ZERO REDUNDANCY DISTRIBUTED LEARNING WITH DIFFERENTIAL PRIVACY

Zhiqi Bu [1] **Justin Chiu** [1] **Ruixuan Liu** [1] **Sheng Zha** [1] **George Karypis** [1]

ABSTRACT

Deep learning using large models has achieved great success in a wide range of domains. However, training these models on billions of parameters is very challenging in terms of training speed, memory cost, and communication efficiency, especially under the privacy-preserving regime with differential privacy (DP). On the one hand, the efficiency of DP optimization is comparable to that of standard non-DP optimization on a single GPU, but existing DP distributed learning is significantly inefficient on multiple GPUs. On the other hand, the Zero Redundancy Optimizer (ZeRO) is a state-of-the-art solution to the standard distributed learning, which can be technically complicated to work compatibly with DP. In this work, we develop a new systematic solution, DP-ZeRO, (I) to scale up the trainable DP model size, e.g. to GPT-100B, (II) to obtain the same computation and communication efficiency as the standard ZeRO, and (III) to enable mixed-precision DP training. Our DP-ZeRO, like the standard ZeRO, has the potential to train models with arbitrary size and exhibits excellent training efficiency on large models. Code at https://github.com/awslabs/fast-differential-privacy.

1 INTRODUCTION

Recent advances in differentially private (DP) deep learning have witnessed the power of large pre-trained models, achieving comparable accuracy to state-of-the-art non-private models across computer vision (De et al., 2022; Bu et al., 2022a; Mehta et al., 2022; Xie et al., 2018) and natural language processing (Yu et al., 2021a; Li et al., 2021; Bu et al., 2023b). Similar to their non-DP counter-parts, it has been observed that larger DP models tend to have better performance. For example, the DP accuracy increases from 83% using RoBERTa-base (123M parameters) to 86% using RoBERTa-large (354M parameters) on GLUE datasets (Li et al., 2021; Bu et al., 2023b; Yu et al., 2021a); the DP BLEU score increases from 61 using GPT2-small (124M parameters) to 64 using GPT2-large (800M parameters) on E2E dataset (Li et al., 2021; Bu et al., 2023b); a similar trend is also observed using ViT (Base/Large/Huge) up to 600M parameters to achieve state-of-the-art DP accuracy on ImageNet, around 81% at $\epsilon = 8$ (Mehta et al., 2022).

Driven by this success and the surge of computational power, it is high time to enable DP deep learning with billions of trainable parameters, at the same scale of the standard non-DP one, e.g., GPT3-175B (Brown et al., 2020) and LLaMA-63B (Touvron et al., 2023a;b). As a consequence, we will select and merge two of the most efficient algorithms – one for (non-DP) distributed solution, and the other for DP algorithm – to build our DP distributed system. Specifically, such a system must have high speed and memory efficiency, low communication cost, and compatibility with general model architectures.

For small to medium models (e.g. with $< 1B$ parameters) that fit within the memory of a single GPU, a range of DP algorithms are feasible, producing the same result at different efficiency. Examples include TensorFlow-privacy (Subramani et al., 2021), Opacus (Yousefpour et al., 2021; Bu et al., 2022b), ghost clipping (GhostClip) (Goodfellow, 2015; Li et al., 2021; Bu et al., 2022a), and Book-Keeping (BK) (Bu et al., 2023c). Specifically, the BK algorithm allows DP optimization to be almost as efficient as the standard one: the time/space complexity of BK algorithm is $1.08\times$ /$1.05\times$ of the standard optimization on ViT-Large (300M parameters, 147 layers) and $1.03\times$ /$1.01\times$ on GPT2-large (800M parameters, 220 layers). We will introduce two key techniques of the BK algorithm in Appendix B, which is also leveraged in our DP-ZeRO solution.

To enable the DP distributed learning for these not-too-large models, one can directly use distributed data parallelism (DDP, (Li et al.)), where each mini-batch of data is partitioned to smaller micro-batches and each GPU computes one micro-batch with a full copy of the DP model. A line of researches (Yousefpour et al., 2021; De et al., 2022; Kurakin et al., 2022) have reported that DDP with DP usually either incurs huge memory cost due to caching the per-sample gradients, or suffers from $2-9\times$ slower training speed than non-DP optimization (De et al., 2022; Bu et al., 2021). We note these attempts have not used more advanced DP algo-

[1]Amazon. Correspondence to: Zhiqi Bu <zhiqibu@amazon.com>.

Proceedings of the 9^{th} MLSys Conference, Bellevue, WA, USA, 2026. Copyright 2026 by the author(s).

Table 1. Summary of selected DP distributed learning.

Distributed solution	Parallelism	Model sharding	Standard version	DP version	Remark
DDP	Data	No	(Li et al.)	(Yousefpour et al., 2021)	unable to fit large model and DP is memory costly
DDP	Data	No	(Frostig et al., 2018)	(De et al., 2022)	unable to fit large model and DP is slow
GPipe	Pipeline	Yes	(Huang et al., 2019)	(He et al., 2022)	pipeline bubble wastes GPU time
ZeRO	Data(&Model)	Yes	(Rajbhandari et al., 2020)	**Ours**	time & memory & communication efficient

rithms, such as BK (Bu et al., 2023c), which can effectively address the efficiency issues. However, the feasibility issue remains insurmountable because DDP cannot train models that exceed the capacity of one GPU.

As the model size further increases beyond a reasonable bound for one GPU (e.g. 40GB memory, which roughly translates to 2B model training with Adam), the model must be partitioned by the pipeline parallelism or the model parallelism, so that each GPU only holds a partial shard of the model (see Figure 2). In (He et al., 2022), DP is combined with pipeline parallelism to fine-tune about 0.1% of GPT3-175B. Yet, the pipeline parallelism can be inefficient due to an issue known as the pipeline bubble, where GPUs are idle while waiting for data to process.

Generally speaking, there are algorithmic challenges to adapt more advanced distributed methods, such as Zero Redundancy Optimizer (Rajbhandari et al., 2020) (**ZeRO**) and mixed-precision training, to DP optimization. These challenges include the incompatibility of per-sample clipping and loss scaling, the interference with inter-GPU communication from DP clipping style, and the conflict between DP modular hooks and ZeRO tensor hooks.

In this work, we develop DP-ZeRO, equipping state-of-the-art distributed learning solution with DP algorithms (see comparison in Table 1), while preserving model accuracy on any number of GPUs (see Appendix F) without altering the mathematics of DP optimization. We summarize our contributions as follows.

1. We propose the zero redundancy distributed learning with differential privacy (DP-ZeRO), achieving the same level of **communication efficiency**, **computation efficiency (speed and memory)**, and **scalability** (e.g. to 100B parameters and 256 GPUs) as the standard ZeRO.

2. One core design of DP-ZeRO is the **locality of DP back-propagation**, which analyzes the interplay between the mathematical and hardware gradient partition, so as to enable DP without interfering the communication orchestra. We overcome this engineering challenge by rewriting the back-propagation of each layer (see Section 5).

3. We propose **DP mixed-precision (bf16 and fp16) training**, specifically not using the loss scaling trick. We show that loss scaling can conflict with DP training even though it is crucial in standard training. This advance allows us to reduce the memory cost or training time by roughly 50%, and enjoy much faster communication in DP distributed learning.

4. We enable DP deep learning with >1B trainable parameters for the first time, including the full GPT2-XL, ViT-Gigantic, ViT-10B and GPT-100B with DP.

5. To support the research community, we will open-source our codebase that automatically applies DP-ZeRO, with one line of code change.

2 PRELIMINARY

2.1 Differential privacy

DP provides a formal privacy guarantee, making it difficult to extract any information from training data. The privacy guarantee is characterized by (ϵ, δ)-DP in Definition 2.1, with smaller (ϵ, δ) representing lower privacy risk.

Definition 2.1 ((Dwork et al., 2006))**.** A randomized algorithm M is (ε, δ)-DP if, for any two neighboring datasets S, S' that differ by one sample and for any event E, we have $\mathbb{P}[M(S) \in E] \leqslant \mathrm{e}^{\varepsilon}\mathbb{P}\left[M\left(S'\right) \in E\right] + \delta$.

In DP deep learning, the privacy is provided by the Gaussian mechanism, which obfuscates the gradient and consequently restricts the influence of any sample onto the model parameters (see more details in Appendix A).

To be specific, the gradients are privatized by post-processing through *per-sample gradient clipping* and *Gaussian noising*: for the m-th group,

$$\mathbf{G}_{[m]} := \sum_i c_{[m],i} \cdot \boldsymbol{g}_{[m],i} + \sigma_{\text{DP}} \cdot \mathcal{N}(0, \mathbf{I}). \tag{1}$$

Here the gradient of all trainable parameters is partitioned into M groups, i.e. $\boldsymbol{g}_{[m],i}$ is the i-th per-sample gradient of the m-th group's parameters, where $m \in \{1 \cdots M\}$ is the group index. c_i is the per-sample gradient clipping factor so that $\|c_{[m],i}\boldsymbol{g}_{[m],i}\| \leq 1/\sqrt{M}$. That is, DP optimization uses the standard optimizers such as stochastic gradient descent (SGD) and Adam (Kingma & Ba, 2014) to update the trainable parameters with the private gradient, instead of the standard gradient $\sum_i \boldsymbol{g}_i$.

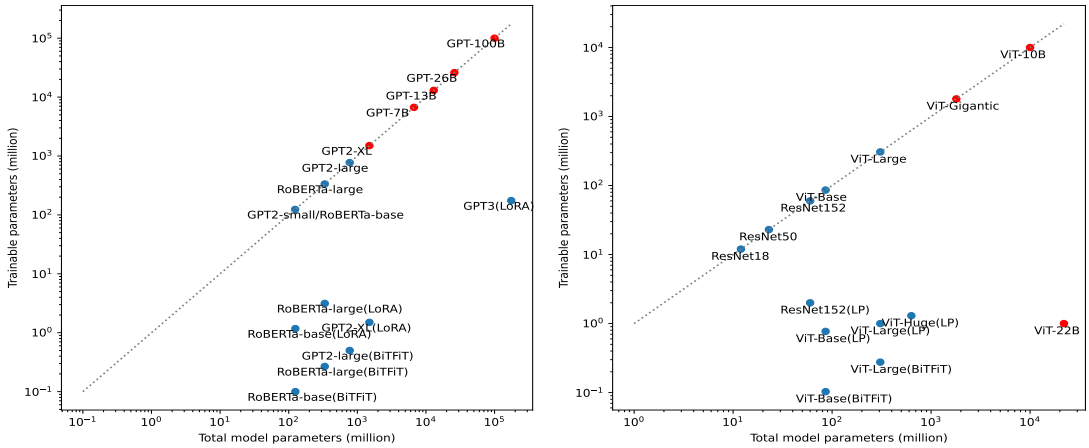

Figure 1. Total/trainable parameters of existing DP models (blue) and ours by DP-ZeRO (red).

