

Improving training time and GPU utilization in geo-distributed language model training

PALAK, Microsoft Research, India

TELLA RAJASHEKHAR REDDY, Microsoft Research, India

BHASKAR KATARIA, Cornell University, USA

ROHAN GANDHI, Microsoft Research, India

KARAN TANDON, Microsoft Research, India

DEBOPAM BHATTACHERJEE, Microsoft Research, India

VENKATA N. PADMANABHAN, Microsoft Research, India

Abstract

The widespread adoption of language models (LMs) has caused a huge surge in demand for GPUs. Training large LMs requires tens of thousands of GPUs and housing them in the same datacenter (DC) is a challenge due to many constraints including availability of peak power. We focus on training such models across multiple DCs connected via the Wide-Area-Network (WAN). We built ATLAS that speeds up the training time using novel workload-aware temporal bandwidth sharing and other design choices. While ATLAS improves the training time, it does not completely eliminate the bubbles (idle GPU cycles). We built BUBBLETEA that runs prefill-as-a-service (part of LM inference) during the bubbles thus improving the GPU utilization without any impact on training. Compared to state-of-the-art designs, ATLAS and BUBBLETEA together achieve up to 17× faster training, and up to 94% GPU utilization. The code will be open-sourced.

1 Introduction

The advent of Generative AI and Language Models (LMs) such as GPT[8], Llama[15], Claude[4] and Mistral[18] has ushered in an era of widespread AI adoption across many industries from healthcare to finance to entertainment. As more sectors realize this newly unleashed potential, AI training is considered one of the larger cloud workloads of today[55].

Training such LMs requires a significant number of GPUs. These LMs have seen a substantial increase in the number of parameters to improve the accuracy[56] and support larger number of tokens. For instance, the largest variant of the Llama model, Llama 4, now boasts 400 billion parameters. Consequently, it is reported that GPT models with trillions of parameters[9] require 10s of thousands of GPUs[5], and Llama models with billions of parameters require thousands of GPUs for months to train[6].

It is desired that all GPUs catering to a training workload are housed in the same data center (DC) and leverage the fast intra-DC interconnects to finish the training quickly[38, 55]. However, it is a non-trivial engineering challenge to deploy a large number of GPUs in the same DC due to space, power, power density, and cooling requirements[11, 12, 74, 76, 81]. Such constraints, coupled with the increasing AI inferencing demand consuming up to 90% of AI compute [69], are already getting reflected in user reports on not being able to provision (book) even a few GPUs[83] per site. This hints at an emerging systems challenge – getting access to thousands of GPUs at a single DC site for training LMs is hard. We here explore the intuitive yet non-trivial solution to this

Authors' Contact Information: Palak, Microsoft Research, India; Tella Rajashekhar Reddy, Microsoft Research, India; Bhaskar Kataria, Cornell University, USA; Rohan Gandhi, Microsoft Research, India; Karan Tandon, Microsoft Research, India; Debopam Bhattacharjee, Microsoft Research, India; Venkata N. Padmanabhan, Microsoft Research, India.

problem – geo-distributing the training jobs across GPUs that span multiple DCs connected via the Wide-Area-Network (WAN). Such an approach has received attention recently[7, 16, 20].

In this paper, we focus on the problem of training individual LMs on GPUs housed across multiple DCs connected via WAN. Training jobs use multiple forms of parallelism such as Data Parallelism (DP), Pipeline Parallelism (PP) and Tensor Parallelism (TP) that all have substantial communication on the critical path (§2). Communication, needed during activation updates, gradient updates, and synchronization, incurs a significantly higher latency over inter-DC WAN than in intra-DC networks. Our experiments for doing PP across DCs show how state-of-the-art schedulers such as Varuna[29] face the following limitations: (a) They end up creating bubbles (idle time; undesired) not only between the forward and backward passes in one training iteration, but also *between microbatches in the same minibatch* (details in §3.2). (b) Similarly, in PP, the DCs running later pipeline stages are idle (bubble) before activations are transferred from the preceding stages. These bubbles are amplified due to slow WAN communication. Consequently, such systems achieve only < 5% GPU utilization and each training iteration is severely elongated. (c) Compared to a hypothetical baseline of housing all such GPUs in a single DC, such geo-distributed training can result in an order of magnitude slower training time. More details are in §3.

We present ATLAS and BUBBLETEA– ATLAS addresses the key limitations in geo-distributed training to substantially improve the training time. However, it does not eliminate all bubbles. We built BUBBLETEA that offers prefill-as-a-service, scheduling the prefill phase of eligible inference requests to further reduce the bubbles. ATLAS and BUBBLETEA are independent systems that complement each other and improve the training time by up to 17× compared to state-of-art and achieve GPU utilization of up to 94% in cross-DC LM training.

1.1 ATLAS design choices

ATLAS addresses above limitations using the following design choices.

Simple idea with large benefits: Prior works[47, 61, 93] paint a bleak picture of WAN bandwidth. E.g., [47] could use only 100 Mbps bandwidth between nodes in Virginia and Sao Paulo on Amazon EC2. Similarly, [61] found an average of 193 Mbps between nodes in Utah and Wisconsin. Note that prior work has extensively focused on improving network designs for GPU to GPU communication on NVLink or InfiniBand[1, 23], but not using Ethernet on WAN resulting in low throughput mentioned above (§3). We found that PyTorch[24] simply used one TCP connection between two nodes on WAN. Our simple idea is to spawn multiple TCP connections between two nodes in the WAN to scale the bandwidth to 5 Gbps compared to 100s of Mbps on a single connection. This simple, intuitive idea of using multiple TCP connections substantially improves the training time.

Coordinating pipelines: We observe that there is no co-ordination among the DP pipelines leaving bubbles in the later stages of PP. To reduce the bubbles in later stages of PP, ATLAS intelligently shares the WAN bandwidth across multiple pipelines and start the processing of microbatches at the later (in the pipeline) DCs sooner. In doing so: (a) ATLAS improves the WAN utilization and improves the training time, and (b) interestingly, such scheduling eliminates the bubbles that arise during the microbatches.

Modeling right combination: While the above ideas substantially improve the training time, we find that it is not always beneficial to use all GPUs. For example, imagine availability of 1000 GPUs in one DC and 10 GPUs in a different DC. It is better to not use the 10 GPUs for the training, as the additional network overhead (slower WAN) results in more bubbles and elongates the training time. Conversely, in a more balanced distribution (1000 GPUs×2), the training time could improve by distributing the job across both DCs. ATLAS uses a heuristic that computes the optimal split of GPUs across DCs for LM training. Additionally, it enables performance modeling, i.e., ‘what-if’

analysis to calculate the cost and training time for a given set of DC locations and the number of GPUs in each DC in simulation so that engineers can choose the right combinations of DCs.

1.2 BUBBLETEA

While ATLAS improves the training time and eliminates bubbles between the microbatches, it does not eliminate bubbles around forward and backward passes (§4.3). Eliminating these bubbles requires either increasing the compute (by processing more layers at a GPU node) or reducing the communication overhead. However, we find that both such approaches do not work well in cross-DC training (details in §5).

To mitigate the wastage of compute during bubbles, we schedule independent workloads during the bubbles. Specifically, we find that *prefill phase from the LM inference workloads* is well suited to consume such bubbles. Inference requests compose of distinct prefill (digesting the prompt) and (auto-regressive) decode phases and recent works have proposed decoupling between such phases during execution[70, 96]. The completion time of the prefill phase is highly deterministic due to known input prompt, which helps BUBBLETEA schedule them effectively during bubbles in training.

BUBBLETEA builds prefill-as-a-service: (a) BUBBLETEA controller receives prefill requests and places them to consume bubbles in the training pipeline. (b) BUBBLETEA controller determines the GPU (in the same DC) where such prefill phase runs and then transfers the KV cache to a different GPU in the same DC (outside BUBBLETEA) for decode phase (following Splitwise architecture[70]). (c) BUBBLETEA reduces the number of GPUs provisioned for inference while improving Time-To-First-Token (TTFT) for large prefills and minor penalty for small prefills.

1.3 Putting it together

We implemented ATLAS and BUBBLETEA that work hand-in-hand to simultaneously reduce the training time and improve GPU utilization in cross-DC LM training. Through testbed experiments and large-scale simulations, we show that: (a) ATLAS can reduce the training time up to 17× compared to state-of-art mechanisms including Varuna[29]. (b) While significant benefits can be attributed to the intuitive idea of using multiple TCP connections, the other idea of intelligently sharing the WAN bandwidth improves the training time by up to 1.82× compared to state-of-art schedulers. (c) ATLAS architecture is effective in using PP across DCs. Such an architecture can scale the throughput with more DCs. (d) Algorithm in ATLAS to calculate the optimal number of GPUs per DC is effective and fast. (e) ATLAS and BUBBLETEA co-ordinate effectively, and BUBBLETEA is able to achieve the GPU utilization up to 94%.

This work does not raise any ethical issues.

2 Background

Most of the LMs such as GPT[8], LLama[15], Claude[4], Mistral[18] use transformer based models[87] with many layers. Each transformer block comprises various components, including attention mechanisms and feedforward neural networks (FFNs), each containing its own neural network. A training job learns the parameters of neural networks.

In each training iteration, a model takes a few samples of a data called *minibatch* and performs a *forward pass* that computes the loss values for the data samples followed by a *backward pass* that computes the gradients. Finally, the model parameters are learnt by applying the negative of the gradients. Recent works save GPU memory by re-running the forward pass just before the backward pass[29], called *recomputation*. Due to the massive sizes of the models (billions to trillions of parameters), the training job is distributed across multiple GPUs. The following parallelism modes distribute the job:

Data parallelism (DP): In this mode, the model (or a subset of its layers) is replicated across GPUs and different minibatches are fed to such replicas. At the end of the iteration (one forward and backward pass), gradients are averaged through *all-reduce*[68] where the communication between replicas is on the critical path.

Pipeline parallelism (PP): While DP helps in speeding up the training time, PP helps in fitting larger models across GPUs. In PP, different layers of the model are assigned to different GPUs. One GPU sends the activations (in forward pass) to the next GPU over network. Similarly, in the backward pass, the gradients are sent between GPUs over network. Furthermore, the minibatch is further split across different *microbatches* that are pipelined in execution. Recent works try to overlap compute of one microbatch with communication of a different microbatch to improve training time[29, 50, 64, 65]. The critical path is shaped by the slowest (due to slower communication, computation, or both) pipeline stage.

Tensor parallelism (TP): Like PP, TP helps in fitting models across GPUs. Unlike PP, TP splits individual layers across different GPUs and uses all-reduce for communication. TP requires significantly higher network bandwidth than DP and PP due to frequent synchronization across shards[39].

