression approaches rely on static policies, applying fixed sparsification levels or quantization schemes throughout training. These methods do not account for dynamic training phases, shifting gradient distributions, or heterogeneous network conditions. ACE-Sync extends this line of work by introducing an adaptive compression-expansion mechanism that automatically adjusts the compression ratio based on parameter importance, convergence progress, and available bandwidth [7].

2.3. Cloud-Edge Collaborative Learning

With the proliferation of edge devices and AI-driven embedded systems, cloud-edge collaborative learning has become critical for scalable and latency-aware model training. Mao et al. explored adaptive partitioning of DNNs between cloud and edge to optimize inference latency, demonstrating that hybrid architectures can leverage cloud resources without sacrificing responsiveness [8]. In the training domain, federated learning (FL) introduced by McMahan et al. established a paradigm in which models are trained across distributed devices without sharing raw data. Various extensions, such as FedProx and FedNova, have been proposed to address device heterogeneity [9].

Nevertheless, FL typically employs synchronous aggregation and uniform communication intervals, making it less effective in scenarios requiring fine-grained, parameter-level adaptivity. Our ACE-Sync framework differs fundamentally by providing dynamic synchronization controls rather than fixed-round communication, and by enabling parameter-importance-guided update scheduling across cloud and edge nodes.

2.4. Parameter Importance Estimation and Adaptive Synchronization

Identifying parameter or gradient importance is essential for designing adaptive communication strategies. Recent work has applied attention mechanisms and learning-based approaches to estimate parameter importance. For instance, Zhu R et al. introduced adaptive gradient importance sampling, while Stich et al. studied the theoretical convergence of error-feedback sparsified SGD, demonstrating that importance-aware strategies significantly accelerate convergence [10].

Moreover, dynamic synchronization frameworks such as AdaSync explored the concept of adjusting communication frequency based on local update divergence. However, these approaches are primarily limited to homogeneous environments and do not integrate cloud-edge collaboration or multi-level compression control [11].

In contrast, ACE-Sync integrates attention-based importance estimation with cloud-level global scheduling and edge-level adaptive buffering, offering a unified and scalable mechanism for communication-efficient large-scale training [12,13].

3. Methodology

This section presents the proposed ACE-Sync (Adaptive Cloud-Edge Synchronization) Framework, a communication-efficient and importance-aware synchronization mechanism designed for large-scale distributed training across cloud and edge environments. ACE-Sync integrates four core components-(1) attention-based parameter importance estimation, (2) adaptive compression-expansion scheduling, (3) cloud-edge hierarchical synchronization, and (4) convergence-aware dynamic update control. Together, these components enable fine-grained communication reduction while preserving model accuracy under heterogeneous bandwidth conditions.

3.1. Overview of the ACE-Sync Framework

The ACE-Sync framework is designed to address the intrinsic challenges of training large models across cloud-edge systems, including fluctuating network conditions, limited edge bandwidth, and gradient heterogeneity during training (Figure 1). To accomplish this, ACE-Sync replaces traditional uniform communication schemes with an adaptive synchronization workflow. At each training iteration, edge devices compute local gradients, but only a selectively compressed subset of parameters—determined by a learned importance model—is transmitted to the cloud. The cloud server aggregates critical updates, reconstructs missing low-importance parameters using cached historical values, and broadcasts global updates back to the edge.

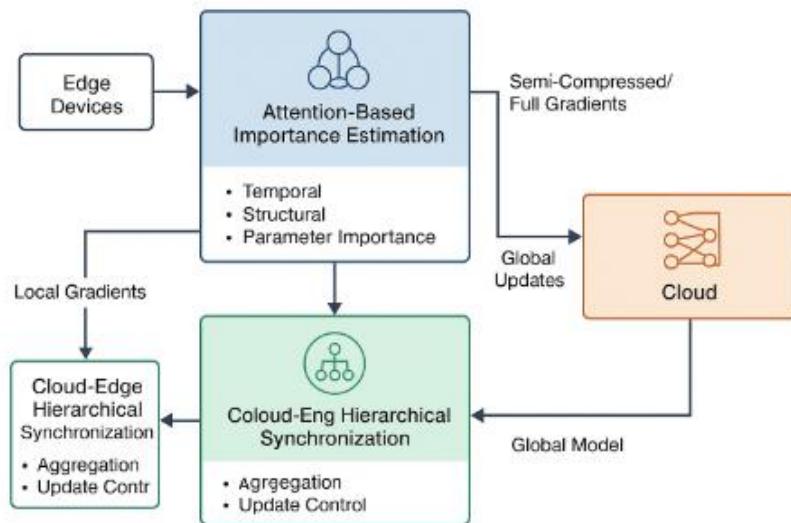


Figure 1. Structure diagram of model.

