



HAL
open science

Cyclic Data Parallelism for Efficient Parallelism of Deep Neural Networks

Louis Fournier, Edouard Oyallon

► **To cite this version:**

Louis Fournier, Edouard Oyallon. Cyclic Data Parallelism for Efficient Parallelism of Deep Neural Networks. 2024. hal-04500345

HAL Id: hal-04500345

<https://hal.science/hal-04500345>

Preprint submitted on 12 Mar 2024

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Cyclic Data Parallelism for Efficient Parallelism of Deep Neural Networks

Louis Fournier¹ and Edouard Oyallon²

¹ ISIR, Sorbonne Université, CNRS, Paris, France
louis.fournier@isir.upmc.fr

² Center for Computational Mathematics, Flatiron Institute, New York 10010, USA

Abstract. Training large deep learning models requires parallelization techniques to scale. In existing methods such as Data Parallelism or ZeRO-DP, micro-batches of data are processed in parallel, which creates two drawbacks: the total memory required to store the model’s activations peaks at the end of the forward pass, and gradients must be simultaneously averaged at the end of the backpropagation step. We propose Cyclic Data Parallelism, a novel paradigm shifting the execution of the micro-batches from simultaneous to sequential, with a uniform delay. At the cost of a slight gradient delay, the total memory taken by activations is constant, and the gradient communications are balanced during the training step. With Model Parallelism, our technique reduces the number of GPUs needed, by sharing GPUs across micro-batches. Within the ZeRO-DP framework, our technique allows communication of the model states with point-to-point operations rather than a collective broadcast operation. We illustrate the strength of our approach on the CIFAR-10 and ImageNet datasets.

Keywords: Data Parallelism · Model Parallelism · Pipeline Parallelism

1 Introduction

Deep learning models have significantly increased in size in recent years, reaching hundreds of billions of parameters and requiring increasingly expensive training [10]. As the cost of training these models continues to rise, there is an increasing need for parallelization strategies. In particular, Data Parallelism (DP) [5, 15] remains the dominant method for training Deep Neural Networks (DNNs) at scale. In DP, the model to be trained is first replicated on multiple workers. Then, during each training step, a mini-batch of data is sliced evenly among the workers in so-called micro-batches. Each worker then executes a forward and backward propagation on each micro-batch, and the locally computed gradients are subsequently averaged across all workers to obtain the gradient with regard to the full mini-batch. The resulting averaged gradient is used to locally update the models, typically following Stochastic Gradient Descent (SGD).

Nevertheless, DP has major drawbacks. First, the communication step between workers is synchronous, as all workers must complete their gradient computations before communicating, leading to idle workers waiting for the slowest

worker [9]. Second, gradients are communicated globally with an all-reduce operation [8], which means that the communication step poses a challenge as the number of workers increases [2]. Last, the total memory used by all workers grows linearly with the number of workers since the model is fully replicated on each one [21]. This requirement can be impractical since modern model sizes exceed the memory capacity of a single device (*e.g.*, a GPU).

In this work, we tackle the memory and communication drawbacks of DP by proposing to change the execution time of workers in DP from simultaneous to sequential. We refer to this process as Cyclic Data Parallelism (CDP). Our modification of standard DP aims to balance both the communication costs and the overall memory usage, by relying heavily on the sequential nature of the execution steps of DNNs during training. More specifically, CDP reduces the cost of gradient communications in DP from collective communications at the end of the training step to point-to-point communications balanced over the entire training step. This balances the total memory used by all workers, at the cost of a small gradient delay. CDP can be combined with standard parallelization implementations for further improvements.

Contributions. We list our specific contributions as follows: **(a)** First, we propose CDP, an alternative to DP that balances gradient communications and overall memory usage across training, at the cost of computing on some delayed gradients. **(b)** We analyze the method by showing that it maintains a low communication and memory overhead during the computation of a mini-batch and allows a much more efficient implementation of mini-batch SGD on one or several GPUs. **(c)** We then particularize the CDP paradigm to state-of-the-art approaches such as Model Parallelism (MP) and Zero Redundancy Optimizer powered DP (ZeRO-DP) [21], showing improvements in all cases. **(d)** Empirically, we show that the gradient delay of CDP leads to equal training of DNNs compared to DP on the large-scale CIFAR-10 and ImageNet datasets.

Paper organization. Our paper is structured as follows. First, we summarize the parallelization techniques used to train DNNs in Sec. 2. We then present standard mini-batch SGD with DP in Sec. 3.1 before introducing our new CDP paradigm and two possible update rules in Sec. 3.2. We then discuss how CDP improves on DP for the implementations of DP (on a single or multiple GPUs), MP and ZeRO-DP in Sec. 4. Finally, we present in Sec. 5 a numerical analysis of the update rules of CDP over the one of DP to train ResNets on CIFAR-10 and ImageNet and show the possible improvement in total memory for a ResNet-50 and a ViT-B/16.

2 Related work

Modern parallelization techniques for DNNs can be separated mainly into either DP, partitioning the mini-batch, or MP, partitioning the DNN model directly.

Data Parallelism. In this class of techniques, the model is replicated on several workers, and each worker computes a micro-batch, a subset of the mini-batch. As noted above, this method suffers from a linear increase in memory as the number of workers increases, while communications costs increase either logarithmically or linearly [2]. To reduce communications, less strict synchronization steps have been introduced, at the cost of updating the parameters with stale gradients [15, 23, 29]. However, these techniques do not take into account the memory required to train DNNs, a point we improve on.

Two notable DP techniques aim to reduce the memory cost of DP: Zero Redundancy Optimizer powered DP (*ZeRO-DP*) [21] and *Fully-Sharded DP* [36]. Rather than fully replicating the entire model across the DP workers, each worker retains only a subset of the model states (*i.e.* optimizer states, gradients, and parameters) corresponding to one stage of the model. The memory required can be further reduced by storing model states on the CPU [22, 24], at the expense of communications between the CPU and the GPUs. Despite the memory gains, the volume of communication increases in ZeRO-DP, which has been addressed by compressing communications or striking a balance with the memory [31, 35]. Nevertheless, the simultaneous nature of the micro-batches executions remains the root cause of ZeRO-DP’s issues, which we tackle in this paper.