Expert parallelism (EP): In the mixture-of-expert (MoE) models [37, 58], the data samples are routed to one or a few *experts*. In this work we focus on the *dense* models (e.g., Llama 3.1) that do not use experts. The principles from this paper apply to MoE models (discussed in Appendix §A).

Sequence parallelism (SP): In PP, we scale sequence dimension and is usually applied at the end of the pre-training to increase model context window[51, 62].

Multiple forms of parallelism: Training large models usually involves all forms of parallelism described above. However, the key is determining how to split the job across such different forms of parallelism.

2.1 Geo-distributed training

Training large language models (with billions of parameters) necessitates a substantial number of GPUs to minimize training latency. Many previous works assume that all GPUs used for training are located within the same datacenter (DC)[29, 31, 50, 54, 65, 72, 91, 95]. These works benefit from the favorable network conditions within a single DC, such as low latency and high bandwidth. However, consolidating all GPUs in a single DC is becoming increasingly challenging as (1) many GPUs (up to 90% of all compute cycles [69]) are being allocated to inference workloads[19] and (2) DCs hit the power draw and cooling thresholds due to high GPU power density[82].




Fig. 1. Experimental setup consisting of 6 GPUs in 3 DCs connected by WAN. WAN offers lower bandwidth compared to intra-DC bandwidth (thicker lines).

A potential solution is to distribute the training across GPUs located in different DCs, which are connected via WANs[45, 53]. Unfortunately, WANs typically do not offer the same network capacities as intra-DC networks. Most WAN links (connecting DCs) range from a few hundred Gbps to a few Tbps[45, 46, 53], and DCs in different countries may experience latencies from tens to hundreds of milliseconds (versus sub-millisecond intra-DC latency).

Private WANs have seen a surge in traffic demand in recent years[13,

Table 1. Bandwidth for single TCP connection for different WAN latencies.

Latency (msec)	10	20	30	40
Bandwidth (Mbps)	1220	600	396	293




Fig. 2. Slow-down in training time in DP as we increase the WAN latency. y -values are multipliers and hence unit-less.




Fig. 3. Slow-down in training time in PP as we increase the WAN latency. y -values are multipliers and hence unit-less.

45, 53]. For example, Meta reported 13 \times increase in WAN traffic in 6 years (not including cross DC training traffic). However, it is challenging to increase the WAN capacity as laying out fiber and (in some cases) under-sea cables can take years [13]. Given this provisioning challenge and the increasing demand, WAN bandwidth is a constrained resource.

However, TP, EP and SP operate within layers and require substantially higher network bandwidth and small latency (that's offered by NVLink or InfiniBand)[39, 58]. As mentioned above, WAN suffers from high latency and low bandwidth compared to networks within DCs. On the other hand, DP and PP require much smaller bandwidth and are good candidates for training over WAN.

3 Limitations of existing designs

It is critical to understand the impact of WAN latency and capacity constraints on LM training workloads that span multiple DC sites. In this section, we quantify the WAN communication overhead on the different modes of parallelism.

Setup: We spawn 6 A100 GPU nodes across 3 DCs ($\times 2$ GPU nodes each). Each node has a single A100 GPU. For controlled experiments, we emulate the WAN latency using tc. Fig.1 shows the experimental setup. We use PyTorch framework for training. We measure the training time of iterations as we vary the WAN latencies. Lower (higher) WAN latencies indicate that the DCs are closer (farther).

Baseline models: We use two baseline models: (a) GPT-A (similar to GPT-3) with context length (L) of 4K and hidden dimension (H) of 4K, each layer is 412 million parameters, (b) GPT-B (bigger model than GPT-3) with L,H = 6K,8K, each layer is 1.2 billion parameters. We limit the number of layers to fit on 6 GPUs.

Note that the individual layer sizes for GPT-B are higher than Llama 3-70B model (such a model has roughly 80 layers resulting in each layer with size of roughly 875 million parameters). We use a smaller number of layers to fit the model on 6 GPUs.

3.1 Impact of WAN latency on DP

In this experiment, all the nodes form a single DP ring, i.e., the model is replicated across all the nodes in the ring. Individual nodes are fed different minibatches, and the gradients are averaged out during the all-reduce communication phase.




Fig. 4. Execution of compute phases (forward, backward, recompute) in PP in Varuna across 6 GPUs (G-1 to G-6) in 3 DCs for $L, H = 6K, 8K$ and WAN latency of 40 msec.

We measure the impact on the training time due to increasing WAN latencies. Fig.2 shows the training time for individual iterations. While compute time stays the same, increase in WAN latency affects the all-reduce time. Now, as shown in Table 1, higher WAN latency translates to lower achievable bandwidth and, hence, higher iteration time.

We make the following observations: (a) As shown in Fig.2, there is a significant slow-down in the training time as the WAN latency increases. For a WAN latency of 40 msec, the training time slows down by more than $15\times$ compared to the same DC baseline. (b) This slowdown is mainly because communication latency increased drastically¹. For WAN latency of 40 msec, 93% and 95% of the iteration time was spent in communication for the two models. (c) We find that all-reduce takes major chunk of the training time. For latency = 40 msec for GPT-B, all-reduce consumed roughly 98% of the training time. (d) PyTorch only uses one TCP connection between one pair of nodes for all-reduce. As the WAN latencies go up, the bandwidth is reduced (Table 1) that in turn increases the training time. (e) Lastly, WAN consists of TCP based interconnections where the data to be communicated is first copied to CPU that runs the TCP connection for communication. That said, time to copy the data between GPU-CPU is minuscule ($<1\%$) and majority of the time is consumed due to communication over the WAN.

3.2 Impact of WAN latency on PP

We evaluate the impact of WAN latencies on the training time when the nodes run PP. In this experiment, we split the number of layers across all 6 nodes. The nodes in the same DC get adjoining layers to minimize the cross-DC traffic (and improve training time). Note that the existing schedulers including GPipe[50] and Varuna[29] try to overlap compute with communication to reduce the job completion time. Fig.3 shows the increase in the training time as we increase the WAN latencies in Varuna scheduler that optimizes the PP scheduling over GPipe and Megatron[65].

Like DP, PP also shows significant slowdown in the training time as WAN latency increases resulting in a reduction in the achievable bandwidth. We also observe that the slowdown in PP is smaller than the slowdown in DP (y -range different for Fig. 2 & 3) at the same WAN latency.

Fig.4 shows the timeline of execution of PP for different microbatches for their forward pass, backward pass, and recomputation in Varuna for WAN latency = 40 msec. The training job starts in GPU G-1 and all microbatches execute back-to-back. The activations are sent to G-2, which is very fast as both the GPUs are in the same DC. However, as G-3 is in a different DC, the activations from G-2 to G-3 are sent over the low-bandwidth WAN. As a result, the activations take significant time (around 2.5 seconds) to transfer from G-2 to G-3². The delay continues to be observed for

¹Analytically, for a model with P parameters, N nodes in the all-reduce ring, and BW bandwidth, communication time per iteration can be calculated as $\frac{4 \cdot P \cdot (N-1)}{N \cdot BW}$ (extra factor 2 is due to fp16).

²The size of activations (and gradients) is $B \cdot L \cdot H$ (B = batch size, L = seq. length, H = hidden size)

subsequent microbatches as well as for the next GPUs (G-3 to G-6). GPU G-6 processes the backward pass at the end of each forward pass and starts sending the gradients for the other GPUs in reverse order (G-6 to G-1). Like activations, gradients also suffer due to lower WAN bandwidth.

We make the following observations: (a) Even with a state-of-the-art scheduler like Varuna, we see significant bubbles between the forward and backward passes and *also between microbatches*. (b) This slowdown is mainly because communication latency increased drastically. For WAN latency of 40 msec, 89.8% and 91% of iteration time was spent in communication for the two models. (c) Nodes G-5 and G-6 are idle for substantial time in the beginning when the activations are transferred between G-2 to G-4. (d) Like DP, PP also used one TCP connection between individual pairs of nodes when run with PyTorch. (e) Bandwidth is *not* shared across activations and gradients in PyTorch. When activations for one microbatch are transferred between two GPUs, activations for subsequent microbatches are queued (instead of parallel transfers). This helps finish transfer of existing microbatches and unblock the next GPUs for further processing. Lastly, as activations and gradients are sent in opposite directions, they do not compete for the same WAN bandwidth.

3.3 Impact of WAN latency on TP

TP requires high performant networks with low latency and very high bandwidth (BW) due to all-reduce phase on the critical path. It is mostly advised to run TP within a node (with multiple GPUs)[42, 88] where GPUs are connected via NVLink offering more than 800 Gbps BW. WAN networks do not provide such high BWs and we do not run TP over WAN. E.g., even with 40 msec WAN latency, the BW is just 293 Mbps (Table 1) (orders of magnitude smaller than NVLink).

4 ATLAS

We quantified in the previous section how lower bandwidth on the WAN substantially elongates training time for different modes of parallelism. We built ATLAS toward improving the training time. We touched upon the summary of design choices in §1.1. We expand on them next.

4.1 Using multiple TCP connections

As mentioned in the previous section, PyTorch[24] simply uses one TCP connection for communication. As shown in Table 1, as latency between the nodes increases, the achievable TCP bandwidth between the nodes shrinks. Recently high bandwidth mechanisms such as NVLink, InfiniBand and RDMA have been improved for longer distances[25]. However, they do not work beyond few hundreds of miles[25]. Most of the cross-region DCs easily fall beyond such distances.




Fig. 5. Bandwidth for single and multiple TCP connections. The server is located in US-East, and we vary the location of the client (X-axis). US-SC denotes US South-central DC. The numbers over the bars denote one-way latencies.

ATLAS enables the use of multiple parallel TCP connections between VMs in different DCs to scale the bandwidth. Fig.5 shows the achievable bandwidth for single and multiple TCP connections as we vary the distance (hence, WAN latency) between DCs. We have a server VM in Azure US-East DC. To vary the WAN latency, we select client VMs in different DCs in US and Asia. We measure the bandwidth using iperf3 and Cubic as TCP transport protocol. The x-axis shows the DC locations for clients. All VMs use Azure F32as_v6 series[3] with 32 cores and 20 Gbps NICs. We also tested with BBR[32]

(delay-based congestion control) flows and observed similar trends as discussed below. We set (multiple) TCP connections to achieve maximum bandwidth.