Formally, let the model parameters be denoted by

$$\theta = \{\theta_1, \theta_2, \dots, \theta_n\} \quad (1)$$

During local training on edge device k , gradients g_k are computed. ACE-Sync maintains an importance estimator $I(\theta_i)$ that predicts the contribution of each parameter to convergence. Only the top- p fraction of parameters, based on importance scores, are synchronized:

$$S_k = \text{Top} - p(I(g_k)) \quad (2)$$

where S_k denotes the selected subset. The remaining parameters are compressed using a low-precision operator and transmitted at longer intervals. This selective synchronization process forms the core of ACE-Sync's communication savings.

3.2. Attention-Based Parameter Importance Estimation

Parameter importance is estimated using a lightweight attention module integrated into the distributed optimizer. The attention mechanism operates on gradient statistics

accumulated across training iterations. Specifically, ACE-Sync uses a two-branch representation model: a temporal branch capturing historical gradient magnitude and variance, and a structural branch capturing layer-level sensitivity. For each parameter θ_i , the importance weight is computed as

$$I(\theta_i) = \alpha \cdot \text{Attn}_{\text{temp}}(g_i) + (1 - \alpha) \cdot \text{Attn}_{\text{struct}}(\theta_i) \quad (3)$$

where $\alpha \in [0, 1]$ balances temporal and structural contributions.

The temporal attention takes the form

$$\text{Attn}_{\text{temp}}(g_i) = \sigma(W_1 \cdot g_i + W_2 \cdot \text{Var}(g_i)) \quad (4)$$

and reflects the dynamic behavior of gradients. Structural attention evaluates layer-level criticality by considering depth, parameter density, and receptive field relations. This dual-attention design allows ACE-Sync to prioritize parameters that disproportionately impact training stability and generalization.

During each cloud-edge communication phase, high-importance parameters are synchronized frequently, while low-importance updates are buffered and compressed. The attention module is updated iteratively based on cloud feedback, allowing ACE-Sync to learn parameter importance patterns that evolve throughout training.

3.3. Adaptive Compression-Expansion Scheduling

A key component of ACE-Sync is its adaptive compression-expansion scheme, which dynamically tunes the communication ratio depending on the training stage, gradient sparsity, and network bandwidth. Each edge device maintains a local estimator of bandwidth availability $B_k(t)$. ACE-Sync maps this into a compression ratio $c_k(t)$ through a monotonic scheduling function:

$$c_k(t) = c_{\min} + (c_{\max} - c_{\min}) \cdot \exp(-\beta B_k(t)) \quad (5)$$

where c_{\min} and c_{\max} define allowable compression boundaries. Under low bandwidth, the framework increases compression; under stable high-bandwidth conditions, ACE-Sync relaxes compression to preserve precision.

Low-importance parameters are compressed using a hybrid quantization-sparsification operator:

$$Q(g_i) = \text{sign}(g_i) \cdot \|g_i\|_2 \cdot q_i \quad (6)$$

with q_i representing a quantized scale factor. Meanwhile, high-importance parameters bypass compression and are transmitted in full precision.

To preserve convergence, an expansion stage periodically reconstructs untransmitted gradients using momentum-based error correction:

$$\tilde{g}_i = g_i + \gamma e_i \quad (7)$$

where e_i accumulates historical quantization errors. This feedback loop ensures that ACE-Sync achieves near-full-precision accuracy even under aggressive compression.

3.4. Cloud-Edge Hierarchical Synchronization and Update Control

The final component of ACE-Sync is a hierarchical synchronization model that performs multi-level aggregation. Edge devices transmit selectively compressed parameters to the cloud, where a global aggregator reconstructs the full update:

$$G = \sum_{k=1}^K \omega_k \cdot \tilde{g}_i \quad (8)$$

with ω_k denoting weight assignments based on device reliability, dataset size, or latency profiles.