Mathematical gradient partition In (1), the trainable parameters and their gradients are mathematically partitioned into M groups, e.g. in all-layer clipping, all parameters form one group ($M = 1$) (Abadi et al., 2016); in layer-wise clipping (McMahan et al., 2018; He et al., 2022), each layer's parameters form a group ($M =$ number of layers). Empirical evidence and theoretical analysis show that different partitions have the same training speed, though a finer partition (e.g. layer-wise) has lighter memory footprint. However, in our DP distributed system, we highlight in Section 3 that the mathematical partition must be subject to the hardware constraints.

Per-sample gradient clipping In (1), a number of clipping functions $c_i = C(\|\boldsymbol{g}_i\|; R)$ are available. Most works (Abadi et al., 2016; Li et al., 2021; Yu et al., 2021a) use the vanilla clipping $c_i = \min(R/\|\boldsymbol{g}_i\|, 1)$. Recently, (Bu et al., 2023b; Yang et al., 2022) advocate the automatic clipping $c_i = 1/(\|\boldsymbol{g}_i\| + 0.01)$ which is hyperparameter-free and comparably accurate. Note if $c_i \equiv 1$, then the clipped gradient reduces to the standard gradient. The main overhead of DP optimization is the computation of per-sample gradient norms. On a single GPU, the mixed ghost clipping (Bu et al., 2023c) has reduced the time complexity to $< 10\%$ on large models like GPT2.

Privacy accounting In (1), adding Gaussian noise to the clipped gradient protects the privacy that is quantifiable by the privacy accounting theory (Abadi et al., 2016; Bu et al., 2020; Dong et al., 2019; Zhu et al., 2022; Gopi et al., 2021; Koskela et al., 2020). The privacy guarantee is increasing in the noise level σ_{DP}, independent of R_m, learning rate, clipping function and model architectures, with $\sigma_{\text{DP}} = 0$ leading to $\epsilon = \infty$ (non-private).

2.2 Zero Redundancy Optimizer (ZeRO)

2.2.1 *Parallel computing*

Parallel computing is necessary to train large-scale models and is critical to the optimization efficiency. For models that fit in a single GPU, data parallelism (DataP) can be used to speed up the training by partition the mini-batch of samples into multiple micro-batches. Then, each GPU (holding a full copy of parameters) executes the forward and backward propagation of one micro-batch, from which the parameter gradients are generated and averaged across GPUs to update the trainable parameters. However, for models that do not fit in a single GPU, the model parameters need to be sharded by alternative solutions such as ZeRO (Rajbhandari et al., 2020), model parallelism (ModelP) and pipeline parallelism (PipeP).

ModelP partitions a model vertically, e.g. using 3 GPUs to store the parameters of one layer. As a consequence, ModelP does not scale efficiently beyond a single node due to fine-grained computation and expensive communication between layers. Implementation-wise, ModelP frameworks usually require heavy code integration that may not be generalizable in model architectures. In contrast, PipeP partitions a model horizontally across layers, e.g. storing 3 layers in each GPU. Each GPU deals with all micro-batches sequentially, though PipeP can be inefficient due to the pipeline bubble, which is overcome by ZeRO (Rajbhandari et al., 2020). ZeRO is an advanced data parallel method that eliminates memory redundancies during the training, and improves the training speed and communication volume proportionally to the number of GPUs. Unlike basic DataP, ZeRO partitions a model's states across GPUs and gather/reduce in a just-in-time manner, thus sustaining the high efficiency of very large model training. Notice that ZeRO can work compatibly with ModelP and optionally offload the model states to CPUs (Ren et al., 2021; Rajbhandari et al., 2021).

2.2.2 *Model state partition*

A major part of the training memory is consumed by the model states[1]. ZeRO has three stages (ZeRO1/2/3) that partition these model states by different levels, with lower level being faster but more memory costly. For instance, in Table 2 of (Rajbhandari et al., 2020), ZeRO1/2/3 at most train 7.6/14.4/128B models on 64 V100 GPUs.

We take an example of mixed-precision Adam optimizer to train a model with Ψ_{model} parameters, which maintains a master copy (fp32) of optimizer states, and the half-precision (fp16/bf16) parameters and gradients.

Optimizer state partition The optimizer states are the (master) parameters, variance and momentum, each taking $4\Psi_{\text{model}}$ memory. ZeRO1 only applies the optimizer state partition and updates the parameters locally, reducing the memory cost of model states from $16\Psi_{\text{model}}$ for basic DataP to $(4+\frac{12}{N_d})\Psi_{\text{model}}$ at each of N_d GPUs.

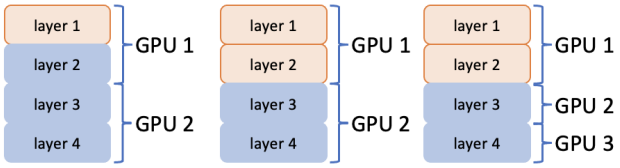

Figure 2. Mathematical partition (left two, same GPU allocation but different convergence) and hardware partition (right two, different GPU allocation but same convergence). Orange and blue are gradient groups $\boldsymbol{g}_{[1]}$ and $\boldsymbol{g}_{[2]}$ in DP optimization (1).

Hardware gradient partition In addition to ZeRO1, during the back-propagation, ZeRO2 (and ZeRO3) further partitions the $2\Psi_{\text{model}}$ gradients into different GPUs, reducing the memory cost to $(2+\frac{14}{N_d})\Psi_{\text{model}}$ at each GPU. Figure 2 illustrates the difference between the hardware partition and the mathematical partition in Section 2.1.

Parameter partition In addition to ZeRO2, ZeRO3 also partitions the $2\Psi_{\text{model}}$ fp16 parameters, further reducing the memory cost to $\frac{16}{N_d}\Psi_{\text{model}}$ at each GPU.

2.3 Mixed-precision training

Mixed-precision training (Micikevicius et al., 2018) performs the forward and backward propagation on the half precision (fp16 or bf16) parameters, activations, and gradients, while performing the model update in full precision (fp32). Compared to the full-precision training, it is capable of saving the memory by $\approx 50\%$ and accelerating the computation by $\approx 2\times$. We note that fp16 has better precision but a more limited range, and vice versa for bf16. Therefore, *loss scaling* is necessary when using fp16 to prevent small gradients from being rounded to zero, and to preserve the accuracy (see Appendix D for details). In contrast, bf16 usually does not need loss scaling since it has the same range as fp32, and it is dominantly popular in large model training. However, bf16 is not as widely supported as fp16, e.g. only available on NVIDIA Ampere GPUs or above (Nvidia).

[1]Another important part of memory consumption is the batch-size-related variables such as the activation tensors, which is instantiated during the forward propagation and independent to DP (which modifies the gradient during the back-propagation).

3 Differentially private ZeRO

3.1 Locality of DP back-propagation

Our DP-ZeRO algorithm introduces the per-sample gradient clipping and noising to the standard ZeRO (Rajbhandari et al., 2020), while maintaining the efficiency. At high level, an iteration of ZeRO consists of the following steps:

$$\Big(\text{all-gather} \to \text{forward} \to \text{all-gather} \to \text{backward} \to \text{reduce}\Big)^{\times L \text{ layers}} \to \text{update(SGD/Adam/...)}$$

The purple operations are **global** and require communication among GPUs, whereas the green operations are **local** and computed within each GPU. To enable DP on different distributed systems, we insist that the DP operations do not interfere with the complicated communication and hence DP back-propagation (red) must be computed locally, similar to the standard back-propagation (black):

$$\begin{aligned}\text{backward} =&(\text{output gradient} \to \text{clipping factor}\\ &\to \text{parameter gradient} \to \text{noising}).\end{aligned}$$

To achieve this, we ensure that the mathematical gradient partition is more fine-grained than the hardware partition[2], which requires the engineering efforts to rewrite the back-propagation for each layer (see details in Section 5.3).

3.2 Overall algorithm

To give more details, we consider training a neural network of linear layers using N_d GPUs. We emphasize that the following procedure is sufficiently generic to cover other layer types, such as convolution, embedding, normalization, and so on, which are all supported by DP-ZeRO. The full algorithm is depicted in Figure 3, where we denote the j-th micro-batched variables like $\boldsymbol{a}_l^{(j)}$, for $1 \le j \le N_d$.