As we spawn more concurrent TCP connections, the cumulative bandwidth grows. Interestingly, we observe that such cumulative bandwidth does not drop below 5 Gbps even when the nodes are at a considerable distance, and does not increase even if we increase the number of concurrent connections. We believe that this is due to the rate-limiting at the hypervisor. AWS also throttles the bandwidth at 5 Gbps[52]. That said, the key take-away is that, ATLAS can help get up to 5 Gbps between two nodes on WAN irrespective of their distance. As a result, ATLAS can cut down the data transfer latency substantially. E.g., instead of using 250 Mbps on a single TCP connection, now ATLAS can get 5 Gbps over multiple connections – cutting down data transfer latency by 20×. ATLAS uses all available bandwidth (e.g., 5 Gbps) between two nodes by spawning multiple TCP connections. We profile the bandwidth for single TCP connection between nodes in different DCs and set the number of TCP connections to consume 5 Gbps between nodes.

We believe that the 5 Gbps limit between two nodes is not unreasonable. The WAN link capacity between two WAN routers typically ranges from a few hundred Gbps to a few Tbps[79] and even 100 nodes at full-throttle could consume 500 Gbps (large fraction of total capacity) – things add up quickly at scale. ATLAS helps consume the available WAN capacity efficiently while not violating any cloud-provider specific rate-limiting enforcement.

Lastly, we leave exploring MPTCP for future work. Currently, we prefer using multiple TCP connections over other alternatives such as MPTCP due to: (a) cloud providers are enforcing bandwidth caps for individual hosts; spawning multiple connections increases the aggregate bandwidth per host linearly with number of TCP connections and are able to reach the bandwidth limit (enforced by cloud provider), (b) MPTCP is primarily used for leveraging multiple interfaces; in our settings, each node has a single NIC, (c) multiple TCP connections are supported by all cloud offerings out of box.

4.2 Multiple forms of parallelism across DCs

Recall that, we use multiple forms of parallelism to speed up training (§2). The key question is how to divide GPUs in different DCs for such forms of parallelism. As mentioned in §3, TP, EP, SP requires significant bandwidth. Thus, ATLAS uses them only across intra-DC nodes connected via NVLink (in the same node), RDMA or Infiniband. This leaves ATLAS with deciding the division of GPUs across DCs for PP and/or DP.

PP across DCs is beneficial: Prior works including [83] have shown that doing PP across DCs and DP within DCs is beneficial in geo-distributed training. Recall that PP tries to overlap compute with communication. Bubbles in compute are formed when communication takes longer than compute. For context length L , hidden size H , and batch size B , communication time is $O(BLH)$ (linear with L and H), while the compute time is $O(BLH^2)$ (for MLP layers) + $O(BHL^2)$ (for attention layers) – quadratic with L and H [65]. Thus, the gap between compute and communication shrinks as model size grows (increasing L , H , or both). On the other hand, we can do DP across DCs where nodes in the same all-reduce ring are in different DCs. However, the all-reduce runs on critical path and its time depends on number of parameters of a layer (that in turn depends on H). It only increases as model size increases. Thus, we choose to do PP across DCs and DP within DCs.

We structure DP+PP parallelism as shown in Fig.6(a) where we use the structure using Varuna (we improve it as detailed in the next section): (a) We have multiple DP pipelines to scale throughput. (b) Each DP pipeline contains PP where subset of layers are assigned to individual GPUs. The PP runs across GPUs in different DCs. For example, let's consider 6 layers. In DP-1, PP is assigned across 6 GPUs in 3 DCs, each GPU is assigned one layer. (c) For different DP pipelines, each layer is

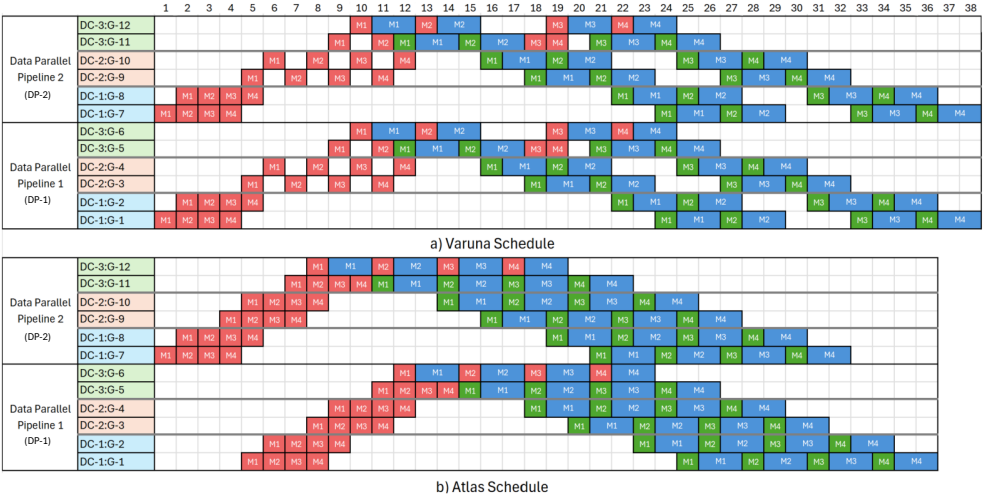


Fig. 6. Bandwidth sharing across DP pipelines in (b) ATLAS versus (a) Varuna. x -tics represent time-slots. M1 to M4 denote microbatches. Red, Green and Blue boxes indicate forward, recompute and backward passes.

assigned in the same DC. For example, layer-1 is assigned to GPUs G-1 and G-7 (in DC-1). Thus, the all-reduce ring (that runs for each layer) only runs across nodes in the same DC. (d) TP, EP, SP is assigned across GPUs on the same node (or nodes in the same DC).

4.3 Improving training time through co-ordination

When running multiform parallelism according to the above architecture, the existing schedulers such as Varuna and Gpipe assume there is no co-ordination between different DP pipelines. As a result, each DP pipeline gets its own share of bandwidth (e.g., 5 Gbps) and roughly follow the same schedule across different pipelines as shown in Fig.6(a). In essence, such sharing is *spatial*, as each pipeline has its own share of bandwidth.

Note that the nodes higher in the PP pipeline (e.g., node G-11 and G-12 in Fig.6(a)) are idle early in the execution as they have not received their microbatches. The delay is exasperated due to slower data transfer over WAN. We observe that there is higher aggregate bandwidth available by using multiple nodes across DP pipelines. In ATLAS, different DP pipelines work in tandem – first, we split the data of one DP pipeline across multiple nodes from other DP pipelines not using their bandwidth at that time. This way, each node gets 5 Gbps across WAN, but we get higher bandwidth using multiple nodes in parallel. To do so, the node that wants to send the data over WAN, first distributes the data to other nodes in different DP pipelines using higher intra-DC bandwidth (100+ Gbps). Next, all such nodes across multiple DP pipelines send the data across WAN in parallel. This way, the DP pipeline can speed up sending its activations to the next nodes in other DCs. We call it *temporal bandwidth sharing*.

We illustrate this key idea in Fig.6(b). Imagine ratio of communication to compute latency (C) to $2\times$. Imagine there are 2 DP pipelines each leveraging 5 Gbps WAN bandwidth between two nodes in different DCs. With schedulers like Varuna, as in Fig.6(a), downstream GPUs in the pipeline (PP) in a different DC (node G-11) need to wait for longer for activation transfer to happen using 5 Gbps WAN bandwidth. with ATLAS, the entire $2 \times 5 = 10$ Gbps bandwidth (across both DP instances) is available to each PP thus speeding up activation transfers to just 1 time-slot instead of 2. As a result, first DP instance progresses faster and makes the microbatches available to node G-11

and G-12 in DC-3 quickly, which then has *cascading effects* to process next batches sooner. These nodes then start their backward pass sooner too. Cumulatively, such temporal sharing improves the training time (Fig.6(b), x -range is shorter ending at time-slot 36 instead of 38). §6.2 shows that such a co-ordination based scheduler can improve the training time up to 1.52 \times compared to Varuna.

Bubble consolidation: The second key benefit of temporal bandwidth sharing is that it reduces the bubbles between microbatches. Notice the bubbles at node G-9, time-slot 5-11 (Fig.6(a)) – these bubbles are formed due to $C > 1$ (communication taking longer than compute). Using temporal bandwidth sharing, we can remove such bubbles between microbatches by setting the number of DP pipelines to C . This way communication takes same time as compute, removing all such bubbles between microbatches³. In other words, we can consolidate bubbles together by using temporal bandwidth sharing. This consolidation leads to continuous bubbles and more opportunities to schedule other workloads during these bubbles as discussed in §5.




Fig. 7. Fluctuations in WAN bandwidth.

Addressing jitters and bandwidth fluctuations, stragglers:

We measure the bandwidth for 24 hours from a VM in US-East DC (node A) to VMs in Southeast Asia (node B) and US-West (node C) DCs in Azure. We found that the bandwidth variations are quite small (Fig.7). Surprisingly, although the distance between nodes A-B is larger, the bandwidth variation is much smaller where coefficient of variance (CoV) is just 0.8% (CoV for A-C is 2.3%). This is because WANs are well provisioned, have

multiple mechanisms for masking network failures[45, 46] and suffer from relatively small packet loss [57]. As such fluctuations are small, ATLAS tries to eliminate bubbles altogether thus ramping up (costly) GPU utilization. Even if there are stragglers, bubbles around microbatches serve as a cushion. For GPU failures, we use the existing checkpointing mechanisms such as asynchronous methods and in-memory approaches[40]. That said, improving checkpointing over WAN is future work.

Lastly, we do not advocate training models using Internet as a transport as Internet suffers substantially higher packet losses[57].

4.4 Schedule for temporal bandwidth sharing

ATLAS uses a heuristic that calculates the schedule for forward and backward passes, as detailed next. The schedule is pre-computed by ATLAS before the training starts, and adjusts to any stragglers.

(1) ATLAS groups DP instances into *DP-cells* during the initialization phase. DP instances within a DP-cell coordinates usage of the aggregate WAN bandwidth. DP-cells operate independent of each other. In Fig.6(b), DP-1 and DP-2 form a single DP-cell. Each DP instance in a DP-cell is assigned a *LocalDPRank* and aggregate WAN bandwidth is shared temporally between these DP instances based on their *LocalDPRanks*. The bubbles between microbatches are because WAN communication is slower than compute. To eliminate such bubbles, we set number of DP pipelines in a DP cell as C , where C is communication to compute ratio.

(2) While scheduling forward passes for a microbatch on any DP pipeline, ATLAS filters only those for which activations/gradients in memory at any point of time is within the peak memory limit. As a result: (a) we are always within the peak memory limit, unlike Varuna, which may exceed memory

³If compute and communication take 10 and 30msec, then by using 3 DP pipelines, we reduce communication to 10msec.

Table 2. Notations used in selecting GPUs for training.