Model Parallelism. Instead of parallelizing computations by slicing a mini-batch, MP [5, 26] aims to partition the components of the DNN. MP can be separated into two categories. First, Intra-layer MP [13, 27], also called tensor parallelism, partitions individual layers across workers but requires a high communication cost. Second, Inter-layer MP [5] splits a DNN into successive stages, and workers are solicited for only one stage, thus requiring less memory. A disadvantage of this method is that only one of the workers computes at a time since the stages are executed sequentially in the forward-backward pass. This limitation was addressed in several ways. Local learning [1, 7, 19] removes the need for a backward pass of the model with local auxiliary objectives at each stage. This allows stages to compute mini-batches simultaneously but at the expense of model performance. Delayed gradient methods [32, 37] allow stages to compute in parallel on micro-batches, by allowing updates with delayed gradients.

Of particular interest is *Pipeline Parallelism* (PP), a form of Inter-layer MP, introduced in GPipe [11], which also avoids this problem. In PP, mini-batches are divided into micro-batches which are passed sequentially to the successive stages. This allows each micro-batch to be propagated to the next stage as soon as it has completed its forward pass, allowing stages to compute micro-batches in parallel. GPipe thus reduces the number of idle workers, while still leaving a "bubble". PipeDream [17] proposes to further reduce this bubble of computation by introducing the one-forward-one-backward (1F1B) pipeline schedule. Within this implementation, a stage alternates between the forward and the backward pass of one micro-batch. However, this comes at a cost as multiple versions of the parameters must be stored by a stage, and potentially storing the activations of the entire batch. Our method similarly introduces a cyclic delay of micro-batches. However, our framework is more general and notably applies to

many other parallelism frameworks, not just PP. GEMS [12] is a MP framework that similarly reduces worker idleness by offsetting computations, like CDP and PP. Still, workers must store parameters of two stages, and GEMS does not extend to other frameworks either, unlike CDP. PipeDream-2BW [18] improves on PipeDream by keeping only two versions of the parameters and accumulating the gradients of the micro-batches to update them. This allows for continuous computations on each stage with no idle workers at the cost of a gradient staleness of one mini-batch. We will show that our more general framework can recover and improve on the update rule of PipeDream-2BW when applied to the particular case of PP. PipeMare [33] keeps a single version of the weights and, like the previous delayed gradient methods, allows asynchronous updates. However, it requires additional hyperparameters, which we do not. DAPPLE [6] improves on PP by using stage replicas to handle micro-batches partitions. This can be combined with our CDP framework. Still, as in PP, GPUs need to store batch activations, not just micro-batches. Finally, note that the parallelization techniques we have discussed are often integrated together, with 3D Parallelism [28] being a notable example that combines DP, tensor parallelism, and PP.

3 The Cyclic Data Parallelism paradigm

3.1 A motivating example: mini-batch SGD with Data Parallelism

The Cyclic Data Parallelism paradigm is best understood from the perspective of the standard DP paradigm. We first describe it before highlighting its caveats.

Mini-batch SGD via DP. A standard DP strategy applied to a mini-batch SGD step is to replicate a model over N workers. Each worker is then fed with mini-batch slices, referred to as micro-batches [11]. At training step t , each replica $i \in [1, N]$ receives a micro-batch of data leading to a training objective f_i parametrized by θ_t , and computes its corresponding gradient $\nabla f_i(\theta_t)$. We assume that the model can be partitioned into N stages, and write $\theta_t = \{\theta_t^j\}_{j \in [1, N]}$ to emphasize the dependency in the stage j . Parameter gradients are computed via a forward and backward pass over the micro-batch. They are averaged over all workers, typically using an all-reduce operation [2]. Finally, each worker locally updates the parameters θ_t with the averaged gradient. With the learning rate γ_t , the standard update rule [25] which results can be written as

$$\theta_{t+1} = \theta_t - \frac{\gamma_t}{N} \sum_{i=1}^N \nabla f_i(\theta_t). \quad (\text{DP})$$

Disadvantages of DP. The execution timeline of DP is illustrated in Fig. 1a, where one time step corresponds to the execution of a forward or backward pass of a stage. We assume that the forward and backward passes of a stage take a similar amount of time for the N workers. A training step t (which indexes the sequence of the parameters) is thus composed of $2N$ time steps. There are