The forward propagation of DP optimization is the same as that of the standard optimization:

$$\boldsymbol{s}_l = \boldsymbol{a}_l \mathbf{W}_l + \mathbf{b}_l, \boldsymbol{a}_{l+1} = \phi_l(\boldsymbol{s}_l). \tag{2}$$

At the l-th layer, $\boldsymbol{a}_l \in \mathbb{R}^{BT_l d_l}$ is the layer's input, also known as the activation, $\boldsymbol{s}_l \in \mathbb{R}^{BT_l p_l}$ is the layer's output, $\mathbf{W}_l \in \mathbb{R}^{d_l p_l}$ is the weight, $\mathbf{b}_l \in \mathbb{R}^{p_l}$ is the bias, and ϕ_l is any inter-layer operation such as ReLU, tanh or pooling. We

[2]The interplay between these two partitions has fundamental impact on the communication and accuracy (see Section 5.1)

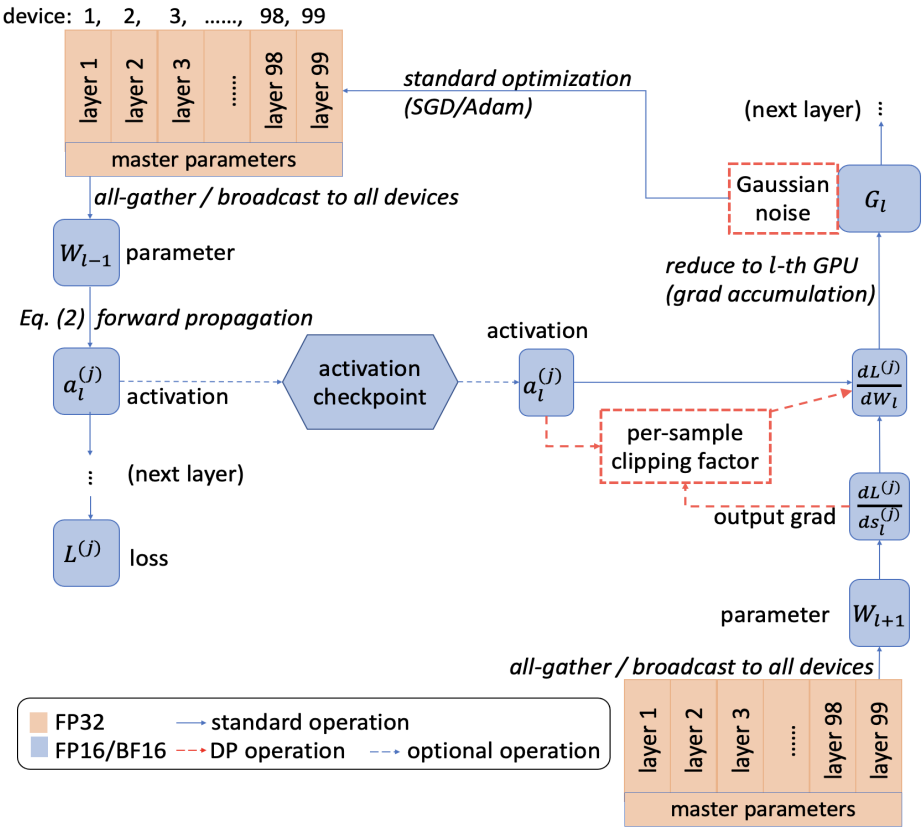

Figure 3. Algorithm of DP-ZeRO with mixed-precision training. We note that the per-sample clipping can be implemented by rewriting the back-propagation on module levels, and that the Gaussian noise is added after the gradient is calculated and before reduction.

denote B as the physical micro-batch size[3] and T_l as the hidden feature dimension (e.g. sentence length or number of pixels). During the forward propagation, the activations $\{\boldsymbol{a}_l\}$ are computed and stored in the computation graph, and the loss $L = \sum_i L_i$ is derived, where L_i is the per-sample losses. During the back-propagation, the output gradient is first computed based on the previous layer.

$$\frac{\partial L}{\partial \boldsymbol{s}_l} = \frac{\partial L}{\partial \boldsymbol{s}_{l+1}} \frac{\partial \boldsymbol{s}_{l+1}}{\partial \boldsymbol{a}_{l+1}} \circ \frac{\partial \boldsymbol{a}_{l+1}}{\partial \boldsymbol{s}_l} = \frac{\partial L}{\partial \boldsymbol{s}_{l+1}} \mathbf{W}_{l+1} \circ \phi_l'(\boldsymbol{s}_l),$$

in which $\circ$ is element-wise multiplication. Specifically, the use of parameter $\mathbf{W}_{l+1}$ necessitates the all-gather operation when the model is partitioned into multiple GPUs, which is not needed in single GPU training. Next, the activation $\boldsymbol{a}_l$ is used together with $\frac{\partial L}{\partial \boldsymbol{s}_l}$ to compute the parameter gradient:

$$\begin{aligned}\text{DP gradient: } & \frac{\partial \sum_i c_i L_i}{\partial \mathbf{W}_l} + \sigma \mathcal{N}(0, \mathbf{I}) \\ =& \boldsymbol{a}_l^\top \operatorname{diag}(c_1, \cdots, c_B) \frac{\partial L}{\partial \boldsymbol{s}_l} + \sigma \mathcal{N}(0, \mathbf{I}).\end{aligned}$$

[3]The micro-batch size B is the number of samples processed by each GPU, which determines the time and memory efficiency, but not the performance. The logical batch size that determines the performance is $B \times N_d \times$ gradient_accumulation_steps.

Note the standard gradient can be viewed as $c_i = 1, \sigma = 0$. Here the per-sample gradient norm (or the clipping factor c_i) can be computed at small cost, as we have discussed in Section 1.

3.3 Time efficiency of DP-ZeRO

The time efficiency of DP-ZeRO consists of two parts: the local computation (including forward and backward propagation) and the global communication (including intra-node and inter-node communication). Given that the only difference between DP-ZeRO and ZeRO is the back-propagation, we claim that DP-ZeRO could enjoy high efficiency on-par with the standard ZeRO when (I) DP back-propagation exhibits a time efficiency comparable to the standard, similar to the single GPU training, and/or (II) the time efficiency of the parts other than back-propagation is not insignificant. We give the time of each part of DP-ZeRO in (3) to illustrate our claim.

$$\begin{aligned}&\frac{\text{DP-ZeRO speed}}{\text{standard ZeRO speed}} \\ =&\frac{\text{standard back-prop} + \text{forward prop} + \text{communication}}{\text{DP back-prop} + \text{forward prop} + \text{communication}}\end{aligned} \tag{3}$$

To be explicit, we summarize the time complexity in Table 2 and refer to Appendix C for details.

In what follows, we analyze the absolute and relative speed (to standard ZeRO) of DP-ZeRO under important settings.

3.3.1 *Number of computation devices*

When scaling from one GPU (zero communication) to one node (multiple GPUs) and to multiple nodes, the communication efficiency decreases sub-linearly (↑communication). On a single node, multiple GPUs can communicate using the high-speed intra-node connections such as NVLink/NVSwitch (Foley & Danskin, 2017; Ishii et al., 2018). On multiple nodes, which are necessary for large models, the inter-node connections are $3 \sim 24\times$ slower than the intra-node connections (Li et al., 2019; Zhang et al., 2022). In short, DP-ZeRO can be as fast as ZeRO by (3) when multiple nodes are employed.

3.3.2 *Memory-efficient distributed learning*

The communication volume is specific to different distributed algorithms, most of which trade the communication or speed for memory, in order to feasibly train very large models. For example, ZeRO3 (but not ZeRO1/2) needs to all-gather the sharded parameters at each iteration, hence suffering from 50% extra communication volume (↑communication). Another example is the activation check-pointing (also known as gradient check-pointing (Chen et al., 2016)), where a second forward propagation re-computes the expensive activations during back-propagation, though at a 33% slower speed (↑forward propagation). These techniques improve the relative speed of DP-ZeRO but worsens the absolute speed.

3.3.3 *Parameter efficient fine-tuning*

Parameter efficient fine-tuning (PEFT), such as LoRA (Hu et al., 2021), Adapter (Houlsby et al., 2019), and BiTFiT (Zaken et al., 2022), optimizes a small fraction (e.g. $\Psi_{\text{train}} = 0.1\%\Psi_{\text{model}}$) of model parameters and thus boosts the efficiency of back-propagation and communication (↓DP back-propagation ↓back-propagation ↓communication). Consequently, (I) the communication volume of the gradient can be reduced possibly by $1000\times$; (II) the local computation can accelerate by 50% (Hu et al., 2021; Bu et al., 2022b), which can be seen by treating Ψ_{train} in Table 2 as almost zero; (III) the memory cost is saved on the non-trainable layers, which translates to larger batch size and faster computation. Hence, both relative and absolute speed of DP-ZeRO improve using PEFT.

[5]Here 0.666 is figurative and dependent on the DP mechanisms. For example, if we denote d, p as a layer's input/output dimensions, then the complexity of BK algorithm (Bu et al., 2023c) is $6BTpd + 2BT^2(p+d) \cdot \mathbb{I}(2T^2 < pd)$, which is non-linear in T but empirically the overall slowdown could be $0.666/6 \approx 11\%$.

3.4 Memory efficiency of DP-ZeRO

We claim that DP-ZeRO is as memory efficient as the standard ZeRO, similar to the single GPU training, when we use (I) the mixed ghost norm trick (Bu et al., 2022a; 2023c), instead of GhostClip (Goodfellow, 2015; Li et al., 2021) or per-sample gradient instantiation (Yousefpour et al., 2021); (II) the layer-wise clipping style instead of the all-layer clipping, so that the book-keeping (Bu et al., 2023c) does not store all output gradients; (III) a large number of GPUs so that the micro-batch size B (i.e. per-GPU batch size) is small: specially, when $B = 1$ in the gradient accumulation, the per-sample gradient is free. We empirically verify our claim in Figures 4 and 7.

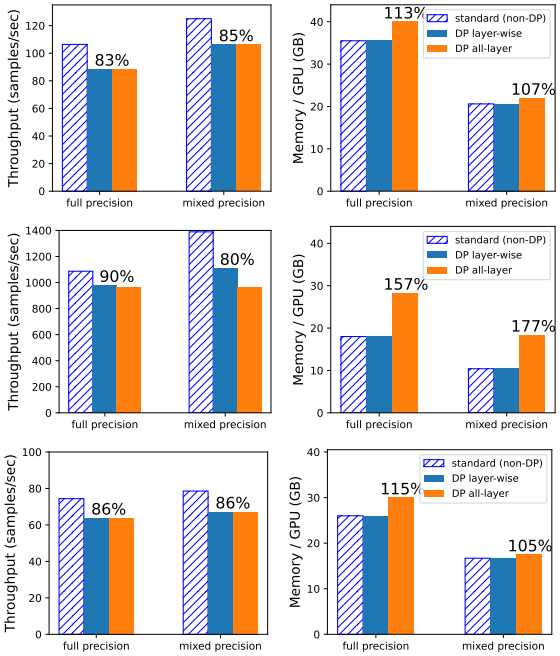

Figure 4. Efficiency on ViT-Gigantic (top row, 1.8B), ResNet152 (middle row) and GPT2-XL (bottom row, 1.5B).