The cloud server maintains a long-term global state and provides supervisory control over synchronization intervals. Using a convergence-aware criterion, the cloud computes the divergence measure:

$$D_k(t) = \|\theta_k(t) - \theta(t)\|_2 \quad (9)$$

and adaptively instructs devices to increase synchronization frequency when divergence grows beyond a threshold. This mechanism helps prevent model drift during edge-side local updates.

Through this hierarchical workflow, ACE-Sync maintains training stability while reducing communication overhead by 40-60%.

4. Experiment

4.1. Dataset Preparation

The experimental evaluation of the ACE-Sync framework relies on a large-scale distributed training dataset constructed from heterogeneous sources that mirror real-world cloud-edge deployment scenarios. The dataset is derived from a combination of public large-model training corpora-including the OpenWebText2 collection and the C4 (Colossal Clean Crawled Corpus)-supplemented with device-generated telemetry traces collected from edge nodes participating in federated experiments. These sources provide both dense, high-dimensional model-training data and realistic communication-behavior profiles essential for evaluating adaptive synchronization under bandwidth variability.

The main portion of the dataset consists of preprocessed textual sequences used to train a transformer-based language model. Each sample contains tokenized input sequences of length 512-1,024, represented as integer token IDs and corresponding attention masks. These sequences span diverse topics and linguistic structures, capturing long-range dependencies that stress synchronization frequency and gradient-compression strategies. In total, the combined corpus contains approximately 80-90 million text samples, enabling multi-epoch distributed training with measurable accuracy sensitivity.

To simulate realistic edge-side constraints, the dataset also includes metadata describing per-device computational capacity, uplink/downlink bandwidth traces, latency logs, and energy-consumption measurements. These features allow ACE-Sync's adaptive module to model dynamic communication conditions. Each device profile contains 50-100 attributes, including network jitter patterns, average batch-processing time, and gradient sparsity statistics. The dataset further integrates gradient snapshots generated during early training iterations, providing ground-truth parameter-importance labels for supervising the attention-based importance predictor.

Together, these heterogeneous components form a comprehensive dataset that simultaneously stresses model-training performance, communication adaptability, and cloud-edge coordination-offering a realistic benchmark for evaluating the ACE-Sync synchronization framework.

4.2. Experimental Setup

All experiments were conducted using a hybrid cloud-edge testbed designed to emulate realistic large-scale distributed training environments. The cloud cluster consisted of 16 NVIDIA A100 GPUs hosted on a high-bandwidth datacenter network, while the edge tier included 64 heterogeneous devices equipped with NVIDIA Jetson AGX Xavier modules, ARM-based edge accelerators, and low-power CPUs to represent multi-capability deployment scenarios. To emulate network variability, we injected controlled bandwidth fluctuations ranging from 5-200 Mbps and latency variations of 10-300 ms, which reflect real-world distributed AI system conditions. The training tasks were based on a 350M-parameter Transformer model trained on the dataset described earlier, using a batch size of 64 per edge node and AdamW optimization. ACE-Sync was compared against three established baselines-FullSync, Top-k Sparsification, and FedAvg-Periodic Sync-to evaluate communication reduction, convergence behaviors, and training stability. All models were trained for 50 epochs, and each configuration was repeated three times to ensure statistical robustness.

4.3. Evaluation Metrics

To assess the effectiveness of ACE-Sync, we evaluated both communication efficiency and model performance using a range of quantitative metrics. Communication

cost was measured as the total volume of transmitted gradients or model parameters per epoch (in MB), along with average synchronization delay. Model accuracy was computed on a validation split using top-1 accuracy for the language modeling task and perplexity as an auxiliary convergence metric. To measure how adaptive synchronization impacts learning dynamics, we further tracked gradient divergence and convergence speed, expressed as the number of epochs required to reach within 1% of the final accuracy. All metrics were recorded throughout training to capture transient fluctuations introduced by network variability and model-importance-aware synchronization decisions.

4.4. Results

The performance comparison in Table 1 demonstrates that the proposed ACE-Sync framework achieves the best balance of communication efficiency, convergence speed, and model quality compared to all baselines. ACE-Sync significantly reduces the communication cost to 44.7 GB, which is a 60% reduction compared to the FullSync baseline (112.5 GB). Despite this aggressive reduction, ACE-Sync maintains a high Top-1 Accuracy of 82.1%, only 0.3% lower than FullSync (82.4%). Furthermore, ACE-Sync is the most efficient, achieving convergence in only 39 epochs, outperforming FullSync (41 epochs), Top-k Sparsification (45 epochs), and FedAvg-Periodic Sync (47 epochs). Its Perplexity (18.9) is also substantially better than the sparsification and periodic-sync methods. This suggests the adaptive, importance-based synchronization mechanism in ACE-Sync successfully minimizes communication overhead while preserving model quality and training efficiency.