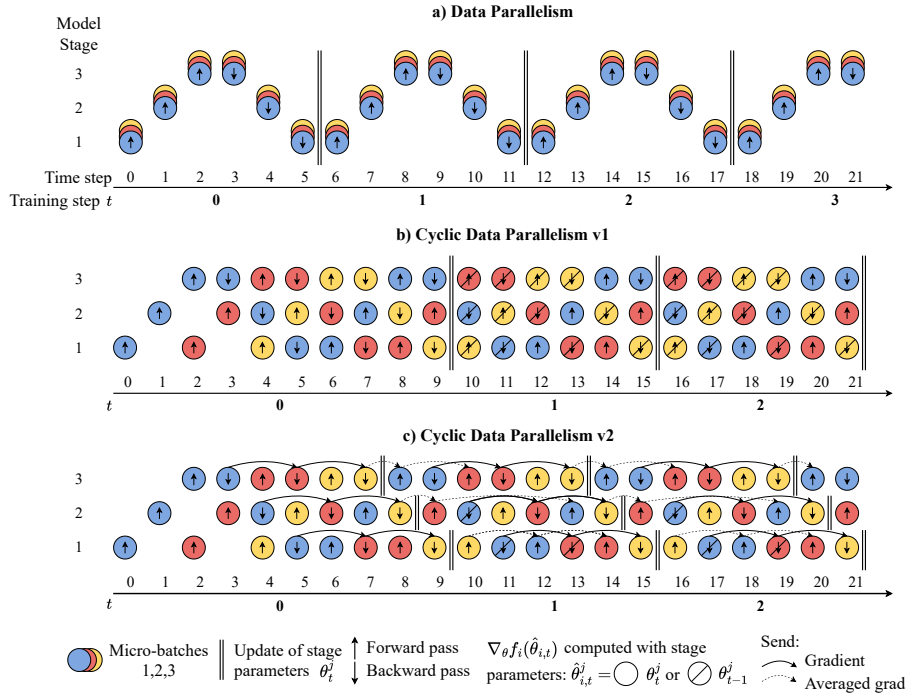


Fig. 1: Timeline of executions for Data Parallelism (DP) and the two versions of Cyclic Data Parallelism (CDP), for $N=3$ workers. (a) DP. The 3 workers begin executing their forward pass simultaneously in DP, and maintain this synchronization throughout the entire forward-backward pass. **(b) CDP-v1.** In CDP, the 3 workers begin executing with an equal delay between them (equal to 2 time steps). For CDP-v1 (see Eq. (CDP-v1)), the parameters of the model are updated with a delay constant equal to one training step. **(c) CDP-v2.** This delay is limited in CDP-v2 (see Eq. (CDP-v2)), by allowing the stages to update and send gradients independently. The communication scheme, balanced across the training step, is indicated. Note that the total complexity of a training step (a forward-backward pass) does not change, but activation memory does not peak in CDP as it does in DP.

several inherent problems with the DP strategy applied to mini-batch SGD. First, the **total number of retained activations** peaks every $2N$ time steps, notably occurring at step 2 in Fig. 1a. This peak results from the intermediate activations that are stored and await release during the backward pass. While this memory overhead can be mitigated by strategies such as pipelining [11], there exists an unavoidable waiting barrier every $2N$ steps to synchronize all gradient computations. This notably occurs between steps 5 and 6 in Fig. 1a. Consequently, this barrier introduces **latency** and leads to **workers idling**. Furthermore, this waiting barrier also leads to a **communication overhead**, as workers must simultaneously communicate their local buffers. While this can be solved by a collective all-reduce operation [2], communication can be an issue in standard centralized frameworks, such as Federated Learning [14].

3.2 Towards delayed mini-batch SGD with Cyclic Data Parallelism

We now propose a method that avoids these problems, caused by the simultaneous execution of computational steps in DP.

Algorithmic description. Contrary to the previous DP approaches, our main idea is to break the synchrony between the forward and backward passes of the N micro-batches. Instead of each step being computed simultaneously, we assume that every micro-batch computation is delayed by an identical number of time steps (*e.g.*, if one worker handles one micro-batch, it will start with a delay). More specifically, the computations of ∇f_i are delayed by a delay of 2 time steps compared to ∇f_{i-1} (where i is taken modulo N , with N the number of stages and micro-batches like before). This results in a cyclic pattern, illustrated in Fig. 1b and 1c. Furthermore, each stage constantly performs either a forward or a backward pass on a single and distinct micro-batch: in particular, it implies that the maximum amount of activations stored at a given time step is nearly constant during training, and in this case, smaller compared to DP. We call this concept **Cyclic Data Parallelism** (CDP).

Update rules. While the forward and backward passes are fully defined assuming an available parameter θ_t , one has to define the corresponding update rule to produce θ_{t+1} , since breaking the synchrony between the backward passes of the micro-batches makes it impossible to follow Eq. (DP). Similarly to [3, 33], at the training step t we introduce an auxiliary variable $\{\hat{\theta}_{i,t}\}_i$, a buffer updated by other concurrent workers. Formally, our algorithm can be written as

$$\begin{aligned} \theta_{t+1} &= \theta_t - \frac{\gamma_t}{N} \sum_{i=1}^N \nabla f_i(\hat{\theta}_{i,t}), & (\text{CDP}) \\ \hat{\theta}_{i,t}^j &= u_{i,j}(\theta_t^j, \theta_{t-1}^j). \end{aligned}$$

where $u_{i,j}(a,b) \in \{a,b\}$ is a rule on the parameters, that depends on both the micro-batch i and the stage j . Thus, $u_{i,j}$ implicitly depends on the time step

of the algorithm but not on the training step in our paradigm. An alternative that would be possible, and would easily allow more complex asynchronous or randomized variants. Note that some rules $u_{i,j}$ are not possible due to the delay between workers, preventing for instance the update rule of DP.

We propose two update rules in particular that follow from Eq. (CDP) with specific rules $u_{i,j}$. They can be considered as the edge cases of the update rule, with maximum and minimum delay respectively, with all other rules $u_{i,j}$ being an intermediary between them. First, we use a simultaneous barrier after the N th batch finishes its computation, which is illustrated in Fig. 1b with a barrier at the time step 9. In this specific case, a consistent update consists of the rule $u_{i,j}(a, b) = b$, using the stored $\hat{\theta}_{i,t} = \theta_{t-1}$. This leads to the delayed update rule

$$\theta_{t+1} = \theta_t - \frac{\gamma_t}{N} \sum_{i=1}^N \nabla f_i(\theta_{t-1}). \quad (\text{CDP-v1})$$

We refer to this update rule as CDP-v1. In the specific case of PP, we recover the update rule of PipeDream-2BW [18]. We now discuss another update rule, which we refer to as CDP-v2, and which is illustrated in Fig. 1c. We fix the rule $u_{i,j}$ such that after the N th micro-batch finishes the gradient computation of a stage, it directly transmits the updated stage parameters to be computed on a micro-batch (the order of communication can be predefined). This gives the rule $u_{i,j}(a, b) = a$ if $j \geq N - i + 1$, and b otherwise, reducing the number of delayed parameters compared to CDP-v1. Note also that we do not need to keep a copy of two parameters, as the parameters in memory are always the freshest available. This update is written as

$$\theta_{t+1} = \theta_t - \frac{\gamma_t}{N} \sum_{i=1}^N \nabla f_i(\theta_{t-1}^1, \dots, \theta_{t-1}^{N-i}, \theta_t^{N-i+1}, \dots, \theta_t^N). \quad (\text{CDP-v2})$$

Note that in both cases, gradient communications need not occur *simultaneously*, instead, results can be communicated at intermediate steps, benefiting the overall communication bandwidth. We propose a possible communication scheme for CDP-v2 in Fig. 1c.