3.5 Mixed-precision training with DP

We now analyze the intricacy in mixed-precision training with DP, which is not unique to DP-ZeRO but present in the general DP optimization. We emphasize that *the per-sample gradient clipping already plays the role of scaling*, and hence DP mixed-precision training must not use loss scaling, as illustrated in Table 3. Specifically, in standard mixed-precision training, there are two steps of scaling: (I) scaling up the loss L_i by $10^3 \sim 10^9$ (and consequently the output gradient $\frac{\partial L_i}{\partial \boldsymbol{s}}$ as well as $\frac{\partial L_i}{\partial \mathbf{W}}$) before the back-propagation, to prevent the underflow where fp16 gradient is too small to be distinguished from 0, and (II) scaling

Table 2. Time complexity (by float-point operations) of one iteration in distributed learning[5]. We denote Ψ_{train} to be the number of trainable parameters ($\Psi_{\text{train}} = \Psi_{\text{model}}$ in full parameter training; $\Psi_{\text{train}} \ll \Psi_{\text{model}}$ when most parameters are frozen, e.g. in parameter efficient fine-tuning), and we have defined B, T below (2).

	forward propagation		back-propagation			communication
	activation	attention	output grad	param grad	DP clipping& noise	
complexity	$2BT\Psi_{\text{model}}$	$O(BT^2)$	$2BT\Psi_{\text{model}}$	$2BT\Psi_{\text{train}}$	$0.666BT\Psi_{\text{train}}$	$O(\Psi_{\text{model}})$

Table 3. Illustration of overflow and underflow issues during mixed-precision training (ghost norm).

loss scale=10^3	activation $\boldsymbol{a}_l$	output grad $\frac{\partial L}{\partial \boldsymbol{s}_l}$ (scaled)	per-sample grad norm		clipping factor	param grad (not scaled down)	param grad (if scaled down)
			$\text{vec}(\boldsymbol{a}_l \boldsymbol{a}_l^\top)$	$\text{vec}(\frac{\partial L}{\partial \boldsymbol{s}_l} \frac{\partial L}{\partial \boldsymbol{s}_l}^\top)$			
standard w/o scaling	$10^{-3} \sim 10^2$	$10^{-8} \sim 10^1$	N/A	N/A	1	$10^{-7} \sim 10^1$	$10^{-7} \sim 10^1$
standard w/ scaling	$10^{-3} \sim 10^2$	$10^{-5} \sim 10^4$	N/A	N/A	1	$10^{-4} \sim 10^4$	$10^{-7} \sim 10^1$
DP w/o scaling	$10^{-3} \sim 10^2$	$10^{-8} \sim 10^1$	$10^2 \sim 10^3$	$10^{-6} \sim 10^0$	$10^{-3} \sim 10^2$	$10^{-7} \sim 10^1$	$10^{-7} \sim 10^1$
DP w/ scaling	$10^{-3} \sim 10^2$	$10^{-5} \sim 10^4$	$10^2 \sim 10^3$	$10^0 \sim 10^6$	$10^{-6} \sim 10^{-1}$	$10^{-7} \sim 10^1$	$10^{-10} \sim 10^{-2}$

down the parameter gradient $\frac{\partial L_i}{\partial \mathbf{W}}$, by the same factor, after the back-propagation to recover the correct magnitude of gradient. However, in DP mixed-precision training, scaling up the loss may cause overflow, while scaling down the gradient incorrectly over-shrinks the gradient and worsens the performance. See Figure 5 for a real example. We dedicate Appendix D to explain this intricacy step-by-step, where we also discuss the *floating-point vulnerability* of DP that is empirically invalidated by (Bu et al., 2024; Yu et al., 2021b; Li et al., 2024).

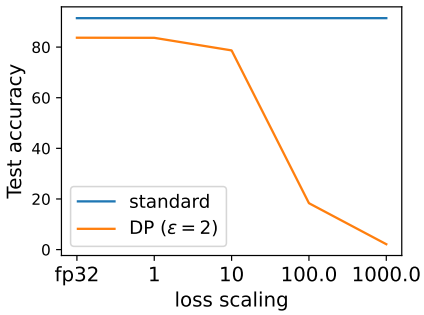

Figure 5. Accuracy of mixed-precision training with loss scaling. ViT-large on CIFAR100.

4 Empirical performance of DP-ZeRO

We evaluate DP-ZeRO on five aspects[6] : model architectures, efficiency, scalability, compatibility with various distributed learning and DP techniques. We use DP-ZeRO to train ResNet (He et al., 2016), ViT (Dosovitskiy et al., 2020; Zhai et al., 2022) and GPT (Radford et al., 2019; Brown et al., 2020), which are workhorses in computer vision and language tasks. We measure the time and memory efficiency of DP-ZeRO under settings such as PEFT and multiple precision formats (fp32 or fp16/bf16). We evaluate the scalability of DP-ZeRO in terms of number of GPUs and number of model parameters. Our experiments scale from single node (8 GPUs) to multiple nodes, up to 256 GPUs, and train models up to 100B trainable parameters. Moreover, DP-ZeRO is compatible with mainstream implementations of ZeRO[7] and with different clipping styles, clipping functions, privacy accountants, and so on. We leave the experimental details in Appendix E. By default, we use AdamW, mixed-precision training, layer-wise clipping style, $B = 4$, and A100 GPU with 40GB memory, unless otherwise stated.

4.1 Generality of DP-ZeRO

DP-ZeRO is generally applicable to different neural network architectures, clipping styles, and precision formats. We test DP-ZeRO1 on single node and observe that different clipping styles are equally fast, but layer-wise clipping is more memory efficient than all-layer clipping. Comparing to the standard ZeRO, our DP-ZeRO enjoys almost the same speed and memory efficiency, while the gap will be further closed as we move to more advanced stages of ZeRO.

4.2 Lighter training of DP-ZeRO

DP-ZeRO can employ low-memory optimizers and train on fewer parameters, therefore vastly reducing the memory and communication cost. On a single node, we demonstrate that DP-ZeRO actually benefits (more than standard ZeRO and single-GPU training) from lighter training, following from

[6]We focus on the system design and less so on the accuracy (except two experiments in Appendix F for ViT and GPT to showcase the consistency of performance across systems and to validate our implementation), because the accuracy depends on hyperparameters like learning rate and number of epochs that has no impact on the efficiency, similar to ZeRO (Rajbhandari et al., 2020) and FSDP (Zhao et al.).

[7]DP-ZeRO is implemented on DeepSpeed (supporting ZeRO1/2/3), FSDP (Zhao et al.) (supporting ZeRO3), MiCS (Zhang et al., 2022) (supporting ZeRO2/3), and any distributed optimizers supported on these libraries.

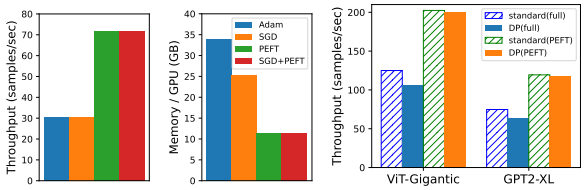

Figure 6. Efficiency of DP-ZeRO with lighter training. Left 2: $> 2\times$ speedup with lighter training; ViT-5B, $B = 1$. Right: DP-ZeRO benefits more than ZeRO from PEFT.

our discussion in Section 3.3.

Low-memory optimizers Low-memory optimizers can boost the training efficiency at the cost of accuracy degradation. For example, SGD requires only 1/3 optimizer states of Adam and significantly saves the memory; 1-bit Adam (Tang et al., 2021) and signSGD compress the gradient and reduce the communication volume up to $32\times$.

Fewer trainable parameters In the fine-tuning phase, as we analyzed in Section 3.3.3, PEFT improves both the local computation and the communication volume. Hence, DP-ZeRO allows PEFT on ViT and GPT to be $\approx 2\times$ faster than full fine-tuning, whereas the single GPU acceleration is $\leq 1.5\times$.

Remark 4.1. We leverage DP-ZeRO3 with SGD to train ViT-10B (full parameters) and ViT-22B (PEFT; 1M trainable parameters) on one node. See Appendix E for configurations.

4.3 Three stages of DP-ZeRO

DP-ZeRO supports all stages of ZeRO under different implementations including DeepSpeed (default) and FSDP. In Figure 7, the efficiency of DP-ZeRO catches up with the standard ZeRO when we move up the stages. For instance of ViT-Gigantic, the throughput increases from 83% by DP-ZeRO1 to 95-97% by DP-ZeRO3. Following (3), we can attribute the relatively fast training of DP-ZeRO to the increase cost of communication, especially in DP-ZeRO3. Additionally, we observe that the throughput of DP-ZeRO1/2 improves to over 95% on 4 nodes, as predicted by Section 3.3.1. Notice that we save DP-ZeRO3 of GPT to Section 4.4 on super-large scale.

4.4 Scalability of DP-ZeRO

We evaluate the scalability of DP-ZeRO3 in terms of large sequence length (2048), large model size ($7 \sim 100$B), and large number of GPUs (up to 256). We use A100 with 80GB memory, as well as the activation check-pointing and ModelP.

In Figure 8 (left), we observe that for a fixed model with 26B trainable parameters, DP-ZeRO is super-linearly scalable to the number of GPUs, achieving $> 95\%$ speed of the standard ZeRO. Here super-linearity is a property of ZeRO (see Figure 3 in (Rajbhandari et al., 2020)) which allows more GPUs to shard the model states (and reduce the per-GPU memory cost) more aggressively, and to train faster since the micro-batch size is larger. Furthermore, in Figure 8 (right), for a fixed number of GPUs, DP-ZeRO is linearly scalable to the model size, achieving the same speed as the standard ZeRO. In short, DP-ZeRO is almost equal to ZeRO in terms of training efficiency in super-large scale.

Remark 4.2. In comparison to DP-ZeRO, DataP (DP or standard) at most fits 5B models (Rajbhandari et al., 2020) in 80GB memory, regardless of the number of GPUs. We cannot compare to DP-PipeP in (He et al., 2022) because the codebase and experiment details (e.g. number of trainable parameters and sequence length) are not publicly available. We believe it suffices to demonstrate the usefulness of DP-ZeRO by comparing ZeRO to PipeP, since DP-ZeRO resembles the efficiency of standard ZeRO.

5 Codebase Design

5.1 Interplay between mathematical partition and hardware partition

In distributed learning where the gradient tensor is too large and must be partitioned on different GPUs, we advocate to avoid the expensive communication between GPUs. The rule of thumb is to enforce that the mathematical partition to be more fine-grained than the hardware partition:

> *Each mathematical group is only materialized on one GPU, so that per-sample gradient clipping does not require cross-group communication among GPUs.*

If we only have one GPU (or multiple GPUs but each holds a full copy of gradient, so no partition, like in ZeRO1/2), then any clipping style (all-layer or layer-wise) satisfies this rule; if we have 10 GPUs and a 10-layer network, and each GPU holds one layer, then all-layer clipping violates this rule. To be more specific, after back-propagating the 10-th layer, its gradient cannot be clipped at the moment because the clipping factor also needs the 1st layer's gradient norm, which is not available yet. Recall in Figure 2, only the middle case satisfies this rule.