Notation	Explanation
D_{max}	Max. number of DP-cells
Num_GPU	Number of GPUs in individual DCs (Map)
C	Communication to compute ratio for PP
P	Max. number of partitions
$total_time$	Total time for different DP-cells (Map)
DCs	Ordered list of DCs
$Partitions$	Partitions assigned to DCs (Map)

limits, and (b) we don't block computation and communication phases on other DP pipeline because of unnecessary utilization of the aggregate WAN bandwidth for transmitting activations/gradient that would have later resulted in exceeding peak memory limit anyway.

(3) ATLAS schedules the compute phase of a micro-batch m in a DP pipeline at time t only when its communication phase could be scheduled immediately next, without overlapping with the communication phase of any other already generated schedule. In case of contention, ATLAS reschedules the compute phase for the micro-batch in a way that the ensuing communication phase does not overlap with any other network communication in the same DP-cell. As a result, DP-2 starts at time 1 and DP-1 starts at time 5 in Fig.6.

(4) If a stage k has both forward and backward tasks ready to be scheduled, ATLAS prioritizes the backward pass to unlock processing at subsequent nodes.

4.5 DC selection

Above sections assume a given topology of GPUs and optimize for such topology. However, it is not always beneficial to use all the GPUs. For example, if 1,000 GPUs are in the same DC, and 10 GPUs are in a different DC, then it's better to forgo the 10 GPUs and avoid slowdown in training time due to poor WAN conditions.

Now, we describe how ATLAS assists in finding the right number of GPUs from individual DCs. The goal here for ATLAS is to determine the right number of GPUs in individual DCs that maximizes training throughput (thus reduces training latency). Users can also run *what-if* analyses to determine the impact of varying number of GPUs in multiple DCs, and choose a set of GPUs to find a balance between the cost and the performance.

Our key observation is that we should (mostly) use all GPUs in a DC or none. It is better to house all the GPUs in the smallest number of DCs possible to improve the training time, as the PP time is better when GPUs are in the same DC compared to having same number of GPUs in different DCs that requires communication over WAN. Thus, ATLAS tries to maximize number of GPUs used in the same DC and tries to minimize the number of DCs used.

Algorithm. 1 calculates the training latency, and the notations are shown in Table 2. The inputs to the algorithm include: (a) an implicit ordering of the DCs (e.g., based on cost of GPUs; the default is based on decreasing order of GPU availability), (b) the number of available GPUs in each DC (Num_GPU), (c) maximum number of DP-cells (D_{max}), (d) communication to compute ratio (C), and (e) number of partitions (P). P is the ratio of total layers in a model to the number of layers fit on a model due to resources on GPU (e.g., GPU memory). Smaller partitions see smaller PP communication overhead. The output of the algorithm is the total time for different values of DP-cells (D). The users then decide D depending on cost, performance, and other metrics. D_{max}

Algorithm 1 Algorithm to calculate latency for different DP-cells.

INPUT: $D_{max}, DCs, Num_GPU, C, P$

OUTPUT: $total_time$

```

1: for  $D$  in  $\{1 \text{ to } D_{max}\}$  do
2:    $part\_left = P$ ;
3:   for  $dc$  in  $DCs$  do
4:      $PP\_GPU = \lfloor \frac{Num\_GPU[dc]}{D \cdot C} \rfloor$ ;
5:      $part\_assigned = \min(part\_left, PP\_GPU)$ ;
6:      $Partitions[dc] = part\_assigned$ ;
7:      $part\_left -= part\_assigned$ ;
8:     if  $part\_left == 0$  then break;
9:   end if
10:  end for
11:  if  $part\_left > 0$  then
12:     $PP\_time = \infty$ ;
13:  else
14:     $PP\_time = get\_latency\_pp(Partitions, D)$ ;
15:     $all\_reduce\_time = get\_latency\_dp(D \cdot C)$ ;
16:  end if
17:   $total\_time[D] = PP\_time + all\_reduce\_time$ ;
18: end for

```

can be set to $\sum \frac{Num_GPUs}{C \cdot P}$. At a high-level, the algorithm calculates the total training time of an iteration for each value (D) of DP-cell $[1, D_{max}]$ that includes the time for running PP and all-reduce in DP. The compute time (including TP, if any) is constant across D that we ignore.

To do so, the algorithm (a) first iterates DCs in an order based on cost, distance, or other metrics (line 3), and then (b) calculates the number of GPUs for the PP stages as $\frac{Num_GPU[dc]}{D \cdot C}$, as there are D DP-cells each with C individual DP pipelines (line 4). It assigns more PP stages to DCs with more GPUs, (c) It assigns the number of partitions that is minimum of partitions left and number of GPUs in PP calculated in (b) (line 5). (d) It stores the GPUs assigned in Partitions map (line 6) and adjusts partitions left (line 7). (e) This routine exits when all partitions are assigned or we run out of GPUs (line 8-12). (f) Lastly, it calculates the total execution time for a given D (line 14-17). (g) $get_latency_pp$ calculates the latency for temporal bandwidth sharing for a DP-cell. Similarly $get_latency_dp$ calculates the latency for the all-reduce phase across DP pipelines.

Calculating throughput: As we show in §6.4, such an algorithm then helps in finding the smallest D (hence GPUs) that provides highest throughput. The throughput is calculated as $\frac{D \cdot C}{total_time[D]}$.

Discussion on selecting DCs: Due to privacy regulations such as GDPR, it may not be possible to move data across jurisdictions. Many techniques can be used for compliance ([14]). Additionally, we can set enough DCs in such jurisdictions to consume the data locally.

Performance and cost modeling: As mentioned above, we can plug-in any combination of DCs and number of GPUs in DCs and the above algorithm can calculate the best *configuration* (DCs and set of GPUs) that provides highest throughput. We can plug-in different combinations and model the cost and performance trade-off *without any deployment*. This way, we can choose the best configuration as per our needs and then deploy it on hardware.

5 BUBBLETEA

As seen already, ATLAS reduces the training time using *temporal bandwidth sharing*. Additionally, it eliminates the bubbles between microbatches and collates them together. However, it does not eliminate the GPU compute bubbles where GPU is idle. We find that despite ATLAS finishing training faster (and improving GPU utilization), it results in up to 55% time spent in bubbles.

Note that achieving 100% GPU utilization (no bubbles) is challenging due to the inherent nature of dependencies between sequential stages of pipeline parallelism[29, 50, 71]. Training over WAN amplifies the bubbles due to longer latencies and smaller bandwidths on WAN.

There are two strawman’s approaches to reduce such residual bubbles, as follows. (a) *Increasing compute*: as compute is shorter than communication in PP, one way is to increase compute to match communication. This can be accomplished by each GPU processing more layers. As GPU memory is limited, the layers could be loaded from CPU (like [73]). However, we find that such a design is bottlenecked by PCIe bandwidth, where the time to load a 1 billion parameter layer from CPU is at least 100 msec even after using pinned-memory [10] (larger than processing the layer itself). (b) *Decreasing communication*: we tried reducing communication through compression that resulted in 2× slowdown in training latency while achieving similar loss rates (§6.7).

We take a novel approach to reduce bubbles. We build BUBBLETEA to schedule a different workload during the bubbles. Specifically, we observe that the prefill phase of the inference workload is well suited for such scheduling. Inference requests have distinct prefill and decode phases[27, 70] – the prefill phase processes the input prompt and the decode phase generates the output tokens in an autoregressive fashion. The execution time of decode phase is unknown at the start as it depends on the number of tokens generated. In contrast, the execution time of prefill phase is highly predictable as it only depends on the input prompt (size) that is known when the request arrives. Leveraging this insight, we build BUBBLETEA that provides *prefill-as-a-service* that we detail next.

5.1 BUBBLETEA design

BUBBLETEA runs prefill for already trained models such as LLama 3.1 while training a different model (e.g., GPT-5). The key goal is to schedule the prefill of inference requests without interfering with the training job. To do so, we do not simultaneously schedule training and prefill. Rather, the prefill kernels are issued only when the training job observes (predictable) bubbles. This way there is *no contention* on compute and memory bandwidth at GPUs. The only overhead comes from storing the models (discuss later). Based on the prefill estimated latency, we choose the GPU so that prefill ends (compute becomes available again) before the training job resumes (does not hamper training job). This approach significantly improves the GPU utilization in training clusters.




Fig. 8. BUBBLETEA Architecture.

The prefill phase requires the model weights stored in the GPU memory. Let’s assume that training uses model $M1$ (could be subset of the layers depending on the memory size), and inference uses model $M2$. One solution could be that we store both the models on CPU and load them to GPU memory before the respective jobs run. However, we observed that such dynamic loading incurs significant delays due to limited PCIe bandwidth. Assuming $M1$ and $M2$ each consume 60 GB memory on a 80 GB A100/H100 GPU, loading the

layers with 64 GBps PCIe bandwidth[21]⁴ takes roughly 0.93 second each. The GPU will be idle for this duration.

Instead with BUBBLETEA, we pre-store subset of layers of $M2$ on GPU, while we store $M1$ as-is. Consequently, $M2$ is sharded across more GPUs and we use PP across different GPUs to run $M2$. This design choice incurs a small overhead in terms of memory capacity. E.g., for serving Llama3-8B using 8 GPUs would require only 2 GB memory (at fp16) per GPU. Transient activations and KV caches use the memory not in use in training. This way, BUBBLETEA incurs very minor overhead. We do so to provide much larger memory for training model ($M1$) so that larger models such as GPT or Llama (405B) can be trained in geo-distributed settings.

BUBBLETEA architecture: Fig.8 shows the architecture in BUBBLETEA. The user request is received by the logically centralized inference controller. The inference controller then chooses the DC (E.g., using locality; its logic is outside the scope of BUBBLETEA), and sends the user prompt, DC and GPU details (for processing the decode) to the BUBBLETEA controller. BUBBLETEA controller then schedules the prefill on one of the ATLAS GPUs with enough capacity. After the completion of the prefill, the KV cache is transferred to the GPU specified by the inference controller for the decode phase (using Splitwise architecture[70]). The decode GPUs are in the same DC to use faster networks to transfer KV cache. In case BUBBLETEA does not have enough capacity to process the prefill phase using training GPUs (e.g., the training is ongoing and there is no bubble at that time), it immediately informs the inference controller accordingly.