Remark on the convergence. We emphasize that the generic update of Eq. (CDP), in the worst case, can be understood as SGD with delayed gradients with a fixed delay of 1. There is a wide body of literature that studies the convergence of delayed first-order methods [16, 30, 34], which guarantee that our rate of convergence is almost equal to the one of SGD. In practice, very large DNNs have been shown to converge similarly with or without a delay of 1 [18, 24], and models up to 530B parameters converge with larger delays with no issues reported [27].

4 Variants of DP, MP and Zero-DP with CDP

We now discuss how CDP reduces the computational overhead of DP (on one or several GPUs), MP, and ZeRO-DP. A schematic representation of the implemen-




Fig. 2: Comparison between parallelism frameworks with and without using CDP, for $N=3$. A device (*e.g.*, a GPU) is represented by a rectangle, and the different micro-batches being computed by the 3 colors. A model stage requires memory, for the parameters used for computation and for the activations retained awaiting the backward pass, indicated with a colored disk and a black circle. Communications are intra or inter-device (thin or thick arrow), collective or point-to-point (double-headed or single-headed arrow). **(a) Single-GPU DP.** This setting corresponds to a high-connectivity device with limited memory. We observe a memory reduction of half. **(b) Multi-GPU DP.** Communications can be drastically reduced when using multiple GPUs with CDP. **(c) DP+MP.** Both the number of required GPUs and the communications are reduced compared to a standard implementation of MP with DP. Only N GPUs are needed in PP, but they require more activation memory, shown with a thicker circle. **(d) ZeRO-DP.** The model states needs to be sent or received by only one worker at each time step, instead of the standard broadcast operation of ZeRO-DP.

tations is presented in Fig. 2, and the theoretical values of the communication and memory requirements are summarized in Tab. 1, which we discuss next.

Analytical comparisons. We will refer to Ψ_P as the memory occupied by the parameters of the entire model (including the optimizer states here). Ψ_A refers to the memory occupied by the activations of one data sample in the entire model. In MP, stages communicate activations at the inputs and outputs of stages. This subset of activations occupies a memory that we note Ψ_A^{int} (and thus $\Psi_A^{\text{int}} \leq \Psi_A$). From now on, N refers to the number of stages of the model as well as the number of micro-batches, which have equal size B . In Tab. 1, we summarize the memory requirements per device, the communication volume, and the maximum number of communication steps required between 2 time steps, for DP, MP, PP, and

Implementation	Memory per GPU		Communications inter-GPUs			
	Activations	Parameters	Volume	Max com. steps	Nb of GPUs	Rule
Single-GPU DP	$NB\Psi_A$	$N\Psi_P$			1	(DP)
+ Cyclic	$\frac{(N+1)}{2}B\Psi_A$	$\frac{(N+1)}{2}\Psi_P$			1	(CDP)
Multi-GPU DP	$B\Psi_A$	Ψ_P	Ψ_P	$\mathcal{O}(\log(N))$	N	(DP)
+ Cyclic	$B\Psi_A$	Ψ_P	Ψ_P	$\mathcal{O}(1)$	N	(CDP)
DP with MP	$B\Psi_A/N$	Ψ_P/N	$\Psi_P+B\Psi_A^{\text{int}}$	$\mathcal{O}(\log(N))$	N^2	(DP)
+ Cyclic	$B\Psi_A/N$	Ψ_P/N	$\frac{1}{2}\Psi_P+B\Psi_A^{\text{int}}$	$\mathcal{O}(1)$	$\frac{1}{2}(N+1)N$	(CDP)
PP	$B\Psi_A$	Ψ_P/N	$B\Psi_A^{\text{int}}$	$\mathcal{O}(1)$	N	(CDP)
ZeRO-DP	$B\Psi_A$	Ψ_P/N	Ψ_P	$\mathcal{O}(\log(N))$	N	(DP)
+ Cyclic	$B\Psi_A$	Ψ_P/N	Ψ_P	$\mathcal{O}(1)$	N	(CDP)

Table 1: Theoretical cost of the parallelism implementations discussed in Sec. 4. All implementations are improved by using CDP (Cyclic) over DP, in terms of memory (per GPU or number of GPUs) or communication. Improvements are noted in bold. The parameter memory of the entire model is noted as Ψ_P , and the activation memory of the entire model for one data sample by Ψ_A , or Ψ_A^{int} for the subset communicated in MP. N indicates both the number of stages and of micro-batches (of size B). The communication volume is the size of the tensors that need to be communicated. ‘Max com. steps’ is the maximum number of steps of communication required between 2 time steps between two workers, which is a minimum of $\mathcal{O}(\log(N))$ steps for a collective operation or a single step for point-to-point communications. The update rule used is specified. Note that the parameter memory required in Single-GPU DP depends on the implementation.

ZeRO-DP with and without CDP. We find the following improvements with CDP compared to DP. CDP halves the memory required to train on DP with a single-GPU device. Multi-GPU DP with CDP doesn’t require an all-reduce operation at the end of the training step but only point-to-point communications at each time step. This improvement in the number of communications required between time steps is also found with MP and ZeRO-DP with CDP. CDP allows MP to halve the number of GPUs needed, thus halving the memory required, as well as the communication of gradients between GPUs. The number of GPUs can be further reduced from MP with CDP to PP, which can be seen as a particular implementation of CDP. However, this requires the GPUs to be able to store the activations of the entire mini-batch. More generally, a disadvantage of MP and PP is that they need to communicate activations, which scale with the batch size B . We will now discuss each setting in more depth.