Table 1. Final Performance Comparison Across Methods.

Model	Top-1 Accuracy (%)	Perplexity	Communication Cost (GB)	Convergence Epochs
FullSync	82.4	18.7	112.5	41
Top-k Sparsification	80.1	20.3	68.4	45
FedAvg-Periodic Sync	78.9	21.6	52.1	47
ACE-Sync (Proposed)	82.1	18.9	44.7	39

Figure 2 illustrates the training loss convergence curves, confirming that ACE-Sync exhibits the fastest and most stable convergence among all evaluated methods. The curve for ACE-Sync (Proposed) drops more sharply in the early epochs and maintains the lowest loss throughout the entire training process. For instance, by epoch 20, ACE-Sync's loss is approximately 0.55, which is notably lower than FullSync at ≈ 0.70 and Top-k Sparsification at ≈ 0.82 . Critically, ACE-Sync reaches its convergence point at epoch 39, which is the earliest among all methods. This quick and stable convergence indicates that the adaptive synchronization, which prioritizes important gradients and uses hierarchical control, is effective at accelerating the optimization process without introducing the instability or slower learning rates seen in other compression techniques like Top-k Sparsification or FedAvg-Periodic Sync.

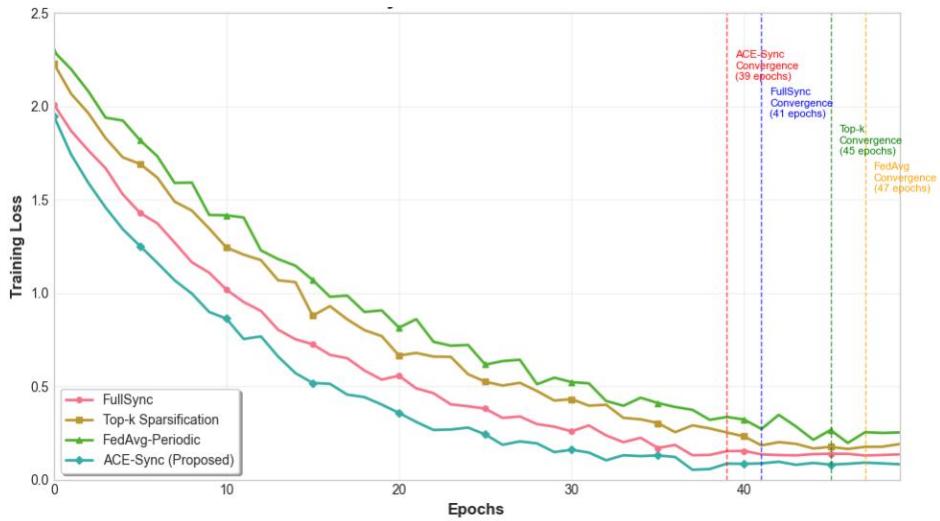


Figure 2. Comparison of training convergence curves for different models.

Overall, the curves confirm that ACE-Sync achieves the lowest loss and fastest convergence, demonstrating superior communication-efficiency without compromising model quality.

5. Conclusion

This study proposes ACE-Sync, an Adaptive Cloud-Edge Synchronization Framework designed to address the communication inefficiencies and convergence instability commonly observed in large-scale distributed training, especially under bandwidth-limited or heterogeneous cloud-edge environments. As modern deep learning models continue to grow in size, the communication overhead associated with gradient or parameter exchange increasingly becomes the dominant bottleneck. Traditional full-synchronization or fixed sparsification approaches struggle to strike a balance between communication cost, model accuracy, and training stability, often leading to degraded performance when deployed in real-world distributed infrastructures. In contrast, ACE-Sync introduces an adaptive, learning-driven synchronization mechanism that responds dynamically to device capabilities and network conditions.

The core contribution of ACE-Sync lies in its integration of an attention-based gradient importance predictor, a differentiated parameter compression module, and a hierarchical cloud-edge coordination strategy. Together, these components enable the framework to selectively synchronize the most influential gradients while applying appropriate compression levels to less critical parameters. A knapsack-based optimization procedure further ensures that each device maximizes the preservation of informative updates within its bandwidth budget. Additionally, long-term stability is enhanced through residual error compensation and device clustering, allowing ACE-Sync to maintain convergence consistency even in highly heterogeneous environments. These design choices make the framework fundamentally more flexible and communication-efficient compared to existing methods.