In fact, it may be beneficial to use the most fine-gradient clipping style, the layer-wise one, so that this rule is always enforced while we can be agnostic of the hardware partition. Additionally, layer-wise clipping may have worse accuracy than all-layer clipping, but empirically the gap is negligible for large models (Bu et al., 2023a).

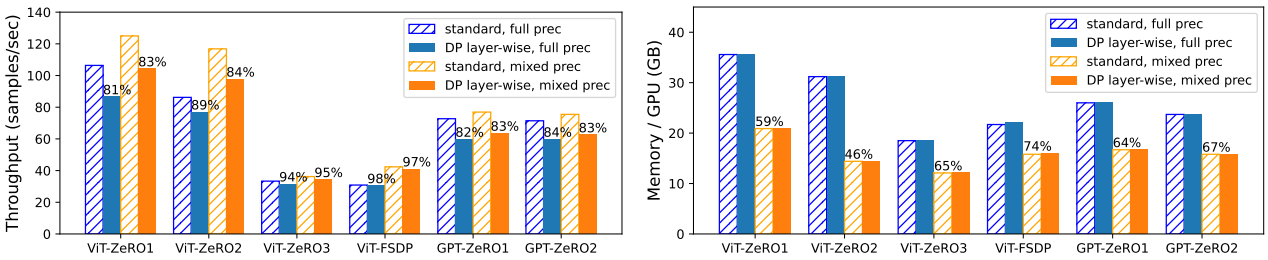

Figure 7. Efficiency of DP-ZeRO on ViT-Gigantic and GPT2-XL under different implementations.

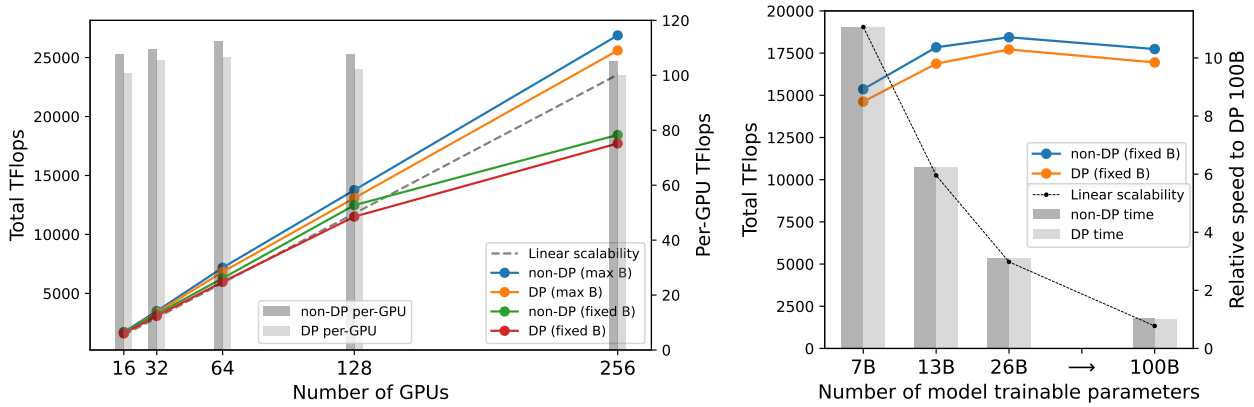

Figure 8. Scalability of DP-ZeRO3 on 26B model (left) and 256 GPUs (right). The colored lines match the left y-axis. The grey bars match the right y-axis. *max B* means we use the maximum micro-batch in each GPU and *fixed B* means $B = 2$.

5.2 DP optimization with forward & backward hooks

Hooks[8] are important functions to enrich the deep learning optimization. To be specific, there are

1. forward modular hook (nn.register_forward_hook),
2. backward modular hook (nn.register_backward_hook),
3. backward tensor hook (tensor.register_hook).

DP libraries including Opacus (Yousefpour et al., 2021), Private-transformers (Li et al., 2021), Private-Vision (Bu et al., 2022a), FastDP (Bu et al., 2023c;b), FastGradClip (Lee & Kifer, 2021) and so on, use modular hooks to modify the standard optimization. However, ZeRO libraries including DeepSpeed and FSDP use tensor hooks. This difference in the types of hooks and many other differences (e.g. both ZeRO libraries and DP libraries modify the optimizer's step function) cause non-trivial problems when combining DP with ZeRO. For example, to keep DP optimization as efficient as the standard, it is necessary to not waste time on computing the non-private gradient. However, if we skip such computation, then ZeRO's tensor hook will not be triggered and the corresponding distributed-learning-related operations cannot carry on. For another example, because DP and ZeRO add different types of hooks, the number of hooks is larger than either optimization and they slows down the training: consider an 100-layer network, each layer with weight and bias (2 tensors), then DP-ZeRO in this subsection needs 100 modular hooks and 200 tensor hooks, adding to a total of 300 hooks. In addition, the Book-Keeping algorithm (in FastDP) in its original form cannot be implemented together with ZeRO3, because all model states are partitioned including the output gradients which are meant to be book-kept. To work around this requires rewriting the distributed solution's communication mechanism, and if successful, still requiring additional communication cost during the second back-propagation. Similar problems are present for Opacus and FastGradClip, which instantiates per-sample gradients that will be partitioned in ZeRO2/3 and requires additional communication cost when gathered to create the privatized gradient.

As a consequence, the hooks are fully supported on DP-ZeRO1 and partially supported on DP-ZeRO2/3 under the layer-wise clipping.

[8] See https://pytorch.org/tutorials/beginner/former_torchies/nnft_tutorial.html#forward-and-backward-function-hooks.

5.3 DP optimization by rewriting back-propagation (without hooks)

Instead of registering hooks on top of the original (non-DP) back-propagation, we can directly modify the back-propagation following Appendix B: given the activation $\boldsymbol{a}_{(l)}$ and output gradient $\frac{\partial L}{\partial \boldsymbol{s}_{(l)}}$, we get the per-sample clipped gradient

$$\frac{\partial C_i L_i}{\partial \mathbf{W}_l} = \frac{\boldsymbol{a}_{(l),i}^\top \frac{\partial L}{\partial \boldsymbol{s}_{(l),i}}}{\sqrt{\text{vec}\left(\frac{\partial L}{\partial \boldsymbol{s}_{(l),i}} \frac{\partial L}{\partial \boldsymbol{s}_{(l),i}}^\top\right) \cdot \text{vec}(\boldsymbol{a}_{(l),i} \boldsymbol{a}_{(l),i}^\top)}}$$

This approach requires rewriting the back-propagation for each base layer (linear, embedding, convolution, normalization, ...) and can be done at different levels of granularity (Pytorch, C++, CUDA kernel).

In this work, we work on Pytorch which builds the models as follows:

$$\text{function} \Longrightarrow \text{base layer} \Longrightarrow \text{advanced layer}.$$

For instance,

`LinearFunction` (modified)
$\longrightarrow$`nn.Linear` (modified)
$\longrightarrow$MoE router, Attention QKV, SWIGLU, LoRA
$\longrightarrow$Mixture-of-experts (MoE; (Fedus et al., 2022)),
$\longrightarrow$Attention QKV (Vaswani et al., 2017),
$\longrightarrow$LoRA (Hu et al., 2021),
$\longrightarrow$SWIGLU (Shazeer, 2020), etc.

Furthermore, non-parameterized modules such as softmax do not require DP-specific handling.

As a consequence, we highlight that

> *Rewriting the back-propagation of base layers is sufficient to directly enable DP training on a wide range of advanced architectures and optimizers.*

We give an example of the standard back-propagation of a linear layer[9].

```
class LinearFunction(Function):
    def forward(ctx, input, weight, bias=None):
        ctx.save_for_backward(input, weight, bias)
        return F.linear(input,weight,bias)

    @staticmethod
    def backward(ctx, grad_output):
        input, weight, bias = ctx.saved_tensors

        grad_input = grad_weight = grad_bias = None

        #-- input gradient
        if ctx.needs_input_grad[0]:
            grad_input = grad_output.mm(weight)

        #-- bias gradient
        if bias is not None and ctx.needs_input_grad[2]:

            grad_bias = grad_output.sum(0)
        #-- weight gradient
        if ctx.needs_input_grad[1]:
            grad_weight = grad_output.t().mm(input)

        return grad_input, grad_weight, grad_bias
```

For comparison, we give the DP back-propagation of a linear layer, which heavily uses `einsum` due to the per-sample operation. *Note the forward function is not modified.* In practice, we apply the mixed ghost clipping for the weight tensor. Only in this section, we instead apply the original ghost clipping for the ease of presentation.

```
class DPLinearFunction(Function):
    def forward(ctx, input, weight, bias=None):
        ctx.save_for_backward(input, weight, bias)
        return F.linear(input,weight,bias)

    @staticmethod
    def backward(ctx, grad_output):
        input, weight, bias = ctx.saved_tensors

        grad_input = grad_weight = grad_bias = None

        #-- input gradient
        if ctx.needs_input_grad[0]:
            grad_input = grad_output.mm(weight)

        #-- bias gradient
        if bias is not None and ctx.needs_input_grad[2]:

            grad_bias = torch.einsum('B...p->Bp',
 grad_output)
            grad_norm2=grad_bias.norm(2, dim=1)**2
        else:
            grad_norm2=torch.zeros(grad_output.shape[0])


        #-- weight gradient
        if ctx.needs_input_grad[1]:
            grad_weight, clip_factor=_linear_w_norm(
 input,grad_output,grad_norm2)
            grad_weight = grad_weight/math.sqrt(
 n_layers)+weight.noise
            if bias is not None and ctx.
 needs_input_grad[2]:
                grad_bias = torch.einsum('B,B...->...',
 clip_factor,grad_bias)
                grad_bias = grad_bias/math.sqrt(
 n_layers) + bias.noise
        return grad_input, grad_weight, grad_bias
```

where the per-sample weight gradients are clipped and aggregated by

```
def _linear_w_norm(input, grad_output, grad_norm2):
    temp1 = torch.bmm(input, input.transpose(-1,-2))
    temp2 = torch.bmm(grad_output, grad_output.
     transpose(-1,-2))
    grad_norm2 = grad_norm2 + (temp1 * temp2).sum(dim
     =(1,2))
    clip_factor = 1/torch.sqrt(grad_norm2.to(input))
    grad_output = torch.einsum('B,B...->B...',
     clip_factor,grad_output)
    return torch.einsum('B...p,B...d->pd',grad_output,
     input), clip_factor
```