4.1 CDP implementation on a single-GPU device

We first explore the setting of a single-GPU device training with mini-batch SGD, assuming that communication within the GPU is cheap and fast, but memory is limited [4].

Standard implementation. We illustrate the training of standard DP on a single-GPU device in Fig. 2.a.i. Then, at the peak of the forward pass, the GPU must

retain the activations of the entire model for N micro-batches, representing the total mini-batch, resulting in a total memory of activations of about $NB\Psi_A$.

Cyclic implementation. Turning our attention to CDP, depicted in Fig. 2.a.ii, we find that its implementation on a single-GPU device requires about half the memory compared to standard DP. Indeed, in this approach, each stage of the model processes one micro-batch at each time step. Consequently, the total memory occupied by activations remains nearly constant across time steps, equal approximately to $\frac{(N+1)}{2}B\Psi_A$, reducing in half the total memory required for a DP implementation. Depending on the implementation, parameters may be shared in the GPU, resulting in no parameter memory improvement. The memory improvement is still significant in this case, as activation memory is generally much larger than parameter memory, peaking at 60GB for a GPT-2 model in [21], compared to 3GB for the fp16 parameters.

4.2 CDP implementation on a multi-GPU device

Here, one GPU processes a single micro-batch, but communications between GPUs may hinder the overall system efficiency.

Standard implementation. We consider the case where DP is implemented on N GPUs that process N micro-batches of data independently, as depicted in Fig. 2.b.i. These GPUs communicate the micro-batches gradients with an all-reduce operation at the end of the training step before proceeding to the next time step. All-reduce is a collective operation that requires at least $\mathcal{O}(\log(N))$ communications steps [2] in a favorable setting, or $\mathcal{O}(N)$ steps for a bandwidth-optimal ring-based implementation [20].

Cyclic implementation. CDP on N GPUs, as depicted in Fig. 2.b.ii, makes better use of the communication bandwidth between GPUs, as the all-reduce operation is fragmented across the training step into point-to-point operations between each time step. Indeed, half of the stages compute a backward pass at any time step, and then send the computed gradient to the following worker (modulo N) for the reduce operation. This communication scheme corresponds to the ring-based all-reduce operation [20], which has optimal bandwidth usage.

4.3 Implementation in a MP paradigm

Standard implementation. Now, in MP, a GPU does not hold a full replica of a model but only of a single stage. During the forward and backward passes, activations and gradients must be communicated between GPUs holding successive stages. A naive implementation of DP and MP results in Fig. 2.c.i., where the N micro-batches require N^2 GPUs to process each slice of the data and the model. Furthermore, note also that only N GPUs are busy at a time, thus most GPUs are idle, waiting for the forward pass of the next batch of data which is a significant inefficiency of MP.

Cyclic implementation. Here, CDP with MP, which we present in Fig. 2.c.ii., once again improves on the DP implementation. In particular, the number of GPUs required for CDP is reduced, and we can show that for each stage j , CDP requires only $N - j + 1$ GPUs to be shared among the N micro-batches. This leads to a "pyramidal" shape of the stage structure. Thus in total, only $\frac{1}{2}N(N+1)$ GPUs holding a stage of the model are needed, halving the number over DP, as for the memory required. Compared to DP, a micro-batch doesn't send the activations of a stage to the same GPU every time for the computation of the next stage. Instead, a GPU in the previous time step will have released the activations stored in its memory after a backward pass. The micro-batch will send the activations to this GPU as its memory is available for computations. Since the number of devices is smaller in DP than in CDP, note also that the total number of gradient communications between devices is reduced. In fact, if a GPU is only able to hold the activation of one micro-batch, the number of GPUs used, $\frac{1}{2}N(N+1)$, is optimal to compute N micro-batches. An in-depth discussion is proposed in the Appendix of this paper.

Pipeline Parallelism implementation. If we assume that one GPU can store the activations of all N micro-batches, then we can further reduce the total number of devices needed to only N , one per stage. Indeed, this implementation of CDP with MP on N devices, depicted in Fig. 2.c.iii., is equivalent to PP as presented in PipeDream [17]. If we follow our first update rule Eq. (CDP-v1), CDP follows the same update rule as PipeDream-2BW [18]. However, our second update rule Eq. (CDP-v2) improves on this update rule by reducing the gradient delay. PP requires fewer GPUs than MP, but only if the GPUs can store the activations of the entire mini-batch. This limits its scaling capacity compared to MP, which is 2 times more GPU-efficient. Indeed, for GPUs with similar memory capacities, MP with CDP requires $\frac{1}{2}(N+1)N$ GPUs to train a batch of size B . Meanwhile, PP requires N GPUs to train a batch of size B/N , or N^2 GPUs to train a batch of size B by combining PP with DP.

4.4 Implementation in a ZeRO-DP paradigm

Standard implementation. ZeRO-DP [21] is a training framework that aims at combining the advantages of DP and MP, which we represent in Fig. 2.d.i.. Rather than replicating the entire model on every one of the N GPUs, ZeRO-DP replicates only the model states of a single stage across the GPUs. The model state refers (in stage 3 of ZeRO-DP) to the model parameters, gradients, and optimizer states. When the workers execute the forward or backward propagation through a stage, the model states of that stage are broadcast from the GPU that stores them to all GPUs. After computation, the model states are deleted to free up the memory, such that a GPU only retains the model states of a maximum of two stages. The communication volume is similar to standard multi-GPU DP, with a 1.5 increase at most. However note that like in PP, the memory taken by the activations on one GPU increases with N since all stages of the model must be retained.