Pipeline parallelism for prefill requests: We do PP for inference model across GPUs. To do so, BUBBLETEA forms a PP pipeline across GPUs of the same rank in individual DP-cells. Fig.6(b) shows a single DP-cell. Imagine another DP-cell with GPUs G-13 to G-24. GPUs G-1, G-13 and GPUs from subsequent DP-cells form a PP where each GPU is assigned a small number of layers. Note that, PP in BUBBLETEA is formed only across GPUs in the same DC to reduce latency. As prefill latency is highly predictable, BUBBLETEA schedules prefill only if there is available capacity. To do so, it finds first available PP pipeline that has bubble across GPUs to accommodate the prefill request scheduled using PP. Importantly, we find that doing PP is beneficial even if the entire inferencing model fits on the same GPU (§6.6) to reduce TTFT (Time-To-First-Token). Such PP improves the TTFT by 67% for larger prefills (§6.6). We discuss this and BUBBLETEA overhead in §6.6. This approach has no impact on TBT (Time-between-Tokens). The TTFT can be further reduced using *chunked* prefills[27] that we leave for future work. No component becomes single point of failure. We discuss BUBBLETEA controller overhead in §6.5.

Scheduling prefills: BUBBLETEA controller gets a rough plan for scheduling the microbatches from ATLAS and identifies the bubbles where prefill can be scheduled. It then waits for signals from the individual GPUs when they process microbatches (using hooks in PyTorch). It then assigns prefills during bubbles based on such signals.

Scheduling decodes: Once the prefill is computed on training GPUs, the KV cache is transferred to decode GPUs using existing designs such as Splitwise[70] and DistServe[96]. As such designs already decouple prefill and decode, they optimize for decodes independently (such as batching). BUBBLETEA uses such designs for scheduling decodes.

LMs where BUBBLETEA helps: We recommend running smaller LMs including small Llama[15], Phi[17], Mistral[18] models. Recall that the training model in ATLAS and inference model in BUBBLETEA are to be stored in the GPU memory. We do not use them simultaneously, so there is no contention on memory bandwidth and compute SMs on GPU. To keep the memory overhead of the inference model small, we only recommend using the small LMs for inference. This way most

⁴PCIe gen5 has bidirectional bandwidth of 128 GBps; one-way bandwidth is 64 GBps.

of the memory is available for the training model and larger models such as GPT or Llama (405B) can be trained in geo-distributed settings. §6.6 shows the overhead in BUBBLETEA.

Lastly, while we propose interleaving prefill (from inference workloads) with the training, there could be other workloads that can be scheduled during bubbles. We detail it in §7.

6 Evaluation

We evaluate ATLAS and BUBBLETEA using testbed experiments and large-scale simulations.

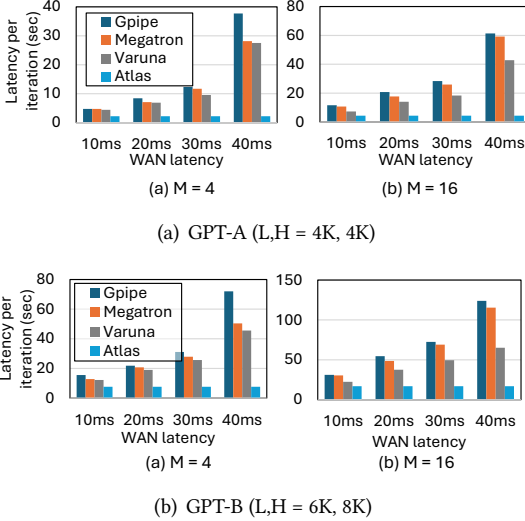


Fig. 9. [Testbed] Training latency (for one iteration) for GPT-A and GPT-B for different WAN latencies and number of microbatches (M).

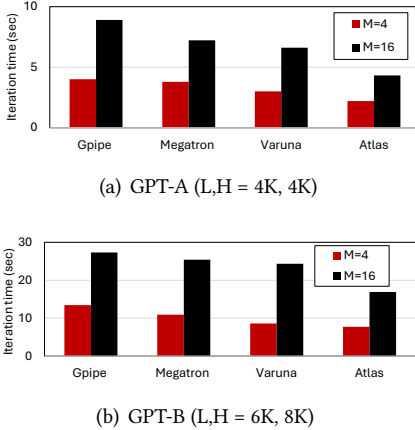


Fig. 10. [Testbed] Iteration time for GPT-A and GPT-B across different number of microbatches (M).

The key observations are: (a) ATLAS, with its optimizations, improves the cross-DC training time by up to $17\times$ compared to baselines like Varuna, GPipe and Megatron. (b) Even when the baselines leverage multiple TCP connections, ATLAS improves training time by up to $1.82\times$ than baselines using temporal bandwidth sharing. (c) ATLAS is able to scale the throughput as we add more DCs. This is in stark contrast with findings in §3 where throughput reduced when using WAN. (d) BUBBLETEA could schedule the prefill requests during bubbles and achieves a GPU utilization of up to 94% in the training cluster, up from 45% when using only ATLAS.

6.1 Setup

Setup for testbed experiments: Our evaluation uses a setup similar to Fig.1 consisting of 12 GPUs in 3 DCs (4×3). We vary the latency between DCs using tc command. We set the maximum cross-DC bandwidth between two nodes to 5 Gbps (§4.1), whereas bandwidth between two nodes in the same DC is capped to 100 Gbps.

All GPUs are A100 80 GB running PyTorch. We compare the training time using ATLAS versus baselines including GPipe[50], Megatron[65] and Varuna[29]. For the baselines, we divide the GPUs among 3 DP pipelines with 4 PP stages. For ATLAS, we have 1 DP-cell with 3 DP pipelines with 4 PP stages as above. While we do not leverage TP for the demonstrations, TP could be used intra-DC. We use GPT-A and GPT-B models (§3) for the evaluation. As mentioned in §3, we set the

number of layers to fit on 4 PP stages. However, individual layer sizes in GPT-B model are higher than Llama 3 80B model. We set the number of microbatches to 4 (inline with the 4 PP stages) and 16.

6.2 ATLAS improves training time

We show that ATLAS is able to improve the training time by up to $17\times$ compared to the baselines using the 2 key ideas of using multiple TCP connections and temporal bandwidth sharing. We compare ATLAS against default baselines (that do not use multiple connections; we give the benefit of multiple TCP connections to the baselines in the next section). We vary the WAN latency from 10 msec to 40 msec. Note that, irrespective of the latency, ATLAS gets same bandwidth due to multiple TCP connections. In contrast, bandwidth for the baselines reduces as WAN latency increases (Table 1).

As shown in Fig.9, ATLAS is able to reduce the training time substantially. We find: (a) Compared to Gpipe, Megatron and Varuna, ATLAS is able to reduce the training time by up to $17\times$, $13\times$, and $12\times$ resp. across different WAN latencies and number of microbatches (M). (b) As expected, ATLAS gains increase as WAN latencies increase. This is because the baselines get lower bandwidth for single TCP connection as WAN latencies increase. That said, even for a small WAN latency of 10 msec, ATLAS is able to improve the training time by up to $2.68\times$. (c) the gains also reduce for $M=16$ compared to $M=4$. This is mainly because the compute takes higher fraction of the total latency for $M=16$. (d) the gains in ATLAS also reduce for GPT-B compared to GPT-A, again due to compute taking higher fraction of total latency.

Temporal bandwidth sharing is crucial In this experiment, we show that just using temporal bandwidth sharing can still substantially improve the training time (by up to $1.82\times$). We continue to use the 12 GPU cluster from previous section. In this experiment, we give the benefit of using multiple TCP connections to all the baselines too alongside ATLAS. This way, we primarily evaluate the benefits of temporal bandwidth sharing (§4.3). Fig.10 shows the training latency per iteration (average) for GPT-A and GPT-B. We also measure such latencies for 4 and 16 microbatches. We found that ATLAS improves the training latency by up to $1.82\times$, $1.72\times$ and $1.52\times$ compared to Gpipe, Megatron and Varuna. Compared to Varuna, we found that ATLAS is able to start the microbatches sooner at the GPUs higher in the PP pipelines that in turn starts the backward passes sooner.

Together, ATLAS achieves GPU utilization to near 45%.

6.3 Efficient use of cross-DC GPUs

We demonstrate now through simulations how ATLAS improves throughput over Varuna by up to 48% (potentially saving 10s of millions of dollars at scale) by using GPUs across different DCs. In the previous experiments, we used 12 GPUs. In this section, we use simulations to evaluate the benefits on ATLAS over Varuna for bigger clusters.

In this experiment, we consider two DC-sets (simulated). DC-set-1 contains 600 GPUs in each DC and we vary the number of DCs from 1 to 5. DC-set-2 contains GPUs, likewise, in 5 DCs as [600, 500, 400, 300, 200]. In this experiment, we give the benefits of using multiple TCP connections to both ATLAS and Varuna. For both ATLAS and Varuna, we assume 5 Gbps bandwidth is available between 2 nodes in different DCs irrespective of the latency. This way, we mainly focus on the benefits due to temporal bandwidth sharing in ATLAS. As mentioned in §4.2, despite multiple TCP connections, communication still takes 3-4 \times compute latency. In this experiment, we consider two communication to compute ratios (C) as 4 \times and 2 \times . We set number of layers (microbatches and PP degree) to 60. In Varuna, we set the number of DP pipelines as per the available GPUs and PP degree (60). In ATLAS, we set the number of DP pipelines calculated using Algorithm. 1 (§4.5). However, we found that such an algorithm used all GPUs. For DC-set-2, it assigned more PP stages

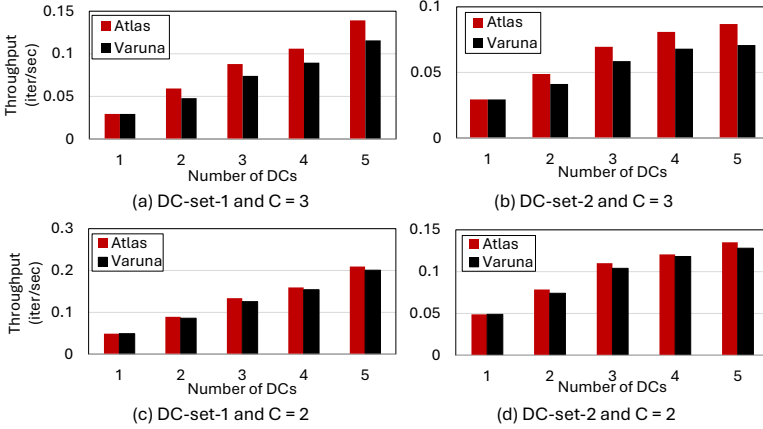


Fig. 11. [Simulation] Throughput in ATLAS and Varuna for different DC-sets and communication to compute ratio (C).

to DCs with more GPUs. In ATLAS, we set the number of DP pipelines in a DP-cell to C as that’s the minimum number of DP pipelines required to overlap communication time with compute.