In addition, we need to rewrite the original non-DP linear module, in order to use this DPLinearFunction. Other layers

[9]This example follows https://pytorch.org/docs/stable/notes/extending.html.

like convolution and normalization require similar modification, though the back-propagation takes different forms for different layers. Specifically, we define the DPLinear module, where the forward function leverages the DPLinearFunction.

```
class DPLinear(Module):
    __constants__ = ['in_features', 'out_features']
    in_features: int
    out_features: int
    weight: Tensor

    def __init__(self, in_features: int, out_features:
     int, bias: bool = True,
                 device=None, dtype=None) -> None:
        factory_kwargs = {'device': device, 'dtype':
     dtype}
        super(DPLinear, self).__init__()
        self.in_features = in_features
        self.out_features = out_features
        self.weight = Parameter(torch.empty((
     out_features, in_features), **factory_kwargs))
        if bias:
            self.bias = Parameter(torch.empty(
     out_features, **factory_kwargs))
        else:
            self.register_parameter('bias', None)
        self.reset_parameters()

    def reset_parameters(self) -> None:
        # Setting a=sqrt(5) in kaiming_uniform is the
     same as initializing with
        # uniform(-1/sqrt(in_features), 1/sqrt(
     in_features)). For details, see
        # https://github.com/pytorch/pytorch/issues
     /57109
        init.kaiming_uniform_(self.weight, a=math.sqrt
     (5))
        if self.bias is not None:
            fan_in, _ = init.
     _calculate_fan_in_and_fan_out(self.weight)
            bound = 1 / math.sqrt(fan_in) if fan_in > 0
      else 0
            init.uniform_(self.bias, -bound, bound)

    def forward(self, input: Tensor) -> Tensor:
        return DPLinearFunction.apply(input, self.
     weight, self.bias)

    def extra_repr(self) -> str:
        return 'in_features={}, out_features={}, bias
     ={}'.format(
            self.in_features, self.out_features, self.
     bias is not None
        )
```

Finally, we replace each non-DP linear layer by our DP linear layer, recursively:

```
def replace_Linear(module):
    for layer_str in dir(module):
        layer = getattr(module, layer_str)
        if type(layer) == torch.nn.Linear and
     requires_grad(layer):
            new_layer = DPLinear(in_features=layer.
     in_features, out_features=layer.out_features,
                                 bias=(layer.bias!=None),
      device=layer.weight.device, dtype=layer.weight.
     dtype)
            new_layer.weight = layer.weight
            new_layer.bias = layer.bias
            del layer

            setattr(module, layer_str, new_layer)

    # iterate through immediate child modules. Note,
     the recursion is done by our code no need to use
     named_modules()
    if hasattr(module,'children'):
        for immediate_child_module in module.children():

            replace_Linear(immediate_child_module)
```

In summary, to convert a standard non-DP model to a DP model, we need to (I) rewrite the function, (II) rewrite the module, and (III) replace each base layer type (linear, convolution, embedding, layer normalization, group normalization, and so on).

5.4 User interface

DP-ZeRO can be enabled by one piece of code: after the model is instantiated,

```
privacy_engine = PrivacyEngine(model, batch_size=256,
     sample_size=50000, epochs=3, target_epsilon=3)
```

The codebase is designed not to modify the optimizer, hence DP-ZeRO can work with arbitrary optimizer.

Consequently, our DP-ZeRO will not distinguish micro-batches when gradient accumulation is used. This is different from Opacus (version ==0.x) and Private-Vision, where only the last micro-batch is processed by "optimizer.step()" with noise addition $\sigma_{\text{DP}}N(0, I)$, but all other micro-batches are processed by "optimizer.virtual_step()" without noising. In constrast, DP-ZeRO adds the noise on each micro-batch equally[10].

6 Discussion

In this work, we develop DP-ZeRO with mixed-precision that enables the optimization of models with 100B trainable parameters, thus allowing DP distributed learning to be as efficient and scalable as the standard distributed learning. We believe this is a significant milestone to pave the path towards DP large models. In particular, our DP-ZeRO is orthogonal to other large-scale training techniques that are not tied to back-propagation: for example, such as activation check-pointing, CPU offloading, weight/activation quantization (Dettmers et al., 2023; Xiao et al., 2023), and tensor parallelism (Narayanan et al., 2021).

Acknowledgement

We thank Dhananjay Ram, Zhenghui Jin and Cole Hawkins for technical support on the AWS cluster, Xinwei Zhang for validating part of our code, Yu-Xiang Wang for the insightful initial discussion, and the hosts of FSDP (especially Min Xu) and DeepSpeed Github repositories for debugging part of our experiments.

[10]Note that the noise level per micro-batch is $\sigma_{\text{DP}}/\sqrt{N_d}$ if a random seed is set across N_d GPUs, or σ_{DP}/N_d otherwise.

REFERENCES

Abadi, M., Chu, A., Goodfellow, I., McMahan, H. B., Mironov, I., Talwar, K., and Zhang, L. Deep learning with differential privacy. In *Proceedings of the 2016 ACM SIGSAC conference on computer and communications security*, pp. 308–318, 2016.

Brown, T., Mann, B., Ryder, N., Subbiah, M., Kaplan, J. D., Dhariwal, P., Neelakantan, A., Shyam, P., Sastry, G., Askell, A., et al. Language models are few-shot learners. *Advances in neural information processing systems*, 33: 1877–1901, 2020.

Bu, Z., Dong, J., Long, Q., and Su, W. J. Deep learning with gaussian differential privacy. *Harvard data science review*, 2020(23), 2020.

Bu, Z., Gopi, S., Kulkarni, J., Lee, Y. T., Shen, H., and Tantipongpipat, U. Fast and memory efficient differentially private-sgd via jl projections. *Advances in Neural Information Processing Systems*, 34, 2021.

Bu, Z., Mao, J., and Xu, S. Scalable and efficient training of large convolutional neural networks with differential privacy. *Advances in Neural Information Processing Systems*, 35:38305–38318, 2022a.

Bu, Z., Wang, Y.-X., Zha, S., and Karypis, G. Differentially private bias-term only fine-tuning of foundation models. In *Workshop on Trustworthy and Socially Responsible Machine Learning, NeurIPS 2022*, 2022b.

Bu, Z., Liu, R., Wang, Y.-X., Zha, S., and Karypis, G. On the accuracy and efficiency of group-wise clipping in differentially private optimization. *arXiv preprint arXiv:2310.19215*, 2023a.

Bu, Z., Wang, Y.-X., Zha, S., and Karypis, G. Automatic clipping: Differentially private deep learning made easier and stronger. *Advances in Neural Information Processing Systems*, 2023b.

Bu, Z., Wang, Y.-X., Zha, S., and Karypis, G. Differentially private optimization on large model at small cost. In *International Conference on Machine Learning*, pp. 3192–3218. PMLR, 2023c.

Bu, Z., Zhang, X., Hong, M., Zha, S., and Karypis, G. Pretraining differentially private models with limited public data. *arXiv preprint arXiv:2402.18752*, 2024.

Chen, T., Xu, B., Zhang, C., and Guestrin, C. Training deep nets with sublinear memory cost. *arXiv preprint arXiv:1604.06174*, 2016.

De, S., Berrada, L., Hayes, J., Smith, S. L., and Balle, B. Unlocking high-accuracy differentially private image classification through scale. *arXiv preprint arXiv:2204.13650*, 2022.

Dettmers, T., Pagnoni, A., Holtzman, A., and Zettlemoyer, L. Qlora: Efficient finetuning of quantized llms. *arXiv preprint arXiv:2305.14314*, 2023.

Dong, J., Roth, A., and Su, W. J. Gaussian differential privacy. *arXiv preprint arXiv:1905.02383*, 2019.

Dosovitskiy, A., Beyer, L., Kolesnikov, A., Weissenborn, D., Zhai, X., Unterthiner, T., Dehghani, M., Minderer, M., Heigold, G., Gelly, S., et al. An image is worth 16x16 words: Transformers for image recognition at scale. In *International Conference on Learning Representations*, 2020.

Dwork, C., McSherry, F., Nissim, K., and Smith, A. Calibrating noise to sensitivity in private data analysis. In *Theory of cryptography conference*, pp. 265–284. Springer, 2006.

Fedus, W., Zoph, B., and Shazeer, N. Switch transformers: Scaling to trillion parameter models with simple and efficient sparsity. *Journal of Machine Learning Research*, 23(120):1–39, 2022.

Foley, D. and Danskin, J. Ultra-performance pascal gpu and nvlink interconnect. *IEEE Micro*, 37(2):7–17, 2017.

Frostig, R., Johnson, M. J., and Leary, C. Compiling machine learning programs via high-level tracing. *Systems for Machine Learning*, 4(9), 2018.

Goodfellow, I. Efficient per-example gradient computations. *arXiv preprint arXiv:1510.01799*, 2015.

Gopi, S., Lee, Y. T., and Wutschitz, L. Numerical composition of differential privacy. *Advances in Neural Information Processing Systems*, 34, 2021.

He, J., Li, X., Yu, D., Zhang, H., Kulkarni, J., Lee, Y. T., Backurs, A., Yu, N., and Bian, J. Exploring the limits of differentially private deep learning with group-wise clipping. *arXiv preprint arXiv:2212.01539*, 2022.

He, K., Zhang, X., Ren, S., and Sun, J. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.

Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., and Gelly, S. Parameter-efficient transfer learning for nlp. In *International Conference on Machine Learning*, pp. 2790–2799. PMLR, 2019.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., and Chen, W. Lora: Low-rank adaptation of large language models. *arXiv preprint arXiv:2106.09685*, 2021.

Huang, Y., Cheng, Y., Bapna, A., Firat, O., Chen, D., Chen, M., Lee, H., Ngiam, J., Le, Q. V., Wu, Y., et al. Gpipe: Efficient training of giant neural networks using pipeline parallelism. *Advances in neural information processing systems*, 32, 2019.