Model	Learning Rule			Model	Learning Rule		
	(CDP-v1)	(CDP-v2)	(DP)		(CDP-v1)	(CDP-v2)	(DP)
ResNet-18	94.1	94.8	94.7	ResNet-18	69.9	70.0	70.1
ResNet-50	94.0	94.5	94.5	ResNet-50	75.8	75.7	75.4

(a) Test accuracy on CIFAR10

(b) Test accuracy on ImageNet

Table 2: Top-1 test accuracy for the three learning rules (CDP-v1), (CDP-v2) and (DP) on the (a) CIFAR10 and (b) ImageNet datasets, by training a ResNet-18 and a ResNet-50. Our results are stable, as the standard deviation over 5 runs is systematically less than 0.08. We observe that on CIFAR10, CDP systematically performs similarly or better than DP, especially CDP-v2. On ImageNet, CDP performs similar to or better than DP.

Cyclic implementation. CDP improves on ZeRO-DP by removing the need for collective communications of the model states, as depicted in Fig. 2.d.ii. Since one stage is computed by a single GPU at a time step, the model states only need to be held by this GPU, with no replication. Then, they only need to be communicated to a single GPU at the next time step. This reduces the communication overhead at each time step, similarly to Multi-GPU DP.

5 Numerical analysis

Hyperparameters. To test our method, we propose to use the standard training pipeline on CIFAR-10 and ImageNet datasets, trained on the ResNet-18 and ResNet-50 architectures, split into 4 stages with similar FLOPs. We simulate this partitioning using the `fvcore` library³ to compute the FLOPs count of each ResNet’s module. For a finer partition of the linear modules, we separate them into weight and bias modules, by approximating the FLOPs of the bias module as the square root of the linear module’s FLOPs. We consider the SGD optimizer with momentum 0.9 and we simulate our delayed activations for DP, CDP-v1, and CDP-v2. We train over 100 epochs with batch size 128 and 90 epochs with batch size 256 respectively for CIFAR-10 and ImageNet. We consider an initial learning rate of 0.05 and 0.1, dropped by a factor of 0.2 and 0.1 on CIFAR-10 and ImageNet respectively, at epochs 30, 60, and 90. The weight decay is equal to 10^{-4} for ImageNet, and 5×10^{-4} or 10^{-3} for CIFAR-10, whether trained on a ResNet-18 or 50. To accommodate the smaller image size of CIFAR-10, as standard, we remove the first max pooling and reduce the kernel size of the first convolutional layer to 3 and the stride to 1. For CIFAR-10, we report the accuracy and its corresponding standard deviation on the test sets for 5 runs. For ImageNet, following standard practice, we report the maximum validation accuracy over the last 10 epochs. Contrary to [33], we did not need any specific tuning of the other hyper-parameters for our update rules. Our source code is available at github.com/fournierlouis/Cyclic_Data_Parallelism.

³ <https://github.com/facebookresearch/fvcore/>

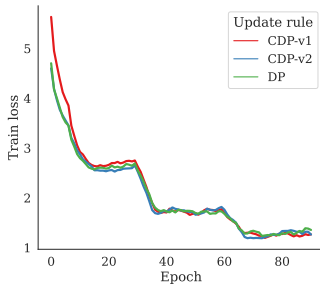


Fig. 3: Training loss of a ResNet-50 trained on ImageNet following the learning rules (DP), (CDP-v1) and (CDP-v2). Values are averaged over a window of 7 epochs for the sake of readability. The loss of CDP-v1 is significantly higher at the beginning of training, which is not the case for CDP-v2. As parameters converge, the effect of the delay disappears and the three losses show a similar training curve, with a small advantage for both CDP-v1 and CDP-v2. This confirms that the delay in our update rules does not affect convergence, even in realistic settings.

Results. We report in Tab. 2a and Tab. 2b the test accuracy of our models for the three rules CDP-v1, CDP-v2, and DP, for 5 runs on CIFAR-10 and 3 runs on ImageNet respectively. The variances reported in our runs are less than 0.08, indicating that our experiments lead to relatively stable results. Both tables show that CDP leads to similar or better performances compared to DP, which is consistent with the experimental findings of [18, 24]. For CIFAR-10, CDP-v2 significantly outperforms CDP-v1, showing our improvement over PipeDream-2BW’s rule.

We also provide in Fig. 3 the training loss of the three concurrent learning rules on the ImageNet dataset: we note that the final loss of DP is slightly higher than both CDP-v1 and CDP-v2, with similar generalization performance, which indicates how close the three methods are, both from an optimization and statistical point of view. This is consistent with the theoretical insights of [16] (see Sec. 3.2).

Activation memory tracking. In Fig. 4, we track the memory used during one forward-backward pass of a ResNet-50 and a ViT-B/16 training on ImageNet. By subtracting the constant value corresponding to the parameter memory, we obtain the activation memory. From this memory curve, which clearly shows the activation memory peaking at the end of the forward pass, we extrapolate for $N = 4, 8,$ and 32 the activation memory usage per worker (total memory across N homogeneous workers divided by N) that an efficient implementation of DP and CDP would use. As expected, the activation memory used varies less during training with CDP, flattening as N increases. Furthermore, the maximum activation memory used by the DP paradigm is significantly higher than that of CDP. The activation memory ratio we find here is approximately 30% for the ResNet-50 and 42% ($= \frac{3.9-2.3}{3.9} \frac{\text{GB}}{\text{GB}}$) for the ViT-B/16, close to the ideal halving




Fig. 4: Activation memory per worker, when training with N workers on ImageNet with an efficient implementation of DP (full) and a CDP (dashed), on a ResNet-50 and a ViT-B/16. An optimal halving of the parameters is represented in black ('Optimal'). The memory required by a forward-backward pass for one work is first tracked, and parameter memory is removed. The figure is extrapolated by mimicking the total memory used by N workers training on DP (*i.e.* simultaneously) or CDP (*i.e.* cyclically), and dividing by N . As N increases, the memory required by CDP flattens, to a value lower than DP. This value reaches close to the reduction in half theorized for the ViT-B/16, with 42%. The heterogeneity of the layers of the ResNet reduces the effectiveness of CDP, only reaching 30% reduction.

theorized. The ResNet reaches a lower memory improvement due to the heterogeneity of its layers, which require varying activation memory for the same execution time (as feature size decreases through depth). This is not an issue for a Transformer-based model, as feature size stays constant across depth. This confirms that CDP will significantly improve the memory used in real implementations when using homogeneous stages and workers.