Scalability with increasing DCs: Fig.11 shows the achieved training throughput (iterations/second) in 4 cases. It can be seen that training throughput improves with more GPUs. With ATLAS, for [DC-set-1, C = 4, 5 DCs], we get roughly 4.7× throughput compared to using a single DC. Similarly for [DC-set-1, C = 2, 5 DCs] the speedup is roughly 4.3×. Similar trends occur for Varuna too. This shows that the multiple forms of parallelism used in ATLAS, where PP is used across DCs, is effective in using GPUs from different DCs.

Next, we consider the benefits of using ATLAS over Varuna. For C = 4, where communication time is 4× compute time, ATLAS substantially outperforms Varuna. For both DC-set-1 and DC-set-2, the improvement in throughput over Varuna is up to 48%. The speedup is mainly because, with temporal bandwidth sharing, ATLAS is able to improve the network utilization and start the forward and backward passes sooner at nodes higher up in pipelines. However, the benefits shrink for C= 2, along the lines of expectation, where ATLAS improves throughput by up to 25% because, at lower C, there are already fewer bubbles and ATLAS has less opportunities to pack the individual passes.

6.4 Cross-DC GPU balancing




Fig. 12. [Simulation] Throughput in ATLAS using 2 DCs.

In the last experiment, we used all the GPUs available. In this experiment, we show that *it’s not always better to use all available GPUs across DCs, and Algorithm 1 is needed to calculate the smallest number of GPUs required for the highest throughput*. In this experiment, we pick two DCs – the first DC has a fixed number (600) of GPUs and

we vary the number of GPUs in the second DC from 0 to 600 in steps of 10% (denoted by F). Through simulations (same setup as last experiment with $C=2$), we calculate the training throughput. We use Algorithm. 1 (§4.5) to calculate the optimal number of GPUs in individual DCs for highest throughput.

Fig.12 shows the improvement in training throughput normalized to using GPUs in a single DC. It can be seen that when F increases to 10%, there is no improvement in throughput. Algorithm. 1, in this case, falls back to only first DC (no GPUs from second DC). This is because, for such F , we need to send the activations from first DC to second DC over the WAN. This inflates the latency to complete one iteration by roughly 11% and, hence, erases any throughput gains by adding the additional 10% GPU compute. We observe similar behavior when increasing F from 20% to 30%. This reinstates the need for Algorithm. 1 that does a sweep across different options and chooses the one that optimizes for a given metric (cost, latency, throughput, etc.).

Lastly, Algorithm. 1 is fast – even for 5 DCs with 600 GPUs each, it took 55 seconds to execute on a single CPU machine. Engineers can quickly do exhaustive *what-if* analysis and estimate the impact of different number of DCs and GPUs.

6.5 BUBBLETEA boosts GPU utilization

In this section, we show how BUBBLETEA can significantly improve the GPU utilization by scheduling prefill stages of inference requests during bubbles in ATLAS. We use the 12 GPU setup and schedule prefills as mentioned in §5.1 running GPT-A. As there is only one DP-cell, we set the PP depth of the inference model to 1. Recall that the BUBBLETEA controller combines (1) the scheduling plan from ATLAS controller and (2) signals from the GPUs when they finish processing training microbatches. Accordingly, BUBBLETEA issues the inference kernels.




Fig. 13. [Testbed] Execution of forward (blue), backward (orange), recompute (green) and *inference prefill* (red) on two GPUs with BUBBLETEA and ATLAS.




Fig. 14. [Testbed] TTFT in using different PP degree in BUBBLETEA for Llama3 8-billion.

Fig.13 shows the GPU utilization for two of the GPUs in a single (training) pipeline. We make the following observations: (a) BUBBLETEA is able to schedule the prefill stages perfectly during the training bubbles. This shows that BUBBLETEA controller can accurately detect the bubbles (through combining scheduling plan and GPU signals). (b) In doing so, BUBBLETEA boosts the GPU utilization to around 94% leaving very small residual bubbles. (c) There is no impact on the training job. We found BUBBLETEA leaves some bubbles between prefill and training

microbatches that helps resume training without delay. (d) BUBBLETEA does not schedule prefill if the bubble is not big enough.

Impact on GPU utilization: Fig.13 shows BUBBLETEA goes one step further than ATLAS and improves the GPU utilization from 45% (ATLAS-only) to 94% using inference prefills.

BUBBLETEA overhead: In the 12-GPU experiments, we could find the prefill bubble to schedule inference request in $<100\mu\text{sec}$. We also do a large scale simulation over 1000 GPUs in a single DC. The GPUs are divided across 50 DP-cells. We replay inference workload as coding dataset from [2]. We found that BUBBLETEA is able to find the training bubble within 200 μsec and the wait time (queuing delay) is within 8 msec.

6.6 BUBBLETEA PP overhead

Recall that, to save the memory overhead of storing inference models in BUBBLETEA, we split the models across different GPUs and use PP. This only impacts TTFT and not TBT as decode runs on dedicated GPUs (§5.1). In this experiment, we use a DGX node using 8 A100 GPUs connected via NVLink. We use the Llama3 8 billion model. We vary the PP degree from 1 to 8 and set the number of GPUs to the PP degree. Fig.14 shows the TTFT for different PP degrees as we vary the prefill length. We make two key observations: (a) Increasing the PP degree negatively impacts the TTFT when prefill length is small. For 512 tokens, PP of degree 8 inflates TTFT by 29%. However, in absolute numbers, the increase is just 16 msec and the TTFT is well under SLOs[96]. The increase is primarily due to communication overhead in PP. (b) As the prefill length increases, the TTFT improves with higher PP degree. In fact, for higher prefill lengths, higher PP degrees actually perform better than lower PP degrees. E.g., for prefill length of 8K tokens, the TTFT for PP=1 is 67% higher than PP=8. We observe that, for larger prefills, the compute becomes saturated even for a subset of the layers. Consequently, the weights are to be swapped in and out, resulting in higher latency. With PP=8, the layers are already partitioned and there is no need to swap, thus improving TTFT. At PP=8, each GPU only uses (small) 2 GB memory for the inferencing model.

LMs where BUBBLETEA helps: As noted previously, we currently limit the memory available to *inference models* to 1-2 GB so that most of the GPU memory is available to the training model (such as GPT or Llama-405B). BUBBLETEA can still support many small/medium LMs from the Llama, Phi, Mistral, and Gemma family for inference. As we relax this constraint, we can support more models that we leave for future work. Additionally, BUBBLETEA can easily support batch (latency insensitive) requests[96]. Note that, we only need to use smaller models for inference (in BUBBLETEA). In ATLAS, we focus on larger LMs such as GPT or Llama-405 for training.

6.7 Comparing with semantics altering methods

ATLAS does not change the semantics (functionalities) of the protocols – it runs *standard* PP and DP to maintain the accuracy. While we tried other designs that change semantics to improve training time at the cost of lower accuracy, our engineers suggest against such coercive approaches.

We tried to reduce the communication time in PP through compression techniques from [30] and SVD [26]. They achieve good compression but suffer from model accuracy loss and/or training time inflation. [30] also notes similar observation for many compression techniques such as Top-K and quantization. Additionally, we found that SVD based compression takes roughly $2\times$ compute time to achieve the same loss rate without using compression. For such reasons, we do not use methods that alter the semantics of the protocols. That said, key ideas in ATLAS are complimentary to semantic altering methods. Using multiple TCP connections and temporal bandwidth sharing can work in tandem and can improve training latency.

7 Related work

ATLAS and BUBBLETEA together improve the training time and GPU utilization when using GPUs from different DCs. We detail the related work below:

Geo-distributed training: There have been much work on geo-distributed training. [47] optimizes for parameter server based training that is orthogonal to ATLAS and does not use PP or DP like parallelism. [61, 86] train for consumer held devices. [61] optimizes for all-reduce traffic. However, they change the semantics that can affect the model accuracy. Similarly, works change the semantics either through quantization[44, 90, 97] or sparsification[36, 60, 77], or a combination of both[89]. Existing works use different strategies for communication frequencies[41, 80, 92] that again can affect the model accuracy. In contrast, ATLAS does not change the semantics preserving the accuracy. Like ATLAS, [93] structures the DP and PP among *heterogeneous* nodes. However, they assign tasks to *all* nodes that may inflate training time. Additionally, such works do not take advantage of DC-like settings that are more homogeneous. [35] focuses on assigning training jobs to individual DCs and do not focus on speeding up training or GPU utilization by making use of GPUs in multiple DCs. Lastly, [85] focuses on searching best configurations that ATLAS can leverage. Similarly, ATLAS is not tied to any one training schedule and can work with schedule computed by other systems such as CrossPipe[34]. ATLAS compliment such systems using temporal bandwidth sharing and using multiple TCP connections.

Mixing training and inference workloads: Like BUBBLETEA and ATLAS, existing works have also looked at combining training and inference workloads. Orion[84] looks at fine grained scheduling of such workloads to reduce interference. Usher[78] focuses on interference due to different inference models. Sarathi[27] mix prefill and decode phases of inference. [48] shows significant interference when scheduling prefills and decodes together. [28] uses other workloads to fill in bubbles. We focus on decoupling training and inference to avoid interference. Lastly, Nvidia supports multiple instances running concurrently on a GPU through MIG[22]. We do not use it as we do not run training and inference concurrently. In fact, we increase the GPU utilization by repurposing idle GPUs for inference.

Improving network communication for training: Recent works have also focused on improving the communication for training workloads. [38] improves RDMA over Ethernet for training. NCCL[23], TACCL[75], TECCL[63] and RCCL[1] optimize communication when GPUs are in the same DC, and do not optimize for TCP used for WAN communication. CASSINI [72] explores traffic interleaving for multiple training jobs while our work improves communication for individual large training jobs in a cross-DC WAN setting leveraging temporal bandwidth sharing as one of the key techniques. Lastly, using multiple TCP connections have been used previously including [33, 66, 94].

Reducing pipeline bubbles: Works such as ZB[71] try to eliminate bubbles during forward and backward passes. However, such works: (a) continue to suffer from bubbles due to lower WAN bandwidth[34], and (b) require higher memory. Such a design is complimentary to temporal bandwidth sharing design in ATLAS. Works such as Hydro[49] use it to try out its trial jobs (not including inference). PipeFischer[67] uses it to schedule parts of the existing job such as scheduling refresh the curvature and inverse matrices. PP in ATLAS is complimentary to PP in Megascala[55] that can further benefit ATLAS.