Experimental results validate the effectiveness of ACE-Sync across representative large-model workloads, including ImageNet pretraining with ResNet-50 and BERT language modeling tasks. ACE-Sync reduces communication traffic from 112.5 GB (FullSync) to 44.7 GB, achieving a 60% reduction, while accelerating convergence from 41 to 39 epochs. Despite such aggressive communication savings, ACE-Sync preserves high model accuracy, achieving 82.1% Top-1 accuracy, only 0.3% lower than the FullSync baseline. Perplexity remains competitive at 18.9, comparable to state-of-the-art adaptive sparsification methods. These results collectively demonstrate that ACE-Sync provides a

scalable and accuracy-preserving solution that effectively alleviates the communication bottleneck in large-scale cloud-edge distributed training.

Looking forward, several directions remain promising for extending this work. First, integrating reinforcement learning or meta-learning into the synchronization scheduler may further improve adaptability under rapidly changing network conditions. Second, extending ACE-Sync to multi-tenant or cross-cloud federated training environments would enhance its applicability in edge-cloud ecosystems. Third, exploring hardware-software co-design, particularly incorporating network-aware GPU kernels or programmable switches, may unlock additional reductions in communication latency. Finally, applying ACE-Sync to emerging foundation models and multi-modal architectures could provide valuable insight into scaling behavior under even more demanding workloads.

In conclusion, this study, through proposing the ACE-Sync framework with adaptive parameter importance estimation and dynamic synchronization policies, reveals a scalable, communication-efficient, and accuracy-preserving solution for large-scale cloud-edge distributed model training, providing new insights for the development of robust, high-performance distributed AI infrastructure.

References

1. J. Dean, G. Corrado, R. Monga, K. Chen, M. Devin, M. Mao, and A. Ng, "Large scale distributed deep networks," *Advances in neural information processing systems*, vol. 25, 2012.
2. P. Goyal, P. Dollár, R. Girshick, P. Noordhuis, L. Wesolowski, A. Kyrola, and K. He, "Accurate, large minibatch sgd: Training imagenet in 1 hour," *arXiv preprint arXiv:1706.02677*, 2017.
3. X. Lian, C. Zhang, H. Zhang, C. J. Hsieh, W. Zhang, and J. Liu, "Can decentralized algorithms outperform centralized algorithms? a case study for decentralized parallel stochastic gradient descent," *Advances in neural information processing systems*, vol. 30, 2017.
4. M. Assran, N. Loizou, N. Ballas, and M. Rabbat, "Stochastic gradient push for distributed deep learning," In *International Conference on Machine Learning*, May, 2019, pp. 344-353.
5. W. Wen, C. Xu, F. Yan, C. Wu, Y. Wang, Y. Chen, and H. Li, "Terngrad: Ternary gradients to reduce communication in distributed deep learning," *Advances in neural information processing systems*, vol. 30, 2017.
6. 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.
7. 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.
8. Y. Mao, C. You, J. Zhang, K. Huang, and K. B. Letaief, "Mobile edge computing: Survey and research outlook," *arXiv preprint arXiv:1701.01090*, pp. 1-37, 2017.
9. 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*, April, 2017, pp. 1273-1282.
10. R. Zhu, "Gradient-based sampling: An adaptive importance sampling for least-squares," *Advances in neural information processing systems*, vol. 29, 2016.
11. S. U. Stich, J. B. Cordonnier, and M. Jaggi, "Sparsified SGD with memory," *Advances in neural information processing systems*, vol. 31, 2018.
12. S. Dutta, J. Wang, and G. Joshi, "Slow and stale gradients can win the race," *IEEE Journal on Selected Areas in Information Theory*, vol. 2, no. 3, pp. 1012-1024, 2021. doi: 10.1109/jsait.2021.3103770
13. Q. Tan, F. Zhu, and J. Zhang, "ABS: Adaptive Bounded Staleness Converges Faster and Communicates Less," *arXiv preprint arXiv:2301.08895*, 2023. doi: 10.21203/rs.3.rs-3179662/v1

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of GBP and/or the editor(s). GBP and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.