6 Conclusion

We introduced Cyclic Data Parallelism (CDP), an alternative framework to Data Parallelism. By executing the forward and backward passes of micro-batches of data cyclically rather than simultaneously, we balance gradient communications and the total memory occupied by activations. We particularize CDP within the context of Data, Model, and Pipeline Parallelism as well as ZeRO-DP, and demonstrate improvements in the total memory required to store activations or in the number of communication steps required between time steps during training. In particular, CDP reduces the number of devices needed in MP and reduces the communication delay in ZeRO-DP. Our results are supported by existing theoretical guarantees in the small delay settings. Finally, our numerical experiments on ImageNet show our update rules achieve similar testing accuracy as standard DP.

In future work, we would like to release a highly efficient implementation compatible with cuDNN frameworks, as well as further relax our update rule to explore the possibility of using asynchronous and random delays.

Acknowledgements

This work was supported by Project ANR-21-CE23-0030 ADONIS, EMERG-ADONIS from Alliance SU, and Sorbonne Center for Artificial Intelligence (SCAI) of Sorbonne University (IDEX SUPER 11-IDEX-0004). This work was granted access to the AI resources of IDRIS under the allocations 2022-AD011013095 and 2023-A0151014526 made by GENCI.

References

1. Belilovsky, E., Eickenberg, M., Oyallon, E.: Decoupled greedy learning of cnns (2020)
2. Chan, E., Heimlich, M., Purkayastha, A., van de Geijn, R.: Collective communication: Theory, practice, and experience. *Concurrency and Computation: Practice and Experience* **19**, 1749–1783 (09 2007). <https://doi.org/10.1002/cpe.v19:13>
3. Chen, C.C., Yang, C.L., Cheng, H.Y.: Efficient and robust parallel dnn training through model parallelism on multi-gpu platform. arXiv preprint arXiv:1809.02839 (2018)
4. Choquette, J., Gandhi, W., Giroux, O., Stam, N., Krashinsky, R.: Nvidia a100 tensor core gpu: Performance and innovation. *IEEE Micro* **41**(2), 29–35 (2021). <https://doi.org/10.1109/MM.2021.3061394>
5. Dean, J., Corrado, G., Monga, R., Chen, K., Devin, M., Mao, M., Ranzato, M.a., Senior, A., Tucker, P., Yang, K., Le, Q., Ng, A.: Large scale distributed deep networks. In: Pereira, F., Burges, C., Bottou, L., Weinberger, K. (eds.) *Advances in Neural Information Processing Systems*. vol. 25. Curran Associates, Inc. (2012), https://proceedings.neurips.cc/paper_files/paper/2012/file/6aca97005c68f1206823815f66102863-Paper.pdf
6. Fan, S., Rong, Y., Meng, C., Cao, Z., Wang, S., Zheng, Z., Wu, C., Long, G., Yang, J., Xia, L., et al.: Dapple: A pipelined data parallel approach for training large models. In: *Proceedings of the 26th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*. pp. 431–445 (2021)
7. Fournier, L., Rivaud, S., Belilovsky, E., Eickenberg, M., Oyallon, E.: Can forward gradient match backpropagation? In: Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., Scarlett, J. (eds.) *Proceedings of the 40th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 202, pp. 10249–10264. PMLR (23–29 Jul 2023), <https://proceedings.mlr.press/v202/fournier23a.html>
8. Goyal, P., Dollár, P., Girshick, R., Noordhuis, P., Wesolowski, L., Kyrola, A., Tulloch, A., Jia, Y., He, K.: Accurate, large minibatch sgd: Training imagenet in 1 hour (2018)
9. Ho, Q., Cipar, J., Cui, H., Kim, J.K., Lee, S., Gibbons, P.B., Gibson, G.A., Ganger, G.R., Xing, E.P.: More effective distributed ml via a stale synchronous parallel parameter server. In: *Proceedings of the 26th International Conference on Neural Information Processing Systems - Volume 1*. p. 1223–1231. NIPS’13, Curran Associates Inc., Red Hook, NY, USA (2013)
10. Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., de Las Casas, D., Hendricks, L.A., Welbl, J., Clark, A., Hennigan, T., Noland, E., Millican, K., van den Driessche, G., Damoc, B., Guy, A., Osindero, S., Simonyan, K., Elsen, E., Rae, J.W., Vinyals, O., Sifre, L.: Training compute-optimal large language models (2022)