8 Conclusion

The widespread adoption of language models (LMs) has caused a huge surge in demand for GPUs. Training large LMs is important to improve accuracy and support larger context lengths. However, training large LMs is increasingly getting challenging as such training requires a large GPU fleet

that’s challenging to house in a single data center (DC). In this paper, we focus on training large LMs when GPUs are distributed across different DCs. We build ATLAS and BUBBLETEA as complimentary tools to improve the training time and GPU utilization. ATLAS improves the training time using multiple design choices including novel temporal bandwidth sharing, but still results in bubbles (idle GPU time). BUBBLETEA schedules prefill phase of inference requests during bubbles to substantially improve GPU utilization. Together, ATLAS and BUBBLETEA improve the training time up to 17× and achieve GPU utilization of up to 94%.

References

- [1] AMD RCCL library. <https://github.com/ROCm/rccl>.
- [2] Azure LLM trace. <https://github.com/Azure/AzurePublicDataset/blob/master/AzureLLMInferenceDataset2024.md>.
- [3] Azure VM. <https://learn.microsoft.com/en-us/azure/virtual-machines/sizes/compute-optimized/fasv6-series?tabs=sizenetwork>.
- [4] Claude AI. <https://claude.ai/>.
- [5] Cost of training GPT models. <https://www.fierceelectronics.com/sensors/chatgpt-runs-10k-nvidia-training-gpus-potential-thousands-more>.
- [6] Cost of training Llama 3 models. <https://news.ycombinator.com/item?id=35008694>.
- [7] Geo-distributed training. <https://datacrunch.io/blog/multi-data-center-training-prime-intellect>.
- [8] GPT-4 model. <https://openai.com/index/gpt-4/>.
- [9] GPT3 and GPT4 sizes. <https://getgenie.ai/gpt-3-vs-gpt-4/>.
- [10] GPT3 and GPT4 sizes. <https://discuss.pytorch.org/t/when-to-set-pin-memory-to-true/19723>.
- [11] GPU power requirements. <https://www.fierce-network.com/cloud/could-gpu-power-levels-break-data-center-ecosystem>.
- [12] GPU power requirements. <https://www.linkedin.com/pulse/heat-challenge-ai-infrastructure-gpu-servers-trgdatacenter-b6vcc>.
- [13] Increase in WAN traffic at Meta. <https://engineering.fb.com/2023/01/27/networking-traffic/optimizing-large-scale-networks-meta-engineers/>.
- [14] Language model compliance. <https://www.linkedin.com/pulse/ensuring-gdpr-compliance-large-language-model-aarush-bhardwaj-m2nye>.
- [15] Llama models. <https://llama.meta.com/>.
- [16] Microsoft geo-distributed training. <https://blogs.microsoft.com/blog/2025/09/18/inside-the-worlds-most-powerful-ai-datacenter/>.
- [17] Microsoft Phi models. <https://azure.microsoft.com/en-us/products/phi>.
- [18] Mistral models. <https://mistral.ai/>.
- [19] Number of GPUs for ChatGPT. <https://www.linkedin.com/pulse/10-mind-blowing-facts-behind-chatgpts-silicon-brain-power-elharony-bvklpf>.
- [20] Nvidia geo-distributed training. <https://developer.nvidia.com/blog/turbocharge-llm-training-across-long-haul-data-center-networks-with-nvidia-nemo-framework/>.
- [21] Nvidia H100 Specs. <https://www.nvidia.com/en-in/data-center/h100/>.
- [22] Nvidia MIG. <https://www.nvidia.com/en-in/technologies/multi-instance-gpu/>.
- [23] Nvidia NCCL library. <https://developer.nvidia.com/nccl>.
- [24] PyTorch. <https://pytorch.org/>.
- [25] RDMA over longer distance. <https://ntt-review.jp/archive/ntttechnical.php?contents=ntr202403ra1.html>.
- [26] Singular Value Decomposition. https://en.wikipedia.org/wiki/Singular_value_decomposition.
- [27] A. Agrawal, N. Kedia, A. Panwar, J. Mohan, N. Kwatra, B. Gulavani, A. Tumanov, and R. Ramjee. Taming {Throughput-Latency} tradeoff in {LLM} inference with {Sarathi-Serve}. In *USENIX OSDI 2024*.
- [28] D. Arfeen, Z. Zhang, X. Fu, G. R. Ganger, and Y. Wang. Pipefill: Using gpus during bubbles in pipeline-parallel llm training. *arXiv preprint arXiv:2410.07192*, 2024.
- [29] S. Athlur, N. Saran, M. Sivathanu, R. Ramjee, and N. Kwatra. Varuna: scalable, low-cost training of massive deep learning models. In *ACM EuroSys 2022*.
- [30] S. Bian, D. Li, H. Wang, E. Xing, and S. Venkataraman. Does compressing activations help model parallel training? *MLSys*, 2024.
- [31] J. Cao, Y. Guan, K. Qian, J. Gao, W. Xiao, J. Dong, B. Fu, D. Cai, and E. Zhai. Crux: Gpu-efficient communication scheduling for deep learning training. In *ACM SIGCOMM 2024*.

- [32] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson. Bbr: congestion-based congestion control. *Communications of the ACM*, 2017.
- [33] L. Chouffournier, A. Ali-Eldin, P. Sharma, P. Shenoy, and D. Towsley. Performance evaluation of multi-path tcp for data center and cloud workloads. In *ACM/SPEC international conference on performance engineering*, 2019.
- [34] T. Chen, A. Kubicek, L. Huang, and T. Hoefler. Crosspipe: Towards optimal pipeline schedules for cross-datacenter training. *arXiv preprint arXiv:2507.00217*, 2025.
- [35] A. Choudhury, Y. Wang, T. Pelkonen, K. Srinivasan, A. Jain, S. Lin, D. David, S. Soleimanifard, M. Chen, A. Yadav, et al. {MAST}: Global scheduling of {ML} training across {Geo-Distributed} datacenters at hyperscale. In *USENIX OSDI 2024*.
- [36] C. Dünner, T. Parnell, and M. Jaggi. Efficient use of limited-memory accelerators for linear learning on heterogeneous systems. *NeurIPS*, 2017.
- [37] W. Fedus, B. Zoph, and N. Shazeer. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *JMLR*, 2022.
- [38] A. Gangidi, R. Miao, S. Zheng, S. J. Bondu, G. Goes, H. Morsy, R. Puri, M. Riftadi, A. J. Shetty, J. Yang, S. Zhang, M. J. Fernandez, S. Gandham, and H. Zeng. Rdma over ethernet for distributed training at meta scale. In *ACM SIGCOMM 2024*.
- [39] R. Gond, N. Kwatra, and R. Ramjee. Tokenweave: Efficient compute-communication overlap for distributed llm inference. *arXiv preprint arXiv:2505.11329*, 2025.
- [40] T. Gupta, S. Krishnan, R. Kumar, A. Vijeve, B. Gulavani, N. Kwatra, R. Ramjee, and M. Sivathanu. Just-in-time checkpointing: Low cost error recovery from deep learning training failures. In *ACM Eurosys*, 2024.
- [41] F. Haddadpour, M. M. Kamani, M. Mahdavi, and V. Cadambe. Local sgd with periodic averaging: Tighter analysis and adaptive synchronization. *NeurIPS*, 2019.
- [42] B. Hanindhito, B. Patel, and L. K. John. Bandwidth characterization of deepspeed on distributed large language model training. In *IEEE ISPASS*, 2024.
- [43] J. He, J. Zhai, T. Antunes, H. Wang, F. Luo, S. Shi, and Q. Li. Fastermoe: modeling and optimizing training of large-scale dynamic pre-trained models. In *ACM Symposium on Principles and Practice of Parallel Programming*, 2022.
- [44] M. Höhfeld and S. E. Fahlman. Probabilistic rounding in neural network learning with limited precision. *Neurocomputing* 1992.
- [45] C.-Y. Hong, S. Kandula, R. Mahajan, M. Zhang, V. Gill, M. Nanduri, and R. Wattenhofer. Achieving high utilization with software-driven wan. In *ACM SIGCOMM 2013*.
- [46] C.-Y. Hong, S. Mandal, M. Al-Fares, M. Zhu, R. Alimi, C. Bhagat, S. Jain, J. Kaimal, S. Liang, K. Mendeleev, et al. B4 and after: managing hierarchy, partitioning, and asymmetry for availability and scale in google’s software-defined wan. In *ACM SIGCOMM 2018*.
- [47] K. Hsieh, A. Harlap, N. Vijaykumar, D. Konomis, G. R. Ganger, P. B. Gibbons, and O. Mutlu. Gaia: geo-distributed machine learning approaching lan speeds. In *USENIX NSDI 2017*.
- [48] C. Hu, H. Huang, L. Xu, X. Chen, J. Xu, S. Chen, H. Feng, C. Wang, S. Wang, Y. Bao, et al. Inference without interference: Disaggregate llm inference for mixed downstream workloads. *arXiv preprint arXiv:2401.11181*, 2024.
- [49] Q. Hu, Z. Ye, M. Zhang, Q. Chen, P. Sun, Y. Wen, and T. Zhang. Hydro-{:Surrogate-Based} hyperparameter tuning service in datacenters. In *USENIX OSDI 2023*.
- [50] Y. Huang, Y. Cheng, A. Bapna, O. Firat, D. Chen, M. Chen, H. Lee, J. Ngiam, Q. V. Le, Y. Wu, et al. Gpipe: Efficient training of giant neural networks using pipeline parallelism. In *NeurIPS 2019*.
- [51] S. A. Jacobs, M. Tanaka, C. Zhang, M. Zhang, S. L. Song, S. Rajbhandari, and Y. He. Deepspeed ulyssees: System optimizations for enabling training of extreme long sequence transformer models. *arXiv preprint arXiv:2309.14509*, 2023.
- [52] P. Jain, S. Kumar, S. Wooders, S. G. Patil, J. E. Gonzalez, and I. Stoica. Skyplane: Optimizing transfer cost and throughput using {Cloud-Aware} overlays. In *USENIX NSDI 2023*.
- [53] S. Jain, A. Kumar, S. Mandal, J. Ong, L. Poutievski, A. Singh, S. Venkata, J. Wanderer, J. Zhou, M. Zhu, J. Zolla, U. Hölzle, S. Stuart, and A. Vahdat. B4: Experience with a globally-deployed software defined wan. In *ACM SIGCOMM 2013*.
- [54] S. Jayaram Subramanya, D. Arfeen, S. Lin, A. Qiao, Z. Jia, and G. R. Ganger. Sia: Heterogeneity-aware, goodput-optimized ml-cluster scheduling. In *ACM SOSP 2023*.
- [55] Z. Jiang, H. Lin, Y. Zhong, Q. Huang, Y. Chen, Z. Zhang, Y. Peng, X. Li, C. Xie, S. Nong, et al. {MegaScale}: Scaling large language model training to more than 10,000 {GPUs}. In *USENIX NSDI 2024*.
- [56] J. Kaplan, S. McCandlish, T. Henighan, T. B. Brown, B. Chess, R. Child, S. Gray, A. Radford, J. Wu, and D. Amodei. Scaling laws for neural language models. *arXiv preprint arXiv:2001.08361*, 2020.
- [57] B. Kataria, P. LNU, R. Bothra, R. Gandhi, D. Bhattacherjee, V. N. Padmanabhan, I. Atov, S. Ramakrishnan, S. Chaturmohta, C. Kotipalli, R. Liang, K. Sueda, X. He, and K. Hinton. Saving private wan: Using internet paths to offload wan traffic in conferencing services. In *ACM CoNEXT 2024*.