Ishii, A., Foley, D., Anderson, E., Dally, B., Dearth, G., Dennison, L., Hummel, M., and Schafer, J. Nvswitch and dgx-2 nvlink-switching chip and scale-up compute server. In *Hot Chips*, 2018.

Katharopoulos, A., Vyas, A., Pappas, N., and Fleuret, F. Transformers are rnns: Fast autoregressive transformers with linear attention. In *International conference on machine learning*, pp. 5156–5165. PMLR, 2020.

Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.

Koskela, A., Jälkö, J., and Honkela, A. Computing tight differential privacy guarantees using fft. In *International Conference on Artificial Intelligence and Statistics*, pp. 2560–2569. PMLR, 2020.

Kurakin, A., Chien, S., Song, S., Geambasu, R., Terzis, A., and Thakurta, A. Toward training at imagenet scale with differential privacy. *arXiv preprint arXiv:2201.12328*, 2022.

Lee, J. and Kifer, D. Scaling up differentially private deep learning with fast per-example gradient clipping. *Proceedings on Privacy Enhancing Technologies*, 2021(1), 2021.

Li, A., Song, S. L., Chen, J., Li, J., Liu, X., Tallent, N. R., and Barker, K. J. Evaluating modern gpu interconnect: Pcie, nvlink, nv-sli, nvswitch and gpudirect. *IEEE Transactions on Parallel and Distributed Systems*, 31(1):94–110, 2019.

Li, Q., Hong, J., Xie, C., Tan, J., Xin, R., Hou, J., Yin, X., Wang, Z., Hendrycks, D., Wang, Z., et al. Llm-pbe: Assessing data privacy in large language models. *Proceedings of the VLDB Endowment*, 17(11):3201–3214, 2024.

Li, S., Zhao, Y., Varma, R., Salpekar, O., Noordhuis, P., Li, T., Paszke, A., Smith, J., Vaughan, B., Damania, P., et al. Pytorch distributed: Experiences on accelerating data parallel training. *Proceedings of the VLDB Endowment*, 13(12).

Li, X., Tramer, F., Liang, P., and Hashimoto, T. Large language models can be strong differentially private learners. *arXiv preprint arXiv:2110.05679*, 2021.

McMahan, H. B., Ramage, D., Talwar, K., and Zhang, L. Learning differentially private recurrent language models. In *International Conference on Learning Representations*, 2018.

Mehta, H., Thakurta, A., Kurakin, A., and Cutkosky, A. Large scale transfer learning for differentially private image classification. *arXiv preprint arXiv:2205.02973*, 2022.

Micikevicius, P., Narang, S., Alben, J., Diamos, G., Elsen, E., Garcia, D., Ginsburg, B., Houston, M., Kuchaiev, O., Venkatesh, G., et al. Mixed precision training. In *International Conference on Learning Representations*, 2018.

Mironov, I. On significance of the least significant bits for differential privacy. In *Proceedings of the 2012 ACM conference on Computer and communications security*, pp. 650–661, 2012.

Narayanan, D., Shoeybi, M., Casper, J., LeGresley, P., Patwary, M., Korthikanti, V., Vainbrand, D., Kashinkunti, P., Bernauer, J., Catanzaro, B., et al. Efficient large-scale language model training on gpu clusters using megatron-lm. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–15, 2021.

Nvidia. The most powerful end-to-end ai and hpc data center platform. https://www.nvidia.com/en-us/data-center/tensor-cores/?ref=blog.paperspace.com.

Radford, A., Wu, J., Child, R., Luan, D., Amodei, D., Sutskever, I., et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.

Rajbhandari, S., Rasley, J., Ruwase, O., and He, Y. Zero: Memory optimizations toward training trillion parameter models. In *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–16. IEEE, 2020.

Rajbhandari, S., Ruwase, O., Rasley, J., Smith, S., and He, Y. Zero-infinity: Breaking the gpu memory wall for extreme scale deep learning. In *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*, pp. 1–14, 2021.

Ren, J., Rajbhandari, S., Aminabadi, R. Y., Ruwase, O., Yang, S., Zhang, M., Li, D., and He, Y. {ZeRO-Offload}: Democratizing {Billion-Scale} model training. In *2021 USENIX Annual Technical Conference (USENIX ATC 21)*, pp. 551–564, 2021.

Shazeer, N. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020.

Shen, Z., Zhang, M., Zhao, H., Yi, S., and Li, H. Efficient attention: Attention with linear complexities. In *Proceedings of the IEEE/CVF winter conference on applications of computer vision*, pp. 3531–3539, 2021.

Subramani, P., Vadivelu, N., and Kamath, G. Enabling fast differentially private sgd via just-in-time compilation and vectorization. *Advances in Neural Information Processing Systems*, 34, 2021.

Tang, H., Gan, S., Awan, A. A., Rajbhandari, S., Li, C., Lian, X., Liu, J., Zhang, C., and He, Y. 1-bit adam: Communication efficient large-scale training with adam's convergence speed. In *International Conference on Machine Learning*, pp. 10118–10129. PMLR, 2021.

Touvron, H., Lavril, T., Izacard, G., Martinet, X., Lachaux, M.-A., Lacroix, T., Rozière, B., Goyal, N., Hambro, E., Azhar, F., et al. Llama: Open and efficient foundation language models. *arXiv preprint arXiv:2302.13971*, 2023a.

Touvron, H., Martin, L., Stone, K., Albert, P., Almahairi, A., Babaei, Y., Bashlykov, N., Batra, S., Bhargava, P., Bhosale, S., et al. Llama 2: Open foundation and fine-tuned chat models. *arXiv preprint arXiv:2307.09288*, 2023b.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, Ł., and Polosukhin, I. Attention is all you need. *Advances in neural information processing systems*, 30, 2017.

Wang, S., Li, B. Z., Khabsa, M., Fang, H., and Ma, H. Linformer: Self-attention with linear complexity. *arXiv preprint arXiv:2006.04768*, 2020.

Xiao, G., Lin, J., Seznec, M., Wu, H., Demouth, J., and Han, S. SmoothQuant: Accurate and efficient post-training quantization for large language models. In *Proceedings of the 40th International Conference on Machine Learning*, 2023.

Xie, L., Lin, K., Wang, S., Wang, F., and Zhou, J. Differentially private generative adversarial network. *arXiv preprint arXiv:1802.06739*, 2018.

Yang, X., Zhang, H., Chen, W., and Liu, T.-Y. Normalized/clipped sgd with perturbation for differentially private non-convex optimization. *arXiv preprint arXiv:2206.13033*, 2022.

Yousefpour, A., Shilov, I., Sablayrolles, A., Testuggine, D., Prasad, K., Malek, M., Nguyen, J., Ghosh, S., Bharadwaj, A., Zhao, J., Cormode, G., and Mironov, I. Opacus: User-friendly differential privacy library in PyTorch. *arXiv preprint arXiv:2109.12298*, 2021.

Yu, D., Naik, S., Backurs, A., Gopi, S., Inan, H. A., Kamath, G., Kulkarni, J., Lee, Y. T., Manoel, A., Wutschitz, L., et al. Differentially private fine-tuning of language models. *arXiv preprint arXiv:2110.06500*, 2021a.

Yu, D., Zhang, H., Chen, W., Yin, J., and Liu, T.-Y. Large scale private learning via low-rank reparametrization. In *International Conference on Machine Learning*, pp. 12208–12218. PMLR, 2021b.

Zaken, E. B., Goldberg, Y., and Ravfogel, S. Bitfit: Simple parameter-efficient fine-tuning for transformer-based masked language-models. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pp. 1–9, 2022.

Zhai, X., Kolesnikov, A., Houlsby, N., and Beyer, L. Scaling vision transformers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12104–12113, 2022.

Zhang, Z., Zheng, S., Wang, Y., Chiu, J., Karypis, G., Chilimbi, T., Li, M., and Jin, X. Mics: near-linear scaling for training gigantic model on public cloud. *Proceedings of the VLDB Endowment*, 16(1):37–50, 2022.

Zhao, Y., Varma, R., Huang, C.-C., Li, S., Xu, M., and Desmaison, A. Introducing pytorch fully sharded data parallel (fsdp) api. https://pytorch.org/blog/introducing-pytorch-fully-sharded-data-parallel-a

Zhu, Y., Dong, J., and Wang, Y.-X. Optimal accounting of differential privacy via characteristic function. In *International Conference on Artificial Intelligence and Statistics*, pp. 4782–4817. PMLR, 2022.

A BACKGROUND OF DIFFERENTIAL PRIVACY

Differential privacy protects the training data by obfuscating the gradient, so that the model memorizes the data less:

$$\textit{data} \longrightarrow \textit{gradient} \xrightarrow{\text{privatization}} \textit{model}$$

Given a mini-batch of per-sample gradients, the Gaussian mechanism adds calculated noise to the per-sample clipped gradients, so as to make the likelihood of the resulting "privatized" gradient nearly indistinguishable from one derived without a specific individual's data.

Here each per-sample gradient is clipped in case the gradient magnitude is too large to be obfuscated by the noise, and more noise equals stronger obfuscation as well as stronger privacy protection.

B THE BOOK-KEEPING (BK) ALGORITHM

B.1 Efficient computation of per-sample gradient norms

The mixed ghost norm (Bu et al., 2022a) is the state-of-the-art technique to compute the per-sample gradient norm of the **weight**, almost for free. It hybridizes two basic techniques – the per-sample gradient instantiation and the ghost norm – to compute the Frobenius norm of weight gradient,

$$\begin{aligned}&\left\|\boldsymbol{a}_{l,i}^\top \frac{\partial L}{\partial \boldsymbol{s}_{l,i}}\right\|^2 \overset{\text{per-sample grad}}{=\!=} \|\frac{\partial L_i}{\partial \mathbf{W}_l}\|_{\text{Fro}}^2 \\ &\overset{\text{ghost norm}}{=\!=} \text{vec}\left(\frac{\partial L}{\partial \boldsymbol{s}_{l,i}}\frac{\partial L}{\partial \boldsymbol{s}_{l,i}}^\top\right)\cdot \text{vec}(\boldsymbol{a}_{l,i}\boldsymbol{a}_{l,i}^\top)\end{aligned} \tag{4}$$

where "vec" flattens the tensor to an one-dimensional vector. In words, Equation (4) gives two equations that are equivalent mathematically, but significantly different in efficiency:

- $\|X\|_{\text{Fro}}^2 = \left\|A^\top B\right\|^2$ firstly computes $A^\top B$ and then its norm.
- $\|X\|_{\text{Fro}}^2 = \text{vec}\left(AA^\top\right)\cdot \text{vec}(BB^\top)$ firstly computes $AA^\top, BB^\top \in \mathbb{R}^{TT}$ and then their dot product.