11. 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)
12. Jain, A., Awan, A.A., Aljuhani, A.M., Hashmi, J.M., Anthony, Q.G., Subramoni, H., Panda, D.K., Machiraju, R., Parwani, A.: Gems: Gpu-enabled memory-aware model-parallelism system for distributed dnn training. In: *SC20: International Conference for High Performance Computing, Networking, Storage and Analysis*. pp. 1–15 (2020). <https://doi.org/10.1109/SC41405.2020.00049>
13. Jia, Z., Zaharia, M., Aiken, A.: Beyond data and model parallelism for deep neural networks. *Proceedings of Machine Learning and Systems* **1**, 1–13 (2019)
14. Konecný, J., McMahan, H.B., Yu, F.X., Richtárik, P., Suresh, A.T., Bacon, D.: Federated learning: Strategies for improving communication efficiency. *ArXiv abs/1610.05492* (2016), <https://api.semanticscholar.org/CorpusID:14999259>
15. Li, M., Andersen, D.G., Park, J.W., Smola, A.J., Ahmed, A., Josifovski, V., Long, J., Shekita, E.J., Su, B.Y.: Scaling distributed machine learning with the parameter server. In: *Proceedings of the 11th USENIX Conference on Operating Systems Design and Implementation*. p. 583–598. OSDI’14, USENIX Association, USA (2014)
16. Mishchenko, K., Bach, F., Even, M., Woodworth, B.: Asynchronous sgd beats minibatch sgd under arbitrary delays (2023)
17. Narayanan, D., Harlap, A., Phanishayee, A., Seshadri, V., Devanur, N.R., Ganger, G.R., Gibbons, P.B., Zaharia, M.: Pipedream: Generalized pipeline parallelism for dnn training. In: *Proceedings of the 27th ACM Symposium on Operating Systems Principles*. pp. 1–15 (2019)
18. Narayanan, D., Phanishayee, A., Shi, K., Chen, X., Zaharia, M.: Memory-efficient pipeline-parallel dnn training. In: *International Conference on Machine Learning*. pp. 7937–7947. PMLR (2021)
19. Nøkland, A., Eidnes, L.H.: Training neural networks with local error signals (2019)
20. Patarasuk, P., Yuan, X.: Bandwidth optimal all-reduce algorithms for clusters of workstations. *Journal of Parallel and Distributed Computing* **69**(2), 117–124 (2009). <https://doi.org/https://doi.org/10.1016/j.jpdc.2008.09.002>, <https://www.sciencedirect.com/science/article/pii/S0743731508001767>
21. Rajbhandari, S., Rasley, J., Ruwase, O., He, Y.: Zero: Memory optimizations toward training trillion parameter models (2020)
22. Rajbhandari, S., Ruwase, O., Rasley, J., Smith, S., He, Y.: Zero-infinity: Breaking the gpu memory wall for extreme scale deep learning (2021)
23. Recht, B., Re, C., Wright, S., Niu, F.: Hogwild!: A lock-free approach to parallelizing stochastic gradient descent. In: Shawe-Taylor, J., Zemel, R., Bartlett, P., Pereira, F., Weinberger, K. (eds.) *Advances in Neural Information Processing Systems*. vol. 24. Curran Associates, Inc. (2011), https://proceedings.neurips.cc/paper_files/paper/2011/file/218a0aefd1d1a4be65601cc6ddc1520e-Paper.pdf
24. Ren, J., Rajbhandari, S., Aminabadi, R.Y., Ruwase, O., Yang, S., Zhang, M., Li, D., He, Y.: Zero-offload: Democratizing billion-scale model training (2021)
25. Robbins, H.E.: A stochastic approximation method. *Annals of Mathematical Statistics* **22**, 400–407 (1951), <https://api.semanticscholar.org/CorpusID:16945044>
26. Shazeer, N., Cheng, Y., Parmar, N., Tran, D., Vaswani, A., Koanantakool, P., Hawkins, P., Lee, H., Hong, M., Young, C., Sepassi, R., Hechtman, B.: Mesh-tensorflow: Deep learning for supercomputers (2018)

27. Shoeybi, M., Patwary, M., Puri, R., LeGresley, P., Casper, J., Catanzaro, B.: Megatron-lm: Training multi-billion parameter language models using model parallelism. arXiv preprint arXiv:1909.08053 (2019)
28. Smith, S., Patwary, M., Norick, B., LeGresley, P., Rajbhandari, S., Casper, J., Liu, Z., Prabhunoye, S., Zerveas, G., Korthikanti, V., et al.: Using deepspeed and megatron to train megatron-turing nlg 530b, a large-scale generative language model. arXiv preprint arXiv:2201.11990 (2022)
29. Stich, S.U.: Local sgd converges fast and communicates little (2019)
30. Stich, S.U., Karimireddy, S.P.: The error-feedback framework: Sgd with delayed gradients. *Journal of Machine Learning Research* **21**(237), 1–36 (2020), <http://jmlr.org/papers/v21/19-748.html>
31. Wang, G., Qin, H., Jacobs, S.A., Holmes, C., Rajbhandari, S., Ruwase, O., Yan, F., Yang, L., He, Y.: Zero++: Extremely efficient collective communication for giant model training (2023)
32. Xu, A., Huo, Z., Huang, H.: On the acceleration of deep learning model parallelism with staleness (2022)
33. Yang, B., Zhang, J., Li, J., Ré, C., Aberger, C., De Sa, C.: Pipemare: Asynchronous pipeline parallel dnn training. *Proceedings of Machine Learning and Systems* **3**, 269–296 (2021)
34. Yang, H., Zhang, X., Khanduri, P., Liu, J.: Anarchic federated learning. In: Chaudhuri, K., Jegelka, S., Song, L., Szepesvari, C., Niu, G., Sabato, S. (eds.) *Proceedings of the 39th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 162, pp. 25331–25363. PMLR (17–23 Jul 2022), <https://proceedings.mlr.press/v162/yang22r.html>
35. Zhang, Z., Zheng, S., Wang, Y., Chiu, J., Karypis, G., Chilimbi, T., Li, M., Jin, X.: Mics: Near-linear scaling for training gigantic model on public cloud (2022)
36. Zhao, Y., Gu, A., Varma, R., Luo, L., Huang, C.C., Xu, M., Wright, L., Shojanazeri, H., Ott, M., Shleifer, S., Desmaison, A., Balioglu, C., Damania, P., Nguyen, B., Chauhan, G., Hao, Y., Mathews, A., Li, S.: Pytorch fsdp: Experiences on scaling fully sharded data parallel (2023)
37. Zhuang, H., Weng, Z., Luo, F., Kar-Ann, T., Li, H., Lin, Z.: Accumulated decoupled learning with gradient staleness mitigation for convolutional neural networks. In: Meila, M., Zhang, T. (eds.) *Proceedings of the 38th International Conference on Machine Learning*. *Proceedings of Machine Learning Research*, vol. 139, pp. 12935–12944. PMLR (18–24 Jul 2021), <https://proceedings.mlr.press/v139/zhuang21a.html>