- [58] D. Lepikhin, H. Lee, Y. Xu, D. Chen, O. Firat, Y. Huang, M. Krikun, N. Shazeer, and Z. Chen. Gshard: Scaling giant models with conditional computation and automatic sharding. *arXiv preprint arXiv:2006.16668*, 2020.
- [59] J. Li, Y. Jiang, Y. Zhu, C. Wang, and H. Xu. Accelerating distributed {MoE} training and inference with lina. In *USENIX ATC*, 2023.
- [60] X. Li, B. Karimi, and P. Li. On distributed adaptive optimization with gradient compression. *arXiv preprint arXiv:2205.05632*, 2022.
- [61] H. Lim, J. Ye, S. Abdu Jyothi, and D. Han. Accelerating model training in multi-cluster environments with consumer-grade gpus. In *ACM SIGCOMM 2024*.
- [62] H. Liu, M. Zaharia, and P. Abbeel. Ring attention with blockwise transformers for near-infinite context. *arXiv preprint arXiv:2310.01889*, 2023.
- [63] X. Liu, B. Arzani, S. K. R. Kakarla, L. Zhao, V. Liu, M. Castro, S. Kandula, and L. Marshall. Rethinking machine learning collective communication as a multi-commodity flow problem. In *ACM SIGCOMM 2024*.
- [64] D. Narayanan, A. Harlap, A. Phanishayee, V. Seshadri, N. R. Devanur, G. R. Ganger, P. B. Gibbons, and M. Zaharia. Pipedream: Generalized pipeline parallelism for dnn training. In *ACM SOSP 2019*.
- [65] D. Narayanan, M. Shoeybi, J. Casper, P. LeGresley, M. Patwary, V. Korthikanti, D. Vainbrand, P. Kashinkunti, J. Bernauer, B. Catanzaro, et al. Efficient large-scale language model training on gpu clusters using megatron-lm. In *ACM SC*, 2021.
- [66] Y. Ni, Z. Zheng, X. Lin, F. Gao, X. Zeng, Y. Liu, T. Xu, H. Wang, Z. Zhang, S. Du, et al. Cellfusion: Multipath vehicle-to-cloud video streaming with network coding in the wild. In *ACM SIGCOMM 2023*.
- [67] K. Osawa, S. Li, and T. Hoefler. Pipefisher: Efficient training of large language models using pipelining and fisher information matrices. *MLSys 2023*.
- [68] P. Patarasuk and X. Yuan. Bandwidth optimal all-reduce algorithms for clusters of workstations. *Journal of Parallel and Distributed Computing*, 2009.
- [69] P. Patel, E. Choukse, C. Zhang, Í. Goiri, B. Warriar, N. Mahalingam, and R. Bianchini. Characterizing Power Management Opportunities for LLMs in the Cloud. In *ACM ASPLOS*, 2024.
- [70] P. Patel, E. Choukse, C. Zhang, A. Shah, Í. Goiri, S. Maleki, and R. Bianchini. Splitwise: Efficient generative llm inference using phase splitting. In *ACM/IEEE ISCA 2024*.
- [71] P. Qi, X. Wan, G. Huang, and M. Lin. Zero bubble pipeline parallelism. *arXiv preprint arXiv:2401.10241*, 2023.
- [72] S. Rajasekaran, M. Ghobadi, and A. Akella. {CASSINI}: {Network-Aware} job scheduling in machine learning clusters. In *USENIX NSDI 2024*.
- [73] S. Rajbhandari, O. Ruwase, J. Rasley, S. Smith, and Y. He. Zero-infinity: Breaking the gpu memory wall for extreme scale deep learning. In *ACM SC*, 2021.
- [74] S. Samsi, D. Zhao, J. McDonald, B. Li, A. Michaleas, M. Jones, W. Bergeron, J. Kepner, D. Tiwari, and V. Gadepally. From words to watts: Benchmarking the energy costs of large language model inference. In *IEEE HPEC 2023*.
- [75] A. Shah, V. Chidambaram, M. Cowan, S. Maleki, M. Musuvathi, T. Mytkowicz, J. Nelson, O. Saarikivi, and R. Singh. {TACCL}: Guiding collective algorithm synthesis using communication sketches. In *USENIX NSDI 2023*.
- [76] P. Shahi, A. Mathew, S. Saini, P. Bansode, R. Kasukurthy, D. Agonafer, et al. Assessment of reliability enhancement in high-power cpus and gpus using dynamic direct-to-chip liquid cooling. *Journal of Enhanced Heat Transfer*, 2022.
- [77] S. Shi, X. Zhou, S. Song, X. Wang, Z. Zhu, X. Huang, X. Jiang, F. Zhou, Z. Guo, L. Xie, et al. Towards scalable distributed training of deep learning on public cloud clusters. *MLSys 2021*.
- [78] S. S. Shubha, H. Shen, and A. Iyer. {USHER}: Holistic interference avoidance for resource optimized {ML} inference. In *USENIX OSDI 2024*.
- [79] R. Singh, M. Ghobadi, K.-T. Foerster, M. Filer, and P. Gill. Radwan: rate adaptive wide area network. In *ACM SIGCOMM 2018*.
- [80] A. Spiridonoff, A. Olshevsky, and Y. Paschalidis. Communication-efficient sgd: From local sgd to one-shot averaging. *NeurIPS*, 2021.
- [81] J. Stojkovic, E. Choukse, C. Zhang, I. Goiri, and J. Torrellas. Towards greener llms: Bringing energy-efficiency to the forefront of llm inference. *arXiv preprint arXiv:2403.20306*, 2024.
- [82] J. Stojkovic, C. Zhang, Í. Goiri, J. Torrellas, and E. Choukse. Dynamollm: Designing llm inference clusters for performance and energy efficiency. In *IEEE HPCA*, 2025.
- [83] F. Strati, P. Elvinger, T. Kerimoglu, and A. Klimovic. ML Training with Cloud GPU Shortages: Is Cross-Region the Answer? In *ACM EuroMLSys 2024*.
- [84] F. Strati, X. Ma, and A. Klimovic. Orion: Interference-aware, fine-grained gpu sharing for ml applications. In *ACM EuroSys 2024*.
- [85] F. Strati, Z. Zhang, G. Manos, I. S. Pérez, Q. Hu, T. Chen, B. Buzcu, S. Han, P. Delgado, and A. Klimovic. Sailor: Automating distributed training over dynamic, heterogeneous, and geo-distributed clusters. *arXiv preprint arXiv:2504.17096*, 2025.
- [86] Z. Tang, Y. Wang, X. He, L. Zhang, X. Pan, Q. Wang, R. Zeng, K. Zhao, S. Shi, B. He, et al. Fusionai: Decentralized training and deploying llms with massive consumer-level gpus. *arXiv preprint arXiv:2309.01172*, 2023.

- [87] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. Attention is all you need. In *NeurIPS 2017*.
- [88] B. Wang, Q. Xu, Z. Bian, and Y. You. Tesseract: Parallelize the tensor parallelism efficiently. In *International Conference on Parallel Processing*, 2022.
- [89] J. Wang, Y. Lu, B. Yuan, B. Chen, P. Liang, C. De Sa, C. Re, and C. Zhang. Cocktailsgd: fine-tuning foundation models over 500mbps networks. In *ICML 2023*.
- [90] 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. *NeurIPS*, 2017.
- [91] W. Xiao, S. Ren, Y. Li, Y. Zhang, P. Hou, Z. Li, Y. Feng, W. Lin, and Y. Jia. {AntMan}: Dynamic scaling on {GPU} clusters for deep learning. In *USENIX OSDI 2020*.
- [92] H. Yu, S. Yang, and S. Zhu. Parallel restarted sgd with faster convergence and less communication: Demystifying why model averaging works for deep learning. In *AAAI*, 2019.
- [93] B. Yuan, Y. He, J. Davis, T. Zhang, T. Dao, B. Chen, P. S. Liang, C. Re, and C. Zhang. Decentralized training of foundation models in heterogeneous environments. In *NeurIPS 2022*.
- [94] J. Zhao, J. Liu, C. Zhang, Y. Cui, Y. Jiang, and W. Gong. Mptcp+: Enhancing adaptive http video streaming over multipath. In *2020 IEEE/ACM 28th International Symposium on Quality of Service (IWQoS)*.
- [95] L. Zheng, Z. Li, H. Zhang, Y. Zhuang, Z. Chen, Y. Huang, Y. Wang, Y. Xu, D. Zhuo, E. P. Xing, et al. Alpa: Automating inter-and {Intra-Operator} parallelism for distributed deep learning. In *USENIX OSDI 2022*.
- [96] Y. Zhong, S. Liu, J. Chen, J. Hu, Y. Zhu, X. Liu, X. Jin, and H. Zhang. DistServe: Disaggregating Prefill and Decoding for Goodput-optimized Large Language Model Serving. In *USENIX OSDI 24*.
- [97] S. Zhou, Y. Wu, Z. Ni, X. Zhou, H. Wen, and Y. Zou. Dorefa-net: Training low bitwidth convolutional neural networks with low bitwidth gradients. *arXiv preprint arXiv:1606.06160*, 2016.

A ATLAS lessons for MoE training

In this paper, we considered training for dense models with all parameters active at a time. However, there are Mixture-of-Expert (MoE) models that route tokens to a subset of experts, reducing number of active parameters at a time. Like TP, the bandwidth heavy all-reduce across different experts is in the critical path[43, 59]. Consequently, BUBBLETEA architecture could be well suited for MoE model training where, like TP, EP (Expert Parallelism) is done within the DC and PP is done across DCs. The key lessons from ATLAS are also applicable to MoE training. E.g., MoE training will benefit from speeding WAN transfer latency using multiple TCP connections as well as temporal bandwidth sharing. Additionally, we can improve the GPU utilization using BUBBLETEA. We leave this for future work.