In summary, the mixed ghost norm always applies the cheaper of two techniques at each layer of a neural network.

Finally, we note that the per-sample gradient norm of the **bias** is computed differently. This is because

$$\frac{\partial L_i}{\partial b_l} = \mathbf{1}^\top \frac{\partial L}{\partial \boldsymbol{s}_{l,i}}$$

is not actually a product of tensors like $X = A^\top B$. In fact, the multiplication with $\mathbf{1}$ turns out to be a summation along the first dimension, and it suffices to use per-sample gradient instantiation for the bias.

B.2 Book-keeping the output gradient

The BK algorithm uses two rounds of back-propagation (though each round only takes half the complexity, hence the total complexity of DP back-propagation matches the non-DP back-propagation). Therefore, output gradients $\frac{\partial L}{\partial \boldsymbol{s}_{l,i}}$ are kept to avoid repeated computation. Notice that the output gradient are relatively cheap to book-keep (see Figure 4 and Figure 5 in (Bu et al., 2023c)).

C COMPONENT-WISE TIME COMPLEXITY OF DP-ZERO

In this section, we explain the time complexity of each part in Table 2, and demonstrate how the complexity can be different under different settings.

Forward propagation: The matrix multiplication during forward propagation results in $2BT\Psi_{\text{model}}$ complexity (see (Bu et al., 2022a)). Notice that, if the activation checkpointing is used, essentially two rounds of forward propagation take place in one iteration. Hence the time complexity doubles and becomes $4BT\Psi_{\text{model}}$.

Back-propagation: This contains two sub-processes: the output gradients are computed at all layers, taking $2BT\Psi_{\text{model}}$ complexity; the parameter gradients are computed only at trainable layers (a few if doing PEFT), taking $2BT\Psi_{\text{train}}$ complexity. Clearly, in full parameter training, the total is $4BT\Psi_{\text{model}}$, and in PEFT, about $2BT\Psi_{\text{model}}$.

Attention: The time complexity of attention is $O(BT^2)$ in (Vaswani et al., 2017), where T is the sequence length (a.k.a. token length). When T is large, e.g. training with long context like $T = 8192$, this cost is prohibitively high. In this regard, a line of researches have proposed linear complexity attention, including but not limited to (Wang et al., 2020; Katharopoulos et al., 2020; Shen et al., 2021).

Communication: For algorithms that don't shard the model, such as data parallelism and ZeRO1/2, the communication is only used to send gradients and optimizer states. Hence the communication volume is proportional to the number of trainable parameters $O(\Psi_{\text{train}})$. Otherwise, for algorithms such as ZeRO3 and tensor parallelism, the communication volume is proportional to the number of total parameters $O(\Psi_{\text{model}})$, because the forward propagation needs to gather the parameters from many GPUs. This makes a big difference in PEFT when $\Psi_{\text{train}} \ll \Psi_{\text{model}}$.

D LOSS SCALING IN MIXED-PRECISION TRAINING

We write the per-sample gradient with loss scaling S as

$$\frac{\partial C_i L_i}{\partial \mathbf{W}_l} = C_i \frac{1}{S}\cdot\left(\boldsymbol{a}_{l,i}^\top (S\cdot \frac{\partial L}{\partial \boldsymbol{s}_{l,i}})\right)$$

training regime	half-precision	loss scaling	steps	remark
Standard	fp16	yes	12356	$\approx$ full-precision standard training
	fp16	no	136	unstable or divergent training
	bf16	no	136	$\approx$ full-precision standard training
DP	fp16	yes	123456	too slow training due to over-shrunk gradient
	fp16	yes	1234^*56	hard to implement and unstable training (Li et al., 2021)
	fp16	no	1346	unstable or divergent training
	bf16	no	1346	$\approx$ full-precision DP training **(our recommendation)**

Table 4. Mixed-precision training with DP or not.

This covers the standard gradient ($C_i = 1$) and DP gradient (e.g. $C_i = 1/||\boldsymbol{g}_i||$, computed by the mixed ghost norm in Appendix B). The key of loss scaling is that $S \gg 1$ enlarges the output gradient to avoid underflow, and $\frac{1}{S}$ shrinks the parameter gradient to the correct magnitude.

Recall that a standard mixed-precision training (with loss scaling[11]) uses steps $1 \to 2 \to 3 \to 5 \to 6$, or $1 \to 3 \to 6$ without loss scaling.

1. Forward propagation (fp16 weights and activations) and get the loss.
2. ~~Scaling up the loss by a factor S.~~
3. Backward propagation on the scaled loss (fp16 parameters and their gradients).
4. Per-sample gradient clipping (sensitivity $= 1$) and noising for DP.
5. ~~Scaling down the parameter gradient by a factor of $1/S$.~~
6. Update the parameters with their gradients.

If we follow the same procedure under the DP regime, say using a hook function to be called after back-propagation creates the gradients like in Opacus (Yousefpour et al., 2021), Private-Transformers (Li et al., 2021), FastDP (Bu et al., 2023c), then the per-sample clipping factor is scaled up S times so as to normalize the gradient. Hence per-sample gradient clipping has already played the role of scaling down. If we scale down the gradient for a second time, the gradient is incorrectly over-shrunk. This is the case in (Yu et al., 2021a) and in the alternative implementation of Appendix T (Li et al., 2021) (see also Figure 5). To be sure, this approach is still DP, but the performance does not match fp32 DP training correctly, and usually degrades too much to be useful.

[11] See https://docs.nvidia.com/deeplearning/performance/mixed-precision-training/index.html#lossscaling.

One walk-around is to prevent per-sample gradient clipping to scale down the gradients and let step 5 do its job, i.e. $1 \to 2 \to 3 \to 4^* \to 5 \to 6$. We note that Appendix T in (Li et al., 2021) follows this path (though no experiment results or codes are available at the time of writing) by modifying step 4: clipping threshold (sensitivity)$= S$ instead of 1, so that the clipped gradient is S times larger than the DP f32 training, to be scaled down by step 5. However, this introduces additional design decisions and does not prevent overflow when using fp16 (due to step 2, see Table 3).

Another walk-around is to delete step 5 and let per-sample gradient clipping do its job, i.e. $1 \to 2 \to 3 \to 4 \to 6$. However, this approach is harder to implement because in the standard process step 2 and 5 are simultaneously enabled or disabled. Also we cannot prevent overflow when using fp16 as we still use step 2.

Therefore, we propose to not use loss scaling (or equivalently we set $S = 1$ statically for all steps) during DP mixed-precision training, i.e. $1 \to 3 \to 4 \to 6$. Although, by not using step 2, we cannot prevent underflow when using fp16, this is much less a problem compared to overflow: underflow (treating small values as 0) makes the training less accurate but does not fail the training like overflow (treating large values as NAN). Lastly, the underflow issue is perfectly mitigated by bf16, which we recommend for DP mixed-precision training whenever possible.

In summary, we recommend using bf16 format and not using loss scaling, although we also support fp16 without loss scaling for DP training. Our implementation is compatible with `torch.amp` package.

Remark D.1. DP may be violated due to the floating-point vulnerability at any precision (Mironov, 2012), but the issues can be more severe at half-precision. Yet, it remains an open discussion whether precision-based attacks can non-trivially break the protection of DP in deep learning, since the model is updated through thousands of iterations which may cover the identifiability of data information. In fact, (Bu et al., 2024; Yu et al., 2021b; Li et al., 2024) have demonstrated that DP mixed-precision training can empirically preserve the protection against various privacy attacks.

E EXPERIMENT SETTINGS

Datasets: To evaluate the efficiency, it suffices to declare the data's dimension (e.g. micro-batch size and feature dimension) without specifying the dataset (though sometimes specifying the dataset means declaring the dimension, e.g. MNIST usually means 28*28 pixels). This is the norm in system papers such as (Rajbhandari et al., 2020; 2021; Zhao et al.). In this work, vision models are trained with 224*224 pixels at ImageNet scale; GPT models are trained with sequence length 100, except in Figure 8 where sequence length is 2048.

Figure 5 and Table 3: We train ViT-large (300M parameters) and CIFAR100, 5 epochs, learning rate 5e-4, logical batch size 1000.

Figure 6: To fit as large a model as possible, we set $B = 1$ and use SGD. We set 48 attention heads, 21 layers, MLP=4*width (also known as embedding dimension), and modify width for all models. For instance, ViT-10B uses width=$768 * 22$, ViT-22B uses width=$768 * 34$.

Figure 8: We train AdamW with layer-wise clipping. DP distributed learning is based on MiCS (ZeRO3) using bf16 mixed-precision training. Most of GPT configuration is the same as (Touvron et al., 2023a) (Table 2) in terms of embedding dimension, attention heads and number of layers. However, GPT-100B uses the configuration from (Brown et al., 2020) (Table 2.1) but a smaller width.

F VERIFICATION OF OUR IMPLEMENTATION

DP-ZeRO is implementing the same DP-SGD (hence the same utility) as in (Abadi et al., 2016), only in a much more scalable and distributed manner than libraries like Private Transformers and Opacus.

To be sure, we provide the results on 5 runs of ViT (image classficiation) and GPT2 (text generation), at $\epsilon = 8$, with standard deviations in the parenthesis. We compare to other libraries (without mixed precision and without distributed learning) and our DP-ZeRO uses mixed precision, in order to validate both (I) the correctness of DP mixed precision and (II) the correctness of our implementation of DP-ZeRO.

	num of GPU	ViT/CIFAR100/accuracy	GPT2/E2E/BLEU
Opacus	1	83.45(0.28)	63.35(0.50)
FastDP	1	83.46(0.19)	64.01(0.54)
GhostClip	1	83.51(0.33)	63.71(0.31)
DP-ZeRO1	8	83.46(0.13)	64.03(0.41)
DP-ZeRO2	8	83.43(0.09)	63.72(0.47)
DP-ZeRO3	8	83.48(0.18)	63.92(0.39)